

AWS Secure Environment Accelerator

Amazon Web Services

Amazon Web Services

None

Table of contents

1. 1. AWS Secure Environment Accelerator	6
1.1 1.1. Overview	6
1.2 1.2. What specifically does the Accelerator deploy and manage?	6
1.2.1 1.2.1. Creates AWS Account	7
1.2.2 1.2.2. Creates Networking	7
1.2.3 1.2.3. Cross-Account Object Sharing	7
1.2.4 1.2.4. Identity	8
1.2.5 1.2.5. Cloud Security Services	8
1.2.6 1.2.6. Other Security Capabilities	8
1.2.7 1.2.7. Centralized Logging and Alerting	8
1.3 1.3. Relationship with AWS Landing Zone Solution (ALZ)	9
1.4 1.4. Relationship with AWS Control Tower	9
1.5 1.5. Accelerator Installation Process (Summary)	10
2. Installation & Upgrades	11
2.1 Accelerator Installation and Upgrades	11
2.2 Installation	12
2.2.1 1. Accelerator Installation Guide	12
2.2.2 1. Accelerator Sample Configurations and Customization	31
2.2.3 1. State Machine Behavior and Inputs	39
2.2.4 1. Multi-file Accelerator Config file and YAML Support Details	42
2.2.5 1. Existing Organizations / Accounts	46
2.2.6 1. How to migrate an AWS Landing Zone (ALZ) account "as is" into an AWS Secure Environment Accelerator (ASEA)	48
2.3 Upgrades	53
2.3.1 1. Accelerator Upgrade Guide	53
2.3.2 1. Accelerator v1.5.x Custom Upgrade Instructions	56
2.4 Functionality	61
2.4.1 Accelerator Service List	62
2.4.2 1. Accelerator Pricing	65
2.4.3 AWS Secure Environment Accelerator Deployment Capabilities	73
2.4.4 1. Accelerator Central Logging Implementation and File Structures	79
2.4.5 Object Naming	84
3. 1. Accelerator Basic Operation and Frequently asked Questions	88
3.1 1.1. Operational Activities	88
3.2 1.2. Existing Accounts / Organizations	101
3.3 1.3. End User Environment	103

3.4	1.4. Upgrades	105
3.5	1.5. Support Concerns	106
3.6	1.6. Deployed Functionality	109
3.7	1.7. Network Architecture	125
4.	Operations & Troubleshooting	129
4.1	Accelerator Operations & Troubleshooting Guide	129
4.2	1. System Overview	130
4.2.1	1.1. Overview	130
4.2.2	1.2. Installer Stack	132
4.2.3	1.3. Initial Setup Stack	137
4.3	1. Troubleshooting	150
4.3.1	1.1. Overview	150
4.3.2	1.2. Components	150
4.3.3	1.3. Examples	159
4.4	1. Common Tasks	163
4.4.1	1.1. Restart the State Machine	163
4.4.2	1.2. Switch To a Managed Account	163
5.	Developer Guide	166
5.1	Accelerator Developer Guide	166
5.2	1. Development Guide	167
5.2.1	1.1. Overview	167
5.2.2	1.2. Project Structure	167
5.2.3	1.3. Installer Stack	167
5.2.4	1.4. Initial Setup Stack	168
5.2.5	1.5. Phase Steps and Phase Stacks	169
5.2.6	1.6. Store outputs to SSM Parameter Store	170
5.2.7	1.7. Libraries & Tools	171
5.2.8	1.8. Workarounds	176
5.2.9	1.9. Local Development	176
5.2.10	1.10. Testing	179
5.3	1. Technology Stack	181
5.3.1	1.1. Overview	181
5.3.2	1.2. TypeScript and NodeJS	181
5.3.3	1.3. CloudFormation	181
5.3.4	1.4. CDK	181
5.4	1. Best Practices	182
5.4.1	1.1. TypeScript and NodeJS	182
5.4.2	1.2. CloudFormation	182

5.4.3	1.3. CDK	183
5.5	1. How to Contribute	187
5.5.1	1.1. General	187
5.5.2	1.2. Adding New Functionality?	187
5.5.3	1.3. Create a CDK Lambda Function with Lambda Runtime Code	187
5.5.4	1.4. Create a Custom Resource	187
5.5.5	1.5. Run All Unit Tests	188
5.5.6	1.6. Accept Unit Test Snapshot Changes	188
5.5.7	1.7. Validate Code with Prettier	188
5.5.8	1.8. Format Code with Prettier	188
5.5.9	1.9. Validate Code with tslint	188
5.6	1. AWS Internal - Accelerator Release Process	189
5.6.1	1.1. Creating a new Accelerator Code Release	189
6.	Sample Sensitive Architecture	190
6.1	Accelerator Sample Sensitive Architecture	190
6.2	1. AWS Secure Environment Accelerator Reference Architecture	191
6.2.1	1.1. Overview	191
6.2.2	1.2. Introduction	191
6.3	1. Account Structure	196
6.3.1	1.1. Overview	196
6.3.2	1.2. Organization structure	196
6.3.3	1.3. Organizational Units	197
6.3.4	1.4. Mandatory Accounts	200
6.3.5	1.5. Functional Accounts	203
6.3.6	1.6. Account Level Security Settings	203
6.3.7	1.7. Private Marketplace	203
6.4	1. Authorization and Authentication	205
6.4.1	1.1. Overview	205
6.4.2	1.2. Relationship to the Organization Management (root) AWS Account	205
6.4.3	1.3. Break Glass Accounts	205
6.4.4	1.4. Multi-Factor Authentication	206
6.4.5	1.5. Control Plane Access via AWS SSO	206
6.4.6	1.6. Root Authorization	208
6.4.7	1.7. Service Roles	208
6.4.8	1.8. Service Control Policies	208
6.5	1. Logging and Monitoring	212
6.5.1	1.1. Overview	212
6.5.2	1.2. CloudTrail	212

6.5.3	1.3. VPC Flow Logs	212
6.5.4	1.4. GuardDuty	212
6.5.5	1.5. Config	212
6.5.6	1.6. CloudWatch Logs	213
6.5.7	1.7. SecurityHub	213
6.5.8	1.8. Systems Manager Session Manager	213
6.5.9	1.9. Systems Manager Inventory	213
6.5.10	1.10. Other Services	213
6.6	1. Networking	214
6.6.1	1.1. Overview	214
6.6.2	1.2. Perimeter	216
6.6.3	1.3. Shared Network	218
6.7	1. Prescriptive Sensitive Sample Architecture Diagrams	225
6.7.1	1.1. Shared VPC Architecture	225
6.7.2	1.2. Spoke VPC Architecture	226
6.7.3	1.3. VPC and Security Group Patterns	226
6.7.4	1.4. Additional Perimeter Patterns	228
7.	Workshops	230
7.1	Accelerator Workshops	230
7.1.1	Accelerator Administrator Immersion Day	230
7.1.2	Accelerator Workload/Application Team Immersion Day	230

The Accelerator is designed to enable customers to upgrade across Accelerator versions while maintaining a customer's specific configuration and customizations, and without the need for any coding expertise or for professional services. Customers have been able to seamlessly upgrade their AWS multi-account environment from the very first Accelerator beta release to the latest release (across more than 50 releases), gaining the benefits of bug fixes and enhancements while having the option to enable new features, without any loss of existing customization or functionality.

Specifically the accelerator deploys and manages the following functionality, both at initial accelerator deployment and as new accounts are created, added, or onboarded in a completely automated but customizable manner:

1.2.1 1.2.1. Creates AWS Account

- Core Accounts - as many or as few as your organization requires, using the naming you desire. These accounts are used to centralize core capabilities across the organization and provide `Control Panel` like capabilities across the environment. Common core accounts include:
 - Shared Network
 - Operations
 - Perimeter
 - Log Archive
 - Security Tooling
- Workload Accounts - automated concurrent mass account creation or use AWS organizations to scale one account at a time. These accounts are used to host a customer's workloads and applications.
- Scalable to 1000's of AWS accounts
- Supports AWS Organizations nested [OU's](#) and importing existing AWS accounts
- Performs 'account warming' to establish initial limits, when required
- Automatically submits limit increases, when required (complies with initial limits until increased)
- Leverages AWS Control Tower

1.2.2 1.2.2. Creates Networking

- Transit Gateways and TGW route tables (incl. inter-region TGW peering)
- Centralized and/or Local (bespoke) VPC's
- Subnets, Route tables, NACLs, Security groups, NATGWs, IGWs, VGWs, CGWs
- **NEW** Outpost, Local Zone and Wavelength support
- VPC Endpoints (Gateway and Interface, Centralized or Local)
- Route 53 Private and Public Zones, Resolver Rules and Endpoints, VPC Endpoint Overloaded Zones
- All completely and individually customizable (per account, VPC, subnet, or OU)
- Layout and customize your VPCs, subnets, CIDRs and connectivity the way you want
- Static or Dynamic VPC and subnet CIDR assignments
- Deletes default VPC's (worldwide)
- AWS Network Firewall

1.2.3 1.2.3. Cross-Account Object Sharing

- VPC and Subnet sharing, including account level re-tagging (Per account security group 'replication')
- VPC attachments and peering (local and cross-account)
- Zone sharing and VPC associations
- Managed Active Directory sharing, including R53 DNS resolver rule creation/sharing
- Automated TGW inter-region peering
- Populate Parameter Store with all `user` objects to be used by customers' IaC

- Deploy and share SSM documents (4 provided out-of-box, ELB Logging, S3 Encryption, Instance Profile remediation, Role remediation)
- customer can provide their own SSM documents for automated deployment and sharing

1.2.4 1.2.4. Identity

- Creates Directory services (Managed Active Directory and Active Directory Connectors)
- Creates Windows admin bastion host auto-scaling group
- Set Windows domain password policies
- Set IAM account password policies
- Creates Windows domain users and groups (initial installation only)
- Creates IAM Policies, Roles, Users, and Groups
- Fully integrates with and leverages AWS SSO for centralized and federated login

1.2.5 1.2.5. Cloud Security Services

- Enables and configures the following AWS services, worldwide w/central designated admin account:
- GuardDuty w/S3 protection
- Security Hub (Enables designated security standards, and disables individual controls)
- Firewall Manager
- CloudTrail w/Insights and S3 data plane logging
- Config Recorders/Aggregator
- Conformance Packs and Config rules (95 out-of-box NIST 800-53 rules, 2 custom rules, customizable per OU)
- Macie
- IAM Access Analyzer
- CloudWatch access from central designated admin account (and setting Log group retentions)

1.2.6 1.2.6. Other Security Capabilities

- Creates, deploys and applies Service Control Policies
- Creates Customer Managed KMS Keys (SSM, EBS, S3), EC2 key pairs, and secrets
- Enables account level default EBS encryption and S3 Block Public Access
- Configures Systems Manager Session Manager w/KMS encryption and centralized logging
- Configures Systems Manager Inventory w/centralized logging
- Creates and configures AWS budgets (customizable per OU and per account)
- Imports or requests certificates into AWS Certificate Manager
- Deploys both perimeter and account level ALB's w/Lambda health checks, certificates and TLS policies
- Deploys & configures 3rd party firewall clusters and management instances (leverages marketplace)
- Gateway Load Balancer w/auto-scaling and VPN IPsec BGP ECMP deployment options
- Protects Accelerator deployed and managed objects
- Sets Up SNS Alerting topics (High, Medium, Low, Blackhole priorities)
- Deploys CloudWatch Log Metrics and Alarms
- Deploys customer provided custom config rules (2 provided out-of-box, no EC2 Instance Profile/Permissions)

1.2.7 1.2.7. Centralized Logging and Alerting

- Deploys an rsyslog auto-scaling cluster behind a NLB, all syslogs forwarded to CloudWatch Logs

- Centralized access to "Cloud Security Service" Consoles from designated AWS account
- Centralizes logging to a single centralized S3 bucket (enables, configures and centralizes)
- VPC Flow logs w/Enhanced metadata fields (also sent to CWL)
- Organizational Cost and Usage Reports
- CloudTrail Logs including S3 Data Plane Logs (also sent to CWL)
- All CloudWatch Logs (includes rsyslog logs)
- Config History and Snapshots
- Route 53 Public Zone Logs (also sent to CWL)
- GuardDuty Findings
- Macie Discovery results
- ALB Logs
- SSM Inventory
- Security Hub findings
- SSM Session Logs (also sent to CWL)
- Resolver Query Logs (also sent to CWL)
- Email alerting for CloudTrail Metric Alarms, Firewall Manager Events, Security Hub Findings incl. GuardDuty Findings
- **NEW** Optionally collect Organization and ASEA configuration and metadata in a new restricted log archive bucket

1.3 1.3. Relationship with AWS Landing Zone Solution (ALZ)

The ALZ was an AWS Solution designed to deploy a multi-account AWS architecture for customers based on best practices and lessons learned from some of AWS' largest customers. The AWS Accelerator draws on design patterns from the Landing Zone, and re-uses several concepts and nomenclature, but it is not directly derived from it, nor does it leverage any code from the ALZ. The Accelerator is a standalone solution with no dependence on ALZ.

1.4 1.4. Relationship with AWS Control Tower

The AWS Secure Environment Accelerator now leverages AWS Control Tower!

With the release of v1.5.0, the AWS Accelerator adds the capability to be deployed on top of AWS Control Tower. Customers get the benefits of the fully managed capabilities of AWS Control Tower combined with the power and flexibility of the Accelerators Networking and Security orchestration.

1.5 1.5. Accelerator Installation Process (Summary)

This summarizes the installation process, the full installation document can be found in the documentation section below.

- Create a config.json (or config.yaml) file to represent your organizations requirements ([several samples provided](#))
- Create a Secrets Manager Secret which contains a GitHub token that provides access to the Accelerator code repository
- Create a unique S3 input bucket in the management account of the region you wish to deploy the solution and place your config.json and any additional custom config files in the bucket
- Download and execute the latest [release](#) installer CloudFormation template in your management accounts preferred 'primary' / 'home' region
- Wait for:
 - CloudFormation to deploy and start the Code Pipeline (~5 mins)
 - Code Pipeline to download the Accelerator codebase and install the Accelerator State Machine (~10 mins)
 - The Accelerator State Machine to finish execution (~1.25 hrs Standalone version, ~2.25 hrs Control Tower Version)
- Perform required one-time [post installation](#) activities (configure AWS SSO, set firewall passwords, etc.)
- On an ongoing basis:
 - Use AWS Organizations to create new AWS accounts, which will automatically be guardrailed by the Accelerator
 - Update the config file in CodeCommit and run the Accelerator State Machine to:
 - deploy, configure and guardrail multiple accounts at the same time (~25 min Standalone, ~50 min/account Control Tower)
 - change Accelerator configuration settings (~25 min)

2. Installation & Upgrades

2.1 Accelerator Installation and Upgrades

This section contains information on the installation and upgrade procedures for ASEA.

- Installation
 - [Installation Guide](#)
 - [Sample Configurations and Customization](#)
 - [State Machine Behavior](#)
 - [Splitting the Config File](#)
 - [Considerations with Existing Organizations](#)
 - [Importing ALZ Accounts](#)
 - [Open Releases](#)
- Upgrades
 - [Upgrade Guide](#)
 - [v1.5.0 Upgrade Instructions](#)
- Functionality
 - [Services](#)
 - [Pricing](#)
 - [Architecture Diagrams](#)
 - [Key Account & Capability Overview](#)
 - [Centralized Logging Details](#)
 - [Accelerator Object Naming](#)
 - [Open Roadmap](#)

2.2 Installation

2.2.1 1. Accelerator Installation Guide

1.1. Overview

We encourage customers installing the Accelerator to get the support of their local AWS account team (SA, TAM, CSM, ProServe) to assist with the installation of the Accelerator, as the Accelerator leverages, deploys, or orchestrates over 50 different AWS services.

Users are strongly encouraged to also read the [Accelerator Operations/Troubleshooting Guide](#) before installation and the [FAQ](#) while waiting for the installation to complete. The Operations/Troubleshooting Guide provides details as to what is being performed at each stage of the installation process, including detailed troubleshooting guidance.

These installation instructions assume one of the prescribed architectures is being deployed.

1.2. Prerequisites

1.2.1. GENERAL

- Management or root AWS Organization account (the AWS Accelerator cannot be deployed in an AWS sub-account)
- No additional AWS accounts need to be pre-created before Accelerator installation
- If required, a limit increase to support your desired number of new AWS sub-accounts (default limit is 10 sub-accounts)
- **recent changes to new AWS account limits are causing accelerator installation failures, please work with your local account team to increase your limits**
- Valid Accelerator configuration file, updated to reflect your requirements (see below)
- Determine your primary or Accelerator `control` or `home` region, this is the AWS region in which you will most often operate
- Government of Canada customers are still required to do a standalone installation at this time, please request standalone installation instructions from your Account SA or TAM
- The Accelerator *can* be installed into existing AWS Organizations - see caveats and notes [here](#)
- Existing AWS Landing Zone Solution (ALZ) customers are required to remove their ALZ deployment before deploying the Accelerator. Scripts are available to assist with this process.
- Changes to the Accelerator codebase are strongly discouraged unless they are contributed and accepted back to the solution. Code customization will block the ability to upgrade to the latest release and upgrades are encouraged to be done between quarterly to semi-annually. The solution was designed to be extremely customizable without changing code, existing customers following these guidelines have been able to upgrade across more than 50 Accelerator releases, while maintaining their customizations and gaining the latest bug fixes, features and enhancements without any developer or professional services based support. Please see [this](#) FAQ for more details.

1.3. Production Deployment Planning

1.3.1. GENERAL

For any deployment of the Accelerator which is intended to be used for production workloads, you must evaluate all these decisions carefully. Failure to understand these choices could cause challenges down the road. If this is a "test" or "internal" deployment of the Accelerator which will not be used for production workloads, you can leave the default config values.

Config file [schema](#) documentation (Draft)

1.3.2. OU STRUCTURE PLANNING

Plan your OU and core account structure carefully. By default, we suggest: Security, Infrastructure, Central, Sandbox, Dev, Test, Prod.

- The Security OU will contain the Security account, the Log Archive account, and the Organization Management account.
- The Infrastructure OU will hold the remainder of the accounts shared or utilized by the rest of the organization (Shared Network , Perimeter , and Operations).
- The remainder of the OUs correspond with major permission shifts in the SDLC cycle and NOT every stage an organization has in their SDLC cycle (i.e. QA or pre-prod would be included in one of the other OUs).
- The Central OU is used to hold accounts with workloads shared across Dev, Test, and Prod environments like centralized CI/CD tooling.
- The v1.5.0+ releases align the Accelerator OU and account structure with AWS multi-account guidance, splitting the core OU into the Security and Infrastructure OUs.

Note: While OUs can be renamed or additional OUs added at a later point in time, deployed AWS accounts CANNOT be moved between top-level OUs (guardrail violation), nor can top-level OUs easily be deleted (requires deleting all AWS accounts from within the OU first).

1.3.3. NETWORK CONFIGURATION PLANNING

If deploying the prescriptive architecture using the Full or Lite sample config files, you will need the following network constructs:

1. Six (6) RFC1918 Class B address blocks (CIDR's) which do not conflict with your on-premise networks (a single /13 block works well)
 - VPC CIDR blocks cannot be changed after installation, this is simply the way the AWS platform works, given everything is built on top of them. Carefully consider your address block selection.
 - one block for each OU, except Sandbox which is not routable (Sandbox OU will use a 7th non-routed address block)
 - the "core" Class B range will be split to support the Endpoint VPC and Perimeter VPC (with extra addresses remaining for future use)
 - Given a shared VPC architecture is leveraged (prevents stranded islands of CIDR blocks and reduces networking costs), we have assigned a class B address block to each VPC to future proof the deployment. Smaller customers can successfully deploy with a half class B CIDR block per shared VPC.
2. Two (2) RFC6598 /23 address blocks (Government of Canada (GC) requirement only)
 - Used for AWS Managed Active Directory (MAD) deployment and perimeter underlay network
 - non-GC customers can replace the RFC6598 address space with the extra unused addresses from the above RFC1918 CIDR range above (the App2 subnets in the Central VPC and the Perimeter VPC address space)
3. BGP ASN's for network routing, one for each of:
 - Transit Gateway (one unique ASN per TGW, multi-region example requires a second ASN)
 - IPSec VPN Firewall Cluster (if deployed)
 - VGW for Direct Connect connectivity (only shown in the config.multi-region-example.json)
 - For example: the Control Tower with Network Firewall example config requires a single BGP ASN for the TGW, the IPSec VPN example requires two BGP ASN's, and the multi-region example requires five unique BGP ASN's.

NOTE: Prior to v1.5.0 CIDR ranges were assigned to each VPC and subnet throughout the config file. This required customers to perform extensive updates across the config file when needing to move to specific IP ranges compatible with a customer's existing on-premise networks.

While this is still supported for those wanting to control exactly what address is used on every subnet, the solution has added support for dynamic CIDR assignments and the sample config files have been updated to reflect. New installs will have CIDR's pulled from CIDR pools, defined in the global-options section of the config file with state maintained in DynamoDB.

The v1.5.0 [custom upgrade guide](#) will provides details on the upgrade process and requirements to migrate to the new CIDR assignment system, if desired. A [script](#) was created to assist with this migration.

1.3.4. DNS, DOMAIN NAME, TLS CERTIFICATE PLANNING

If deploying the prescriptive architecture, you must decide on:

1. A unique Windows domain name (`organizationaws / organization.aws` , `organizationcloud / organization.cloud` , etc.). Given this is designed as the primary identity store and used to domain join all cloud hosted workloads, changing this in future is difficult. Pick a Windows domain name that does NOT conflict with your on-premise AD domains, ensuring the naming convention conforms to your organizations domain naming standards to ensure you can eventually create a domain trust between the MAD and on-premise domains/forests
2. DNS Domain names and DNS server IP's for on-premise private DNS zones requiring cloud resolution (can be added in future)
3. DNS Domain for a cloud hosted public zone "public": ["organization.cloud-nuage.canada.ca"] (can be added in future)
4. DNS Domain for a cloud hosted private zone "private": ["organization.cloud-nuage.gc.ca"] (can be added in future)
5. Wildcard TLS certificate for each of the 2 previous zones (can be added/changed in future)

1.3.5. EMAIL ADDRESS PLANNING

1. While you require a minimum of 6 **unique** email addresses (1 per sub-account being created), we recommend at least 20 **unique** email ALIASES associated with a single mailbox, never used before to open AWS accounts, such that you do not need to request new email aliases every time you need to create a new AWS account and they can all be monitored via a single mailbox. These email addresses can **never** have been used to previously open an AWS account.
2. You additionally require email addresses for the following additional purposes (these can be existing monitored mailboxes and do not need to be unique):
 - Accelerator execution (state machine) notification events (1 address)
 - High, Medium and Low security alerts (3 addresses if you wish to segregate alerts)
 - Budget notifications

1.3.6. CENTRALIZED INGRESS/EGRESS FIREWALLS

As of v1.5.0 the Accelerator offers multiple automated firewall deployment options:

a) AWS Network Firewall (native AWS Cloud service)

- Defined in the config file as part of a VPC

b) 3rd party firewalls interconnected to the cloud tenancy via IPSec VPN (Active/Active using BGP + ECMP)

- Defined in the config file under deployments w/TGW VPN attachments
- this was the only automated option prior to v1.5.0
- a sample Fortinet Fortigate configuration is provided (both PAYGO and BYOL supported)
- For Fortinet BYOL, requires minimum 2 valid license files (evaluation licenses adequate) (can be added in future)

c) 3rd party firewalls interconnected to the cloud tenancy via Gateway Load Balancer (GWLB) in an auto-scaling group

- Defined in the config file under both deployments and load balancers
- a sample Checkpoint CloudGuard configuration is provided (both PAYGO and BYOL supported)

d) Customer gateway (CGW) creation, to enable connectivity to on-premises firewalls or manually deployed cloud firewalls

- Defined in the config file under deployments w/TGW VPN attachments (but without an AMI or VPC association)

Examples of each of the firewall options have been included as variants of the Lite config file [example](#).

Note: While we only provide a single example for each 3rd party implementation today, the implementations are generic and should be usable by any 3rd party firewall vendor, assuming they support the required features and protocols. The two examples were driven by customer demand and heavy lifting by the 3rd party vendor. We look forward to additional vendors developing and contributing additional sample configurations. For new 3rd party integrations, we encourage the use of the GWLB approach.

1.3.7. OTHER

1. We recommend installing with the default Accelerator Name (ASEA) and Accelerator Prefix (ASEA-), but allow customization. Prior to v1.5.0 the defaults were (PBMM) and (PBMMAcceL-) respectively.
 - the Accelerator name and prefix **CANNOT** be changed after the initial installation;
 - the Accelerator prefix including the mandatory dash cannot be longer than 10 characters.
2. New installations, which now leverage Control Tower, require the `organization-admin-role` be set to `AWSControlTowerExecution`. Existing standalone installations will continue to utilize their existing role name for the `organization-admin-role`, typically `OrganizationAccountAccessRole`, as this role is used by AWS Organizations by default when no role name is specified while creating AWS accounts through the AWS console.
 - the Accelerator leverages this role name to create all new accounts in the organization;
 - this role name, as defined in the config file, **MUST** be utilized when manually creating all new sub-accounts in the Organization;
 - existing installs wishing to change the role name are required to first deploy a new role with a trust to the root account, in all accounts in the organization.

1.4. Accelerator Pre-Install Steps

1.4.1. GENERAL

Before installing, you must first:

1. Login to the Organization **Management (root) AWS account** with `AdministratorAccess`.
 2. **Set the region to your desired home region** (i.e. `ca-central-1`)
 3. Install AWS Control Tower:
 - Government of Canada customers are *required* to skip this step
 - OU and account names can ONLY be customized during initial installation. These values MUST match with the values supplied in the Accelerator config file.
 - a. Go to the AWS Control Tower console and click `Set up Landing zone`
 - b. Select your `home region` (i.e. `ca-central-1`) - the Accelerator home region must match the Control Tower home region
 - c. Leave the Region deny setting set to `Not enabled` - the Accelerator needs a customized region deny policy
 - d. Select *all* regions for `Additional AWS Regions for governance`, click `Next`
 - The Control Tower and Accelerator regions MUST be properly aligned
 - If a region is not `governed` by Control Tower, it must NOT be listed in `control-tower-supported-regions`
 - To manage a region requires the region:
 - be enabled in Control Tower (if supported)
 - added to the config file `control-tower-supported-regions` list (if supported)
 - added to the config file `supported-regions` list (even if not supported by Control Tower, as the Accelerator can manage regions not yet supported by Control Tower, but only when NOT listed in `control-tower-supported-regions`)
 - While we highly recommend guardrail deployment for all AWS enabled by default regions, at minimum
 - the home region MUST be enabled in Control Tower and must be listed in `control-tower-supported-regions`
 - both the home-region and `#{GBL*REGION}` must be listed in `supported-regions`
 - e. For the `Foundational OU`, leave the default value `Security`
 - f. For the `Additional OU` provide the value `Infrastructure`, click `Next`
 - g. Enter the email addresses for your `Log Archive` and `Audit` accounts, change the `Audit` account name to `Security`, click `Next` - OU and account names can ONLY be customized during initial installation. OU names, account names and email addresses `_must*` match identically with the values supplied in the Accelerator config file.
 - h. Select `Enabled` for AWS CloudTrail configuration (if not selected), click `Next`
 - i. Click `Set up Landing zone` and wait ~60 minutes for the Control Tower installation to complete
 - j. Select `Add or register organizational units`, Click `Add an OU`
 - k. Type `Dev`, click `Add`, wait until the OU is finished provisioning (or it will error)
 - l. Repeat step 9 for each OU (i.e. `Test`, `Prod`, `Central`, `Sandbox`)
 - m. Select `Account factory`, Edit, Subnets: 0, Deselect all regions, click `Save`
 - n. In AWS Organizations, move the Management account from the `root OU` into the `Security OU`
4. Verify:
 - a. AWS Organizations is enabled in `ALL features` mode
 - if required, navigate to AWS Organizations, click `Create Organization`, `Create Organization`
 - b. Service Control Policies are enabled
 - if required, in Organizations, select `Policies`, `Service control policies`, `Enable service control policies`
 5. Verify the Organization Management (root) account email address
 - In AWS Organizations, Settings, "[Send Verification Request](#)"
 - Once it arrives, complete the validation by clicking the validation link in the email

6. Create a new KMS key to encrypt your source configuration bucket (you can use an existing key)
 - AWS Key Management Service, Customer Managed Keys, Create Key, Symmetric, and then provide a key name (`ASEA-Source-Bucket-Key`), Next
 - Select a key administrator (Admin Role or Group for the Organization Management account), Next
 - Select key users (Admin Role or Group for the Organization Management account), Next
 - Validate an entry exists to "Enable IAM User Permissions" (critical step if using an existing key)
 - `"arn:aws:iam::123456789012:root"` , where `123456789012` is your **Organization Management** account ID.
 - Click Finish
 - Select the new key, Select Key Rotation , Automatically rotate this CMK every year , click Save.
7. Enable "Cost Explorer" (My Account, Cost Explorer, Enable Cost Explorer)
 - With recent platform changes, Cost Explorer *may* now be auto-enabled (unable to confirm)
8. Enable "Receive Billing Alerts" (My Account, Billing Preferences, Receive Billing Alerts)
9. It is **extremely important** that **all** the account contact details be validated in the Organization Management (root) account before deploying any new sub-accounts.
 - This information is copied to every new sub-account on creation.
 - Subsequent changes to this information require manually updating it in **each** sub-account.
 - Go to My Account and verify/update the information lists under both the Contact Information section and the Alternate Contacts section.
 - Please ESPECIALLY make sure the email addresses and Phone numbers are valid and regularly monitored. If we need to reach you due to suspicious account activity, billing issues, or other urgent problems with your account - this is the information that is used. It is CRITICAL it is kept accurate and up to date at all times.

1.4.2. CREATE GITHUB PERSONAL ACCESS TOKEN AND STORE IN SECRETS MANAGER

As of v1.5.0, the Accelerator offers deployment from either GitHub or CodeCommit:

GitHub (recommended)

1. You require a GitHub access token to access the code repository
2. Instructions on how to create a personal access token are located [here](#).
3. Select the scope `public_repo` underneath the section `repo: Full control over private repositories` .
4. Store the personal access token in Secrets Manager as plain text. Name the secret `accelerator/github-token` (case sensitive).
 - Via AWS console
 - Store a new secret, and select Other type of secrets , Plaintext
 - Paste your secret with no formatting no leading or trailing spaces (i.e. completely remove the example text)
 - Select the key you created above (`ASEA-Source-Bucket-Key`),
 - Set the secret name to `accelerator/github-token` (case sensitive)
 - Select Disable rotation

CodeCommit (alternative option)

Multiple options exist for downloading the GitHub Accelerator codebase and pushing it into CodeCommit. As this option is only for advanced users, detailed instructions are not provided.

1. In your AWS Organization Management account, open CodeCommit and create a new repository named `aws-secure-environment-accelerator`
2. Go to GitHub and download the repository Source code zip or tarball for the [release](#) you wish to deploy
 - Do NOT download the code off the main GitHub branch, this will leave you in a completely unsupported state (and with beta code)
3. Push the extracted codebase into the newly created CodeCommit repository, maintaining the file/folder hierarchy
4. Set the default CodeCommit branch for the new repository to main
5. Create a branch following the Accelerator naming format for your release (i.e. `release/v1.5.5`)

1.4.3. AWS INTERNAL (EMPLOYEE) ACCOUNTS ONLY

If deploying to an internal AWS employee account and installing the solution with a 3rd party firewall, you need to enable Private Marketplace (PMP) before starting:

1. In the Organization Management account go here: <https://aws.amazon.com/marketplace/privatemarketplace/create>
2. Click `Create a Private Marketplace`, and wait for activation to complete
3. Go to the "Account Groups" sub-menu, click `Create account group`
4. Enter an Account Group Title (i.e. `Default`) and Add the Management (root) account number in `Associate AWS account`
5. Associate the default experience `New Private Marketplace`, then click `Create account group` and wait for it to create
6. Go to "Experiences" sub-menu, select `New Private Marketplace`
7. Select the "Settings" sub-tab, and click the `Not Live` slider to make it `Live` and wait for it to complete
8. Ensure the "Software requests" slider is set to `Requests off` and wait for it to complete
9. Change the name field (i.e. append `-PMP`) and change the color, so it is clear PMP is enabled for users, click `Update`
10. Go to the "Products" sub-tab, then select the `ALL AWS Marketplace products` nested sub-tab
11. Search Private Marketplace for the Fortinet or Checkpoint products and select
 - `Fortinet FortiGate (BYOL) Next-Generation Firewall` and
 - `Fortinet FortiManager (BYOL) Centralized Security Management` **or**
 - `CloudGuard Network Security for Gateway Load Balancer - BYOL` and
 - `Check Point Security Management (BYOL)`
12. Select "Add" in the top right
 - Due to PMP provisioning delays, this sometimes fails when attempted immediately following enablement of PMP or if adding each product individually - retry after 20 minutes.
13. While not used in this account, you must now subscribe to the two subscriptions and accept the EULA for each product (you will need to do the same in the perimeter account, once provisioned below)
 - To subscribe, select the "Approved products" tab
 - Click on the product you want to subscribe, in this case `Fortinet FortiGate (BYOL) Next-Generation Firewall` and `Fortinet FortiManager (BYOL) Centralized Security Management` **or** `CloudGuard Network Security for Gateway Load Balancer - BYOL` and `Check Point Security Management (BYOL)`
 - Click on "Continue to Subscribe"
 - Click on "Accept Terms" and wait for subscription to be completed
 - If you are deploying in any region except `ca-central-1` or wish to switch to a different license type, you need the new AMI IDs. After successfully subscribing, continue one more step and click the "Continue to Configuration". When you get the below screen, select your region and version (**Fortinet v6.4.7**, **Checkpoint Mgmt** `R81.10-335.883` and **CloudGuard** `R80.40-294.374` recommended at this time). Marketplace will provide the required AMI ID. Document the two AMI IDs, as you will need to update them in your `config.json` file below.

{Vendor and product name appear here}

Continue to Launch

[< Product Detail](#) [Subscribe](#) [Configure](#)

Configure this software

Choose a fulfillment option below to select how you wish to deploy the software, then enter the information required to configure the deployment.

Delivery Method

64-bit (x86) Amazon Machine Image (AMI) ▼

Software Version

6.4.2 (Sep 02, 2020) ▼

Region

Asia Pacific (Sydney) ▼

Ami Id: ami-01a19dcdcc11b06d4

Product code: dlaioq277sglm5mw1y1dmeuqa

[Release notes \(updated September 2, 2020\)](#)

Pricing information

This is an estimate of typical software and infrastructure costs based on your configuration. Your actual charges for each statement period may differ from this estimate.

Software Pricing

Fortinet	\$0/hr
FortiGate (BYOL)	
Next-Generation Firewall	
BYOL	
running on c5.xlarge	

Infrastructure Pricing

EC2:	1 * c5.xlarge
Monthly Estimate:	\$160.00/month

1.5. Basic Accelerator Configuration

1. Select a sample config file as a baseline starting point

- **IMPORTANT: Use a config file from the GitHub code branch you are deploying from, as valid parameters change over time. The main branch is NOT the current release and often will not work with the GA releases.**
- sample config files can be found in [this](#) folder;
- descriptions of the sample config files and customization guidance can be found [here](#);
- unsure where to start, use the `config.Lite-CTNFW-example.json`, where CTNFW is for Control Tower w/NFW;
- These configuration files can be used, as-is, with only minor modification to successfully deploy the sample architectures;
- On upgrades, compare your deployed configuration file with the latest branch configuration file for any new or changed parameters;

2. At minimum, you MUST update the AWS account names and email addresses in the sample file:

- For existing accounts, they *must* match identically to both the account names and email addresses defined in AWS Organizations (including the management account);
- For new accounts, they must reflect the new account name/email you want created;
- All new AWS accounts require a unique email address which has never before been used to create an AWS account;
- When updating the budget or SNS notification email addresses within the sample config, a single email address for all is sufficient;
- Update the IP address in the "alarm-not-ip" variable with your on-premise IP ranges (used for the AWS-SSO-Authentication-From-Unapproved-IP alarm);
- If deploying the Managed AD, update the dns-domain, netbios-domain, log-group-name, as well as the AD users and groups that will be created;
- For a test deployment, the remainder of the values can be used as-is;
- While it is generally supported, we recommend not adding more than 1 or 2 workload accounts to the config file during the initial deployment as it will increase risks of hitting a limit. Once the Accelerator is successfully deployed, add the additional accounts to the config file and rerun the state machine.
- More information [here](#) on the fields in the config file that need to be updated.

- 20/230 -

3. A successful deployment of the prescriptive architecture requires VPC access to 9 AWS endpoints, you cannot remove both the perimeter firewalls (all public endpoints) and the 9 required central VPC endpoints from the config file (ec2, ec2messages, ssm, ssmessages, cloudformation, secretsmanager, kms, logs, monitoring).
4. When deploying to regions other than `ca-central-1`, you need to modify your config file as follows (for Canada Central 1, the AMI IDs are pre-populated for you):
 - a. Update the firewall and firewall manager AMI IDs to reflect your home regions regional AMI IDs (see 2.3.3, item 13), making sure you select the right version and region per the recommendations.
 - b. Validate all the Interface Endpoints defined in your config file are supported in your home region (i.e. Endpoint VPC). Remove unsupported endpoints from the config file, add additional endpoints as available.
 - c. If you are installing into a home region which is explicitly named in any of the replacements\addl_regions_x, remove it from the list. If deploying in us-east-1, remove `#{GBL_REGION}`.
5. Create an S3 bucket in your Organization Management account `your-bucket-name`
 - you must supply this bucket name in the CFN parameters *and* in the config file (`global-options\central-bucket`)
 - the bucket name *must* be the same in both spots
 - the bucket must have versioning enabled
 - the bucket must be `S3-KMS` encrypted using the `ASEA-Source-Bucket-Key` created above
6. Place your customized config file(s), named `config.json` (or `config.yaml`), in your new bucket
7. If required, place the firewall configuration and license files in the folder and path defined in the config file
 - For AWS Network Firewall: `nfw/nfw-example-policy.json`
 - For Fortinet: `firewall/firewall-example.txt`, `firewall/license1.lic` and `firewall/license2.lic`
 - We have made a sample available [here](#): `./reference-artifacts/Third-Party/`
 - the sample configures an active / active firewall pair with two tunnels per firewall
 - If you updated your perimeter VPC subnet names, you must also make these changes in your `firewall-example.txt` file
 - If you don't have any license files, update the config file with an empty array (`"license": []`). Do NOT use the following: `[""]`.
 - The basic Checkpoint configuration is stored directly in `config.json`
8. Place any defined certificate files in the folder and path defined in the config file
 - i.e. `certs/example1-cert.key`, `certs/example1-cert.crt`
 - Sample available [here](#): `./reference-artifacts/Certs-Sample/*`
 - Ideally you would generate real certificates using your existing certificate authority
 - Should you wish, instructions are provided to aid in generating your own self-signed certificates (Self signed certificates are NOT secure and simply for demo purposes)
 - Use the examples to demonstrate Accelerator TLS functionality only
9. Detach **ALL** SCPs (except `FullAWSAccess` which remains in place) from all OU's and accounts before proceeding
 - For Control Tower based installs do NOT detach Control Tower SCPs (i.e. `aws-guardrails-xxxxxx`)
 - Installation **will fail** if this step is skipped

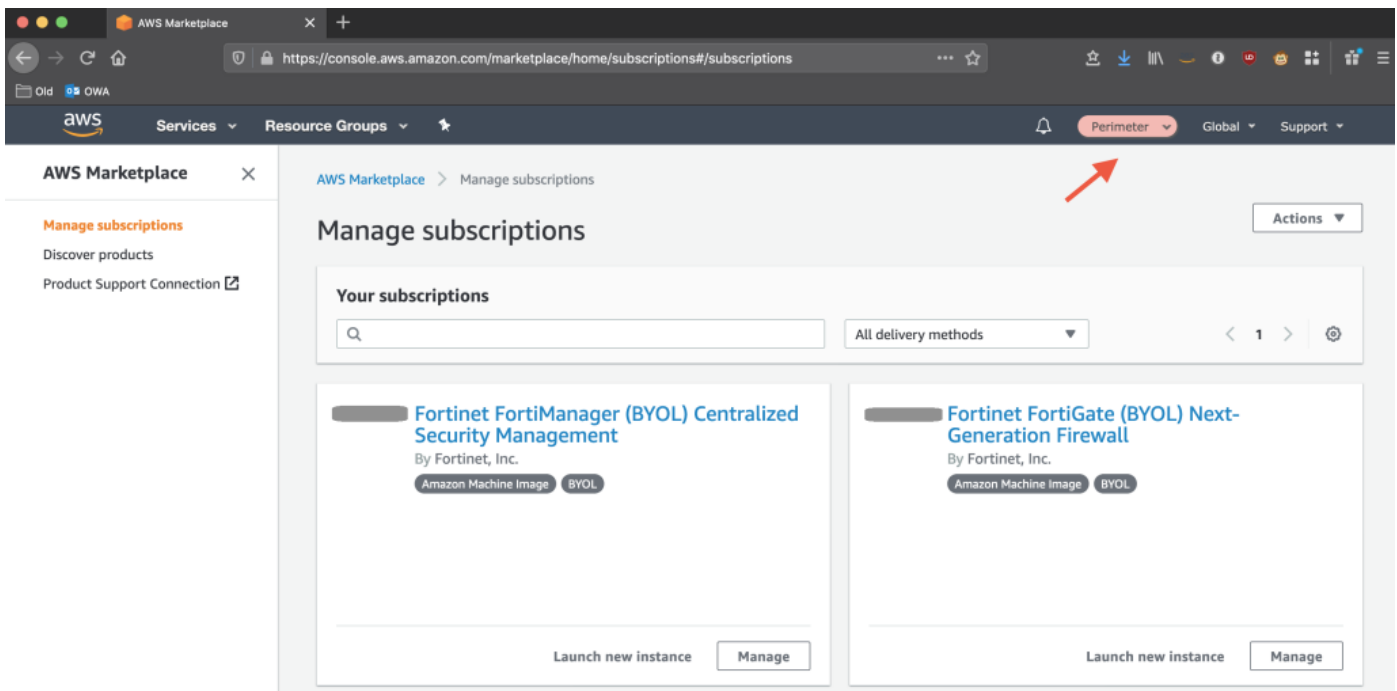
1.6. Installation

1. You can find the latest release in the repository [here](#).
 - We only support new installations of v1.5.5 or above (older releases continue to function)
2. Download the CloudFormation (CFN) template for the release you plan to install (either `AcceleratorInstallerXXX.template.json` for GitHub or `AcceleratorInstallerXXX-CodeCommit.template.json` for CodeCommit)
3. Use the provided CloudFormation template to deploy a new stack in your Management (root) AWS account
 - As previously stated we do not support installation in sub-accounts
4. Login to your Organization Management account and **make sure you are in your desired home region** (i.e. `ca-central-1`) (your desired primary or control region)

5. Navigate to **CloudFormation** in the AWS Console and click `Create stack with new resources (standard)`, then
 - Select "Template is ready"
 - For the "Specify template" select "Upload a template file"
 - Select the `*.template.json` file you downloaded in step 2 above
 - Click Next
6. Fill out the required parameters - **LEAVE THE DEFAULTS UNLESS SPECIFIED BELOW**
 - Specify `Stack Name` STARTING with `ASEA-` (case sensitive) suggest a suffix of `orgname` or `username`
 - Change `ConfigS3Bucket` to the name of the bucket you created above `your-bucket-name`
 - Add an `Email` address to be used for State Machine Status notification
 - The `GitHub Branch` should point to the release you selected
 - if upgrading, change it to point to the desired release
 - the latest stable branch is currently `release/v1.5.5`, case sensitive
 - click `Next`
7. Finish deploying the stack
 - Apply a tag on the stack, `Key= Accelerator`, `Value= ASEA` (case sensitive).
 - **ENABLE STACK TERMINATION PROTECTION** under `Stack creation options`
 - Click `Next`, Acknowledge resource creation, and click `Create stack`
 - The stack typically takes under 5 minutes to deploy.
8. Once deployed, you should see a CodePipeline project named `ASEA-InstallerPipeline` in your account. This pipeline connects to GitHub, pulls the code from the prescribed branch and deploys the Accelerator state machine.
 - if the CloudFormation fails to deploy with an `Internal Failure`, or, if the pipeline fails connecting to GitHub, then:
 - fix the issue with your GitHub secret created in section 2.3.2, then delete the Installer CloudFormation stack you just deployed, and restart at step 3 of this section.
9. For new stack deployments, when the stack deployment completes, the Accelerator state machine will automatically execute (in Code Pipeline). When upgrading you must manually `Release Change` to start the pipeline.
10. **While the pipeline is running:**
 - review the list of [Known Installation Issues](#) in the section below
 - review the Accelerator Basic Operation and Frequently Asked Questions ([FAQ](#)) [Document](#)
11. Once the pipeline completes (~10 mins), the main state machine, named `ASEA-MainStateMachine_sm`, will start in Step Functions
12. The state machine time is dependent on the quantity of resources being deployed. On an initial installation of a more complex sample configuration files, it takes approximately 2 hours to execute (depending on the configuration file). Timing for subsequent executions depends entirely on what resources are changed in the configuration file, but often takes as little as 20 minutes.
 - While you can watch the state machine in Step Functions, you will also be notified via email when the State Machine completes (or fails). Successful state machine executions include a list of all accounts which were successfully processed by the Accelerator.
13. The configuration file will be automatically moved into Code Commit (and deleted from S3). From this point forward, you must update your configuration file in CodeCommit.
14. You will receive an email from the State Machine SNS topic and the 3 SNS alerting topics. Please confirm all four (4) email subscriptions to enable receipt of state machine status and security alert messages. Until completed, you will not receive any email messages (must be completed within 7-days).
15. If the state machine **fails**:
 - Refer to the [Troubleshooting Guide](#) for instructions on how to inspect and retrieve the error
 - You can also refer to the [FAQ](#) and [Known Installation Issues](#)
 - Once the error is resolved, re-run the step function `ASEA-MainStateMachine_sm` using `{"scope": "FULL", "mode": "APPLY"}` as input

16. If deploying a prescriptive architecture with 3rd party firewalls, after the perimeter account is created in AWS Organizations, but before the Accelerator reaches Stage 2:

- a. NOTE: If you miss the step, or fail to execute it in time, no need to be concerned, you will simply need to re-run the main state machine (`ASEA-MainStateMachine_sm`) to deploy the firewall (no SM input parameters required)
- b. Login to the **perimeter** sub-account (Assume your `organization-admin-role`)
- c. Activate the 3rd party vendor firewall and firewall manager AMI's in the AWS Marketplace
 - Navigate back to your private marketplace
 - Note: Employees should see the private marketplace, including the custom color specified in prerequisite step 4 above.
 - Select "Discover products" from the side bar, then in the "Refine Results" select "Private Marketplace => Approved Products"
 - Subscribe and Accept the Terms for each product (firewall and firewall manager)
 - When complete, you should see the marketplace products as subscriptions in the **Perimeter account**:



17. If deploying the prescriptive architecture, once the main state machine (`ASEA-MainStateMachine_sm`) completes successfully, confirm the status of your perimeter firewall deployment

- If you have t2.micro ec2 instances running in any account which had the account-warming flag set to true, they will be removed on the next state machine execution;
- If your perimeter firewalls were defined but not deployed on first run, you will need to rerun the state machine. This happens when:
 - a. you were unable to activate the firewall AMI's before stage 2 (step 16)
 - b. we were not able to fully activate your account before we were ready to deploy your firewalls. This case can be identified by a running EC2 micro instance in the account, or by looking for the following log entry 'Minimum 15 minutes of account warming required for account'.
 - c. In these cases, simply select the `ASEA-MainStateMachine_sm` in Step Functions and select `Start Execution` (no SM input parameters required)

1.6.1. KNOWN INSTALLATION ISSUES

Current Issues:

- If dns-resolver-logging is enabled, VPC names containing spaces are not supported at this time as the VPC name is used as part of the log group name and spaces are not supported in log group names. By default in many of the sample config files, the VPC name is auto-generated from the OU name using a variable. In this situation, spaces are also not permitted in OU names (i.e. if any account in the OU has a VPC with resolver logging enabled and the VPC is using the OU as part of its name)
- On larger deployments we are occasionally seeing state machine failures when `Creating Config Recorders`. Simply rerun the state machine with the input of `{"scope": "FULL", "mode": "APPLY"}`.
- Occasionally CloudFormation fails to return a completion signal. After the credentials eventually fail (1 hr), the state machine fails. Simply rerun the state machine with the input of `{"scope": "FULL", "mode": "APPLY"}`
- If the State Machine fails on an initial execution of NEW-ACCOUNTS, it must be re-run to target the failed accounts (i.e. with `{"scope": "FULL", "mode": "APPLY"}`) or the new sub-accounts will fail to be properly guardrailed

Issues in Older Releases:

- New installs to releases prior to v1.5.5 are no longer supported.
- Upgrades to releases prior to v1.5.5 are no longer supported.
- Upgrades to v1.3.9 in preparation for an upgrade to v1.5.5 may be possible with manual workarounds.
- FROM 2022-08-07 to 2022-10-12: An issue with the version of cfn-init in the "latest" AWS standard Windows AMI will cause the state machine to fail during a new installation when deploying an RDGW host. RDGW hosts in existing deployments will fail to fully initialize if the state machine is or has been recently run and the auto-scaling group subsequently refreshes the host (default every 7 days).
- To temporarily workaround this issue, assume an administrative role in your `operations` account, open Systems Manager Parameter Store, and `Create parameter` with a Name of `/asea/windows-ami` and a value of `ami-0d336ea070bc06fb8` (which is the previous good AMI in ca-central-1), accepting the other default values. Update your config file to point to this new parameter by changing `image-path` (under `\deployments\mad`) to `/asea/windows-ami` instead of `/aws/service/ami-windows-latest/Windows_Server-2016-English-Full-Base`. Rerun your state machine. If you have an existing RDGW instance it should be terminated to allow the auto-scaling group to redeploy it. In other regions you will need to lookup the previous working ami-id (you cannot use `ami-0d336ea070bc06fb8`)
- **This issue was resolved with the 2022-10-12 Windows AMI [release](#). Customers that implemented this workaround must revert the above config file entry and rerun their state machines (the above AMI has been deprecated).**

1.7. Post-Installation

The Accelerator installation is complete, but several manual steps remain:

1. Enable and configure AWS SSO in your `home` region (i.e. ca-central-1)

- **NOTE: AWS SSO has been renamed to AWS IAM Identity Center (IdC). The IdC GUI has also been reworked. The below steps are no longer click-by-click accurate. An update to the below documentation is planned, which will also include instructions to delegate AWS IdC administration to the Operations account enabling connecting IdC directly to MAD, rather than through an ADC.**
- Login to the AWS Console using your Organization Management account
- Navigate to AWS Single Sign-On, click `Enable SSO`
- Set the SSO directory to AD ("Settings" => "Identity Source" => "Identity Source" => click `Change` , Select Active Directory, and select your domain from the list)
- Under "Identity Source" section, Click `Edit` beside "Attribute mappings", then set the `email` attribute to: `${dir:email}` and click `Save Changes`
- Configure Multi-factor authentication, we recommend the following minimum settings:
 - Every time they sign in (always-on)
 - Security key and built-in authenticators
 - Authenticator apps
 - Require them to provide a one-time password sent by email to sign in
 - Users can add and manage their own MFA devices
 - Create all the default permission sets and any desired custom permission sets
 - e.g. Select `AWS accounts` from the side bar, select "Permission sets" tab then `Create permission set`
 - Use an existing job function policy => `Next`
 - Select job function policy `AdministratorAccess`
 - Add Tags, Review and Create
 - repeat for each default permission set and any required custom permission sets
- For Control Tower based installations, remove the orphaned Permission Sets from each AWS accounts (select the account, expand Permission Sets, click Remove for each)
- Map MAD groups to permission sets and accounts
 - Select `AWS accounts` from the side bar and select `AWS organization` tab
 - Select the accounts you want to map to each MAD group and click `Assign users`
 - Select your DNS domain e.g. `example.local` , and search for the group you would like to assign (e.g. `aws-` for the pre-created groups) and click `Search connected directory`
 - Select the desired group `aws-log-archive-View`
 - Select the permission set you would like to assign to the MAD group to (e.g. `ViewOnlyAccess`)
 - Click `Finish` (Note: if it fails during provisioning, simply select the failed accounts and click on "Retry changes")
- AWS SSO should be the primary mechanism for all access to all AWS accounts in your Organization, to determine or update the login page for your organization:
 - Click on `Dashboard` within the AWS SSO console and note the `User portal URL`
 - Share this url with all your users
- NOTE: the url prefix can only be changed ONCE (ever) using the settings tab, so update with caution
- Any pre-created AWS MAD users passwords are available in secrets manager in the AWS management account. To reset these passwords login to the Operations account through AWS SSO, then:
 - Navigate to "Directory Service" , select the directory => Actions => Reset user password
 - Users can change their passwords from any MAD domain connected instance

2. Configure the new alb-forwarding feature (added in v1.5.0)

- AWS ALB's are published using DNS names which resolve to backing IP's which could silently change at any time due to a scaling event, maintenance, or a hardware failure. While published as a DNS name, ALB's can only target IP addresses. This presents a challenge as we need the ALB's in the perimeter account to target ALB's in the various back-end workload accounts;
- Alb-forwarding solves this problem by executing a small snippet of code every 60 seconds which updates managed ALB listeners with any IP changes, ensuring any managed flows do not go offline. This removes the requirement to leverage a 3rd party appliance to perform NAT to a DNS name;
- Alb-forwarding is enabled on the perimeter VPC in the perimeter account in many of the sample configs and can be used with Accelerator created or customer created ALB's. This setting causes a new DynamoDB table named `ASEA-ALB-IP-Forwarding-vpc-*` to be created in the perimeter account. Additionally the parameter store in the perimeter account contains the details of all Accelerator deployed ALB's across the environment for easy reference;
- Steps to configure:
 - First you need to manually create a listener on the front-end ALB (without a target group), multiple listeners are supported;
 - Next, for each application that needs to be published, a record needs to be added to the DynamoDB table, see sample below;
 - Records can be added to the table for any ALB in the account running the alb-forwarding tool. Records can be added at any time. DDB change logs will trigger the initial creation of the appropriate target group(s) and IP addresses will be verified and updated every 60 seconds thereafter.

Sample DynamoDB JSON to add an entry to the table: ▼

```
{
  "id": "App1",
  "targetAlbDnsName": "internal-Core-mydevacct1-alb-123456789.ca-central-1.elb.amazonaws.com",
  "targetGroupDestinationPort": 443,
  "targetGroupProtocol": "HTTPS",
  "vpcId": "vpc-0a6f44a80514daaa",
  "rule": {
    "sourceListenerArn": "arn:aws:elasticloadbalancing:ca-central-1:123456789012:listener/app/Public-DevTest-perimeter-alb/b1b12e7a0c412bf3/ef9b022a4 added8bdf",
    "condition": {
      "paths": ["/img/*", "/myApp2"],
      "hosts": ["aws.amazon.com"],
      "priority": 30
    }
  }
}
```

- where `id` is any unique text, `targetAlbDnsName` is the DNS address for the backend ALB for this application (found in parameter store), `vpcId` is the vpc ID containing the front-end ALB (in this account), `sourceListenerArn` is the arn of the listener of the front-end ALB, `paths` and `hosts` are both optional, but one of the two must be supplied. Finally, `priority` must be unique and is used to order the listener rules. Priorities should be spaced at least 40 apart to allow for easy insertion of new applications and forwarder rules.

- the provided `targetAlbDnsName` must resolve to addresses within a [supported](https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html) IP address space.

3. On a per role basis, you need to enable the CWL Account Selector in the Security and the Operations accounts, in each account:

- Go to CloudWatch, Settings, Under `Cross-account cross-region` select `Configure`, Under `View cross-account cross-region` select `Edit`, choose `AWS Organization account selector`, click `Save changes`

4. Configure central Ingress/Egress firewalls, if deployed

- Layer 3/4 `appliance` based inspection is an optional feature

General

- If deployed, login to any 3rd party firewalls and firewall manager appliances and update any default passwords;
- Tighten security groups on the 3rd party firewall instances (using the Accelerator configuration file), further limiting access to firewall management interfaces to a set of designated and controlled CIDR ranges;
- Update the firewall configuration per your organization's security requirements and best practices;
- Diagrams reflecting perimeter traffic flows when NFW and/or GWLB are used can be found [here](#) on slides 6 through 9.

AWS Network Firewall

- The AWS Network Firewall policies and rules deployed by the Accelerator, can only be updated using the Accelerator. Customers wishing to manage the AWS Network Firewall from the console GUI, must create a new policy with new rules created through the console and then manually associate this new policy to the Accelerator deployed Network Firewall. Customers can choose either option, but they cannot be mixed to ensure that Accelerator updates do not overwrite console based updates.

Fortinet

- Manually update firewall configuration to forward all logs to the Accelerator deployed NLB addresses fronting the rsyslog cluster
- login to each firewall, select `Log Settings`, check `Send logs to syslog`, put the NLB FQDN in the `IP Address/FQDN` field (stored in parameter store of perimeter account)
- Manually update the firewall configuration to connect perimeter ALB high port flows through to internal account ALB's
- Note: while this option is still available, a new alb-forwarding mechanism is available in v1.5.0+ (see section 2 above) which simplifies and eliminates this more complicated "NAT to DNS name" option;
- login to each firewall, switch to `FG-traffic vdom`, select `Policies & Objects`, select `Addresses`, `Expand Addresses`
- Set `Prod1-ALB-FQDN` to point to a reliable sub-account ALB FQDN, this is used for full-path health checks on **all** ALB's
- Set additional `DevX-ALB-FQDN`, `TestX-ALB-FQDN` and `ProdX-ALB-FQDN` to point to workload account ALB FQDNs
- Two of each type of ALB FQDN records have been created, when you need more, you need to create BOTH an additional FQDN and a new VIP, per ALB
- Each new VIP will use a new high port (i.e. 7007, 7008, etc.), all of which map back to port 443
- Detailed steps can be read [here](#).

Checkpoint

- Download and install Checkpoint SmartConsole client on your PC (only available for Windows, if you have a Mac, you can use Workspaces to install the client)
 - Retrieve the firewall and firewall manager SSL certificate from secrets manager in the perimeter account and save to a pem file (convert to ppk on Windows)
 - Login to the `AWS Console` for the `Perimeter` account
 - Navigate to `Secrets Manager` and select `Secrets`
 - Look for Secret name `ASEA/keypair/ASEA-Appliances` and click on it
 - Go to the `Secret value` section and click on `Retrieve secret value`
 - Make sure you have the `Plaintext` tab selected and copy the content to a `.pem` file using your preferred text editor. (the file will need to be converted to ppk format on Windows)
 - Wait approx. 25 min. after the managers "Launch time" and then SSH into the Firewall Manager using the SSL certificate (.pem file retrieved above) and login user `admin`
 - The firewall manager EC2 instance name is `ASEA-Checkpoint-FirewallMgr`
 - Once you SSH successfully, execute the following commands:
 - `set user admin password`
 - `set expert-password`
 - `set user admin shell /bin/bash`
 - `save config`
 - The following commands are useful for troubleshooting (in expert mode):
 - `autoprov_cfg -v` (check cme at Take 155 or greater)
 - `autoprov_cfg show all` (check cme configuration)
 - `cat /var/log/aws-user-data.log` (validate bootstrap, file should end with "Publish operation" succeeded (100%))
 - `tail -f /var/log/CPcme/cme.log` (watch to ensure it finds the instances, establishes SIC and adds the nodes)
 - Login to SmartConsole, and update the firewall policy per your organizations security requirements
 - An outbound rule allowing http and https should exist
 - From the RDGW host in Operations, test to see if outbound web browsing is enabled
 - NOTES:
 - No best practice or security configuration has been configured on the Checkpoint firewalls. These firewalls have been configured to work with GWLB, but otherwise have the default/basic Checkpoint out-of-box configuration installed
 - Do NOT reboot the Checkpoint appliances until bootstrap is complete (~25 minutes for the manager), or you will be required to redeploy the instance
5. Recover root passwords for all sub-accounts and apply strong passwords
- Process documented [here](#)
6. Enable MFA for **all** IAM users and **all** root account users, recommendations:
- Yubikeys provide the strongest form of MFA protection and are strongly encouraged for all account root users and all IAM users in the Organization Management (root) account
 - the Organization Management (root) account requires a dedicated Yubikey (if access is required to a sub-account root user, we do not want to expose the Organization Management accounts Yubikey)
 - every ~50 sub-accounts requires a dedicated Yubikey (minimize the required number of Yubikeys and the scope of impact should a Yubikey be lost or compromised)
 - each IAM breakglass user requires a dedicated Yubikey, as do any additional IAM users in the Organization Management (root) account. While some CSPs do not recommend MFA on the breakglass users, it is strongly encouraged in AWS
 - all other AWS users (AWS SSO, IAM in sub-accounts, etc.) can leverage virtual MFA devices (like Google Authenticator on a mobile device)

7. Customers are responsible for the ongoing management and rotation of all passwords on a regular basis per their organizational password policy. This includes the passwords of all IAM users, MAD users, firewall users, or other users, whether deployed by the Accelerator or not. We do NOT automatically rotate any passwords, but strongly encourage customers do so, on a regular basis.
8. During the installation we request required limit increases, resources dependent on these limits will not be deployed
 - a. Limit increase requests are controlled through the Accelerator configuration file `"limits":{}` setting
 - b. The sample configuration file requests increases to your EIP count in the perimeter account and to the VPC count and Interface Endpoint count in the `shared-network` account
 - c. You should receive emails from support confirming the limit increases
 - d. On the next state machine execution, resources blocked by limits should be deployed (i.e. additional VPC's and Endpoints)
 - e. If more than 2 days elapses without the limits being increased, on the next state machine execution, they will be re-requested
9. Note: After a successful install the Control Tower `Organizational units'` dashboard will indicate `2 of 3` in the `Accounts enrolled` column for the Security OU, as it does not enable enrollment of the management account in guardrails. The Accelerator compliments Control Tower and enables guardrails in the management account which is important to high compliance customers.

1.8. Other Operational Considerations

- The Organization Management (root) account does NOT have any preventative controls to protect the integrity of the Accelerator codebase, deployed objects or guardrails. Do not delete, modify, or change anything in the Organization Management (root) account unless you are certain as to what you are doing. More specifically, do NOT delete, or change *any* buckets in the Organization Management (root) account.
- While generally protected, do not delete/update/change S3 buckets with `cdk-asea-`, or `asea-` in *any* sub-accounts.
- ALB automated deployments only supports Forward and not redirect rules.
- AWS generally discourages cross-account KMS key usage. As the Accelerator centralizes logs across an entire organization as a security best practice, this is an exception/example of a unique situation where cross-account KMS key access is required.
- Only 1 auto-deployed MAD in any mandatory-account is supported today.
- VPC Endpoints have no Name tags applied as CloudFormation does not currently support tagging VPC Endpoints.
- If the Organization Management (root) account coincidentally already has an ADC with the same domain name, we do not create/deploy a new ADC. You must manually create a new ADC (it won't cause issues).
- 3rd party firewall updates are to be performed using the firewall OS based update capabilities. To update the AMI using the Accelerator, you must first remove the firewalls and then redeploy them (as the EIP's will block a parallel deployment), or deploy a second parallel FW cluster and de-provision the first cluster when ready.
- When adding more than 100 accounts to an OU which uses shared VPC's, you must *first* increase the Quota `Participant accounts per VPC` in the shared VPC owner account (i.e. `shared-network`). Trapping this quota before the SM fails has been added to the backlog.
- The default limit for Directory Sharing is 125 accounts for an Enterprise Managed Active Directory (MAD), a quota increase needs to be manually requested through support from the account containing the MAD before this limit is reached. Standard MAD has a sharing limit of 5 accounts (and only supports a small quota increase). The MAD sharing limit is not available in the Service Quota's tools.

2.2.2 1. Accelerator Sample Configurations and Customization

1.1. Summary

- Sample config files can be found in [this](#) folder
- Most of the examples reflect a medium security profile (NIST, ITSG, FEDRAMP)
- Unsure where to start, use [config.lite-CTNFW-example.json](#) (CT w/NFW variant of option 2)
- Frugal and want something comprehensive to experiment with, use [config.test-example.json](#) (option 5)
- Config file [schema](#) documentation (Draft)
- Estimated monthly [pricing](#) for sample configurations

1.2. Sample Configuration Files with Descriptions

1.2.1. FULL CONFIGURATION ([CONFIG.EXAMPLE.JSON](#))

- The full configuration file was based on feedback from customers moving into AWS at scale and at a rapid pace. Customers of this nature have indicated that they do not want to have to upsize their perimeter firewalls or add Interface endpoints as their developers start to use new AWS services. These are the two most expensive components of the deployed architecture solution.
- Default settings:
- AWS Control Tower: No
- Firewall: IPSec VPN with Active/Active Fortinet cluster (uses BGP+ECMP)

1.2.2. LITE WEIGHT CONFIGURATION FILES

- Four variants with differing central ingress/egress firewalls
- *Variant 1: **Recommended starting point*** ([config.lite-CTNFW-example.json](#))
- Default Settings:
- AWS Control Tower: Yes
- Firewall: AWS Network Firewall
- *Variant 2: **Recommended for new GC PBMM customers*** ([config.lite-VPN-example.json](#))
- requires 3rd party licensing (BYOL or PAYGO)
- Default Settings:
- AWS Control Tower: No
- Firewall: IPSec VPN with Active/Active Fortinet cluster (uses BGP+ECMP)
- *Variant 3: ([config.lite-NFW-example.json](#))*
- Same as *Variant 1* config without AWS Control Tower
- Default Settings:
- AWS Control Tower: No
- Firewall: AWS Network Firewall
- *Variant 4: ([config.lite-GWLB-example.json](#))*
- requires 3rd party licensing (BYOL or PAYGO)
- Default Settings:
- AWS Control Tower: No
- Firewall: Gateway Load Balancer with Checkpoint firewalls in an autoscaling group

- To reduce solution costs and allow customers to grow into more advanced AWS capabilities, we created these lite weight configurations that does not sacrifice functionality, but could limit performance. These config files:
- only deploys the 9 required centralized Interface Endpoints (removes 50 from full config). All services remain accessible using the AWS public endpoints, but require traversing the perimeter firewalls
- removes the perimeter VPC Interface Endpoints
- reduces the Fortigate instance sizes from c5n.2xl to c5n.xl (VM08 to VM04) in *Variant 2: IPSec VPN with Active/Active Fortinet cluster option*
- removes the UnClass ou and VPC
- AWS Control Tower can be implemented in all sample configs using *Variant 1: AWS Control Tower with AWS Network Firewall* as an example (new installs only).
- The Accelerator allows customers to easily add or change this functionality in future, as and when required without any impact

1.2.3. ULTRA-LITE SAMPLE CONFIGURATION

- *Variant 1:* ([config.ultralite-CT-example.json](#))
- AWS Control Tower: Yes
- Firewall: None
- Networking: None
- *Variant 2:* ([config.ultralite-example.json](#))
- AWS Control Tower: No
- Firewall: None
- Networking: None
- This configuration file was created to represent an extremely minimalistic Accelerator deployment, simply to demonstrate the art of the possible for an extremely simple config. This example is NOT recommended as it violates many AWS best practices. This config has:
 - no `shared-network` or `perimeter` accounts
 - no networking (VPC, TGW, ELB, SG, NACL, endpoints) or `route53` (zones, resolvers) objects
 - no Managed AD, AD Connector, rsyslog cluster, RDGW host, or 3rd party firewalls
 - only enables/deploys AWS security services in 2 regions (ca-central-1, us-east-1) (Not recommended)
 - only deploys 2 AWS config rules w/SSM remediation
 - renamed log-archive (Logs), security (Audit) and operations (Ops) account names

1.2.4. MULTI-REGION SAMPLE CONFIGURATION ([CONFIG.MULTI-REGION-EXAMPLE.JSON](#))

- This configuration file was created to represent a more advanced multi-region version of the Full configuration file from configuration 1 above. This config:
 - adds a TGW in us-east-1, peered to the TGW in ca-central-1
 - adds TGW static routes, including several dummy sample static routes
 - adds a central Endpoint VPC in us-east-1 with us-east-1 endpoints configured
 - adds a shared VPC for all UnClass OU accounts in us-east-1, connected to the us-east-1 TGW (accessible through ca-central-1)
 - creates additional zones and resolver rules
 - Sends us-east-1 CloudWatch Logs to the central S3 log-archive bucket in ca-central-1
 - Deploys SSM documents to us-east-1 and remediates configured rules in UnClass OU
 - adds a local account specific VPC, in us-east-1, in the account MyUnClass and connects it to the us-east-1 TGW (i.e. shares TGW)
 - local account VPC set to use central endpoints, associates appropriate centralized hosted zones to VPC (also creates 5 local endpoints)
 - adds a VGW for DirectConnect to the perimeter VPC
 - adds the 3rd AZ in ca-central-1 (MAD & ADC in AZ a & b)

- Default Settings:
- AWS Control Tower: No
- Firewall: IPSec VPN with Active/Active Fortinet cluster (uses BGP+ECMP)

1.2.5. TEST CONFIGURATION ([CONFIG.TEST-EXAMPLE.JSON](#)) (USE FOR TESTING OR LOW SECURITY PROFILES)

- Further reduces solution costs, while demonstrating full solution functionality (NOT recommendend for production). This config file:
- uses the Lite weight configuration as the starting point (NFW variant)
- consolidates Dev/Test/Prod OU to a single Workloads OU/VPC
- only enables Security Hub, Config and Macie in ca-central-1 and us-east-1
- removes the Fortigate firewall cluster (per NFW variant)
- removes the rsyslog cluster
- reduces the RDGW instance sizes from t2.large to t2.medium
- reduces the size of the MAD from Enterprise to Standard edition
- removes the on-premise R53 resolvers (hybrid dns)
- reduced various log retention periods and the VPCFlow log interval
- removes the two example workload accounts
- adds AWS Network Firewall (NFW) and AWS NATGW for centralized ingress/egress (per NFW variant)
- Default Settings:
- AWS Control Tower: No
- Firewall: AWS Network Firewall

1.3. Deployment Customizations

1.3.1. [MULTI-FILE CONFIG FILE AND YAML FORMATTING OPTION](#)

- The sample configuration files are provided as single, all encompassing, json files. The Accelerator also supports both splitting the config file into multiple component files and configuration files built using YAML instead of json. Details can be found in the linked document.

1.3.2. [SAMPLE SNIPPETS](#)

- The sample configuration files do not include the full range of supported configuration file parameters and values, additional configuration file parameters and values can be found in the sample snippets document.

1.3.3. [THIRD PARTY FIREWALL EXAMPLE CONFIGS](#)

- The Accelerator is provided with a sample 3rd party configuration file to demonstrate automated deployment of 3rd party firewall technologies. Given the code is vendor agnostic, this process should be able to be leveraged to deploy other vendors firewall appliances. When and if other options become available, we will add them here as well.
- Automated [firewall configuration customization](#) possibilities
- Sample Fortinet Fortigate firewall config [file](#)

1.4. Other Configuration File Hints and Tips

- It is critical that all accounts that are leveraged by other accounts (i.e. accounts that any workload accounts are dependant on), are included in the mandatory-accounts section of the config file (i.e. shared-network, log-archive, operations)
- Account pointers within the config file point to the account key (i.e. (mandatory-account-configs\account-key) and NOT the account name field (mandatory-account-configs\account-key\account-name: "account name"). This allows for easy account names, duplicate account names, and no requirement to update account pointers during account renames.
- If any of the account pointers within `global-options` does not point to a valid mandatory account key, the State Machine will fail with the error `EnvironmentVariable value cannot be null before starting CodeBuild Phase -1`

- You cannot supply (or change) configuration file values to something not supported by the AWS platform
- For example, CWL retention only supports specific retention values (not any number)
- Shard count - can only increase/reduce by half the current limit. i.e. you can change from 1 - 2 , 2 - 3 , 4 - 6
- Always add any new items to the END of all lists or sections in the config file, otherwise
- Update validation checks will fail (VPC's, subnets, share-to, etc.)
- To skip, remove or uninstall a component, you can often simply change the section header, instead of removing the section
- change "deployments"/"firewalls" to "deployments"/"xxfirewalls" and it will uninstall the firewalls and maintain the old config file settings for future use
- Objects with the parameter `deploy: true`, support setting the value to false to remove the deployment
- As you grow and add AWS accounts, the Kinesis Data stream in the log-archive account will need to be monitored and have its capacity (shard count) increased by setting `"kinesis-stream-shard-count"` variable under `"central-Log-services"` in the config file
- Updates to NACL's requires changing the rule number (`100` to `101`) or they will fail to update
- When adding a new subnet or subnets to a VPC (including enabling an additional AZ), you need to:
 - increment any impacted NACL id's in the config file (`100` to `101` , `32000` to `32001`) (CFN does not allow nacl updates)
 - make a minor change to any impacted route table names (`MyRouteTable` to `MyRouteTable1`) (CFN does not allow updates to route table associated ids)
- The sample VPN firewall configuration uses an instance with 4 NIC's, make sure you use an instance size that supports 4 ENI's
- Firewall names, CGW names, TGW names, MAD Directory ID, account keys, and OU's must all be unique throughout the entire configuration file (also true for VPC names given NACL and security group referencing design)
- The configuration file *does* have validation checks in place that prevent users from making certain major unsupported configuration changes
- **The configuration file does NOT have extensive error checking. It is expected you know what you are doing. We eventually hope to offer a config file, wizard based GUI editor and add the validation logic in this separate tool. In most cases the State Machine will fail with an error, and you will simply need to troubleshoot, rectify and rerun the state machine.**
- You cannot move an account between top-level OU's. This would be a security violation and cause other issues. You can move accounts between sub-ou. Note: The Control Tower version of the Accelerator does NOT support sub-ou's.
- When using YAML configuration files, we only support the subset of yaml that converts to JSON (we do not support anchors)
- Security Group names were designed to be identical between environments, if you want the VPC name in the SG name, you need to do it manually in the config file
- Adding more than approximately 50 *new* VPC Interface Endpoints across *all* regions in any one account in any single state machine execution will cause the state machine to fail due to Route 53 throttling errors. If adding endpoints at scale, only deploy 1 region at a time. In this scenario, the stack(s) will fail to properly delete, also based on the throttling, and will require manual removal.
- We do not support Directory unsharing or ADC deletion, delete methods were not implemented. We only support ADC creation in mandatory accounts.
- If `use-central-endpoints` is changed from true to false, you cannot add a local VPC endpoint on the same state machine execution (add the endpoint on a prior or subsequent execution)
- If you update the 3rd party firewall names, be sure to update the routes and alb's which point to them. Firewall licensing occurs through the management port, which requires a VPC route back to the firewall to get internet access and validate the firewall license.
- Removing the AWS NFW requires 2 state machine executions, in the first you must remove all routes that reference the NFW, and in the second you can remove or xx out the NFW (also true for the GWLB implementation).

1.5. Config file and Deployment Protections

- The config file is moved to AWS CodeCommit after the first execution of the state machine to provide strong configuration history, versioning and change control
- After each successful state machine execution, we record the commit id of the config file used for that execution in secrets manager
- On **every** state machine execution, before making any changes, the Accelerator compares the latest version of the config file stored in CodeCommit with the version of the config file from the last successful state machine execution (after replacing all variables)

- If the config file includes any changes we consider to be significant or breaking, we immediately fail the state machine
- if a customer somehow accidentally uploads a different customers config file into their Accelerator CodeCommit repository, the state machine will fail
- if a customer makes what we consider to be a major change to the config file, the state machine will fail
- if a customer makes a change that we believe has a high likelihood to cause a deployment failure, the state machine will fail
- If a customer believes they understand the full implications of the changes they are making (and has made any required manual changes to allow successful execution), we have provided protection override flags. These overrides should be used with extremely caution:
- To provide maximum protection we have provided scoped override flags. Customers can provide a flag or flags to only bypass specific type(s) of config file validations or blocks. If using an override flag, we recommend customers use these scoped flags in most situations.
- If a customer is purposefully making extensive changes across the config file and wants to simply override all checks with a single override flag, we also have this option, but discourage it use.
- The various override flags and their format can be found in [here](#).

1.6. Summary of Example Config File Minimum Changes for New Installs

At a minimum you should consider reviewing the following config file sections and make the required changes.

1.6.1. GLOBAL OPTIONS

- S3 Central Bucket
- `global-options/central-bucket` : "AWSDOC-EXAMPLE-BUCKET"
- replace with `your-bucket-name` as referenced in the Installation Guide [Step #5](#)
- Central Log Services SNS Emails
- `global-options/central-log-services/sns-subscription-emails` : "myemail+notifyT-xxx@example.com"
- update the 3 email addresses (high, medium and low) as required. Each address will receives alerts or alarms of the specified level. The same email address can be used for all three.
- The default dynamic CIDR pools (`global-options/cidr-pools`) listed below are used to assign ranges based on the subnet mask set in each VPC and subnet throughout the configuration file.
- `global-options/cidr-pools/0/cidr` : "10.0.0.0/13"
- The main address pool used to dynamically assign CIDR ranges for most VPCs
- `global-options/cidr-pools/1/cidr` : "100.96.252.0/23"
- Address pool used to dynamically assign CIDR ranges for the Managed Active Directory subnets in the Ops account
- `global-options/cidr-pools/2/cidr` : "100.96.250.0/23"
- Address pool used to dynamically assign CIDR ranges for the Perimeter VPC
- `global-options/cidr-pools/3/cidr` : "10.249.1.0/24"
- A non-routable pool of addresses used to dynamically assign CIDR ranges for the Active Directory Connector subnets in the Organization Management/root account

1.6.2. MANDATORY ACCOUNT CONFIGS

- All mandatory accounts specific to your config file, that are present under the `mandatory-account-config` section require you to assign a unique email address for each account listed below. Replace the email values in the JSON config file for these accounts with unique email addresses.
- `mandatory-account-configs/shared-network/email` : "myemail+aseaT-network@example.com-----REPLACE-----"
- `mandatory-account-configs/operations/email` : "myemail+aseaT-operations@example.com-----REPLACE-----"
- `mandatory-account-configs/perimeter/email` : "myemail+aseaT-perimeter@example.com-----REPLACE-----"
- `mandatory-account-configs/management/email` : "myemail+aseaT-management@example.com-----REPLACE-----" (Note: This is the email of your root account)
- `mandatory-account-configs/log-archive/email` : "myemail+aseaT-log@example.com-----REPLACE-----"
- `mandatory-account-configs/security/email` : "myemail+aseaT-sec@example.com-----REPLACE-----"
- Budget Alerts email addresses need to be replaced with an email address in your organization. It can be the same email address for all budget alerts. Config located at the following path (Multiple exist for different thresholds, update all under each account):
- `mandatory-account-configs/shared-network/budget/alerts/emails` : "myemail+aseaT-budg@example.com"
- `mandatory-account-configs/perimeter/budget/alerts/emails` : "myemail+aseaT-budg@example.com"
- `mandatory-account-configs/management/budget/alerts/emails` : "myemail+aseaT-budg@example.com"
- For the `shared-network` account, review and update the following (or delete the sections):
- `mandatory-account-configs/shared-network/vpc/on-premise-rules/zone` : "on-premise-privatedomain1.example.ca" (qty 2)
- `mandatory-account-configs/shared-network/vpc/zones/private` : "cloud-hosted-privatedomain.example.ca"
- `mandatory-account-configs/shared-network/vpc/zones/public` : "cloud-hosted-publicdomain.example.ca"
- For the `operations` account, review and update the following:
- `mandatory-account-configs/operations/deployments/mad/dns-domain` : "example.local"
- `mandatory-account-configs/operations/deployments/mad/netbios-domain` : "example"
- `mandatory-account-configs/operations/deployments/mad/log-group-name` : "/\${ACCELERATOR_PREFIX_ND}/MAD/example.local" (replace example.local)
- `mandatory-account-configs/operations/deployments/mad/ad-users` (update user, email and group of each user as required)
- do not remove or change permissions on the `adconnector-usr`

- For `perimeter` account, review and update the following:
 - `mandatory-account-configs/perimeter/certificates/priv-key` : "certs/example1-cert.key"
 - `mandatory-account-configs/perimeter/certificates/cert` : "certs/example1-cert.crt"
- If you are using VPN config:
 - `mandatory-account-configs/perimeter/deployments/firewalls/image-id` : "ami-0d8e2e78e928def11"
- Update AMI with the AMI collected from the Marketplace for **Fortinet FortiGate (BYOL) Next-Generation Firewall**
 - `mandatory-account-configs/perimeter/deployments/xxfirewall-manager/image-id` : "ami-0e9f45c3ec34c3a9a"
- Update AMI with the AMI collected from the Marketplace for **Fortinet FortiManager (BYOL) Centralized Security Management**
- NOTE: Default config of "xxfirewall-manager" will prevent the firewall manager from being deployed. To deploy the firewall manager remove the "xx" to set the parameter to "firewall-manager"
- `mandatory-account-configs/perimeter/deployments/firewalls/license` : ["firewall/license1.lic", "firewall/license2.lic"]
- Two Fortinet FortiGate firewall licenses, if you don't have any license files, update the config file with an empty array ("license": []). Do NOT use the following: [""]
- Place files in a folder (eg. firewall) in the same S3 bucket in your Organization Management account as the deployment configuration file.
 - `mandatory-account-configs/perimeter/deployments/firewalls/config` : "firewall/firewall-example.txt"
- The Fortinet configuration file to initially configure the firewalls. Sample configuration files can be found in the [reference-artifacts/Third-Party](#) folder
- Place file in a folder (eg. firewall) in the same S3 bucket in your Organization Management account as the deployment configuration file.
- If you are using GWLB config:
 - `mandatory-account-configs/perimeter/deployments/firewalls/Checkpoint-Firewall` - `image-id` : "ami-0217611bf09d5b4c1"
- Update AMI with the AMI collected from the Marketplace for **CloudGuard Network Security for GWLB - BYOL**
 - `mandatory-account-configs/perimeter/deployments/firewall-manager/image-id` : "ami-0071a3b4ef9ac766a"
- Update AMI with the AMI collected from the Marketplace for **Checkpoint Security Management**
 - `mandatory-account-configs/perimeter/deployments/firewall-manager/version` : "R8110BYOLMGMT"
- Update version based on the selected ami-id version from the Private Marketplace
- For `management` , review and update the following:
 - `mandatory-account-configs/management/account-name` : "ASEA-Main"
- Update this field with your Organization Management (root) account name, if it is not set to ASEA-Main.
 - `mandatory-account-configs/management/iam/users`
- the names of your break-glass and ASEA operation users

1.6.3. WORKLOAD ACCOUNT CONFIGS

- As mentioned in the Installation Guide, we recommend not adding more than 1 or 2 workload accounts to the config file during the initial deployment as it will increase risks of hitting a limit. Once the Accelerator is successfully deployed, add the additional accounts back into the config file and rerun the state machine.
- Review the workload accounts in the config that you selected and change the name and email as desired
- Modify `mydevacct1` with the account name of your choosing
- Modify `mydevacct1/account-name` : "MyDev1" with the account name
- Modify `mydevacct1/email` : "myemail+aseaT-dev1@example.com-----REPLACE-----" with a unique email address for the account
- Modify `mydevacct1/description` : "This is an OPTIONAL SAMPLE workload account..." with a description relevant to your account
- Modify `mydevacct1/ou` : "Dev" with the OU that you would like the account to be attached to

1.6.4. ORGANIZATION UNITS

- For all organization units, update the budget alerts email addresses:
- `organizational-units/core/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- `organizational-units/Central/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- `organizational-units/Dev/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- `organizational-units/Test/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- `organizational-units/Prod/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- `organizational-units/Sandbox/default-budgets/alerts/emails` : "myemail+aseaT-budg@example.com"
- For organization units with `certificates` , review the certificates and update as you see fit. These certificates are used in the `alb` section under `alb/cert-name` of each OU

2.2.3 1. State Machine Behavior and Inputs

1.1. State Machine Behavior

Accelerator v1.3.0 makes a significant change to the manner in which the state machine operates. These changes include:

1. Reducing the `default scope` of execution of the state machine to only target newly created AWS accounts and AWS accounts listed in the mandatory accounts section of the config file.
 - `default scope` refers to running the state machine without any input parameters;
 - This new default scope disallows any changes to the config file outside new accounts;
 - NOTE: it is critical that accounts for which others are dependant upon, MUST be located within the `mandatory-account-configs` section of the config file (i.e. management, log-archive, security, operations, shared-network, perimeter, etc.).
2. The state machine now accepts a new input parameter, `scope`, which accepts the following values: `FULL` | `NEW-ACCOUNTS` | `GLOBAL-OPTIONS` | `ACCOUNT` | `OU`.
 - when the `scope` parameter is supplied, you must also supply the `mode` parameter. At this time `mode` only accepts the value `APPLY`. To be specific `"mode": "APPLY"` is mandatory when running the state machine with the `"scope":` parameter.
3. Starting the state machine with `{"scope": "FULL", "mode": "APPLY"}` makes the state machine execute as it did in v1.2.6 and below.
 - The state machine targets all AWS accounts and allows changes across any section of the config file;
 - The blocks and overrides described in section 1.4 above remain valid;
 - `FULL` mode must be run at least once immediately after any Accelerator version upgrade. Code Pipeline automatically starts the state machine with `{"scope": "FULL", "mode": "APPLY"}`. If the state machine fails for any reason after upgrade, the state machine must be restarted with these parameters until a successful execution of the state machine has completed.
4. Starting the state machine with `{"scope": "NEW-ACCOUNTS", "mode": "APPLY"}` is the same as operating the state machine with the `default scope` as described in the first bullet
5. Starting the state machine with `{"scope": "GLOBAL-OPTIONS", "mode": "APPLY"}` restricts changes to the config file to the `global-options` section.
 - If any other portion of the config file was updated or changed, the state machine will fail;
 - The global options scope executes the state machine on the entire managed account footprint.
6. Starting the state machine with `{"scope": "OU", "targetOus": [X], "mode": "APPLY"}` restricts changes to the config file to the specified `organizational-units` section(s) defined by `targetOus`.
 - When `scope=OU`, `targetOus` becomes a mandatory parameter;
 - X can be any one or more valid OU names, or the value `"ALL"`;
 - When `["ALL"]` is specified, the state machine targets all AWS accounts, but only allows changes to the `organizational-units` section of the config file;
 - When OUs are specified (i.e. `["Dev", "Test"]`), the state machine only targets mandatory accounts plus accounts in the specified OUs (Dev, Test), and only allows changes to the specified OUs sections (Dev, Test) of the config file;
 - If any other portion of the config file was updated or changed, the state machine will fail.
7. Starting the state machine with `{"scope": "ACCOUNT", "targetAccounts": [X], "mode": "APPLY"}` restricts changes to the config file to the specified `xxx-account-configs` section(s) defined by `targetAccounts`.
 - When `scope=ACCOUNT`, `targetAccounts` becomes a mandatory parameter;
 - X can be any one or more valid account numbers, the value `"NEW"`, or the value `"ALL"`;
 - When `["ALL"]` is specified, the state machine targets all AWS accounts, but only allows changes to the `xxx-account-configs` sections of the config file;
 - When specific accounts and/or `NEW` is specified (i.e. `["NEW", "123456789012", "234567890123"]`), the state machine only targets mandatory accounts plus the listed accounts and any newly created accounts. It also only allows changes to the specified accounts sections (New, 123456789012, 234567890123) of the config file;
 - If any other portion of the config file was updated or changed, the state machine will fail.

Starting in v1.3.0, we recommend running the state machine with the parameters that most tightly scope the state machines execution to your planned changes and minimizing the use of `FULL` scope execution.

- should you accidentally change the wrong section of the config file, you will be protected;
- as you grow and scale to hundreds or thousands of accounts, your state machine execution time will remain fast.

NOTE 1: The `scope` setting has no impact on SCP application, limit requests, custom tagging, or directory sharing.

NOTE 2: All comparisons for config file changes are assessed AFTER all replacements have been made. Changing variable names which result in the same end outcome do NOT appear as a change to the config file.

1.2. Accelerator State Machine Inputs

1.2.1. REBUILD DYNAMODB TABLE CONTENTS

With the exception of the Outputs table, the contents of the Accelerator DynamoDB tables are rebuilt on every state machine execution. We recently started depending on the Outputs DynamoDB tables to ensure the parameters in parameter store are consistently maintained in the same order as objects are created and deleted. Should the CONTENTS of the tables be destroyed or corrupted, customers can force a rebuild of the CloudFormation Outputs in DynamoDB by starting the state machine with the parameter:

```
{ "storeAllOutputs": true }
```

This should be completed BEFORE running the state machine with a corrupt or empty DynamoDB table or the Accelerator is likely to reorder a customers parameters. If the DynamoDB tables were completely destroyed, they must be recreated before running the state machine with this parameter.

1.2.2. BYPASS ALL CONFIG FILE VALIDATION CHECKS

This parameter should be specified with extreme caution, as it bypasses all config file validation. The state machine typically has protections enabled preventing customers from making breaking changes to the config file. Under certain conditions with the support of a trained expert, bypassing these checks is required. Start the state machine with the parameter:

```
{ "overrideComparison": true }
```

Customers are encouraged to use the specific override variables below, rather than the all-inclusive override, to ensure they only bypasses intended config changes.

1.2.3. BYPASSING SPECIFIC CONFIG FILE VALIDATION CHECKS

Providing any one or more of the following flags will only override the specified check(s):

```
{
  "configOverrides": {
    "ov-global-options": true,
    "ov-del-accts": true,
    "ov-ren-accts": true,
    "ov-acct-email": true,
    "ov-acct-ou": true,
    "ov-acct-vpc": true,
    "ov-acct-subnet": true,
    "ov-acct-vpc-optin": true,
    "ov-tgw": true,
    "ov-mad": true,
    "ov-ou-vpc": true,
    "ov-ou-subnet": true,
    "ov-share-to-ou": true,
    "ov-share-to-accounts": true,
    "ov-naci": true,
    "ov-nfw": true
  }
}
```

1.2.4. GENERATE VERBOSE LOGGING WITHIN STATE MACHINE

- Added "verbose": "1" state machine input options
- parameter is optional

- parameter defaults to 0

```
{ "scope": "FULL", "mode": "APPLY", "verbose": "1" }
```

1.2.5. STATE MACHINE SCOPING INPUTS

Summary of inputs, per section 1.1 above:

```
{ "scope": "FULL", "mode": "APPLY" }
```

```
{ "scope": "NEW-ACCOUNTS", "mode": "APPLY" }
```

```
{ "scope": "GLOBAL-OPTIONS", "mode": "APPLY" }
```

```
{ "scope": "OU", "targetOus": ["ou-name", "ou-name"], "mode": "APPLY" }
```

```
{ "scope": "ACCOUNT", "targetAccounts": ["123456789012", "234567890123"], "mode": "APPLY" }
```

1.2.6. EXAMPLE OF COMBINED INPUTS

```
{  
  "scope": "FULL",  
  "mode": "APPLY",  
  "configOverrides": { "ov-ou-vpc": true, "ov-ou-subnet": true, "ov-acct-vpc": true }  
}
```

2.2.4 1. Multi-file Accelerator Config file and YAML Support Details

1.1. Customers would like the ability to specify their configuration in YAML. This facilitates

- commenting out entire sections, which is unavailable in standard JSON
- annotating aspects of configuration (e.g. cidr: "10.100.0.0/16" # We chose this for \reason.)
- aligning the Accelerator with CloudFormation, which supports JSON/YAML as input format

1.2. Customers would like the configuration file split into multiple files

- one file for Global options + Mandatory accounts
- one file per OU
- one file for every approx. 2000 lines of workload accounts (Code Commit diff stops working at 3000 lines, allow for adding to each file)

1.3. Benefits

1. Easier cut/paste/comparison of OU configurations
2. Allow CodeCommit diff functionality to function (File currently too large)
3. Allow easier updates to workload accounts (simple append)
4. Smaller scoped updates (de-risk accidentally changing the wrong section)
5. Both a customer request and something the team thought was a good idea

1.4. Steps FOR YAML

- The loadAcceleratorConfig functionality should no longer assume config.json as the config filename in the config repo and/or S3, instead it should look for config.yaml and config.json
- Check for the existence of config.yaml and config.json (initially in S3, but also in CodeCommit on future executions)
- If both files exist, fail with an error message
- **Infer the file type from the extension, and parse accordingly**
- Any other failure should also be an error, fail with an error message
- The accelerator will continue to use JSON formatting internally, if a yaml file is supplied, we are simply converting it to JSON for use by the Accelerator
- All examples throughout this document use config.json as the example, but also apply to YAML
- Both JSON and YAML input files will be equally supported
- Only one file format is supported across all config files, either JSON or YAML, customers can NOT mix YAML and JSON file formats

1.5. Steps For File Split

- When the `__LOAD` keyword is encountered, search relatively (from the same location as root config file) for the file, and insert into the config tree, recursively following `__LOAD` if necessary (to max depth of 2). Any file referenced in `__LOAD` must parse successfully in one of the two formats, otherwise FAIL.

```
{
  "core": {
    "__LOAD": "ous/core.json"
  }
}
```

Note that while we will provide sensible examples, there is no prescriptive requirement for file organization within a customer's configuration, customers can use the feature to break-out sections as is most effective for their deployment. Breaking out large repeatable sections like security groups is a good example and could be included off the main file, an account file, or off an ou file:

```
"security-groups": [ "__LOAD": "global/security-groups.json" ]
```

Examples:

1. All in one (single file like today):

```
.
├── config.json
```

1. Split along major sections:

```
.
├── config.json
├── ou
│   ├── core.json
│   ├── dev.json
│   └── test.json ---> could be one per ou, could only be for some ou's as determined by customer
├── accounts
│   ├── workload-accounts-group1.json
│   ├── workload-accounts-group2.json
│   ├── my-workload-accounts.json
│   └── more-accounts.json ---->we will encourage each file being as close to 2000 lines as possible (not one per account, not all in one file)
├── global
│   ├── global-options.json
│   └── security-groups.json
etc
```

- Max depth of 2 means config.json can load ou/dev.json, which can load global/security-groups.json.
- security-groups.json CANNOT load another sub-file (unless security-groups.json was only directly loaded from config.json).

1.6. Dealing with Accelerator Automatic Config File Updates

When customers create AWS accounts directly through AWS Organizations, the Accelerator automatically updates the config file, adding these new accounts. If a customer renames an OU we automatically update the config file. With multi-part files, how do we know what source file to update? We require two mechanisms:

1. Add the following new parameters to the global-options section of the config file

```
"workloadaccounts-param-filename": "accounts/more-accounts2.json",
"workloadaccounts-prefix" : "accounts/more-accounts",
"workloadaccounts-suffix" : 3,
```

- filename is set to config.json, and prefix to config in a single file configuration scenario (suffix is not used)
- While OU contents can be moved into __LOADED sub-files, it was decided the OU object itself must remain in the main config file
- The above parameters:
 - are required to be in the main config file and cannot be __LOAD 'ed
 - Must be present or SM fails
 - Are used to decide where to add new accounts to the config file

2. Add the following new parameter to each mandatory and workload account config

```
"src-filename": "accounts/my-workload-accounts.json",
```

1.7. Accelerator Internal Operations

- when updating an account in the config file, we use the "src-filename" parameters to find and update an accounts ou, ou-path, account-name, and email parameters

- When creating new accounts (inserting into config file):
- if the update is not going to make the file larger than 2000 lines, insert the new account into the config file `"workloadaccounts-param-filename"`
- if the insert will push the file over 2000 lines:
- create the next unused filename for the given prefix in Code Commit (`{"workloadaccounts-prefix"}{"workloadaccounts-suffix"}.{customer file format}`), i.e. `"accounts/more-accounts3.json"`
- insert the new account into the new file in it's entirety
- update `"workloadaccounts-param-filename"` to: `{"workloadaccounts-prefix"}{"workloadaccounts-suffix"}.{customer file format}`
- add a new load stmt to the workload-accounts section of the config file with the name `{"workloadaccounts-prefix"}{"workloadaccounts-suffix"}.{customer file format}`
- update `"workloadaccounts-suffix"` to: `{"workloadaccounts-suffix"} + 1`
- be careful with comma's between files (JSON sections) when appending/connecting

1.8. Example

The entire main config file could be reduced to this:

```
{
  "global-options": {
    "workloadaccounts-param-filename": "accounts/more-accounts2.json",
    "workloadaccounts-prefix": "accounts/more-accounts",
    "workloadaccounts-suffix": 3,
    "__LOAD": "global/global-options.json"
  },
  "mandatory-account-configs": {
    "__LOAD": "accounts/mandatory-accounts.json"
  },
  "workload-account-configs": {
    "__LOAD": [
      "accounts/workload-accounts1.json",
      "accounts/my-other-accounts.json",
      "accounts/workload-accounts2.json"
    ]
  },
  "organizational-units": {
    "core": {
      "__LOAD": "ous/core.json"
    },
    "Central": {
      "__LOAD": "ous/central.json"
    },
    "Dev": {
      "__LOAD": "ous/dev.json"
    },
    "Test": {
      "__LOAD": "ous/test.json"
    },
    "Prod": {
      "__LOAD": "ous/prod.json"
    },
    "UnClass": {
      "__LOAD": "ous/unclass.json"
    },
    "Sandbox": {
      "__LOAD": "ous/sandbox.json"
    }
  }
}
```

1.9. Acceptance Criteria

- A new customer may start an Accelerator deployment with a config.json or config.yaml, and have it deploy as expected so long as the file is semantically correct according to structure and expected keys (and of course syntactically correct in either YAML or JSON)
- Accelerator should continue to function as it does today o i.e. on startup creates repo and copies all referenced config files, not just config.json to repo (json or YAML) o leverages config files in CodeCommit repo from this point forward (json or YAML as provided by customer) o SM runs against the commit id of each file at the start of the SM (i.e. don't allow changes to any file during execution)
- Accelerator leverages multiple config files to receive the same input parameters it previously did from one file
- All accelerator functionality both ALZ and Standalone versions continue to function as previously defined

- Customer can successfully provides multiple config files with the same result as the current one file

2.2.5 1. Existing Organizations / Accounts

1.1. Considerations: Importing existing AWS Accounts / Deploying Into Existing AWS Organizations

- The Accelerator *can* be installed into existing AWS Organizations
- our early adopters have all successfully deployed into existing organizations
- Existing AWS accounts *can* also be imported into an Accelerator managed Organization
- Caveats:
 - Per AWS Best Practices, the Accelerator deletes the default VPC's in all AWS accounts, worldwide. The inability to delete default VPC's in pre-existing accounts will fail the installation/account import process. Ensure default VPC's can or are deleted before importing existing accounts. On failure, either rectify the situation, or remove the account from Accelerator management and rerun the state machine
 - The Accelerator will NOT alter existing (legacy) constructs (e.g. VPC's, EBS volumes, etc.). For imported and pre-existing accounts, objects the Accelerator prevents from being created using preventative guardrails will continue to exist and not conform to the prescriptive security guidance
 - Existing workloads should be migrated to Accelerator managed VPC's and legacy VPC's deleted to gain the full governance benefits of the Accelerator (centralized flow logging, centralized ingress/egress, no IGW's, Session Manager access, existing non-encrypted EBS volumes, etc.)
 - Existing AWS services will be reconfigured as defined in the Accelerator configuration file (overwriting existing settings)
 - We do NOT support *any* workloads running or users operating in the Organization Management (root) AWS account. The Organization Management (root) AWS account MUST be tightly controlled
 - Importing existing *workload* accounts is fully supported, we do NOT support, recommend and strongly discourage importing mandatory accounts, unless they were clean/empty accounts. Mandatory accounts are critical to ensuring governance across the entire solution
 - We've tried to ensure all customer deployments are smooth. Given the breadth and depth of the AWS service offerings and the flexibility in the available deployment options, there may be scenarios that cause deployments into existing Organizations to initially fail. In these situations, simply rectify the conflict and re-run the state machine.
 - If the Firewall Manager administrative account is already set for your organization, it needs to be unset before starting a deployment.

1.2. Process to import existing AWS accounts into an Accelerator managed Organization

- Newly invited AWS accounts in an Organization will land in the root ou
- Unlike newly created AWS accounts which immediately have a Deny-All SCP applied, imported accounts are not locked down as we do not want to break existing workloads (these account are already running without Accelerator guardrails)
- In AWS Organizations, select ALL the newly invited AWS accounts and move them all (preferably at once) to the correct destination OU (assuming the same OU for all accounts)
- In case you need to move accounts to multiple OU's we have added a 2 minute delay before triggering the State Machine
- Any accounts moved after the 2 minute window will NOT be properly ingested, and will need to be ingested on a subsequent State Machine Execution
- This will first trigger an automated update to the config file and then trigger the state machine after a 2 minute delay, automatically importing the moved accounts into the Accelerator per the destination OU configuration
- As previously documented, accounts CANNOT be moved between OU's to maintain compliance, so select the proper top-level OU with care
- If you need to customize each of the accounts configurations, you can manually update the configuration file either before or after you move the account to the correct ou
- if before, you also need to include the standard 4 account config file parameters, if after, you can simply add your new custom parameters to the account entry the Accelerator creates
- if you add your imported accounts to the config file, moving the first account to the correct ou will trigger the state machine after a 2 minutes delay. If you don't move all accounts to their correct ou's within 2 minutes, your state machine will fail. Simply finish moving all accounts to their correct OU's and then rerun the state machine.

- If additional accounts are moved into OUs while the state machine is executing, they will not trigger another state machine execution, those accounts will only be ingested on the next execution of the state machine
- customers can either manually initiate the state machine once the current execution completes, or, the currently running state machine can be stopped and restarted to capture all changes at once
- Are you unsure if an account had its guardrails applied? The message sent to the state machine Status SNS topic (and corresponding email address) on a successful state machine execution provides a list of all successfully processed accounts.
- The state machine is both highly parallel and highly resilient, stopping the state machine should not have any negative impact. Importing 1 or 10 accounts generally takes about the same amount of time for the Accelerator to process, so it may be worth stopping the current execution and rerunning to capture all changes in a single execution.
- We have added a 2 min delay before triggering the state machine, allowing customers to make multiple changes within a short timeframe and have them all captured automatically in the same state machine execution.

1.3. Deploying the Accelerator into an existing Organization

- As stated above, if the ALZ was previously deployed into the Organization, please work with your AWS account team to find the best mechanism to uninstall the ALZ solution
- Ensure all existing sub-accounts have the role name defined in `organization-admin-role` installed and set to trust the Organization Management (root) AWS Organization account
- prior to v1.2.5, this role must be named: `AWSCloudFormationStackSetExecutionRole`
- if using the default role (`AWSCloudFormationStackSetExecutionRole`) we have provided a CloudFormation stack which can be executed in each sub-account to simplify this process
- As stated above, we recommend starting with new AWS accounts for the mandatory functions (shared-network, perimeter, security, log-archive accounts).
- To better ensure a clean initial deployment, we also recommend the installation be completed while ignoring most of your existing AWS sub-accounts, importing them post installation:
- create a new OU (i.e. `Imported-Accounts`), placing most of the existing accounts into this OU temporarily, and adding this OU name to the `global-
options\ignored-ous` config parameter;
- any remaining accounts must be in the correct ou, per the Accelerator config file;
- install the Accelerator;
- import the skipped accounts into the Accelerator using the above import process, paying attention to the below notes
- NOTES:
- Do NOT move any accounts from any `ignored-ous` to the root ou, they will immediately be quarantined with a Deny-All SCP, they need to be moved directly to their destination ou
- As stated above, when importing accounts, there may be situations we are not able to fully handle
- If doing a mass import, we suggest you take a quick look and if the solution is not immediately obvious, move the account which caused the failure back to `ignored-ous` and continue importing the remainder of your accounts. Once you have the majority imported, you can circle back and import outstanding problem accounts with the ability to focus on each individual issue
- The challenge could be as simple as someone has instances running in a default VPC, which may require some cleanup effort before we can import (coming soon, you will be able to exclude single account/region combinations from default VPC deletion to gain the benefits of the rest of the guardrails while you migrate workloads out of the default VPC)

2.2.6 1. How to migrate an AWS Landing Zone (ALZ) account "as is" into an AWS Secure Environment Accelerator (ASEA)

1.1. Overview

This document describes the steps to migrate an existing linked account from an AWS Landing Zone (ALZ) to an AWS Secure Environment Accelerator (ASEA).

1.2. Prerequisites / Setup

1.2.1. CONFIRM ASEA SSO AND OU CONFIGURATION

On the ASEA, setup and run initial tests with SSO and permission sets with an account under the OU where the linked account will be migrated to. Confirm that SSO is properly configured with permissions required for the team members whose account is being migrated. This would include configuration of the ASEA's AWS Managed Active Directory (MAD) which should align with how the team migrating their account has their AWS SSO and MAD configured today.

1.2.2. SWITCH THE ALZ LINKED ACCOUNT PAYMENT METHOD TO INVOICING

If working with your AWS account team (TAM/SA) they will reach out to an internal team within AWS to have the linked account payment method switched to invoicing. This way the customer doesn't have to enter a credit card when making the account standalone in the upcoming steps.

1.2.3. CONFIRM CONSOLE ACCESS TO THE ALZ LINKED ACCOUNT AND ALSO TO THE EMAIL ACCOUNT

Confirm you have access to login as root to the ALZ linked account AWS console. Confirm you have access to the email account associated to the ALZ linked account. The upcoming steps will first make the account standalone (remove from ALZ organizations) so you need to make sure you have root access to the account. If required, you can reset the password following: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_passwords_change-root.html

1.2.4. IF AN ENTERPRISE SUPPORT (ES) CUSTOMER, THEN CONFIRM ES IS ENABLED ON THE ALZ LINKED ACCOUNT

If the ALZ management account is on Enterprise Support (ES), then make sure ES is enabled on the linked account being migrated to the ASEA. If its not, then raise a support case to activate ES on the linked account. This is to make sure an ES support case can be created and escalated during step 2 if any unforeseen issue occurs.

1.2.5. CONFIRM THE ALZ CODEPIPELINE IS EXECUTING SUCCESSFULLY

Make sure the ALZ CodePipeline is still running successfully. Execute the ALZ CodePipeline from the management account to make sure it runs successfully.

- AWS Console -> CodePipeline
- Select "AWS-Landing-Zone-CodePipeline"
- Select "Release Change"
- Click on the pipeline and confirm it successfully runs through to completion

1.2.6. CONFIRM CLI ACCESS AND SETUP PYTHON AND THE AWS PYTHON SDK (BOTO3)

Confirm SSO temporary command line access from the management account with AdminAccess.

- SSO login → Select linked account → "Command line or programmatic access"
- Select Option 2 and add to your AWS credentials file under "[default]"
- This is required as the python script in step 3 takes a "profile" parameter
- Confirm you have the AWS CLI tool installed.
- <https://aws.amazon.com/cli/>
- Confirm by running a command such as "aws s3 ls"
- Confirm you have python3 and the AWS python library (boto3) installed which is required in step 2 to confirm the account has been disassociated from the landing zone correctly.
- <https://boto3.amazonaws.com/v1/documentation/api/latest/guide/quickstart.html>

1.3. Landing Zone - Disassociate the account from the ALZ

- Login to the ALZ management account, and go to "Service Catalog" -> "Provisioned products"
- Select "Access Filter" -> "Account" to see a list of the account products

1.3.1. SELECT THE PRODUCT FOR THE SPECIFIC LINKED ACCOUNT

- Put the linked account name in the provisioned products search bar
- This will narrow down the list and show a product name "AWS-Landing-Zone-Account-Vending-Machine" with a name "*lzapplicaitons**"
- Select that product and then "Actions->Terminate"

1.3.2. CONFIRM THE PRODUCT SUCCESSFULLY TERMINATES

- The provisioned product entry will show a status of "Under change"
- You can also verify by going to CloudFormation → Stacks and you will see "DELETE IN PROGRESS" for the AVM Template stack being deleted.
- Go to the Resources tab to see the deleted resources associated to this stack.
- Once the provisioned product no longer says "Under change" move to the next step.
- Please note, this can take 1-2 hours.

1.3.3. GO TO THE LINKED ACCOUNT (ASSUME ROLE)

- From the management account, assume the role "AWSCloudFormationStackSetExecutionRole" to the linked account
- or optionally, SSO with console access to that account

1.3.4. UNDER "CLOUDFORMATION" VERIFY THAT THE ALZ STACKS (STACKSETS FROM ALZ MGMT) WERE DELETED

- There should be no stack left in the linked account with the prefix "StackSet-AWS-Landing-Zone-Baseline*". For example:
- StackSet-AWS-Landing-Zone-Baseline-CentralizedLoggingSpoke-
- StackSet-AWS-Landing-Zone-Baseline-EnableConfigRules-
- StackSet-AWS-Landing-Zone-Baseline-EnableNotifications-
- StackSet-AWS-Landing-Zone-Baseline-EnableConfigRulesGlobal-
- StackSet-AWS-Landing-Zone-Baseline-EnableConfig-
- StackSet-AWS-Landing-Zone-Baseline-ConfigRole-
- StackSet-AWS-Landing-Zone-Baseline-IamPasswordPolicy-
- StackSet-AWS-Landing-Zone-Baseline-SecurityRoles-
- StackSet-AWS-Landing-Zone-Baseline-EnableCloudTrail-

1.3.5. VERIFY THAT THE ACCOUNT IS READY TO BE INVITED AND BASELINED BY THE ASEA

- You need to ensure that resources don't exist in the default VPC, there is no config recorder channel, no CloudTrail Trail and STS is active in all regions.
- This can be done manually, but ideally use this python script that can be run as well to automate the verification
- https://github.com/paulbayer/Inventory_Scripts/blob/mainline/ALZ_CheckAccount.py
- mkdir test; cd test
- git clone https://github.com/paulbayer/Inventory_Scripts.git
- python3 ALZ_CheckAccount.py -a LINKED_ACCOUNT_HERE -p default

- It will run through 5 steps and output the following. If you were to run this script before the "terminate" step above is complete you would have warnings in steps 2 and 3 below.
- Step 0 completed without issues
- Checking account 111122223333 for default VPCs in any region
- Step 1 completed with no issues
- Checking account 111122223333 for a Config Recorders and Delivery Channels in any region
- Step 2 completed with no issues
- Checking account 111122223333 for a specially named CloudTrail in all regions
- Step 3 completed with no issues
- Checking account 111122223333 for any GuardDuty invites
- Step 4 completed with no issues
- Checking that the account is part of the AWS Organization.
- Step 5 completed with no issues
- We've found NO issues that would hinder the adoption of this account ****

1.4. Landing Zone (ALZ) - Remove the account from the ALZ organizations and make standalone

Removing the account from the ALZ organizations and making it standalone is required so it can be invited into the ASEA organization.

1.4.1. READ THE FOLLOWING SUMMARY/CONSIDERATIONS

- <https://aws.amazon.com/premiumsupport/knowledge-center/organizations-move-accounts/>

1.4.2. VERIFY ACCESS

- As stated in the previous sections, verify you have a mechanism to access the account post leaving the ALZ organization
- Former SSO roles will no longer function nor will the "AWSCloudFormationStackSetExecutionRole" role as it will have a trust relationship to the ALZ management account.
- Confirm the root credentials have been recovered and are usable
- As an alternative, confirm access with a new role/IAM user with Admin permissions on the account

1.4.3. VERIFY BILLING FLIPPED TO INVOICING

- As stated in the previous sections, verify the account payment method has been flipped to "invoicing" to avoid having to enter a Credit Card when going standalone. This can be done working with your AWS account team who will coordinate internally, or by raising a support case describing the use case.

1.4.4. REMOVE THE ACCOUNT FROM THE ORGANIZATIONS AND MAKE STANDALONE

- Follow the instructions on the following link to remove the account
- The short version is select the account from the ALZ mgmt account Organizations and select "remove"
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_remove.html
- <https://aws.amazon.com/blogs/security/aws-organizations-now-supports-self-service-removal-of-accounts-from-an-organization>
- Note, when moving the account standalone do not select Enterprise Support. You shouldn't get a popup dialog asking for a Credit Card and the Support level since the account should have been moved to invoicing. Support can be reenabled on the linked account once it's invited into the ASEA organization.

1.5. Accelerator - Invite the account into its organization

1.5.1. FROM THE ASEA MGMT ACCOUNT, SEND AN INVITE TO THE STANDALONE ACCOUNT

- Follow the instructions on the following link to invite the account

- The short version is go to the ASEA mgmt account organizations and select "Add an account" -> "Invite existing account" -> "enter the linked account account ID"

• https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invites.html

1.5.2. IN THE FORMER ALZ ACCOUNT, ACCEPT THE INVITATION

• https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invites.html#orgs_manage_accounts_accept-decline-invite

1.5.3. KEEP THE LINKED ACCOUNT AT THE ROOT LEVEL OF THE ORGANIZATIONS

- Verify access to the linked account using your root login credentials
- If you had created an IAM role/user with Admin permissions, then verify access as well

1.5.4. ACTIVATE ENTERPRISE SUPPORT (ES) ON THIS LINKED ACCOUNT

- If ES is enabled on the ASEA management account, open a support case to enable ES on this linked account
- Go to the Support center and create a billing support case with "Account" and "Activation"
- Subject "Requesting ES enablement on linked account"
- Body "Requesting ES enablement on linked account "
- Your AWS TAM can escalate the case with the support team if it's time sensitive.
- This is to make sure an ES support case can be created and escalated during the next steps if any unforeseen issue occurs.

1.5.5. UPDATE (OR ADD) THE ORGANIZATION ADMINING ROLE SO ONE CAN ASSUME THE ROLE INTO THE LINKED ACCOUNT

- Login to the linked account which just joined the organization.
- Create a new Organization Admin role, as defined in the customers config file: "organization-admin-role": "OrganizationAccountAccessRole".
- With newer customers the default is "OrganizationAccountAccessRole, with older customers it is "AWSCloudFormationStackSetExecutionRole".
- If "AWSCloudFormationStackSetExecutionRole" then you can edit the trust relationship directly
- Go to IAM -> Role -> AWSCloudFormationStackSetExecutionRole
- Update the trust relationship to have the management account ID of the ASEA (instead of the account ID of the previous ALZ)
- Verify that you can assume this role from the management account into the linked account

1.6. Accelerator - Move the linked account from the top level root OU into the appropriate OU managed by the ASEA

1.6.1. PLAN WHAT OU THIS ACCOUNT WILL BE MOVED INTO

- Option 1 - Create a new OU and move the account into that OU
- Before the migration, the team would have created a new OU (ie-similar to the sandbox OU).
- This would be needed if they need to isolate this account from TGW attachments/Networking and want to keep it isolated.
- The state machine will run and start to baseline the account.
- It will create a new VPC and deploy resources using CFN such as Config, CloudTrail, etc.
- Note, if the OU is setup similar to the sandbox OU it does not provide access to the shared VPCs that have the TGW attachments.
- Creating a new OU also requires adding that new OU and the OU persona to the config file in advance of the next state machine execution.
- Option 2 - Move account into an existing OU (ie-prod)
- The state machine will run and start to baseline the account.
- It will create a new VPC and deploy resources using CFN such as Config, CloudTrail, etc.
- The customers existing VPC will remain, as a 2nd DETACHED VPC.
- Mote. if it is non-compliant to security rules, it remains non-compliant and needs to be cleaned up and brought into compliance
- If the VPC is compliant and it has unique IP addresses, it could be attached to the TGW.

1.6.2. MOVE THE ACCOUNT FROM THE ROOT OU TO THE CORRECT OU

- THIS CANNOT BE EASILY UNDONE - MAKE SURE YOU MOVE TO THE CORRECT OU

- Follow the instructions on the following link to move the account to the correct OU
- The short version is go to the ASEA management account organizations and "select the account" -> "actions" -> "move" -> "select the correct OU"
- NOTE: The ASEA state machine will automatically start within 1-2 minutes of the account being moved into the OU
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_ous.html#move_account_to_ou
- Verify that the ASEA main state machine (under AWS->Step Functions) is triggered and runs cleanly (~30-45 minutes)

1.7. Accelerator (ASEA) - Verify access control with roles, SSO, etc

- Update and verify SSO and permission sets for the linked account now part of the ASEA
- Verify you still have access to the linked account via root (or other mechanisms)
- Verify you still can assume the operations role into the linked account

1.8. Landing Zone - Close down the ALZ core accounts and then the management account

Once all workloads have been migrated from the ALZ to the ASEA, then you may decide to shutdown your ALZ.

1.8.1. CLOSE DOWN THE ALZ LINKED ACCOUNTS

- Close all the linked accounts "as is" without making them standalone
- This will be the ALZ core linked accounts, but you might have some remaining workload accounts you decided not to migrate to the ASEA.
- <https://aws.amazon.com/premiumsupport/knowledge-center/close-aws-account>
- The management account will remain with organizations and the core accounts will show as suspended for 90 days.

1.8.2. CLOSE DOWN THE ALZ MANAGEMENT ACCOUNT

- After 90 days, the suspended linked accounts will be completely closed
- Go to the root account and turn off Organizations and then close the root account

2.3 Upgrades

2.3.1 1. Accelerator Upgrade Guide

1.1. General Upgrade Considerations

- Due to some breaking dependency issues, customers can only upgrade to v1.3.8 or above (older releases continue to function, but cannot be installed).
- While an upgrade path is planned, customers with a standalone Accelerator installation can upgrade to v1.5.x but need to continue with a standalone installation until the Control Tower upgrade option becomes available.
- Always compare your configuration file with the config file from the release you are upgrading to in order to validate new or changed parameters or changes in parameter types / formats.
- do NOT update to the latest firewall AMI - see the last bullet in section [1.8. Other Operational Considerations](#) of the installation guide
- do NOT update the `organization-admin-role` - see item 2 in section [1.3.7. Other](#)
- do NOT update account-keys (i.e. existing installations cannot change the internal values to `management` from `master`)
- do NOT make changes outside those required for the upgrade (those stated in the release notes or found through the comparison with the sample config file(s)). Customers wishing to change existing Accelerator configuration should either do so before their upgrade, ensuring a clean/ successful state machine execution, or after a successful upgrade.
- The Accelerator name and prefix **CANNOT** be changed after the initial installation
- Customers which customized any of the Accelerator provided default configuration files (SCPs, rsyslog config, ssm-documents, iam-policies, etc.) must manually merge the latest Accelerator provided updates with deployed customizations:
- it is important customers assess the new defaults and integrate them into their custom configuration, or Accelerator functionality could break or Accelerator deployed features may be unprotected from modification
- if customers don't take action, we continue to utilize the deployed customized files (without the latest updates)
- The below release specific considerations need to be cumulatively applied (an upgrade from v1.2.3 to v1.2.5 requires you to follow both v1.2.4 and v1.2.5 considerations)

1.2. Release Specific Upgrade Considerations:

- Upgrades to `v1.5.1-a` and above from `v1.5.0` or `v1.5.1`:
- Do not add the parameter: `"ssm-inventory-collection": true` to OUs or accounts which already have SSM Inventory configured or the state machine will fail
- Follow the standard upgrade steps detailed in section 1.3 below
- `v1.5.1` was replaced by `v1.5.1-a` and is no longer supported for new installs or upgrades
- Upgrades to `v1.5.0` and `v1.5.1-a` and above from `v1.3.8` through `v1.3.9`:
- We recommend upgrading directly to `v1.5.1-a`
- Due to the size and complexity of this upgrade, we require all customers to upgrade to `v1.3.8` or above before beginning this upgrade
- While `v1.5.0` supports Control Tower for *NEW* installs, existing Accelerator customers **CANNOT** add Control Tower to their existing installations at this time (planned enhancement for 22H1)
- Attempts to install Control Tower on top of the Accelerator will corrupt your environment (both Control Tower and the Accelerator need minor enhancements to enable)
- The **v1.5.x custom upgrade guide can be found [here](#)**

- Upgrades to v1.3.9 and above from v1.3.8-b and below :
- All interface endpoints containing a period must be removed from the config.json file either before or during the upgrade process
- i.e. ecr.dkr, ecr.api, transfer.server, sagemaker.api, sagemaker.runtime in the full config.json example
- If you remove them on a pre-upgrade State Machine execution, you can put them back during the upgrade, if you remove them during the upgrade, you can put them back post upgrade.
- Upgrades to v1.3.3 and above from v1.3.2 and below :
- Requires mandatory config file schema changes as documented in the [release notes](#).
- These updates cause the config file change validation to fail and require running the state machine with the following input to override the validation checks on impacted fields: `{"scope": "FULL", "mode": "APPLY", "configOverrides": {"ov-ou-vpc": true, "ov-ou-subnet": true, "ov-acct-vpc": true }}`
- Tightens VPC interface endpoint security group permissions and enables customization. If you use VPC interface endpoints that requires ports/ protocols other than TCP/443 (such as email-smtp), you must customize your config file as described [here](#)
- Upgrades from v1.3.0 and below :
- Please review the [Release Specific Upgrade Considerations](#) from ASEA v1.5.0 or below, they were removed from this release.

1.3. Summary of Upgrade Steps (all versions except [v1.5.0](#))

1. Login to your Organization Management (root) AWS account with administrative privileges
2. Either: a) Ensure a valid Github token is stored in secrets manager ([per the installation guide](#)), or b) Ensure the latest release is in a valid branch of CodeCommit in the Organization Management account
3. Review and implement any relevant tasks noted in the General Upgrade Considerations [section](#) above
4. Update the config file in CodeCommit with new parameters and updated parameter types based on the version you are upgrading to (this is important as features are iterating rapidly)
 - An automated script is available to help convert config files to the new v1.5.0 format
 - Compare your running config file with the sample config file from the latest release
 - Review the [Config file changes](#) section of the [release notes](#) for **all** Accelerator versions since your current deployed release
5. If you customized any of the other Accelerator default config files by overriding them in your S3 input bucket, merge the latest defaults with your customizations before beginning your upgrade
6. Download the latest installer template (`AcceleratorInstallerXYZ.template.json` or `AcceleratorInstallerXXX-CodeCommit.template.json`) from the [Assets](#) section of the latest [release](#)
7. Do **NOT** accidentally select the `ASEA-InitialSetup` CloudFormation stack **below**
8. If you are replacing your GitHub Token:
 - Take note of the `AcceleratorName`, `AcceleratorPrefix`, `ConfigS3Bucket` and `NotificationEmail` values from the Parameters tab of your deployed Installer CloudFormation stack (`ASEA-what-you-provided`)
 - Delete the Installer CloudFormation stack (`ASEA-what-you-provided`)
 - Redeploy the Installer CloudFormation stack using the template downloaded in step 6, providing the values you just documented (changes to `AcceleratorName` or `AcceleratorPrefix` are not supported)
 - The pipeline will automatically run and trigger the upgraded state machine

9. If you are using a pre-existing GitHub token, or installing from CodeCommit:

- Update the Installer CloudFormation stack using the template downloaded in step 5, updating the `GitHubBranch` to the latest release (eg. `release/v1.5.1-a`)
- Go to AWS CloudFormation and select the stack: `ASEA-what-you-provided`
- Select Update, select Replace current template, Select Upload a template file
- Select Choose File and select the template you downloaded in step 6 (`AcceleratorInstallerXYZ.template.json` or `AcceleratorInstallerXXX-CodeCommit.template.json`)
- Select Next, Update `GitHubBranch` parameter to `release/vX.Y.Z` where X.Y.Z represents the latest release
- Click Next, Next, I acknowledge, Update
- Wait for the CloudFormation stack to update (`Update_Complete` status) (Requires manual refresh)
- Go To Code Pipeline and Release the ASEA-InstallerPipeline

2.3.2 1. Accelerator v1.5.x Custom Upgrade Instructions

1.1. Overview

The upgrade from v1.3.8/v1.3.9 to v1.5.x is generally the same as any previous Accelerator upgrades, with a couple of key differences:

- the magnitude of this release has resulted in a requirement for significant updates to the config file
- we have provided a script to assist with this process. A manual verification of the changes and customer custom updates are often still required.
- we are re-aligning the OU structure with AWS guidance and that of AWS Control Tower (optional, but highly recommended)
- the core OU is being split into a "Security" OU and an "Infrastructure" OU
- we've added the capability to manage your IP addresses in DynamoDB, rather than with the config file
- this includes the ability to dynamically allocate CIDR ranges to VPCs and subnets
- more information on this features design can be found on this [ticket](#)
- the config file conversion script will:
 - update your config file in a manner that supports both CIDR management schemes (but continues to leverage the previous mechanism)
 - copy your currently configured CIDR ranges into the appropriate DynamoDB tables (optional, but recommended)
 - you can change your IP address mechanism for any VPC at any time
 - customers can mix and match IP address management mechanisms as they choose (`provided` , `lookup` , and `dynamic`)

1.2. Upgrade Caveats

1. **While an upgrade path is planned, customers with a Standalone Accelerator installation can upgrade to v1.5.x but need to continue with a Standalone installation until the Control Tower upgrade option becomes available.**
2. The script to assist with config file conversion and DynamoDB population only supports single file json based config files, customers that leverage YAML and/or multi-part config files, have several options:
 - manually update your yaml or multi-part json config file to reflect the config file format for the latest release (similar to all previous upgrades)
 - use the config.json file found in the `raw` folder of your CodeCommit repo to run the conversion script
 - this version of the config file has resolved all variables with their final values, all variables will be removed from config.json in this scenario
 - the new config file can be converted back to json/multi-part format before being placed back into your CodeCommit repository
 - or it could be used to simply validate the changes you made using option a
 - do not manually update the config file in the `raw` folder, as it will be overwritten based on the json or yaml file in the root of your repository
 - use a 3rd party tool to manually convert your yaml / multi-part config files to a single file json file to run the conversion script
 - the new config file can be converted back to json/multi-part format before being placed back into your CodeCommit repository
3. Config files which are significantly different than the example config files may not be properly converted. This includes config files which use different mandatory account keys or renamed the core OU.
4. This guide and its examples assume the existing accelerator deployment uses the `PBMMAccel-` accelerator prefix, if a different prefix is used on the existing installation, it is important it is specified when execution section 1.6 below.

1.3. Config File Conversion

- You must first upgrade to Accelerator v1.3.8 or v1.3.9
- Login to your AWS Organization Management account
- Pull your current config.json file from CodeCommit and save as a text file
- Locate the python conversion script and review its readme [here](#)
- To convert your configuration file execute: (completely offline process)


```
python update.py --Region ca-central-1 --LoadConfig --ConfigFile config.json
```

- This will output a new config file named: `update-config.json`
- Save *both* the original v13.8 and the new v1.5.0 config files for future reference/use
- After conversion, we recommend running the updated config file back prettier to simplify file comparisons
- While the conversion script often does much of the heavy lifting, we still require customers to manually verify the changes and make manual adjustments as appropriate:
- If you use a relatively standard config file you MAY not need to make any changes manually
- Ensure the value of `account-name` for the Organization Management account matches the actual account name of the Organization management account (the account key is generally either `management` or `master`).
- we recommend you change your `rdgw-instance-type` and `rsyslog-instance-type` from `t2` to `t3`. (they will auto-replace on the next instance refresh) (Optional).
- optionally remove the `"API_GW_EXECUTION_LOGGING_ENABLED"` config rule throughout, as it overlaps with an identical Security Hub config rule.
- we added the capability to deploy a Config aggregator in any of the central services accounts (i.e. Log-archive, Security, Operations), by adding `"config-aggr": true` to *either*: `central-security-services`, `central-operations-services`, or `central-log-services`. The existing aggregator in the Org management account will remain. Do **not** set it in all 3 sections, as AWS only supports a maximum of 3 config aggregators.
- the optional attribute `endpoint-port-overrides` has been properly renamed to `endpoint-port-overrides`. If you have the `endpoint-port-overrides` in your config file you must rename it to `endpoint-port-overrides`.
- the new example config files also introduced several new internally resolvable variables (`${CONFIG::OU_NAME}` and `${CONFIG::VPC_NAME}`), which when used thoughtfully along with the new dynamic CIDR feature, enables multi-part config file customers to define the VPCs for multiple OU's in a single shared nested config file. These new variables should be ignored during an upgrade.
- the accelerator supports 3 types of CIDR ranges `provided`, `lookup`, and `dynamic`. The upgrade script sets the `cidr-src` to `provided`, meaning it uses the CIDR ranges provided in the config file, as per the previous release. The upgrade script also adds the additional required fields (`pool` and `size`) to every CIDR range defined in the config file to leverage the `lookup` type, but when set to `provided` these fields are NOT required and could be removed. They were added by the script for the sole purpose of making it easy to switch from `provided` to `lookup` in future. Once a customer switches to `lookup`, the `cidr\value` field is no longer used and can be removed from the config file. The `cidr-src` for should remain set at `provided` during upgrade.
- do **not** add the `cidr-pools` section to the config file during or before the upgrade, this section is only used for new installations.
- New description fields have been added to the config file to help provide context to certain objects. These will be used by a future GUI that is under development, and serve no functional purpose at this time. Customers can alter this text as they please.
- Most of the example config files have been converted to `dynamic` `cidr-src` as it provides simpler CIDR management for new customers. Two example config files ending in `-oldIP.json` have been maintained to aid upgrading customers in config file comparison.
- Be advised - in v1.5.0 we restructured the SCPs based on a) customer requests, and b) the addition of Control Tower support for new installs.
- customers are responsible for reviewing the SCPs to ensure they have not been altered in a manner that no longer meets an organizations security requirements;
- we reorganized and optimized our SCP's from 4 SCP files down to 3 SCP files, without removing any protections or guardrails;
- these optimizations have resulted in minor enhancements to the SCP protections and in some cases better scoping;
- the first two SCP files (Part-0 and Part-1) contain the controls which protect the integrity of the Accelerator itself;
- the third file (Sensitive, Unclass, Sandbox) contains customer data protection specific guardrails, which may change based on workload data classification or customer profiles and requirements;
- this frees the fourth SCP for use by Control Tower, or for use by customers for custom guardrails (Standalone installs only). As Control Tower leverages 2 SCP files on the Security OU, we have moved some of our SCP's to the account level (Control Tower installations only).
- The script and upgrade instructions above do not include the new config file parameters added in v1.5.1+. These new parameters can be added either during or after the upgrade. New parameters include: `"rdgw-enforce-imdsv2": true`, `"rsyslog-enforce-imdsv2": true`, `"ssm-inventory-collection": true` on each ou, and `"dynamic-s3-log-partitioning": [{"values}]`

1.4. Upgrade process

- Before proceeding with your upgrade please review the General and Release Specific Upgrade Considerations in the [Upgrade Guide](#)
- upgrades directly from v1.3.8 need to ensure they include the extra step required for v1.3.9 upgrades (removal of endpoints with periods)
- Login to your AWS Organization Management account, in your home or default region
- Place your *updated and validated* config file back in the root folder of your CodeCommit repository
- json, yaml and multi-part config files continue to be supported
- Change to the us-east-1 region and open CloudWatch
- Select Rules under events in the left navigation pane
- Select the `PBMMAccel-MoveAccount_rule`, select `actions`, select `Disable`
- Select the `PBMMAccel-PolicyChanges_rule`, select `actions`, select `Disable`
- Open AWS Organizations
- Select the `core` OU and rename it to `Security` (case sensitive) (i.e. Actions, Rename)
- In the `Root` ou, create a new OU named `Infrastructure` (case sensitive) (i.e. Actions, Create new)
- Expand the `Security` OU, select all EXCEPT the `security`, `log-archive`, and `Organization Management` account
- (i.e. the `Operations`, `Perimeter`, and `SharedNetwork` accounts)
- Select Actions, Move, then select the newly created `Infrastructure` OU (note these accounts temporarily missing guardrails)
- NOTE: the key to this being a possible move/change, is the new `Infrastructure` OU is defined identically to the old `core` OU. Typically you CANNOT move accounts between OU's without breaking something and violating security guardrails.
- Select Policies from the left navigation pane, then Service Control Policies
- Click the `PBMMAccel-Guardrails-Part-2` policy, and Select Targets
- Detach the policy from ALL OUs
- Change to the us-east-1 region and open CloudWatch
- Select Rules under events in the left navigation pane
- Select the `PBMMAccel-MoveAccount_rule`, select `actions`, select `Enable`
- Select the `PBMMAccel-PolicyChanges_rule`, select `actions`, select `Enable`
- Follow the Standard Upgrade instructions from the section `Summary of Upgrade Steps (all versions)` of the Installation and Upgrade guide, repeated verbatim below for ease of reference

1.5. "Summary of Upgrade Steps (all versions)" (Copied from upgrade guide)

1. Login to your Organization Management (root) AWS account with administrative privileges
2. Either:
 - a) Ensure a valid Github token is stored in secrets manager, or
 - b) Ensure the latest release is in a valid branch of CodeCommit in the Organization Management account. See this [\(section\)](#) of the installation guide for more details.
3. Review and implement any relevant tasks noted in the upgrade consideration sections (sections 1.1 and 1.2) of the [Upgrade Guide](#)
4. Update the config file in CodeCommit with new parameters and updated parameter types based on the version you are upgrading to (this is important as features are iterating rapidly)
 - An automated script is available to help convert config files to the new v1.5.0 format
 - Compare your running config file with the sample config file from the latest release
 - Review the `Config file changes` section of the [release notes](#) for **all** Accelerator versions since your current deployed release
5. If you customized any of the other Accelerator default config files by overriding them in your S3 input bucket, merge the latest defaults with your customizations before beginning your upgrade
6. Download the latest installer template (`AcceleratorInstallerXYZ.template.json` or `AcceleratorInstallerXXX-CodeCommit.template.json`) from the `Assets` section of the latest [release](#)

7. Do **NOT** accidentally select the `PBMMAccel-InitialSetup` CloudFormation stack **below**

8. If you are replacing your GitHub Token:

- Take note of the `AcceleratorName`, `AcceleratorPrefix`, `ConfigS3Bucket` and `NotificationEmail` values from the Parameters tab of your deployed Installer CloudFormation stack (`PBMMAccel-what-you-provided`)
- Delete the Installer CloudFormation stack (`PBMMAccel-what-you-provided`)
- Redeploy the Installer CloudFormation stack using the template downloaded in step 6, providing the values you just documented (changes to `AcceleratorName` or `AcceleratorPrefix` are not supported)
- The pipeline will automatically run and trigger the upgraded state machine

9. If you are using a pre-existing GitHub token, or installing from CodeCommit:

- Update the Installer CloudFormation stack using the template downloaded in step 5, updating the `GithubBranch` to the latest release (eg. `release/v1.5.0`)
- Go to AWS CloudFormation and select the stack: `PBMMAccel-what-you-provided`
- Select Update, select Replace current template, Select Upload a template file
- Select Choose File and select the template you downloaded in step 6 (`AcceleratorInstallerXYZ.template.json` or `AcceleratorInstallerXXX-CodeCommit.template.json`)
- Select Next, Update `GithubBranch` parameter to `release/vX.Y.Z` where X.Y.Z represents the latest release
- Click Next, Next, I acknowledge, Update
- Wait for the CloudFormation stack to update (`Update_Complete` status) (Requires manual refresh)
- Go To Code Pipeline and Release the `PBMMAccel-InstallerPipeline`

1.6. Post Upgrade Follow-up Tasks for v1.5.x Upgrade

- Once the State Machine finishes in AWS Organizations, delete the `PBMMAccel-Guardrails-Part-2` SCP, as we do not remove policies from Organizations given they could be used by customers for other purposes.
- Optionally, populate DynamoDB with your allocated CIDR ranges:
- You *must* perform the following steps if you plan to:
 - move from `"cidr-src": "provided"` to `"cidr-src": "lookup"` for your existing cidr blocks; or
 - dynamically assign CIDR ranges to new VPCs using pools that overlap with existing VPC's currently defined in your config file.
- this step populates the DynamoDB tables (`PBMMAccel-cidr-vpc-assign` and `PBMMAccel-cidr-subnet-assign`) with your EXISTING utilized CIDR ranges using the upgrade script a second time *after* the upgrade is complete.
- Using the converted version of your config file (`update-config.json`)
- Validate you are happy with the `pool` names assigned to each VPC and subnet throughout the config file. Update as appropriate, pool names can be any alpha-numeric string, but a subnets pool must match one of its VPCs pools.
- Locate the python conversion script and review its readme [here](#)
- To load DynamoDB with your CIDR ranges, execute: (online, requires credentials to the Organization Management account)


```
python update.py --Region ca-central-1 --LoadDB --ConfigFile update-config.json --AcceleratorPrefix PBMMAccel-
```
- Note the use of the `--LoadDB` switch, the UPDATED config filename, and that the Accelerator prefix is specified
- If you run this command using your v1.3.8 or v1.3.9 config file, the `\cidr\pool:` values will be improperly assigned in DynamoDB.
- In the Organization Management account, in DynamoDB, select Items, and Verify that the tables have been properly populated.
- the script should have populated the following two DynamoDB tables: `PBMMAccel-cidr-vpc-assign` and `PBMMAccel-cidr-subnet-assign` with all your existing vpc and subnet assignments.
- if you plan to dynamically assign CIDR ranges for any new VPCs, you need to manually create the CIDR pools by adding new item(s) to the DynamoDB Table `PBMMAccel-cidr-pool`. The `PBMMAccel-cidr-pool` table stores CIDR ranges to select from for new CIDR assignments. This table works together with the other two DynamoDB tables to track, assign and maintain non-overlapping CIDR ranges based on a pool name and region.

Example DynamoDB JSON to add an entry to the `PBMMAccel-cidr-pool` table:

```
{
  "id": {
    "S": "1"
  },
  "cidr": {
    "S": "10.0.0.0/13"
  },
  "region": {
    "S": "ca-central-1"
  },
  "pool": {
    "S": "main"
  }
}
```

- where `id` is any unique text, `cidr` is the main cidr block from which VPC cidrs are taken. `region` is the AWS region where the pool is used. `pool` is the name of the pool

NOTES:

- You can populate the `cidr-pools` section of the config file/DynamoDB with values that overlap with the existing assigned ranges in your config file. In this situation, it is CRITICAL that you execute this entire process, to avoid issuing duplicate or overlapping CIDR ranges with those already issued. Alternatively, leverage new unique ranges when populating the `cidr-pools`.
- `cidr-pools` only needs to be populated when a VPC has a `cidr-src` set to `dynamic`.
- Optionally, change all the `cidr-src` values throughout your config file to `lookup`, and remove all the `cidr\value` fields. Once changed, CIDR values will be provided by DynamoDB. Switching to `lookup` requires completion of the previous optional step to first load DynamoDB.
- run the state machine with the input parameters `{"scope": "FULL", "mode": "APPLY", "verbose": "0"}`
- during the state machine execution, the Accelerator will compare the values returned by DynamoDB with the values from the previous successful state machine execution. If the DynamoDB values were incorrectly populated, the state machine will catch it with a comparison failure message and gracefully fail.

2.4 Functionality

2.4.1 Accelerator Service List

Services

This table indicates whether services are leveraged and/or orchestrated by the Accelerator.

CATEGORY	SERVICE	LEVERAGED	ORCHESTRATED
Compute			
	AWS Lambda	X	
	Amazon Elastic Compute Cloud (EC2)		X
Monitoring & Alerts			
	Amazon CloudTrail		X
	AWS Config		X
	Amazon CloudWatch	X	X
	Amazon EventBridge	X	X
	Amazon Simple Notification Service (SNS)	X	
	AWS Budgets		X
	Systems Manager Inventory		X
Infrastructure			
	AWS CodeCommit	X	
	AWS CodeBuild	X	
	AWS CodePipeline	X	
	AWS CloudFormation	X	
	AWS Cloud Development Kit (CDK) / Software Development Kit (SDK)	X	
	AWS Step Functions	X	
	Amazon Kinesis Data Stream	X	
	Amazon Kinesis Data Firehose	X	
	Amazon Simple Queue Service (SQS)	X	
Data			
	Amazon Simple Storage Service (S3)	X	X
	Amazon DynamoDB	X	
	Amazon Elastic Container Registry (ECR) (incl. ECR Public)	X	
	Systems Manager Parameter Store	X	X
	AWS Secrets Manager	X	
Networking			
	Amazon Virtual Private Cloud (VPC)		X
	AWS Transit Gateway		X
	AWS PrivateLink		X
	Elastic Load Balancer (ELB) (incl. ALB, NLB, GWLB)		X
	Route53		X
	Route53 Resolver		X

CATEGORY	SERVICE	LEVERAGED	ORCHESTRATED
Management			
	AWS Organizations	X	X
	AWS Resource Access Manager (RAM)		X
	AWS Identity and Access Management (IAM)	X	X
	AWS Single Sign-On (SSO)	X	
	AWS Directory Service (incl. AWS Managed AD and AD Connector)		X
	AWS Control Tower	X	X
	AWS IAM Access Analyzer		X
	AWS Cost and Usage Reports		X
	AWS Service Quotas		X
Security			
	AWS GuardDuty		X
	AWS Security Hub		X
	Amazon Macie		X
	Systems Manager Automation		X
	Systems Manager Session Manager		X
	AWS Key Management Service (KMS)	X	X
	AWS Security Token Service (STS)	X	
	AWS Firewall Manager		X
	AWS Network Firewall		X
	AWS Certificate Manager (ACM)		X
Third-Party			
	Fortinet FortiGate and FortiManager (Firewall & Mgmt)		X
	Checkpoint CloudGuard and Manager (Firewall & Mgmt)		X
	rsyslog on Amazon Linux 2		X
	Windows Remote Desktop Gateway Bastion		X

If we missed a service, let us know!

2.4.2 1. Accelerator Pricing

1.1. Overview

The AWS Secure Environment Accelerator (ASEA) is available free of charge as an open source solution on GitHub. **You are responsible for the cost of the AWS services enabled, configured, and deployed by the solution.**

The ASEA solution enables, configures and deploys two types of AWS [services](#): services leveraged by the ASEA itself to deliver its capabilities; and services orchestrated by the ASEA to help create a secure multi-account AWS foundation for your users and workloads.

The pricing for services leveraged by the ASEA are relatively consistent and small. The pricing for services orchestrated by the ASEA can vary dramatically based on the underlying architecture, services and features selected by a customer through the customizable configuration file.

Most of the provided example ASEA configuration files (except ultra-lite) build a highly available and scalable multi-datacenter environment with hyperscale routing and enterprise grade security worldwide, something that would cost tens of millions of dollars on-premises and still not achieve the same results.

As shown below, different configuration files can dramatically change the monthly cost of running the solution from \$30/month, to \$1500/month, to \$2400/month, to over \$3700/month. The price of the deployed solution is 100% dependent on what the customer deploys, and not on the Accelerator automation engine itself. While the example deployment(s) may appear expensive when used solely for testing in a personal account, they typically only represent a very small percentage of a production customers AWS spend. The examples were designed to minimize costs as a customer scales.

This document is designed to assist customers in understanding the pricing associated with operating the example ASEA configuration files. For full pricing details, please refer to each services [pricing page](#).

1.2. Example Configuration File Pricing

The pricing found in this document is provided as an example only. Pricing represents reasonably steady state, minimal activity or traffic flows, and only includes sample workload accounts when they exist in the example config files.

Pricing is based on the ca-central-1 region, a month with 31 days (744 hours), on-demand pricing and Bring Your Own Licensing (BYOL) for any 3rd party firewalls. This is estimated pricing, the solution is regularly updated and pricing is dependent on the actual version and configuration used to implement the solution.

Any changes to the example configuration file will impact the pricing. These estimates do not include any customer workloads, workloads must be independently priced.

1.2.1. PRICING BY CONFIGURATION FILE

The following table provides the estimated monthly pricing based on the example configuration. Additional information on each of the example config files can be found [here](#).

Example Configuration	Description	Estimated Monthly Pricing
Ultra-Lite	This configuration file was created to represent an extremely minimalistic Accelerator deployment, to demonstrate the art of the possible for an extremely simple config. This example is NOT recommended as it violates many AWS best practices.	\$30
Test	Designed to reduce solution costs, while demonstrating full solution functionality (Use for testing Full/Lite configurations or Low Security Profiles). Based on Lite Config w/AWS Network Firewall.	\$1,500
Lite	Same as Full Config with the following changes: 1) Reduces the FortiGate instance sizes from c5n.2xl to c5n.xl (VM08 to VM04); 2) Only deploys the 9 required centralized Interface Endpoints (removes 50). All services remain accessible using the AWS public endpoints, but require traversing the perimeter firewalls; 3) Removes the perimeter VPC Interface Endpoints; 4) Removes the Unclass ou and VPC.	\$2,575
		\$2,550
		\$2,450 +FW lic.
		\$2475 +FW lic.
	Four variants of the lite configuration file are provided: - AWS Control Tower w/AWS Network Firewall instead of IPSEC VPN Firewalls (recommended starting point) - AWS Network Firewall instead of IPSEC VPN Firewalls - IPSEC VPN integrated 3rd party firewalls - AWS Gateway Load Balancer integrated 3rd party firewalls	
Full	Large IPSEC VPN Firewalls w/Endpoints - The full configuration file was based on feedback from customers moving into AWS at scale and at a rapid pace. Customers of this nature have indicated that they do not want to have to upsize their perimeter firewalls or add Interface endpoints as their developers start to use new AWS services. These are the two most expensive components of the deployed architecture solution.	\$4,200

1.2.2. PRICING BY AWS ACCOUNT (ALL CONFIGURATIONS)

The following table provides the estimated monthly pricing per AWS account for each of the example configuration files.

AWS Account	Description	Ultra Lite	Test	Lite	Full
Management	This is the organization management or root account. This account aggregates organization wide billing, and is used to manage the Accelerator, AWS SSO and SCPs. Access to this account must be highly restricted. This account should not contain any customer resources or workloads.	\$10	\$75	\$140	\$140
Operations	This Account is used for centralized IT operational resources (MAD, rsyslog, ITSM, etc.) which need to be made available to all accounts in the organization and would generally be used and managed by the Cloud Operations team.	-	\$275	\$680	\$680
Security	The security account is generally used and managed by the customers security and compliance teams, and contains an organizations security tooling and consoles. This account functions as the organization administrative account for Security Hub, GuardDuty, Macie, Firewall Manager, and Access Analyzer. This account also has the ability to assume a view-only role in every account in the organization to conduct security investigations.	\$5	\$10	\$25	\$25
Log Archive	The log archive account provides a central aggregation and secure long-term storage location for all logs created within the AWS organization. Logs created in every account in the organization are centralized to an S3 bucket in this account.	\$15	\$35	\$55	\$55
Perimeter	This account is used as the centralized internet facing ingress/egress point and contains edge security services for the organizations IaaS based workloads.	-	\$590	\$410- \$700	\$1,200
Shared Network	This account is used for centralized or shared networking resources and will typically contain a transit gateway to enable routing between different AWS based and on-premises networks. If a centralized or shared VPC architecture is deployed, this account will also contain VPCs (i.e. Dev, Test, Prod) which are shared via RAM sharing to accounts within designated OUs in the organization. If a spoke architecture is used, the Transit gateway is instead shared to the accounts within the organization.	-	\$515	\$825- \$995	\$1,950
MyDev1	This is an optional sample workload account which lives in the Dev organizational unit. Dev accounts have a full set of security guardrails similar to a production accounts and are designed to be used by developers. These accounts leverage either local or centralized networking and are connected to the organizations network via the centralized transit gateway, which is used to access the internet via the perimeter security account or on-premises networks.	-	-	\$80	\$80
TheFunAccount	This is an optional sample workload account that is created in Sandbox organizational unit. Sandbox accounts are designed for experimentation only, as they have the fewest guardrails, and provide the most cloud native experience. These accounts leverage localized networking and are fully isolated from all other organization networks, with no transit gateway connectivity and direct internet access via a local internet gateway.	-	-	\$70	\$70

AWS Account	Description	Ultra Lite	Test	Lite	Full
TOTAL	Estimated Monthly Pricing	\$30	\$1500	\$2,450 - \$2,575	\$4,200

1.2.3. DETAILED PRICING BY AWS SERVICE (LITE CONFIG – IPSEC VPN ACTIVE/ACTIVE FIREWALLS)

We picked a single example configuration file to provide detailed pricing per service.

The following table provides the estimated monthly pricing per AWS services provisioned by the Accelerator, across all accounts, for the Lite – IPsec VPN configuration.

AWS service	Quantity	Estimated Monthly Pricing
CloudTrail (All Regions)		\$28
CloudWatch (All Regions)		\$35
CloudWatch Events (All Regions)		\$0
CodeBuild		\$2
CodeCommit		\$0
CodePipeline		\$0
Config (All Regions)		\$85
Data Transfer		\$0
Directory Service	- Managed Active Directory (2 domain controllers) - Shared Directory (2 accounts) - Small AD Connector (1)	\$444
DynamoDB		\$0
EC2 Container Registry (ECR)		\$0.2
Elastic Compute Cloud (EC2)	- NAT Gateway (1) - Remote Desktop Gateway (1 x Windows t3.large) - rsyslog Servers (2 x Linux t3.large) - Fortinet Firewalls (2 x Linux c5n.xlarge) - EBS Volumes (30 GB x 3 instances, 100 GB x 2 instances)	\$669
Elastic Load Balancing	- Application Load Balancing (2) - Network Load Balancing (rsyslog) (1)	\$55
GuardDuty (All Regions)		\$41
Key Management Service (All Regions)		\$44
Kinesis		\$12
Kinesis Firehose		\$2
Lambda (All Regions)		\$0
Macie (All Regions)		\$4
Route 53	- HostedZones (11) - Resolver Network Interfaces (4)	\$378
Secrets Manager		\$5
Security Hub (All Regions)		\$97
Simple Notification Service (All regions)		\$0
Simple Queue Service (All Regions)		\$0
Simple Storage Service (All regions)		\$6
Step Functions		\$1
Systems Manager		\$0

AWS service	Quantity	Estimated Monthly Pricing
Virtual Private Cloud	- VPC Endpoints (18) - VPN Connections (2) - Transit Gateway VPC Attachments (5) - Transit Gateway VPN Attachments (2)	\$542
TOTAL	Estimated Monthly Pricing	\$2,450

2.4.3 AWS Secure Environment Accelerator Deployment Capabilities

Overview

Deploys, creates, manages and updates the following objects across a multi-region, multi-account AWS environment

TASK	Accelerator - What happens, WHERE, under what condition, on each state machine execution
AWS Accounts	
- Creates mandatory accounts (accounts which other accounts are dependent on)	organization management (root) account, global scope
- Creates workload accounts (individually or in bulk), base personality determined by ou placement	organization management (root) account, global scope
- Supports native AWS Organization account and OU activities (OU and account rename, move account between OU's, create accounts, etc.)	organization management (root) account, global scope
- Applies a Deny All SCP on any newly created account(s) until successfully guardrailed	organization management (root) account, new account scope (failure to apply guardrails fails the Accelerator and leaves account blocked until remediated)
- Allows bulk parallel* account creation, configuration, updates and guardrail application	creates, guardrails and configures new accounts and regions in parallel per defined personas, organization management (root) account. Control Tower account ingestion is sequential at this time.
- Performs 'account warming' to establish initial limits, when required	state Machine region only, defined accounts (per region potential)
- Checks limit increases, when required (complies with initial limits until increased)	per account, per region (supported limits only)
- Automatically submits limit increases, when required	state Machine region only, defined accounts (per region potential)
- Leverages AWS Control Tower	Accelerator and Control Tower home regions must match, the Accelerator supports all on-by-default regions and will require a standalone install in regions not yet supported by Control Tower
Networking	
- Creates Transit Gateways and TGW route tables incl. static routes and inter-region TGW peering	in the defined region(s), defined account(s)
- Creates centralized and/or local account (bespoke) VPC's	in the defined region(s), defined account(s)
...all completely and individually customizable (per account, VPC, subnet, or OU), Static or Dynamic VPC and subnet CIDR assignments	
- Creates Subnets, Route tables, NACLs, Security groups, NATGWs, IGWs, VGWs, CGWs (per customer specs)	part of any VPC, in the defined region(s), defined account(s) - allows detailed CIDR allocation, and cross-account security group referencing
- Deletes default VPC's (worldwide)	in all regions, in all accounts, can disable regions (all accounts or specific account)
- Creates VPC Endpoints (Gateway and Interface)	part of any VPC, in the defined region(s), defined account(s)
- Configures centralized endpoints (R53 zones populated, shared and attached to local and cross-account VPC's)	configures regional central endpoints (only one 'central' VPC per region)
- Creates Route 53 Private and Public Zones	in the defined account(s), defined region(s), defined VPC(s), global scope
- Creates Resolver Rules and Resolver (inbound/outbound) Endpoints	part of a specific VPC(s), in the defined region(s), defined account(s) (i.e. per region possible)
...including MAD R53 DNS resolver rule creation	created in same region as MAD only, shared to same region VPC's when use-central-endpoints set
- Automatically creates R53 VPC Endpoint Overloaded Zones	same region(s), same account(s) as the endpoint and VPC(s)
- Deploys and configures AWS Network Firewall	on any VPC, any region, any account

TASK	Accelerator - What happens, WHERE, under what condition, on each state machine execution
Cross-Account Object Sharing	
- VPC and Subnet sharing, including account level retagging/naming (and per account security group 'replication')	VPC's are shared to accounts within the SAME REGION as the source VPC only An OU could have additional VPC's defined for additional regions and would be shared to the appropriate accounts in the same additional regions
- VPC peering and TGW attachments (local and cross-account)	in the defined region, no cross-region attachments or peering supported
- Managed Active Directory sharing	state machine region only (consider same region as the MAD only)(unshare method not implemented)
- Automated TGW inter-region peering	cross-region, cross-account or same-account
- Shares SSM remediation documents	from defined account(s), to defined OU's, in defined regions
Zone sharing and VPC associations	
- Public Hosted Zones	no sharing, no association required (any account, any VPC, any region)
- Private Hosted Zones - i.e. Cloud DNS domains	associated worldwide to all VPCs with use-central-endpoints
- Endpoint Private Hosted Zones	associate within region, for all VPC use-central-endpoints (including cross-account)
- On-premise resolver rules	associate within region, for all VPC use-central-endpoints (including cross-account)
- MAD resolver rule association	same region as the MAD resolver only, assoc. w/all VPC use-central-endpoints
Identity	
- Creates Directory services (Managed Active Directory and Active Directory Connectors)	in a specific VPC, in the defined region, defined account - only 1 per account, therefore can't have a second region in the same account (ADC creation only supported in mandatory accounts)
- Creates Windows admin bastion host auto-scaling group	once per above MAD (once per account), same region as MAD
- Set Windows domain password policies (initial installation only)	once per above MAD (once per account), same region as MAD
- Set IAM account password policies	once per account, global scope
- Creates Windows domain users and groups (initial installation only)	once per above MAD (once per account), same region as MAD
- Creates IAM Policies, Roles, Users, and Groups	once per account, global scope
Cloud Security Services	
- Enables and configs the following AWS services, worldwide w/central specified admin account:	(each service can have specified regions disabled)
- GuardDuty w/S3 protection	enabled all regions, all accounts, admin account per region
- Security Hub (Enables specified security standards, and disables specified individual controls)	enabled all regions, all accounts, admin account per region
- Firewall Manager	enabled once per account (global scope), single admin account
- CloudTrail w/Insights and S3 data plane logging	enabled all regions (using Organization trail, stored in Organization Management account)
- Config Recorders/Aggregator	enabled all regions, all regions include global events, aggregator set to specified region in Organization Management account

TASK	Accelerator - What happens, WHERE, under what condition, on each state machine execution
- Macie	enabled all regions, admin account per region
- IAM Access Analyzer	enabled once per account (global scope), single admin account
- Enables CloudWatch access from central specified admin account	enabled once per account (global scope), two admin accounts (Ops & Security)
- Deploys customer provided SSM remediation documents (four provided out-of-box today)	customized per OU, defined regions, defined accounts
...remediates S3 buckets without KMS CMK encryption and ALB's without centralized logging	customized per OU, all regions, integrated w/SSM remediation, when desired
- Deploys AWS Config rules (managed and custom) including AWS Conformance packs (NIST 800-53 deployed by default + 2 custom)	customized per OU, all regions, all accounts integrated w/SSM remediation, when desired
Other Security Capabilities	
- Creates, deploys and applies Service Control Policies	at the top OU level only, sub-ou's managed directly through AWS Organizations
- Creates Customer Managed KMS Keys w/automatic key rotation (SSM, EBS, S3)	SSM and EBS keys are created if a VPC exists in the region, S3 if we need an Accelerator bucket in the region, per account
- Enables account level default EBS KMS CMK encryption	set if a VPC exists in the region, per account
- Enables S3 Block Public Access	once per account, global scope
- Configures Systems Manager Session Manager w/KMS CMK encryption and centralized logging	set if a VPC exists in the region, per account
- Imports or requests certificates into AWS Certificate Manager	State Machine region only (per region potential, required for ALB deployments)
- Deploys both perimeter and account level ALB's w/Lambda health checks, certs & TLS policies	State Machine region only (per region potential)
- Deploys & configures 3rd party firewall clusters and management instances	in the defined region(s), defined account(s)
...Gateway Load Balancer w/auto-scaling (NEW) and VPN IPSec BGP ECMP deployment options	
- Configuration is fully managed and maintained in AWS CodeCommit - full multi-account configuration history	organization management (root) account
...breaking configuration changes block Accelerator execution	Idempotent - extensive error handling and failure cleanup - Accelerator can be stopped, started, and rerun without implication
Centralized Logging	
- Deploys an rsyslog auto-scaling cluster behind an NLB, all syslogs forwarded to CWL	State Machine region only (per region potential)
- Centralizes logging to a single centralized S3 KMS CMK encrypted bucket (enables, configures and centralizes) incl:	Sets S3 ownership flag, sets bucket retentions
- VPC Flow logs (w/Enhanced metadata fields and optional CWL destination)	part of a specific VPC, in the defined region, defined account (to local account bucket in state machine region, replicated to log-archive primary region)
- Organizational Cost and Usage Reports	once per organization, global scope (to local account bucket in state machine region, replicated to log-archive primary region)
- CloudTrail Logs including S3 Data Plane Logs (also sent to CWL)	directly back to log-archive, specified primary region

TASK	Accelerator - What happens, WHERE, under what condition, on each state machine execution
- All CloudWatch Logs (includes rsyslog logs) (and setting Log group retentions)	State machine region, plus configured regions
- Config History and Snapshots	directly back to log-archive account specified primary region
- Route 53 Public Zone Logs, DNS Resolver Query Logs	to CloudWatch Logs in us-east-1 (which are sent to S3)
- GuardDuty Findings	directly back to log-archive, specified primary region
- Macie Discovery results	directly back to security, specified primary region, replicated to log-archive
- ALB Logs	State Machine region only (same as ALB deployment)
- SSM Session Logs (also sent to CWL)	All regions currently send back to central region, log-archive account
Extensibility	
- Populates each accounts Parameter Store with the Accelerator deployed objects (allows customer IaC to extend/leverage)	each account, defined regions (all ELB's across the environment are populated in specified accounts, i.e. perimeter, to enable automated end-to-end plumbing)
- Every execution outputs the execution status and a list of successfully guardrailed accounts to a SNS topic	allows 3rd party framework to execute after every Accelerator execution by hooking to SNS topic
...which emails a customer defined email address	...or hooking to the email alert
- Deploys roles with customized access (read-only,write) to the log-archive buckets (enabling customer SIEM deployments, SSM, EC2 CWL)	defined account, global scope
- Designed for Day 1, 2 and day 10. Customers get new features without any customization effort no matter the deployed architecture	Upgradable from any version to any version, no customization or professional services required (Customer production proven across multiple releases)
Alerting	
- Deploys global High, Medium, Low, Ignore priority SNS topics and email subscriptions	in the defined account, org accessible regional topics, each region subscribed to a single defined central region which has the email subscriptions
- Deploys customer defined CloudWatch Log Metrics and Alarms w/prioritized alarms (19 out-of-box)	all accounts, home region only, as this is where the Org/account CloudTrail exists
- Creates and configures AWS budgets w/alerting (customizable per OU and per account)	once per account, global scope
- Configures email alerting for CloudTrail Metric Alarms, Firewall Manager Events, Security Hub Findings incl. GuardDuty Findings	

General

- "defined" region, "defined" account, means "customer defined", either at installation, upgrade, or any time they decide to reconfigure
- all items are created per customer defined parameters and configurations and are fully customizable without changing a single line of code
- security services are enabled and deployed globally, but, each service can be disabled per region. A single region deployment is possible.
- customer can enable/disable features, or change the configuration of each feature in the Accelerator config file
- customers can evolve their configurations over time, as they evolve and as their requirements change, without the requirement for code changes or professional services


Region support

- All AWS commercial regions are supported. Lack of availability of CodeBuild, CodeCommit, or AWS Organizations in the Accelerator primary or installation region will prevent installation directly in that region. In these cases, customers can select a different installation region and the Accelerator can remotely deploy configurations and guardrails to that unsupported installation region.
- Prior to v1.2.5, we utilized a single StackSet, which blocked several additional installation regions. The Accelerator no longer leverages any StackSets, unblocking installing directly in several additional regions.
- As most features can be toggled on/off (per region), we expect most regions should be supportable both as a primary (or installation) region with the three above noted exceptions, and in these cases should still be fully supported as a managed (or secondary) region.
- Opt-in regions are not yet supported, but given enough demand, could easily be added.

[...Return to Accelerator Table of Contents](#)

2.4.4 1. Accelerator Central Logging Implementation and File Structures

The following diagram details the ASEA central logging implementation:



1.1. Accelerator Central Logging Buckets

Bucket Type	Bucket Name	Purpose
AES Encrypted Bucket	pbummaccel-logarchive-phase0-aescentral1-1py9vr4cdwuxu	ALB Logs - ALB's do not support logging to a KMS bucket
KMS Encrypted Bucket	pbummaccel-logarchive-phase0-cacentral1-1tr23emhncdzo	All other AWS Accelerator initiated logs
AES or KMS Encrypted	aws-controltower-logs-123456789012-ca-central-1	All Control Tower initiated logs
AES or KMS Encrypted	aws-controltower-s3-access-logs-123456789012-ca-central-1	S3 Access logs for the Control Tower logs bucket

1.1.1. NOTES

- Every customer has two Accelerator logging buckets
- Control Tower installations have an additional two Control Tower logging buckets
- Customers could use any account name for their central logging account

- Bucket name format is: {Accel-Prefix}-{Account-Name}-{Accel-Phase}-xxx{Region}-{Random}
- {Accel-Prefix} defaults to 'asea' (previously 'pbmmacel' for Canada)
- {Accel-Phase} should always be 'phase0'
- {region} should always be 'cacentral1' for Canada
- {account} is likely to be 'log-archive'
- xxx is either "aes" or "" (nothing)

1.2. Accelerator Bucket Folders

Log Type	Folder Path	Example
ELB (in AES bucket)	{account#}/elb- {elbname}/AWSLogs/ {account#}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-aescacental1-1py9vr4ucwuxu/123456789012/elb-Core-mydevacct1-alb/AWSLogs/123456789012/ELBAccessLogTestFile s3://pbmmacel-logarchive-phase0-aescacental1-1py9vr4ucwuxu/123456789013/elb-Public-Prod-perimeter-alb/AWSLogs/123456789013/ELBAccessLogTestFile
VPC Flow Logs	{account#}/{vpc- name}/AWSLogs/ {account#}/ vpcflowlogs/{region}/ {year}/{month}/{day}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/123456789012/Test-East-icl/AWSLogs/123456789012/vpcflowlogs/us-east-1/2020/08/31/123456789012_vpcflowlogs_us-east-1_fl-04af3543c74402594_20200831T1720Z_73d3922a.log.gz
Macie Reports	{account#}/ maciestestobject	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/123456789014/macie-test-object
Cost and Usage Reports	{account#}/cur/Cost- and-Usage-Report/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/123456789015/cur/Cost-and-Usage-Report/*
Config History*	AWSLogs/{account#}/ Config/{region}/{year}/ {month}/{day}/ ConfigHistory/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/AWSLogs/123456789016/Config/ca-central-1/2020/8/31/ConfigHistory/123456789016_Config_ca-central-1_ConfigHistory_AWS::CloudFormation::Stack_20200831T011226Z_20200831T025845Z_1.json.gz
Config Snapshot*	AWSLogs/{account#}/ Config/{region}/{year}/ {month}/{day}/ ConfigSnapshot/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/AWSLogs/123456789016/Config/ca-central-1/2020/8/30/ConfigSnapshot/123456789016_Config_ca-central-1_ConfigSnapshot_20200830T193058Z_5d173149-e6d0-41e4-af7f-031ff736f8c8.json.gz
GuardDuty	AWSLogs/{account#}/ GuardDuty/{region}/ {year}/{month}/{day}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/AWSLogs/123456789014/GuardDuty/ca-central-1/2020/09/02/294c9171-4867-3774-9756-f6f6c209616f.jsonl.gz
CloudWatch Logs****	CloudWatchLogs/ {year}/{month}/{day}/ {hour}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/CloudWatchLogs/2020/08/30/00/PBMMACcel-Kinesis-Delivery-Stream-1-2020-08-30-00-53-33-35aeea4c-582a-444b-8afa-848567924094
CloudTrail Digest***	{org-id}/AWSLogs/ {org-id}/{account#}/ CloudTrail-Digest/ {region}/{year}/ {month}/{day}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/o-fxozgwu6rc/AWSLogs/o-fxozgwu6rc/123456789016/CloudTrail-Digest/ca-central-1/2020/08/30/123456789016_CloudTrail-Digest_ca-central-1_PBMMACcel-Org-Trail_ca-central-1_20200830T190938Z.json.gz
CloudTrail Insights**	{org-id}/AWSLogs/ {org-id}/{account#}/ CloudTrail-Insights/ {region}/{year}/ {month}/{day}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/o-fxozgwu6rc/AWSLogs/o-fxozgwu6rc/123456789015/CloudTrail-Insight/ca-central-1/2020/09/23/123456789015_CloudTrail-Insight_ca-central-1_20200923T0516Z_KL5e9VCV2SS7IqzB.json.gz
CloudTrail***	{org-id}/AWSLogs/ {org-id}/{account#}/ CloudTrail/{region}/ {year}/{month}/{day}/*	<ul style="list-style-type: none"> s3://pbmmacel-logarchive-phase0-cacental1-1tr23emhncdzo/o-fxozgwu6rc/AWSLogs/o-fxozgwu6rc/123456789016/CloudTrail/ca-central-1/2020/08/30/123456789016_CloudTrail_ca-central-1_20200830T0115Z_3YQJxwt5qUaOzMtL.json.gz
CT S3 Access Logs	{no folders}	<ul style="list-style-type: none"> s3://aws-controltower-s3-access-logs-123456789012-ca-central-1/2021-04-26-18-11-21-8647E1080048E5CB
SSM Inventory	ssm-inventory/{ssm- inventory-type}/ accountid={account#}/ region={region}/ resourcetype={rt}/*	<ul style="list-style-type: none"> s3://asea-logarchive-phase0-cacental1-1tr23emhncdzo/ssm-inventory/AWS:Application/accountid=123456789012/region=ca-central-1/resourcetype=ManagedInstanceInventory/i-001188b4e152aecaf.json

1.2.1. NOTES

* Located in Control Tower bucket when installed, Control Tower adds the {org-id} (i.e. o-h9ho05hcxl/) as the top level folder

** Only available in Accelerator Standalone deployments

*** CloudTrail control plane logs located in Control Tower bucket when installed, Control Tower drops the {org-id} (i.e. o-h9ho05hcxl/) from the middle of the folder path. This may change when Control Tower migrates to Organization Trails. CloudTrail data plane logs remain in the Accelerator bucket.

**** v1.5.1 introduces the capability to split CloudWatch log groups starting with specific prefixes out into customer named subfolders. The folder/file structure is otherwise identical. The v1.5.1 example config files separate out MAD, RQL, Security Hub, NFW, rsyslog, and SSM logs by default.

Example: Security Hub logs will be in the following structure: `CloudWatchLogs/security-hub/{year}/{month}/{day}/{hour}/`

- Account number is sometimes duplicated in path because logs replicated from another account always need to start with the source account number
- Macie reports will only appear in the {account#} for the central security account, and only if a customer schedules PII discovery reports
- All CloudWatch Logs from all accounts are mixed in the same folder, the embedded log format contains the source account information as documented here: <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/ValidateLogEventFlow.html>
- With the exception of CloudWatch Logs, all logs are in the original format provided by the log source/service.

2.4.5 Object Naming

ACCELERATOR OBJECT NAMING

- Resources will have the 'Name' tag assigned, where Name={name}{suffix}
- No prefix or suffix will be applied to DNS records/zones (as that breaks them)
- When _ is not supported, a - will be used
- Stacks/stacksets/functions and **non-end user** accessed objects deployed in all accounts will also start with the {AcceleratorPrefix} prefix (i.e. "**PBMMAccel-**" or "**ASEA-**")
- The prefix does not apply to objects like VPC's, subnets, or TGW's which customers need to directly access. This is for objects deployed to build the customer accessible objects
- This prefix will be protected by SCP's so customers don't break 'managed' features
- Resources will have the tag 'Accelerator={AcceleratorName}' assigned when tags are supported
- Stacks will have the tag 'AcceleratorName={AcceleratorName}' assigned, which will often (but not always) be inherited by objects created by the stack (due to TGW duplicate tag issue)

DEFAULTS

- the default {AcceleratorName} is 'PBMM' before v1.5.0 and 'ASEA' after v1.5.0
- the default {AcceleratorPrefix} is 'PBMMAccel-' before v1.5.0 and 'ASEA-' after v1.5.0

SUFFIX'S

suffix	object type
_vpc	VPC
_azN_net	Subnet
_azN_rt	RouteTable
_tgw	Transit Gateway
-key	KMS key
_pcx	Peering Connection
_sg	Security Group
_nacl	NACL
_alb	Application Load Balancer
_nlb	Network Load Balancer
_agw	Appliance Gateway
_vpce	VPC Endpoint
_AMI	AMI
_dhcp	DHCP option set
_snap	snapshot
_ebs	Block storage
_igw	internet gateway
_lgw	Local gateway
_nat	NAT gateway
_vpg	Virtual private gateway
_cgw	Customer gateway
_vpn	VPN Connection
_sm	Step Functions
_rule	CW Event Rule
_pl	CodeBuild

NO SUFFIX

suffix	object type
None	Stacks
None	CFN_Stack_Sets
None	Lambda
None	Cloud Trails
None	CWL Groups
None	Config Rules
None	OU
None	Service Control Policy

3. 1. Accelerator Basic Operation and Frequently asked Questions

3.1 1.1. Operational Activities

1.1.1. How do I add new AWS accounts to my AWS Organization?

How do I add new AWS accounts to my AWS Organization?

We offer three options and all can be used in the same Accelerator deployment. All options work with AWS Control Tower, ensuring the account is both ingested into Control Tower and all Accelerator guardrails are automatically applied.

No matter the mechanism you choose, new accounts will automatically be blocked from use until fully guardrailed, the Accelerator will automatically execute, and accounts will automatically be ingested into AWS Control Tower (if deployed).

Option 1

Users can simply add the following five lines to the configuration file `workload-account-configs` section and rerun the state machine. The majority of the account configuration will be picked up from the OU the AWS account has been assigned. You can also add additional account specific configuration, or override items like the default OU budget with an account specific budget. This mechanism is often used by customers that wish to programmatically create AWS accounts using the Accelerator and allows for adding many new accounts at one time.

```
"fun-acct": {
  "account-name": "TheFunAccount",
  "email": "myemail+aseaT-funacct@example.com",
  "src-filename": "config.json",
  "ou": "Sandbox"
}
```

Option 2

We've heard consistent feedback that our customers wish to use native AWS services and do not want to do things differently once security controls, guardrails, or accelerators are applied to their environment. In this regard, simply create your new AWS account in AWS Organizations as you did before**, either by a) using the AWS Console or b) by using standard AWS account creation API's, CLI or 3rd party tools like Terraform.

- **** IMPORTANT:** When creating the new AWS account using AWS Organizations, you need to specify the role name provided in the Accelerator configuration file `global-options\organization-admin-role`, otherwise we cannot bootstrap the account. In Control Tower installations, this **MUST** be set to `AWSControlTowerExecution`, for customers who installed prior to v1.2.5 this value is `AWSCloudFormationStackSetExecutionRole` and after v1.2.5 we were recommending using the role `OrganizationAccountAccessRole` as this role is used by default by AWS Organizations if no role name is specified when creating AWS accounts through the AWS console or cli.
- On account creation we will apply a quarantine SCP which prevents the account from being used by anyone until the Accelerator has applied the appropriate guardrails
- Moving the account into the appropriate OU triggers the state machine and the application of the guardrails to the account, once complete, we will remove the quarantine SCP.
- **NOTE:** Accounts CANNOT be moved between OU's to maintain compliance, so select the proper top-level OU with care
- In AWS Organizations, select ALL the newly created AWS accounts and move them all (preferably at once) to the correct destination OU (assuming the same OU for all accounts)
- In case you need to move accounts to multiple OU's we have added a 2 minute delay before triggering the State Machine
- Any accounts moved after the 2 minute window will NOT be properly ingested, and will need to be ingested on a subsequent State Machine Execution.

Option 3

Create your account using Account Factory in the AWS Control Tower console.

1.1.2. I tried to enroll a new account via Control Tower but it failed?**I tried to enroll a new account via Control Tower but it failed?**

or "The state machine failed during the 'Load Organization Configuration' step with the error 'The Control Tower account: ACCOUNT_NAME is in a failed state ERROR'"

If account enrollment fails within Control Tower, you will need to follow the troubleshooting steps [here](#). A common reason for this is not having the `ControlTowerExecution` role created in the account you are trying to enroll. Even after you successfully enroll the account, it is possible the state machine will fail at `Load Organization Configuration`. If you look at the CloudWatch logs you will see the error message:

```
There were errors while loading the configuration: The Control Tower account: ACCOUNT_NAME is in a failed state ERROR.
```

This is because the Accelerator checks that there are no errors with Control Tower before continuing. In some cases Control Tower can leave an orphaned Service Catalog product in an **Error** state. You need to cleanup Control Towers Service Catalogs Provisioned Products so there are no products remaining in an error or tainted state before you can successfully re-run the state machine.

1.1.3. Can I use AWS Organizations for all tasks I currently use AWS Organizations for?

Can I use AWS Organizations for all tasks I currently use AWS Organizations for?

In AWS Organizations you can continue to:

- create and rename AWS accounts
- move AWS accounts between OU's
- create, delete and rename OU's, including support for nested OU's
- create, rename, modify, apply and remove SCP's

What can't I do:

- modify Accelerator or Control Tower controlled SCP's
- add/remove SCP's on top-level OU's (these are Accelerator and/or Control Tower controlled)
- users can change SCP's on non-top-level OU's and non-Accelerator controlled accounts as they please
- add/remove SCP's on specific accounts that have Accelerator controlled SCPs
- move an AWS account between top-level OU's (i.e. `Sandbox` to `Prod` is a security violation)
- moving between `Prod/sub-ou-1` to `Prod/sub-ou2` or `Prod/sub-ou2/sub-ou2a/sub-ou2ab` is fully supported
- create a top-level OU (need to validate, as they require config file entries)
- remove quarantine SCP from newly created accounts
- we do not support forward slashes (/) in OU names, even though the AWS platform does

More details:

- If an AWS account is renamed, an account email is changed, or an OU is renamed, on the next state machine execution, the config file will automatically be updated.
- If you edit an Accelerator controlled SCP through Organizations, we will reset it per what is defined in the Accelerator configuration files.
- If you add/remove an SCP from a top-level OU or Accelerator controlled account, we will put them back as defined in the Accelerator configuration file.
- If you move an account between top-level OU's, we will put it back to its original designated top-level OU.
- The Accelerator fully supports nested OU's, customers can create any depth OU structure in AWS Organizations and add/remove/change SCP's *below* the top-level as they desire or move accounts between these OU's without restriction. Users can create OU's to the full AWS OU structure/depth
- Except for the Quarantine SCP applied to specific accounts, we do not 'control' SCP's below the top level, customers can add/create/customize SCP's
- as of v1.3.3 customers can optionally control account level SCP's through the configuration file

1.1.4. How do I make changes to items I defined in the Accelerator configuration file during installation?

How do I make changes to items I defined in the Accelerator configuration file during installation?

Simply update your configuration file in CodeCommit and rerun the state machine! In most cases, it is that simple.

If you ask the Accelerator to do something that is not supported by the AWS platform, the state machine will fail, so it needs to be a supported capability. For example, the platform does not allow you to change the CIDR block on a VPC, but you can accomplish this as you would today by using the Accelerator to deploy a new second VPC, manually migrating workloads, and then removing the deprecated VPC from the Accelerator configuration.

Below we have also documented additional considerations when creating or updating the configuration file.

It should be noted that we have added code to the Accelerator to block customers from making many 'breaking' or impactful changes to their configuration files. If someone is positive they want to make these changes, we also provide override switches to allow these changes to be attempted forcefully.

1.1.5. Can I update the config file while the State Machine is running? When will those changes be applied?

Can I update the config file while the State Machine is running? When will those changes be applied?

Yes. The state machine captures a consistent input state of the requested configuration when it starts. The running Accelerator instance does not see or consider any configuration changes that occur after it has started. All configuration changes occurring after the state machine is running will only be leveraged on the *next* state machine execution.

1.1.6. What if I really mess up the configuration file?

What if I really mess up the configuration file?

The Accelerator is designed with checks to compare your current configuration file with the version of the config file from the previous successful execution of the state machine. If we believe you are making major or breaking changes to the config file, we will purposefully fail the state machine. See [Config file and Deployment Protections](#) for more details.

With the release of v1.3.0 we introduced state machine scoping capabilities to further protect customers, detailed [here](#).

1.1.7. What if my State Machine fails? Why? Previous solutions had complex recovery processes, what's involved?

What if my State Machine fails? Why? Previous solutions had complex recovery processes, what's involved?

If your main state machine fails, review the error(s), resolve the problem and simply re-run the state machine. We've put a huge focus on ensuring the solution is idempotent and to ensure recovery is a smooth and easy process.

Ensuring the integrity of deployed guardrails is critical in operating and maintaining an environment hosting protected data. Based on customer feedback and security best practices, we purposely fail the state machine if we cannot successfully deploy guardrails.

Additionally, with millions of active customers each supporting different and diverse use cases and with the rapid rate of evolution of the AWS platform, sometimes we will encounter unexpected circumstances and the state machine might fail.

We've spent a lot of time over the course of the Accelerator development process ensuring the solution can roll forward, roll backward, be stopped, restarted, and rerun without issues. A huge focus was placed on dealing with and writing custom code to manage and deal with non-idempotent resources (like S3 buckets, log groups, KMS keys, etc.). We've spent a lot of time ensuring that any failed artifacts are automatically cleaned up and don't cause subsequent executions to fail. We've put a strong focus on ensuring you do not need to go into your various AWS sub-accounts and manually remove or cleanup resources or deployment failures. We've also tried to provide usable error messages that are easy to understand and troubleshoot. As new scenarios are brought to our attention, we continue to adjust the codebase to better handle these situations.

Will your state machine fail at some point in time, likely. Will you be able to easily recover and move forward without extensive time and effort, YES!

1.1.8. How do I update some of the supplied sample configuration items found in reference-artifact, like SCPs and IAM policies?

How do I update some of the supplied sample configuration items found in reference-artifact, like SCPs and IAM policies?

To override items like SCP's or IAM policies, customers simply need to provide the identically named file in their input bucket. As long as the file exists in the correct folder in the customers input bucket, the Accelerator will use the customers supplied version of the configuration item, rather than the Accelerator version. Customer SCP's need to be placed into a folder named `scp` and IAM policies in a folder named `iam-policy` (case sensitive).

The Accelerator was designed to allow customers complete customization capabilities without any requirement to update code or fork the GitHub repo. Additionally, rather than forcing customers to provide a multitude of config files for a standard or prescriptive installation, we provide and auto-deploy with Accelerator versions of most required configuration items from the reference-artifacts folder of the repo. If a customer provides the required configuration file in their Accelerator S3 input bucket, we will use the customer supplied version of the configuration file rather than the Accelerator version. At any time, either before initial installation, or in future, a customer can place new or updated SCPs, policies, or other supported file types into their input bucket and we will use those instead of or in addition to Accelerator supplied versions. Customer only need to provide the specific files they wish to override, not all files.

Customers can also define additional SCPs (or modify existing SCPs) using the name, description and filename of their choosing, and deploy them by referencing them on the appropriate organizational unit in the config file.

Prior to v1.2.5, if we updated the default files, we overwrote customers customizations during upgrade. Simply updating the timestamp *after* upgrade on the customized versions and then rerunning the state machine re-instates customer customizations. In v1.2.5 we always use the customer customized version from the S3 bucket. Its important customers assess newly provided defaults during an upgrade process to ensure they are incorporating all the latest fixes and improvements. If a customer wants to revert to Accelerator provided default files, they will need to manually copy it from the repo into their input bucket.

NOTE: Most of the provided SCPs are designed to protect the Accelerator deployed resources from modification and ensure the integrity of the Accelerator. Extreme caution must be exercised if the provided SCPs are modified. In v1.5.0 we restructured the SCPs based on a) customer requests, and b) the addition of Control Tower support for new installs.

- we reorganized and optimized our SCP's from 4 SCP files down to 3 SCP files, without removing any protections or guardrails;
- these optimizations have resulted in minor enhancements to the SCP protections and in some cases better scoping;
- the first two SCP files (Part-0 and Part-1) contain the controls which protect the integrity of the Accelerator itself;
- the third file (Sensitive, Unclass, Sandbox) contains customer data protection specific guardrails, which may change based on workload data classification or customer profiles and requirements;
- this freed the fourth SCP for use by Control Tower. As Control Tower leverages 2 SCP files on the Security OU, we have moved some of our SCP's to the account level.

1.1.9. I deployed AWS Managed Active Directory (MAD) as part of my deployment, how do I manage Active Directory domain users, groups, and domain policies after deployment?

I deployed AWS Managed Active Directory (MAD) as part of my deployment, how do I manage Active Directory domain users, groups, and domain policies after deployment?

Customers have clearly indicated they do NOT want to use the Accelerator to manage their Active Directory domain or change the way they manage Active Directory on an ongoing basis. Customer have also indicated, they need help getting up and running quickly. For these reasons, the Accelerator only sets the domain password policy, and creates AD users and groups on the initial installation of MAD. After the initial installation, customers must manage Windows users and groups using their traditional tools. A bastion Windows host is deployed as a mechanism to support these capabilities. Passwords for all newly created MAD users have been stored, encrypted, in AWS Secrets Manager in the Management (root) Organization AWS account.

To create new users and groups:

- RDP into the ASEA-RDGW bastion host in the Ops account
- Run ADUC and create users and groups as you please under the NETBIOSDOMAIN (example) tree
- Or run the appropriate PowerShell command
- Go to AWS SSO and map the Active Directory group to the appropriate AWS account and permission set

The Accelerator will not create/update/delete new AD users or groups, nor will it update the domain password policy after the initial installation of Managed Active Directory. It is your responsibility to rotate these passwords on a regular basis per your organizations password policy. (NOTE: After updating the admin password it needs to be stored back in secrets manager).

1.1.10. How do I suspend an AWS account?

How do I suspend an AWS account?

Suspending accounts is blocked via SCP and purposely difficult, two options exist:

1. Modify SCP method (not desired)
2. Leverage the UnManaged OU
 - validate your config file contains the value: `"ignored-ous": ["UnManaged"]`
 - the state machine must be executed at least once after this value is added to the config file
 - In AWS Organizations create an OU named `UnManaged` in the root of the OU tree, if it does not exist
 - Change to the `us-east-1` region and open CloudWatch and navigate to Rules
 - Select the `PBMMAccel-MoveAccount_rule`, select actions, select `Disable`
 - In Organizations move the account to be suspended to the `UnManaged` OU
 - Change to the `us-east-1` region and open CloudWatch and navigate to Rules
 - Select the `PBMMAccel-MoveAccount_rule`, select actions, select `Enable`
 - login to the account to be suspended as the account root user
 - suspend the account through `My Account`
 - Run the state machine (from the Organization management account), the account will:
 - have a `deleted=true` value added to the config file
 - be moved to the suspended OU (OU value and path stays the same in the config file)
 - `deleted=true` causes OU validation to be skipped on this account on subsequent SM executions
 - If the AWS account was listed in the mandatory-accounts section of the config file the SM will fail (expected)
 - after the above tasks have been completed, remove all references to the suspended mandatory account from the config file
 - rerun the state machine, specifying: `{ "overrideComparison": true }`
 - Deleted accounts will continue to appear under the `Suspended` OU for 90-days

1.1.11. I need a new VPC, where shall I define it?

I need a new VPC, where shall I define it?

You can define a VPC in one of four major sections of the Accelerator configuration file:

- within an organization unit (this is the recommended and preferred method);
- within an account in mandatory-account-configs;
- within an account in workload-account-configs;
- defined within an organization unit, but opted-in within the account config.

We generally recommend most items be defined within organizational units, such that all workload accounts pickup their persona from the OU they are associated and minimize per account configuration. Both a local account based VPC (as deployed in the Sandbox OU accounts), or a central shared VPC (as deployed in the Dev/Test/Prod OU accounts in many of the example configs) can be defined at the OU level.

As mandatory accounts often have unique configuration requirements, for example the centralized Endpoint VPC, they must be configured within the account's configuration. Customers can define VPC's or other account specific settings within any account's configuration, but this requires editing the configuration file for each account configuration.

Prior to v1.5.0, local VPC's defined at the OU level were each deployed with the same CIDR ranges and therefor could not be connected to a TGW. Local VPC's requiring centralized networking (i.e. TGW connectivity) were required to be defined in each account config, adding manual effort and bloating the configuration file.

The addition of `dynamic` and `lookup` CIDR sources in v1.5.0 resolves this problem. Local VPCs can be defined in an OU, and each VPC will be dynamically assigned a unique CIDR range from the assigned CIDR pool, or looked up from the DynamoDB database. Customers can now ensure connected, templated VPCs are consistently deployed to every account in an OU, each with unique IP addresses.

v1.5.0 also added a new opt-in VPC capability. A VPC is defined in an OU and a new config file variable is added to this VPC `opt-in: true`. When opt-in is set to true, the state machine does NOT create the VPC for the accounts in the OU, essentially ignoring the VPC definition. Select accounts in the OU can then be opted-in to the VPC(s) definition, by adding the value `accountname\opt-in-vpcs: ["opt-in-vpc-name1", "opt-in-vpc-name2", "opt-in-vpc-nameN"]` to the specific accounts which need the VPC(s). A VPC definition with the specified name (i.e. `opt-in-vpc-name1`) and the value `opt-in: true`, must exist in the OU config for the specified account. When these conditions apply, the VPC will be created in the account per the OU definition. Additional opt-in VPCs can be added to an account, but VPC's cannot be removed from the `opt-in-vpcs` array. VPC's can be TGW attached, assuming `dynamic cidr-src` is utilized, or DynamoDB is prepopulated with the required CIDR ranges using `lookup` mode. `cidr-src` provided is suitable for disconnected Sandbox type accounts.

The Future: While Opt-In VPCs are powerful, we want to take this further. Why not deploy an AWS Service Catalog template which contains the names of all the available opt-in VPCs for the accounts OU, inside each account. An account end user could then request a new VPC for their account from the list of available opt-in patterns. A user's selection would be sent to a centralized queue for approval (w/auto-approval options), which would result in the `opt-in-vpc` entry in that account being updated with the end users requested VPC pattern and the personalized VPC being created in the account and attached to the centralized TGW (if part of the pattern). This would ensure all VPC's conformed to a set of desirable design patterns, but also allow the end-user community choices based on their desired development and app patterns. If you like this idea, please +1 [this](#) feature request.

1.1.12. How do I modify and extend the Accelerator or execute my own code after the Accelerator provisions a new AWS account or the state machine executes?

How do I modify and extend the Accelerator or execute my own code after the Accelerator provisions a new AWS account or the state machine executes?

Flexibility:

- The AWS Secure Environment Accelerator was developed to enable extreme flexibility without requiring a single line of code to be changed. One of our primary goals throughout the development process was to avoid making any decisions that would result in users needing to fork or branch the Accelerator codebase. This would help ensure we had a sustainable and upgradable solution for a broad customer base over time.
- Functionality provided by the Accelerator can generally be controlled by modifying the main Accelerator configuration file.
- Items like SCP's, rsyslog config, PowerShell scripts, and iam-policies have config files provided and auto-deployed as part of the Accelerator to deliver on the prescriptive architecture (these are located in the \reference-artifacts folder of the GitHub repo for reference). If you want to alter the functionality delivered by any of these additional config files, you can simply provide your own by placing it in your specified Accelerator bucket in the appropriate sub-folder. The Accelerator will use your provided version instead of the supplied repo reference version.
- As SCP's and IAM policies are defined in the main config file, you can simply define new policies, pointing to new policy files, and provide these new files in your bucket, and they will be used.
- While a sample firewall config file is provided in the \reference-artifacts folder, it must be manually placed in your S3 bucket/folder on new Accelerator deployments
- Any/all of these files can be updated at any time and will be used on the next execution of the state machine
- Over time, we predict we will provide several sample or reference architectures and not just the current single PBMM architecture (all located in the \reference-artifacts\SAMPLE_CONFIGS folder).

Extensibility:

- Every execution of the state machine sends a state machine status event to a state machine SNS topic
- These status events include the Success/Failure status of the state machine, and on success, a list of all successfully processed AWS accounts
- While this SNS topic is automatically subscribed to a user provided email address for user notification, users can also create additional SNS subscriptions to enable triggering their own subsequent workflows, state machines, or custom code using any supported SNS subscription type (Lambda, SQS, Email, HTTPS, HTTPS)
- Additionally, objects deployed within an account have been populated in Parameter Store, see answer 1.3.2 for details

Example:

- One of our early adopter customers has developed a custom user interface which allows their clients to request new AWS environments. Clients provide items like cost center, budget, and select their environment requirements (i.e. Sandbox, Unclass or full sensitive SDLC account set). On appropriate approval, this pushes the changes to the Accelerator configuration file and triggers the state machine.
- Once the state machine completes, the SNS topic triggers their follow-up workflow, validates the requested accounts were provisioned, updates the customer's account database, and then executes a collection of customer specific follow-up workflow actions on any newly provisioned accounts.

1.1.13. How can I easily access my virtual machines or EC2 instances?

How can I easily access my virtual machines or EC2 instances?

The preferred and recommended method to connect to instances within the Accelerator is by using AWS Systems Manager Session Manager. Session Manager allows access to instances without the need to have any open firewall ports. Session Manager allows for Command line access to instances (both Windows and Linux) directly through the AWS console, with all activity logged to CloudWatch Logs. Session Manager enables customers to connect to Windows instances with a native RDP client and Linux instances with a native SSH client, if desired. Customers can gain quick access to instances through the AWS console, or using their preferred clients.

General

- Both the RDGW and rsyslog instances deployed in the Ops account are properly configured for Systems Manager Session Manager
- We have implemented automation such that all instances are also automatically configured for Session Manager (i.e. configured with the appropriate launch role, has a recent session manager agent installed (most amazon AMI's do), has access to an SSM endpoint)

NEW - System Manager Fleet Manager is now available, which allows connecting graphically to Windows desktops directly from the AWS console without the need for any command line access or tools, and without any requirement for an RDSH/ RDP client. As the Accelerator configures Systems Manager, this feature "simply works" for most instances deployed within the Accelerator.

- Navigate to Systems Manager, open Fleet Manager, select an instance, click "Node Actions" and select "Connect with Remote Desktop". If you are logging in via AWS SSO, you can seamlessly connect. IAM users must supply valid credentials.

Connecting to an Instance

- From the AWS Console
- Go to: EC2, Instances, select the instance (i.e. ASEA-RDGW), click "Connect", select Session Manager, Connect
- Ideal for Linux or Windows PowerShell users
- Everything is fully logged for future reference
- Directly through local RDP client using Session Managers tunnel capability:
- Provides direct access to your instances/host without any open internet ports on the RDGW instance, using a local/fat client tool
- Install AWS CLI v2 on your PC - available [here](#) (uninstall CLIV1 first, if installed)
- Install the SSM plugin on your PC - available [here](#)
- Get AWS temporary credentials from AWS SSO for the account your workload resides (i.e. Ops account when accessing the ASEA-RDGW instance) by selecting "Command line or programmatic access" instead of "Management Console" and paste them into a command prompt
- i.e. via logging in here: <https://xxxxxxx.awsapps.com/start> or
- This [blog](#) describes the process to use SSO to get credentials for the AWS CLI directly without the GUI
- Then enter: `aws ssm start-session --target "i-12345678901234567" --document-name AWS-StartPortForwardingSession --parameters portNumber="3389",localPortNumber="56789"--region ca-central-1`
- Command syntax is slightly different on Linux/Mac
- Replace i-1111adddce582b23c with the instance id of your RDGW instance
- A tunnel will open
- As these are tunnels to proprietary protocols (i.e. RDP/screen scraping) session content is not logged.
- Run mstsc/rdp client and connect to 127.0.0.1:56789
- By replacing 3389 with a new port for another applications (i.e. SSH running on a Linux instance), you can connect to a different application type
- You can change the local port by changing 56789 to any other valid port number (i.e. connecting to multiple instances at the same time)
- Login with the windows credentials discussed above in the format NETBIOSDOMAIN\User1 (i.e. example\user1)
- Your Netbios domain is found here in your config file: "netbios-domain": "example",
- Connect to your desktop command line to command line interface of remote Windows or Linux servers, instead of through console (i.e. no tunnel):
- `aws ssm start-session --target "i-090c25e64c2d9d276""--region ca-central-1`
- Replace i-xxx with your instance ID
- Everything is fully logged for future reference
- If you want to remove the region from your command line, you can:
- Type: "aws configure" from command prompt, hit {enter} (key), {enter} (secret), enter: ca-central-1, {enter}

1.1.14. I ran the state machine but it failed when it tried to delete the default VPC?

I ran the state machine but it failed when it tried to delete the default VPC? The state machine cannot delete the default VPC (Error: VPC has dependencies and cannot be deleted)

You need to ensure that resources don't exist in the default VPC or else the state machine won't be able to delete it. If you encounter this error, you can either delete the resources within the VPC or delete the default VPC manually and run the state machine again.

3.2 1.2. Existing Accounts / Organizations

1.2.1. How do I import an existing AWS account into my Accelerator managed AWS Organization (or what if I created a new AWS account with a different Organization trust role)?

How do I import an existing AWS account into my Accelerator managed AWS Organization (or what if I created a new AWS account with a different Organization trust role)?

- Ensure you have valid administrative privileges for the account to be invited/added
- Add the account to your AWS Organization using standard processes (i.e. Invite/Accept)
- this process does NOT create an organization trust role
- imported accounts do NOT have the quarantine SCP applied as we don't want to break existing workloads
- Login to the account using the existing administrative credentials
- Execute the Accelerator provided CloudFormation template to create the required Accelerator bootstrapping role - in the GitHub repo here: `reference-artifacts\Custom-Scripts\Import-Account-CFN-Role-Template.yml`
- add the account to the Accelerator config file and run the state machine
- If you simply created the account with an incorrect role name, you likely need to take extra steps:
- Update the Accelerator config file to add the parameter: `global-options\ignored-ous = ["UnManagedAccounts"]`
- In AWS Organizations, create a new OU named `UnManagedAccounts` (case sensitive)
- Move the account to the `UnManagedAccounts` OU
- You can now remove the Quarantine SCP from the account
- Assume an administrative role into the account
- Execute the Accelerator provided CloudFormation template to create the required Accelerator bootstrapping role

1.2.2. Is it possible to deploy the Accelerator on top of an AWS Organization that I have already installed the AWS Landing Zone (ALZ) solution into?

Is it possible to deploy the Accelerator on top of an AWS Organization that I have already installed the AWS Landing Zone (ALZ) solution into?

Existing ALZ customers are required to uninstall their ALZ deployment before deploying the Accelerator. Please work with your AWS account team to find the best mechanism to uninstall the ALZ solution (procedures and scripts exist). It is often easier to migrate AWS accounts to a new Accelerator Organization, per the process detailed in the next FAQ question. Additionally, please reference the following [section](#) of the Installation Guide for additional considerations.

1.2.3. What if I want to move an account from an AWS Organization that has the ALZ deployed into an AWS Organization running the Accelerator?

What if I want to move an account from an AWS Organization that has the ALZ deployed into an AWS Organization running the Accelerator?

Before removing the AWS account from the source organization, terminate the AWS Service Catalog product associated with the member account that you're interested in moving. Ensuring the product terminates successfully and that there aren't any remaining CloudFormation stacks in the account that were deployed by the ALZ. You can then remove the account from the existing Organization and invite it into the new organization. Accounts invited into the Organization do NOT get the `Deny ALL` SCP applied, as we do not want to break existing running workloads. Moving the newly invited account into its destination OU will trigger the state machine and result in the account being ingested into the Accelerator and having the guardrails applied per the target OU persona.

For a detailed procedure, please review this [document](#).

3.3 1.3. End User Environment

1.3.1. Is there anything my end users need to be aware of? Why do some of my end users struggle with CloudWatch Log groups errors?

Is there anything my end users need to be aware of? Why do some of my end users struggle with CloudWatch Log groups errors?

CloudWatch Log group deletion is prevented for security purposes and bypassing this rule would be a fundamental violation of security best practices. This protection does NOT exist solely to protect ASEA logs, but ALL log groups. Users of the Accelerator environment will need to ensure they set CloudFormation stack Log group retention type to RETAIN, or stack deletes will fail when attempting to delete a stack (as deleting the log group will be blocked) and users will encounter errors. As repeated stack deployments will be prevented from recreating the same log group name (as it already exists), end users will either need to check for the existence of the log group before attempting creation, or include a random hash in the log group name. The Accelerator also sets log group retention for all log groups to value(s) specified by customers in the config file and prevents end users from setting or changing Log group retentions. When creating new log groups, end users must either *not* configure a retention period, or set it to the default `NEVER expire` or they will also be blocked from creating the CloudWatch Log group. If applied by bypassing the guardrails, customer specified retention periods on log group creation will be overridden with the Accelerator specified retention period.

While a security best practice, some end users continue to request this be changed, but you need to ask: Are end users allowed to go in and clean out logs from Windows Event Viewer (locally or on domain controllers) after testing? Clean out Linux kernel logs? Apache log histories? The fundamental principal is that all and as many logs as possible will be retained for a defined retention period (some longer). In the "old days", logs were hidden deep within OS directory structures or access restricted by IT from developers - now that we make them all centralized, visible, and accessible, end users seem to think they suddenly need to clean them up. Customers need to establish a usable and scalable log group naming standard/convention as the first step in moving past this concern, such that they can always find their active logs easily. As stated, to enable repeated install and removal of stacks during test cycles, end user CloudFormation stacks need to set log groups to RETAIN and leverage a random hash in log group naming (or check for existence, before creating).

The Accelerator provided SCPs (guardrails/protections) are our recommendations, yet designed to be fully customizable, enabling any customer to carefully override these defaults to meet their individual requirements. If insistent, we'd suggest only bypassing the policy on the Sandbox OU, and only for log groups that start with a very specific prefix (not all log groups). When a customer wants to use the delete capability, they would need to name their log group with the designated prefix - i.e. opt-in to allow CloudWatch log group deletes.

1.3.2. How can I leverage Accelerator deployed objects in my IaC? Do I need to manually determine the arn's and object id's of Accelerator deployed objects to leverage them in my IaC?

How can I leverage Accelerator deployed objects in my IaC? Do I need to manually determine the arn's and object id's of Accelerator deployed objects to leverage them in my IaC?

Objects deployed by the Accelerator which customers may need to leverage in their own IaC have been populated in parameters in AWS parameter store for use by the IaC tooling of choice. The Accelerator ensures parameters are deployed consistently across accounts and OUs, such that a customer's code does not need to be updated when it is moved between accounts or promoted from Dev to Test to Prod.

Objects of the following types and their associated values are stored in parameter store: VPC, subnet, security group, ELB (ALB/NLB w/DNS address), IAM policy, IAM role, KMS key, ACM cert, SNS topic, and the firewall replacement variables.

Additionally, setting "populate-all-elbs-in-param-store": true for an account will populate all Accelerator wide ELB information into parameter store within that account. The sample PBMM configuration files set this value on the perimeter account, such that ELB information is available to configure centralized ingress capabilities.

1.3.3. How do I deploy AWS Elastic Beanstalk instances?

How do I deploy AWS Elastic Beanstalk instances?

If your deployed environment contains an SCP enforcing volume encryption of EC2 instances, your Elastic Beanstalk deployment will fail.

The SCP will contain an entry like this:

```
{
  "Sid": "EBS1",
  "Effect": "Deny",
  "Action": "ec2:RunInstances",
  "Resource": "arn:aws:ec2:*:*:volume/*",
  "Condition": {
    "Bool": {
      "ec2:Encrypted": "false"
    }
  }
},
```

A solution is to encrypt the root volume of the AMI that Elastic Beanstalk uses for your selected platform, and perform a custom AMI deployment of your Elastic Beanstalk application.

You can gather the AMI that Elastic Beanstalk uses via CLI with the following command:

```
aws elasticbeanstalk describe-platform-version --region <YOUR_REGION> --platform-arn <ARN_EB_PLATFORM>
```

Once you have gathered the AMI ID successfully, go to the EC2 console and:

- Click on the 'AMIs' option in the left navigation pane
- Search for your AMI after selecting 'Public Images' from the dropdown list.
- Select the AMI
- Go to Actions and Copy AMI
- Click on the checkbox to enable 'Encryption' and then select "Copy AMI".

Once the AMI is successfully copied, you can use this AMI to specify a custom AMI in your Elastic Beanstalk environments with root volume encrypted.

3.4 1.4. Upgrades

1.4.1. Can I upgrade directly to the latest release, or must I perform upgrades sequentially?

Can I upgrade directly to the latest release, or must I perform upgrades sequentially?

Yes, currently customers can upgrade from whatever version they have deployed to the latest Accelerator version. There is no requirement to perform sequential upgrades. In fact, we strongly discourage sequential upgrades.

Given the magnitude of the v1.5.0 release, we have added a one-time requirement that all customers upgrade to a minimum of v1.3.8 before attempting to upgrade to v1.5.0.

1.4.2. Why do I get the error "There were errors while comparing the configuration changes:" when I update the config file?

Why do I get the error "There were errors while comparing the configuration changes:" when I update the config file?

In v1.3.0 we added protections to allow customers to verify the scope of impact of their intended changes to the configuration file. In v1.3.0 and above, the state machine does not allow changes to the config file (other than new accounts) without providing the `scope` parameter. Please refer to the [State Machine behavior and inputs Guide](#) for more details.

3.5 1.5. Support Concerns

1.5.1. The Accelerator is written in CDK and deploys CloudFormation, does this restrict the Infrastructure as Code (IaC) tools that I can use?

The Accelerator is written in CDK and deploys CloudFormation, does this restrict the Infrastructure as Code (IaC) tools that I can use?

No. Customers can choose the IaC framework or tooling of their choice. The tooling used to deploy the Accelerator has no impact on the automation framework customers use to deploy their applications within the Accelerator environment. It should be noted that the functionality deployed by the Accelerator is extremely platform specific and would not benefit from multi-platform IaC frameworks or tooling.

1.5.2. What happens if AWS stops enhancing the Accelerator?

What happens if AWS stops enhancing the Accelerator?

The Accelerator is an open source project, should AWS stop enhancing the solution for any reason, the community has access to the full codebase, its roadmap and history. The community can enhance, update, fork and take ownership of the project, as appropriate.

The Accelerator is an AWS CDK based project and synthesizes to native AWS CloudFormation. AWS sub-accounts simply contain native CloudFormation stacks and associated custom resources, when required. The Accelerator architecture is such that all CloudFormation stacks are native to each AWS account with no links or ties to code in other AWS accounts or even other stacks within the same AWS account. This was an important initial design decision.

The Accelerator codebase can be completely uninstalled from the organization management (root) account, without any impact to the deployed functionality or guardrails. In this situation, guardrail updates and new account provisioning reverts to a manual process. Should a customer decide they no longer wish to utilize the solution, they can remove the Accelerator codebase without any impact to deployed resources and go back to doing things natively in AWS as they did before they deployed the Accelerator. By adopting the Accelerator, customers are not locking themselves in or making a one-way door decision.

1.5.3. What level of Support will the ASEA have from AWS Support?

What level of Support will the ASEA have from AWS Support?

The majority of the solution leverages native AWS services which are fully supported by AWS Support. Additionally, the Accelerator is an AWS CDK based project and synthesizes to native AWS CloudFormation. AWS sub-accounts simply contain native CloudFormation stacks and associated custom resources (when required). The Accelerator architecture is such that all CloudFormation stacks are native to each AWS account with no direct links or ties to code in other AWS accounts (no stacksets, no local CDK). This was an important project design decision, keeping deployed functionality in independent local CloudFormation stacks and decoupled from solution code, which allows AWS support to effectively troubleshoot and diagnose issues local to the sub-account.

As the Accelerator also includes code, anything specifically related to the Accelerator codebase will be only supported on a "best effort" basis by AWS support, as AWS support does not support custom code. The first line of support for the codebase is typically your local AWS team (your SA, TAM, ProServe and/or AWS Partner). As an open source project, customers can file requests using GitHub Issues against the Accelerator repository or open a discussion in GitHub discussions. Most customer issues arise during installation and are related to configuration customization or during the upgrade process.

1.5.4. What does it take to support the Accelerator?

What does it take to support the Accelerator?

We advise customers to allocate a 1/2 day per quarter to upgrade to the latest Accelerator release.

Customers have indicated that deploying the Accelerator reduces their ongoing operational burden over operating in native AWS, saving hours of effort every time a new account is provisioned by automating the deployment of the persona associated with new accounts (guardrails, networking and security). The Accelerator does NOT alleviate a customer's requirement to learn to effectively operate in the cloud (like monitoring security tooling/carrying out Security Operation Center (SOC) duties). This effort exists regardless of the existence of the Accelerator.

1.5.5. Is the Accelerator only designed and suitable for Government of Canada or PBMM customers?

Is the Accelerator only designed and suitable for Government of Canada or PBMM customers?

No. The Accelerator is targeted at **any AWS customer** that is looking to automate the deployment and management of a comprehensive end-to-end multi-account environment in AWS. It is ideally suited for customers interested in achieving a high security posture in AWS.

The Accelerator is a sophisticated deployment framework that allows for the deployment and management of virtually any AWS multi-account "Landing Zone" architecture without any code modifications. The Accelerator is actually delivering two separate and distinct products which can each be used on their own:

1. the Accelerator the tool, which can deploy virtually any architecture based on a provided config file (no code changes), and;
2. the Government of Canada (GC) prescriptive PBMM architecture which is delivered as a sample configuration file and documentation.

The tooling was purposely built to be extremely flexible, as we realized that some customers may not like some of the opinionated and prescriptive design decisions we made in the GC architecture. Virtually every feature being deployed can be turned on/off, not be used or can have its configuration adjusted to meet your specific design requirements.

We are working on building a library of sample config files to support additional customer needs and better demonstrate product capabilities and different architecture patterns. In no way is it required that the prescriptive GC architecture be used or deployed. Just because we can deploy, for example, an AWS Managed Active Directory, does not mean you need to use that feature of the solution. Disabling or changing these capabilities also requires zero code changes.

While the prescriptive sample configuration files were originally developed based on GC requirements, they were also developed following AWS Best Practices. Additionally, many security frameworks around the world have similar and overlapping security requirements (you can only do security so many ways). The provided architecture is applicable to many security compliance regimes around the world and not just the GC.

3.6 1.6. Deployed Functionality

1.6.1. I wish to be in compliance with the 12 GC TBS Guardrails, what don't you cover with the provided sample architecture?

I wish to be in compliance with the 12 GC TBS Guardrails, what don't you cover with the provided sample architecture?

The AWS SEA allows for a lot of flexibility in deployed architectures. If used, the provided PBMM sample architecture was designed to help deliver on the technical portion of *all* 12 of the GC guardrails, when automation was possible.

What don't we cover? Assigning MFA to users is a manual process. Specifically, you need to procure Yubikeys for your root/break glass users, and enable a suitable form of MFA for *all* other users (i.e. virtual, email, other). The guardrails also include some organizational processes (i.e. break glass procedures, or signing an MOU with CCCS) which customers will need to work through independently.

While AWS is providing the tools to help customer be compliant with the 12 PBMM guardrails (which were developed in collaboration with the GC) - it's up to each customer's ITSec organization to assess and determine if the deployed controls actually meet their security requirements.

Finally, while we started with a goal of delivering on the 12 guardrails, we believe we have extended well beyond those security controls, to further help customers move towards meeting the full PBMM technical control profile (official documentation is weak in this area at this time).

1.6.2. Does the ALB perform SSL offloading?

Does the ALB perform SSL offloading?

As configured - the perimeter ALB decrypts incoming traffic using its certificate and then re-encrypts it with the certificate for the back-end ALB. The front-end and back-end ALB's can use the same or different certs. If the Firewall needs to inspect the traffic, it also needs the backend certificate be manually installed.

1.6.3. What is the recommended approach to manage the ALB certificates deployed by the Accelerator?

What is the recommended approach to manage the ALB certificates deployed by the Accelerator?

The Accelerator installation process allows customers to provide their own certificates (either self-signed or generated by a CA), to enable quick and easy installation and allowing customers to test end-to-end traffic flows. After the initial installation, we recommend customers leverage AWS Certificate Manager (ACM) to easily provision, manage, and deploy public and private SSL/TLS certificates. ACM helps manage the challenges of maintaining certificates, including certificate rotation and renewal, so you don't have to worry about expiring certificates.

The Accelerator provides 3 mechanisms to enable utilizing certificates with ALB's:

- **Method 1** - IMPORT a certificate into AWS Certificate Manager from a 3rd party product
- When using a certificate that does not have a certificate chain (usually this is the case with Self-Signed)

```
"certificates": [
  {
    "name": "My-Cert",
    "type": "import",
    "priv-key": "certs/example1-cert.key",
    "cert": "certs/example1-cert.crt"
  }
]
```

- When using a certificate that has a certificate chain (usually this is the case when signed by a Certificate Authority with a CA Bundle)

```
"certificates": [
  {
    "name": "My-Cert",
    "type": "import",
    "priv-key": "certs/example1-cert.key",
    "cert": "certs/example1-cert.crt",
    "chain": "certs/example1-cert.chain"
  }
]
```

- this mechanism allows a customer to generate certificates using their existing tools and processes and import 3rd party certificates into AWS Certificate Manager for use in AWS
- Self-Signed certificates should NOT be used for production (samples were provided simply to demonstrate functionality)
- both a `.key` and a `.crt` file must be supplied in the customers S3 input bucket
- "cert" must contain only the certificate and not the full chain
- "chain" is an optional attribute that contains the certificate chain. This is generally used when importing a CA signed certificate
- this will create a certificate in ACM and a secret in secrets manager named `accelerator/certificates/My-Cert` in the specified AWS account(s), which points to the newly imported certificates ARN

- **Method 2** - REQUEST AWS Certificate Manager generate a certificate

```
"certificates": [
  {
    "name": "My-Cert",
    "type": "request",
    "domain": "*.example.com",
    "validation": "DNS",
    "san": ["www.example.com"]
  }
]
```

- this mechanism allows a customer to generate new public certificates directly in ACM
- both `DNS` and `EMAIL` validation mechanisms are supported (DNS recommended)
- this requires a **Public** DNS zone be properly configured to validate you are legally entitled to issue certificates for the domain
- this will also create a certificate in ACM and a secret in secrets manager named `accelerator/certificates/My-Cert` in the specified AWS account(s), which points to the newly imported certificates ARN
- this mechanism should NOT be used on new installs, skip certificate and ALB deployment during initial deployment (removing them from the config file) and simply add on a subsequent state machine execution
- Process:
 - you need a public DNS domain properly registered and configured to publicly resolve the domain(s) you will be generating certificates for (i.e. example.com)
 - domains can be purchased and configured in Amazon Route53 or through any 3rd party registrar and DNS service provider
 - in Accelerator phase 1, the cert is generated, but the stack does NOT complete deploying (i.e. it waits) until certificate validation is complete
 - during deployment, go to the AWS account in question, open ACM and the newly requested certificate. Document the authorization CNAME record required to validate certificate generation
 - add the CNAME record to the zone in bullet 1 (in Route53 or 3rd party DNS provider) (documented [here](#))
 - after a few minutes the certificate will validate and switch to `Issued` status
 - Accelerator phase 1 will finish (as long as the certificate is validated before the Phase 1 credentials time-out after 60-minutes)
 - the ALB will deploy in a later phase with the specified certificate
- **Method 3** - Manually generate a certificate in ACM
 - this mechanism allows a customer to manually generate certificates directly in the ACM interface for use by the Accelerator
 - this mechanism should NOT be used on new installs, skip certificate and ALB deployment during initial deployment (removing them from the config file) and simply add on a subsequent state machine execution
 - Process:
 - go to the AWS account for which you plan to deploy an ALB and open ACM
 - generate a certificate, documenting the certificates ARN
 - open Secrets manager and generate a new secret of the format `accelerator/certificates/My-Cert` (of type `Plaintext` under `Other type of secrets`), where `My-Cert` is the unique name you will use to reference this certificate
 - In all three mechanisms a secret will exist in Secrets Manager named `accelerator/certificates/My-Cert` which contains the ARN of the certificate to be used.

- In the Accelerator config file, find the definition of the ALB for that AWS account and specify `My-Cert` for the ALB `cert-name`

```
"alb": [
  {
    "cert-name": "My-Cert"
  }
]
```

- The state machine will fail if you specify a certificate in any ALB which is not defined in Secrets Manager in the local account.

We suggest the most effective mechanism for leveraging ACM is by adding CNAME authorization records to the relevant DNS domains using Method 2, but may not appropriate right for all customers.

1.6.4. Why do we have rsyslog servers? I thought everything was sent to CloudWatch?

Why do we have rsyslog servers? I thought everything was sent to CloudWatch?

The rsyslog servers are included to accept logs for appliances and third party applications that do not natively support the CloudWatch Agent from any account within a customers Organization. These logs are then immediately forwarded to CloudWatch Logs within the account the rsyslog servers are deployed (Operations) and are also copied to the S3 immutable bucket in the log-archive account. Logs are only persisted on the rsyslog hosts for 24 hours. The rsyslog servers are required to centralize the 3rd party firewall logs (Fortinet Fortigate).

1.6.5. Can you deploy the solution without Fortinet Firewall Licenses?

Can you deploy the solution without Fortinet Firewall Licenses?

Yes, if license files are not provided, the firewalls will come up configured and route traffic, but customers will have no mechanism to manage the firewalls/change the configuration until a valid license file is added. If invalid licence files are provided, the firewalls will fail to load the provided configuration, will not enable routing, will not bring up the VPN tunnels and will not be manageable. Customers will need to either remove and redeploy the firewalls, or manually configure them. If performing a test deployment, please work with your local Fortinet account team to discuss any options for temporary evaluation licenses.

Additionally, several additional firewall options are now available, including using AWS Network Firewall, a native AWS service.

1.6.6. I installed additional software on my Accelerator deployed RDGW / rsyslog host, where did it go?

I installed additional software on my Accelerator deployed RDGW / rsyslog host, where did it go?

The RDGW and rsyslog hosts are members of auto-scaling groups. These auto-scaling groups have been configured to refresh instances in the pool on a regular basis (7-days in the current sample config files). This ensures these instances are always clean. Additionally, on every execution of the Accelerator state machine the ASG are updated to the latest AWS AMI for the instances. When the auto-scaling group refreshes its instances, they will be redeployed with the latest patch release of the AMI/OS. It is recommended that the state machine be executed monthly to ensure the latest AMI's are always in use.

Customers wanting to install additional software on these instances should either a) update the automated deployment scripts to install the new software on new instance launch, or b) create and specify a custom AMI in the Accelerator configuration file which has the software pre-installed ensuring they are also managing patch compliance on the instance through some other mechanism.

At any time, customers can terminate the RDGW or rsyslog hosts and they will automatically be re-created from the base images with the latest patch available at the time of the last Accelerator State Machine execution.

1.6.7. Some sample configurations provide NACLs and Security Groups. Is that enough?

Some sample configurations provide NACLs and Security Groups. Is that enough?

Security group egress rules are often used in 'allow all' mode (0.0.0.0/0), with the focus primarily being on consistently allow listing required ingress traffic (centralized ingress/egress controls are in-place using the perimeter firewalls). This ensures day to day activities like patching, access to DNS, or to directory services access can function on instances without friction.

The Accelerator provided sample security groups in the workload accounts offer a good balance that considers both security, ease of operations, and frictionless development. They allow developers to focus on developing, enabling them to simply use the pre-created security constructs for their workloads, and avoid the creation of wide-open security groups. Developers can equally choose to create more appropriate least-privilege security groups more suitable for their application, if they are skilled in this area. It is expected as an application is promoted through the SDLC cycle from Dev through Test to Prod, these security groups will be further refined by the extended customers teams to further reduce privilege, as appropriate. It is expected that each customer will review and tailor their Security Groups based on their own security requirements. The provided security groups ensures day to day activities like patching, access to DNS, or to directory services access can function on instances without friction, with the understanding further protections are providing by the central ingress/egress firewalls.

The use of NACLs are general discouraged, but leveraged in this architecture as a defense-in-depth mechanism. Security groups should be used as the primary access control mechanism. As with security groups, we encourage customers to review and tailor their NACLs based on their own security requirements.

1.6.8. Can I deploy the solution as the account root user?

Can I deploy the solution as the account root user?

No, you cannot install as the root user. The root user has no ability to assume roles which is a requirement to configure the sub-accounts and will prevent the deployment. As per the [installation instructions](#), you require an IAM user with the `AdministratorAccess` policy attached.

1.6.9. Is the Organizational Management root account monitored similarly to the other accounts in the organization?

Is the Organizational Management root account monitored similarly to the other accounts in the organization?

Yes, all accounts including the Organization Management or root account have the same monitoring and logging services enabled. When supported, AWS security services like GuardDuty, Macie, and Security Hub have their delegated administrator account configured as the "security" account. These tools can be used within each local account (including the Organization Management account) within the organization to gain account level visibility or within the Security account for Organization wide visibility. For more information about monitoring and logging refer to [architecture documentation](#).

1.6.10. How are the perimeter firewall configurations and licensing managed after deployment?

How are the perimeter firewall configurations and licensing managed after deployment?

While you deploy the perimeter firewalls with the Accelerator you will continue to manage firewall updates, configuration changes, and license renewals from the respective firewall management interface and not from the Accelerator config file. As these changes are not managed by the Accelerator you do not need to rerun the state machine to implement or track any of these changes. You can update the AMI of the 3rd party firewalls using the Accelerator, you must first remove the existing firewalls and redeploy them (as the Elastic IP's (EIP's) will block a parallel deployment) or deploy a second parallel firewall cluster and de-provision the first cluster when ready.



1.6.11. Can the Fortinet Firewall deployments use static private IP address assignments?

Can the Fortinet Firewall deployments use static private IP address assignments?

Yes, the `"port"` stanza in the configuration file can support a private static IP address assignment from the AZ and subnet. Care must be exercised to assure the assigned IP address is within the correct subnet and availability zone. Consideration must also be given to the Amazon reserved IP addresses (first three addresses, and the last) within subnets when choosing an IP Address to assign.

Using the `config.example.json` as a reference, static IP Assignments would look like this in the `ports:` stanza of the firewall deployment.

```
"ports": [
  {
    "name": "Public",
    "subnet": "Public",
    "create-eip": true,
    "create-cgw": true,
    "private-ips": [
      {
        "az": "a",
        "ip": "100.96.250.4"
      },
      {
        "az": "b",
        "ip": "100.96.250.132"
      }
    ]
  },
  {
    "name": "OnPremise",
    "subnet": "OnPremise",
    "create-eip": false,
    "create-cgw": false,
    "private-ips": [
      {
        "az": "a",
        "ip": "100.96.250.68"
      },
      {
        "az": "b",
        "ip": "100.96.250.196"
      }
    ]
  }
  ...
],
```

Where `private-ips` are not present for the subnet or availability zone an address will be assigned automatically from available addresses when the firewall instance is created.

1.6.12. I've noticed CloudTrail logs and in certain situation VPC flow logs are stored in the centralized log-archive account logging bucket twice?

I've noticed CloudTrail logs and in certain situation VPC flow logs are stored in the centralized log-archive account logging bucket twice?

Yes. CloudTrail is configured to send its logs directly to S3 for centralized immutable log retention. CloudTrail is also configured to send its logs to a centralized Organizational CloudWatch Log group such that the trail can be a) easily queried online using CloudWatch Insights across all AWS accounts in the organization, and b) to enable alerting based on undesirable API activity using CloudWatch Metrics and Alarms. All CloudWatch Log groups are also configured to be sent, using Amazon Kinesis, to S3 for centralized immutable log retention.

VPC flow log destinations can be configured in the config file. The example config files are set to send the VPC flow logs to both S3 and CloudWatch Logs by default for the same reasons as CloudTrail.

To reduce the duplicate long-term storage of these two specific CloudWatch Log types, customers can set `cwl-glbl-exclusions` under `central-log-services` to: `["/${ACCELERATOR_PREFIX_ND}/flowlogs/*", "${ACCELERATOR_PREFIX_ND}/CloudTrail*"]` to prevent these specifically named log groups from being stored on S3. This setting also prevents the Accelerator from setting the customer desired log group retention period defined in the config file, once implemented, for those log groups. Therefore, we do not recommend this exception be applied during the initial installation, as the retention setting on these CWL groups will remain the default (infinite). If `cwl-glbl-exclusions` is set after initial install, the defined retention will be configured during install and will remain set to the value present when the exception was applied to those log groups. This allows logs to be stored in CloudWatch Logs for quick and easy online access (short-retention only), and stored in S3 for long-term retention and access.

Side note: CloudTrail S3 data plane logs are enabled at the Organizational level, meaning all S3 bucket access is logged. As CloudTrail is writing to a bucket within the Organization, CloudTrail itself is accessing the bucket, seemingly creating a cyclical loop. As CloudTrail writes to S3 in 5-10min batches, CloudTrail will actually only cause one extra log 'entry' every 5-10minutes and not per S3 event, mitigating major concerns. Today, with an Organization trail logging data plane events for all buckets - there is no way to exclude any one bucket. But - having clear view of who accessed/changed logs, including AWS services, is important.

1.6.13. I need a Route53 Private Hosted Zone in my workload account. How shall I proceed?

I need a Route53 Private Hosted Zone in my workload account. How shall I proceed?

The workload account requires creating a temporary local VPC before creating the Private Hosted Zone (PHZ). Creating a PHZ in Route53 requires association with a VPC. You cannot specify a shared VPC when creating the PHZ, hence the need for this workaround.

Create the temporary workload account VPC

You can create the temporary VPC during AWS account creation via the ASEA config (preferred way). Insert the "vpc" JSON object like shown below when using the ASEA config to create an AWS account.

If you don't use the ASEA config you will need to assume the proper ASEA elevated IAM role in the workload account in order to create the VPC manually.

```
"mydevacct": {
  "account-name": "MyDev1",
  "email": "dev1-main@super-corp.co",
  "src-filename": "config.json",
  "ou": "dev",
  "vpc": [
    {
      "deploy": "local",
      "name": "Local",
      "description": "This VPC Temp VPC to create the local hosted zone.",
      "cidr-src": "provided",
      "cidr": [
        {
          "value": "192.168.100.0/24"
        }
      ],
      "region": "${HOME_REGION}"
    }
  ]
}
```

Create in the workload account a Private Hosted Zone

Using an IAM role assumed in the workload account:

List the VPCs.

```
aws ec2 describe-vpcs
```

Then retrieve the VpcId attribute for the newly created VPC as well as the Id for the shared VPC.

Create the Private Hosted Zone

```
aws route53 create-hosted-zone --name <MY_DOMAIN> --hosted-zone-config PrivateZone=true --vpc VPCRegion=<VPC_REGION>,VPCId=<VPC_ID> --caller-reference <YOUR_REFERENCE_ID>
```

Insert the proper values for:

- <MY_DOMAIN>
- <VPC_REGION>
- <VPC_ID> (id of new the local VPC)
- <YOUR_REFERENCE_ID> (can be any value)

Take note of the newly created hosted zone id by looking at the output of the command. The Id is the value after / hostedzone/ from the Id attribute. For example, the value is Z0123456NWOWQ4HNN40U from "Id": "/hostedzone/Z0123456NWOWQ4HNN40U".

Create an authorization to associate with this new zone

While still in the workload account; you need to create an association request authorization to allow the shared VPC to associate with this newly created Route53 PHZ.

```
aws route53 create-vpc-association-authorization --hosted-zone-id <ZONE_ID> --vpc VPCRegion=<SHARED_VPC_REGION>,VPCId=<SHARED_VPC_ID>
```

Insert the proper values for:

- <ZONE_ID>
- <SHARED_VPC_REGION>
- <SHARED_VPC_ID>

Confirm the association request for the shared VPC

After switching to an IAM role in the SharedNetwork account associate the Private Hosted Zone from the workload account.

```
aws route53 associate-vpc-with-hosted-zone --hosted-zone-id <ZONE_ID> --vpc VPCRegion=<SHARED_VPC_REGION>,VPCId=<SHARED_VPC_ID>
```

Insert the proper values for:

- <ZONE_ID>
- <SHARED_VPC_REGION>
- <SHARED_VPC_ID>

Validate Association and clean-up

Back in the workload account and assuming its IAM role, validate the association using the below command. You should see two VPCs attached. The local VPC and the shared VPC.

```
aws route53 get-hosted-zone --id <ZONE_ID>
```

Insert the proper values for:

- <ZONE_ID>

You can now dissociate the local VPC from the zone.

```
aws route53 disassociate-vpc-from-hosted-zone --hosted-zone-id <ZONE_ID> --vpc VPCRegion=<VPC_REGION>,VPCId=<VPC_ID>
```


Insert the proper values for:

- <ZONE_ID>
- <VPC_REGION>
- <VPC_ID>

You can now delete the local VPC. We recommend you leverage the ASEA configuration file. Simply remove the `vpc` section from the workload account:

```
"mydevacct": {
  "account-name": "MyDev1",
  "email": "dev1-main@super-corp.co",
  "src-filename": "config.json",
  "ou": "dev"
}
```

and rerun the State Machine.

1.6.14. How do I create a role which has read access to the log-archive bucket to enabling log forwarding to my favorite SIEM solution?

How do I create a role which has read access to the log-archive bucket to enabling log forwarding to my favorite SIEM solution?

You can update the ASEA config file to provision an IAM role that has cross-account access to the Log Archive S3 Buckets. Attempting to do this outside the ASEA config file is blocked by security guardrails. Additionally, even if the guardrails are bypassed, it is likely the ASEA will revert any manual changes on subsequent State Machine executions. The below example creates a Lambda role which is provided permissions to Amazon OpenSearch, S3 Read Only, Lambda VPC Execution, the Log Archive S3 buckets and the KMS key. Update the below example with the least-privilege policies needed to meet the requirements of your chosen SIEM solution.

The primary trick, is the use of the `"ssm-log-archive-read-only-access": true` flag.

As we generally recommend the SIEM be deployed into the Operations account, add the following to the roles array within the **Operations** account section in the ASEA config file:

```
{
  "role": "SIEM-Lambda-Processor",
  "type": "lambda",
  "ssm-log-archive-read-only-access": true,
  "policies": [
    "AmazonOpenSearchServiceFullAccess",
    "service-role/AWSLambdaVPCAccessExecutionRole",
    "AmazonS3ReadOnlyAccess"
  ],
  "boundary-policy": "Default-Boundary-Policy"
}
```

1.6.15. How do I create a role for use by Azure Sentinel using the new S3 Connector method?

How do I create a role for use by Azure Sentinel using the new S3 Connector method?

This process is very similar to FAQ #1.6.14, except we need to allow for a cross-cloud role assumption. This will be done in the Log Archive account, instead of the Operations account.

The following config snippet should be added to the roles array within the **Log Archive** account section in the ASEA config file:

```
{
  "role": "MicrosoftSentinelRole",
  "type": "account",
  "ssm-log-archive-read-only-access": true,
  "policies": [
    "AmazonSQSReadOnlyAccess",
    "service-role/AWSLambdaSQSQueueExecutionRole",
    "AmazonS3ReadOnlyAccess"
  ],
  "boundary-policy": "Default-Boundary-Policy",
  "trust-policy": "sentinel-trust-policy.json",
  "source-account": "log-archive",
  "source-account-role": "OrganizationAccountAccessRole"
}
```

- The value of the `source-account-role` above needs to be replaced with the value of `organization-admin-role` from your config file (`OrganizationAccountAccessRole`, `AWSCloudFormationStackSetExecutionRole`, or `AWSControlTowerExecution`).

The above role uses a custom trust policy, and also requires a file of the name `sentinel-trust-policy.json` be placed into the `iam-policy` folder of the customers S3 input bucket. This file must contain the following text:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::197857026523:root"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "{CUSTOMER-VALUE-HERE}"
        }
      }
    }
  ]
}
```

- The IAM account number listed above is a value provided by Microsoft in their documentation (hard-coded to the same value for all customers).
- The value of `sts:ExternalId`, shown as `{CUSTOMER-VALUE-HERE}` above, must be replaced with the ID of the Log Analytics Workspace in your Azure tenant.
- This information is based on the requirements published [here](#) as of 2022-03-10.

1.6.16. Does the ASEA include a full SIEM solution?

Does the ASEA include a full SIEM solution?

We've found a diverse set of differing customer needs and requirements across our customer base. The ASEA:

- enables AWS security services like Amazon GuardDuty (a Cloud native IDS solution) and centralizes the consoles of these tools in the Security account;
- audits the entire environment for compliance and consolidates findings from AWS security services in the Security Hub console in the Security account;
- sends prioritized email alerts for Security Hub Findings, Firewall Manager alerts and customizable CloudWatch Alarms;
- centralizes logs across the environment in a central bucket in the Log Archive account;
- in addition, retains logs locally in CloudWatch Logs for simple query using CloudWatch Insights.

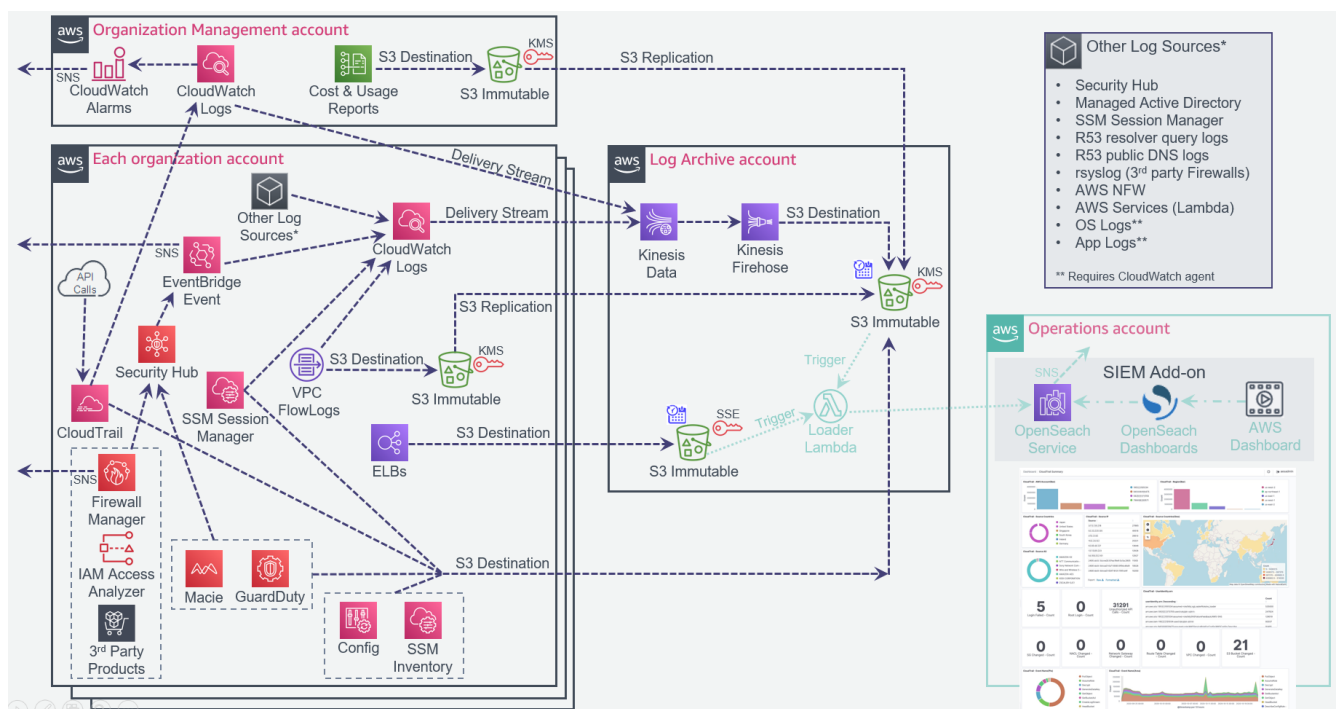
This makes it extremely simple to layer a customer's preferred SIEM solution on top of the ASEA, enabling easy consumption of the comprehensive set of collected logs and security findings.

Customers ask for examples of what this integration looks like. We've also had a number of customers ask for a reasonably functional and comprehensive open source SIEM-like solution to provide more advanced dashboarding, log correlation and search capabilities.

While not a part of the ASEA, we've made the [SIEM on Amazon OpenSearch Service](#) available as an ASEA **Add-on** to satisfy these requirements.

This independent solution can easily and quickly be deployed on top of the ASEA by following the documentation and using the scripts available [here](#). This process takes less than an hour.

The overall logging architecture is represented in this diagram:



1.6.17. Why are only select interface endpoints provisioned in the sample configuration files?

Why are only select interface endpoints provisioned in the sample configuration files?

For economic reasons, most of the sample configuration files only include the following minimum set of required interface endpoints:

"ec2", "ec2messages", "ssm", "ssmmessages", "secretsmanager", "cloudformation", "kms", "logs", "monitoring"

The full sample configuration file included all interface endpoints that existed in the Canada (Central) region at the time the configuration file was originally developed:

"access-analyzer", "acm-pca", "application-autoscaling", "appmesh-envoy-management", "athena", "autoscaling", "autoscaling-plans", "awsconnector", "cassandra", "clouddirectory", "cloudformation", "cloudtrail", "codebuild", "codecommit", "codepipeline", "config", "datasync", "ebs", "ec2", "ec2messages", "ecr.api", "ecr.dkr", "ecs", "ecs-agent", "ecs-telemetry", "elasticbeanstalk", "elasticbeanstalk-health", "elasticfilesystem", "elasticloadbalancing", "elasticmapreduce", "email-smtp", "events", "execute-api", "git-codecommit", "glue", "kinesis-firehose", "kinesis-streams", "kms", "license-manager", "logs", "macie2", "monitoring", "notebook", "sagemaker.api", "sagemaker.runtime", "secretsmanager", "servicecatalog", "sms", "sns", "sqs", "ssm", "ssmmessages", "states", "storagegateway", "sts", "synthetics", "transfer", "transfer.server", "workspaces"

Since that time these additional endpoints have been launched in the ca-central-1 region and can be optionally added to customer configuration files to make them accessible from private address space:

"airflow.api", "airflow.env", "airflow.ops", "app-integrations", "appstream.api", "appstream.streaming", "auditmanager", "backup", "backup-gateway", "batch", "cloudhsmv2", "codedeploy", "codedeploy-commands-secure", "codestar-connections.api", "comprehend", "comprehendmedical", "databrew", "dms", "elasticache", "emr-containers", "finspace", "finspace-api", "fis", "fsx", "greengrass", "imagebuilder", "inspector2", "iot.data", "iot.fleethub.api", "iotsitewise.api", "iotsitewise.data", "kendra", "lakeformation", "lambda", "memory-db", "mgn", "models-v2-lex", "nimble", "panorama", "profile", "qldb.session", "rds", "rds-data", "redshift", "redshift-data", "rekognition", "runtime-v2-lex", "sagemaker.featurestore-runtime", "securityhub", "servicecatalog-appregistry", "ssm-contacts", "ssm-incidents", "sync-states", "textextract", "transcribe", "transcribestreaming", "translate", "xray"

The aws.sagemaker.ca-central-1.studio interface endpoint was also launched, but cannot be auto-deployed by the Accelerator at this time as it does not utilize standardized naming and requires a code update to enable deployment.

Additional endpoints may exist in other AWS regions. Any endpoint can be added to any Accelerator configuration file, as long as it follows the standardized endpoint naming convention (e.g. com.amazonaws.{region}.{service}).

3.7 1.7. Network Architecture

1.7.1. We want to securely connect our on-premises networks/datacenters to our AWS Cloud PBMM tenancy, what does AWS you recommend?

We want to securely connect our on-premises networks/datacenters to our AWS Cloud PBMM tenancy, what does AWS you recommend?

We recommend customers create a new AWS sub-account in your organization in the Infrastructure OU to “own” the Direct Connect (DX), segregating Direct Connect management and billing from other organization activities. Once provisioned you would create a Public VIF on the DX in this account. You can also create additional Private VIF’s when and if required, and share them directly with any sub-account that needs to consume them.

We recommend customers then inter-connect directly to the Transit Gateway, in the Shared Network sub-account, from your on-premises network/datacenters.

- Initiate IPsec VPN tunnels from on-premises to the TGW using BGP w/ECMP to scale and balance the traffic. Equal Cost Multi-Pathing (ECMP) is used to balance the traffic across the available VPN tunnels.
- You need to create as many VPN attachments to the TGW as is required to meet your bandwidth requirements or DX capacity. Today IPsec attachments are limited to 1.25 Gbps each (10 Gbps would require 8 attachments) and is scalable to 50 Gbps.
- Each VPN attachment would comprise two tunnels (active/passive), each connecting to a different on-premises firewall/VPN appliance.

The VPN attachments would then be connected to an appropriately configured route table on the TGW. TGW route tables provide VRF like segregation capabilities, allowing customers to control which of their cloud based networks are allowed to communicate on-premises, or visa-versa.

This architecture is fully managed and easy to manage, highly available, scalable, cost effective, and enables customers to reserve all their 3rd party Perimeter firewall capacity for public or internet facing traffic.

(This guidance will be updated once MACSEC is broadly available across AWS transit centers)

1.7.2. Does this configuration violate PBMM / ITSG-22/38/33 principals?

Does this configuration violate PBMM / ITSG-22/38/33 principals?

No. Data center interconnects are not zoning boundaries (or ZIPs). Additionally, in many cases the on-premises VPN termination device used to interconnect to the cloud either contains, or is placed in-line with firewall and/or inspection devices. Customers insistent on placing a firewall between datacenters can enable the appropriate filtering or inspection on these on-premise devices. Enabling the same capabilities inside AWS would mean a customer is inspecting both ends of the same wire, a pointless activity. The TGW approach is being used by several gov’t PBMM customers.

Additionally, it should be noted that workloads in all the AWS accounts are fully protected using AWS Security Groups (stateful firewalls) wrapped around each and every instance comprising a workload.

1.7.3. Why do you NOT recommend using a VGW on the perimeter VPC?

Why do you NOT recommend using a VGW on the perimeter VPC?

The VGW solution was not designed to support an enterprise cloud environment – it was designed to provide single VPC connectivity. The VGW solution offers lower availability than other options as it relies on VPC route tables to steer traffic, which need to be updated using custom scripts in the event the failure of an appliance or availability zone. The VGW solution is typically harder to maintain and troubleshoot. The VGW solution has limited scalability, as the VGW only supports a single active connection and does not support BGP or ECMP (i.e. supports a maximum bandwidth of 1.25Gbps). Most customers providing enterprise cloud connectivity have switched away from this approach. This approach is highly discouraged.

1.7.4. Why do you NOT recommend connecting directly to the 3rd party firewall cluster in the perimeter account? (not GWLB, not NFW)

Why do you NOT recommend connecting directly to the 3rd party firewall cluster in the perimeter account? (not GWLB, not NFW)

This approach was common with AWS customers before the TGW was introduced, with many customers upgrading or considering upgrading to the TGW approach. We also have some customers using this architecture based on a very specific limitation of the customer's Direct Connect architecture, these customers would also like to migrate to the TGW approach, if they could.

While viable, this approach adds unneeded complexity, reduces cloud availability, is expensive to scale, and reduces bandwidth to internet facing workloads. This solution doubles the IPsec VPN tunnels using BGP w/ECMP requirements as it needs tunnels on both sides of the firewall. In this configuration each firewall appliance typically only provides a single pair of IPsec connections supporting marginally more bandwidth than the TGW VPN attachments. Adding tunnels and bandwidth requires adding firewall appliances. Stateful capabilities typically need to be disabled due to performance and asymmetric routing challenges. This typically means a very expensive device is being deployed inside AWS simply to terminate a VPN tunnel.

1.7.5. What if I really want to inspect this traffic inside AWS, but like the TGW architecture?

What if I really want to inspect this traffic inside AWS, but like the TGW architecture?

Customers who insist on inspecting the ground to cloud traffic inside AWS *can* do this with the proposed TGW architecture. The TGW route tables can be adjusted to hairpin the traffic through either a dedicated Inspection VPC, or to the Perimeter account firewall cluster for inspection. The Inspection VPC option could leverage 3rd party firewalls in an autoscaling group behind a Gateway Load Balancer, or leverage AWS network firewall to inspection traffic. To maximize internet throughput, the Inspection VPC option is generally recommended. While we do not feel inspection is needed in this situation, it is possible.

1.7.6. What does the traffic flow look like for an application running in a workload account?

What does the traffic flow look like for an application running in a workload account?


The perimeter (ingress/egress) account typically contains two ALB's, one for production workloads and another for Dev/Test workloads. The Dev/Test ALB should be locked to restrict access to on-premises users (using a security group) or have authentication enabled to prevent Dev/Test workloads from being exposed to the internet. Additionally, each workload account (Dev/Test/Prod) contains a local (back-end) ALB.

AWS Web Application Firewall (WAF) should be enabled on both front-end and back-end ALB's. The Front-end WAF would contain rate limiting, scaling and generic rules. The back-end WAF would contain workload specific rules (i.e. SQL injection). As WAF is essentially a temporary fix for broken applications before a developer can fix the issue, these rules typically require the close involvement of the application team. Rules can be centrally managed across all WAF instances using AWS Firewall Manager from the Security account.

The front-end ALB is then configured to target the back-end ALB using the process described in the [Post Installation](#) section of the installation guide, step 2 (Configure the new alb-forwarding feature (added in v1.5.0)). This enables configuring different DNS names and/or paths to different back-end ALB's using the ASEA's alb-forwarder. We recommend moving away from the NAT to DNS mechanism used in previous released as it was too complex, does not work with bump-in-the-wire inspection devices (NFW, GWLB), and only available on a limited number of 3rd party firewalls.

This implementation allows workload owners to have complete control of workloads in a local account including the ELB configuration, and allow site names and paths to be defined and setup at sub-account creation time (instead of during development) to enable publishing publicly or on-premises in a rapid agile manner.

This overall flow is depicted in this diagram:



NOTE1: Distinct route tables required per AZ, which targets the local AZ's nfw, gwlb or ngw endpoint

1.7.7. How does CloudFront and API Gateway fit with the answer from question 1.7.6?**How does CloudFront and API Gateway fit with the answer from question 1.7.6?**

The perimeter account is focused on protecting legacy IaaS based workloads. Cloud Native applications including CloudFront and API Gateway should be provisioned directly in the same account as the workload and should NOT traverse the perimeter account.

These services must still be appropriately configured. This includes ensuring both WAF and logging are enabled on each endpoint.

The GC guidance on Cloud First patterns and anti-patterns can be downloaded [here](#).

4. Operations & Troubleshooting

4.1 Accelerator Operations & Troubleshooting Guide

This document is targeted at individuals installing or executing the AWS Secure Environment Accelerator. It is intended to guide individuals who are executing the Accelerator by providing an understanding as to what happens at each point throughout execution and to assist in troubleshooting state machine failures and/or errors. This is one component of the provided documentation package and should be read after the Installation Guide, but before the Developer Guide.

- [System Overview](#)
- [Troubleshooting](#)
- [Common Tasks](#)

4.2 1. System Overview

This document is targeted at individuals installing or executing the AWS Secure Environment Accelerator. It is intended to guide individuals who are executing the Accelerator by providing an understanding as to what happens at each point throughout execution and to assist in troubleshooting state machine failures and/or errors. This is one component of the provided documentation package and should be read after the Installation Guide, but before the Developer Guide.

4.2.1 1.1. Overview

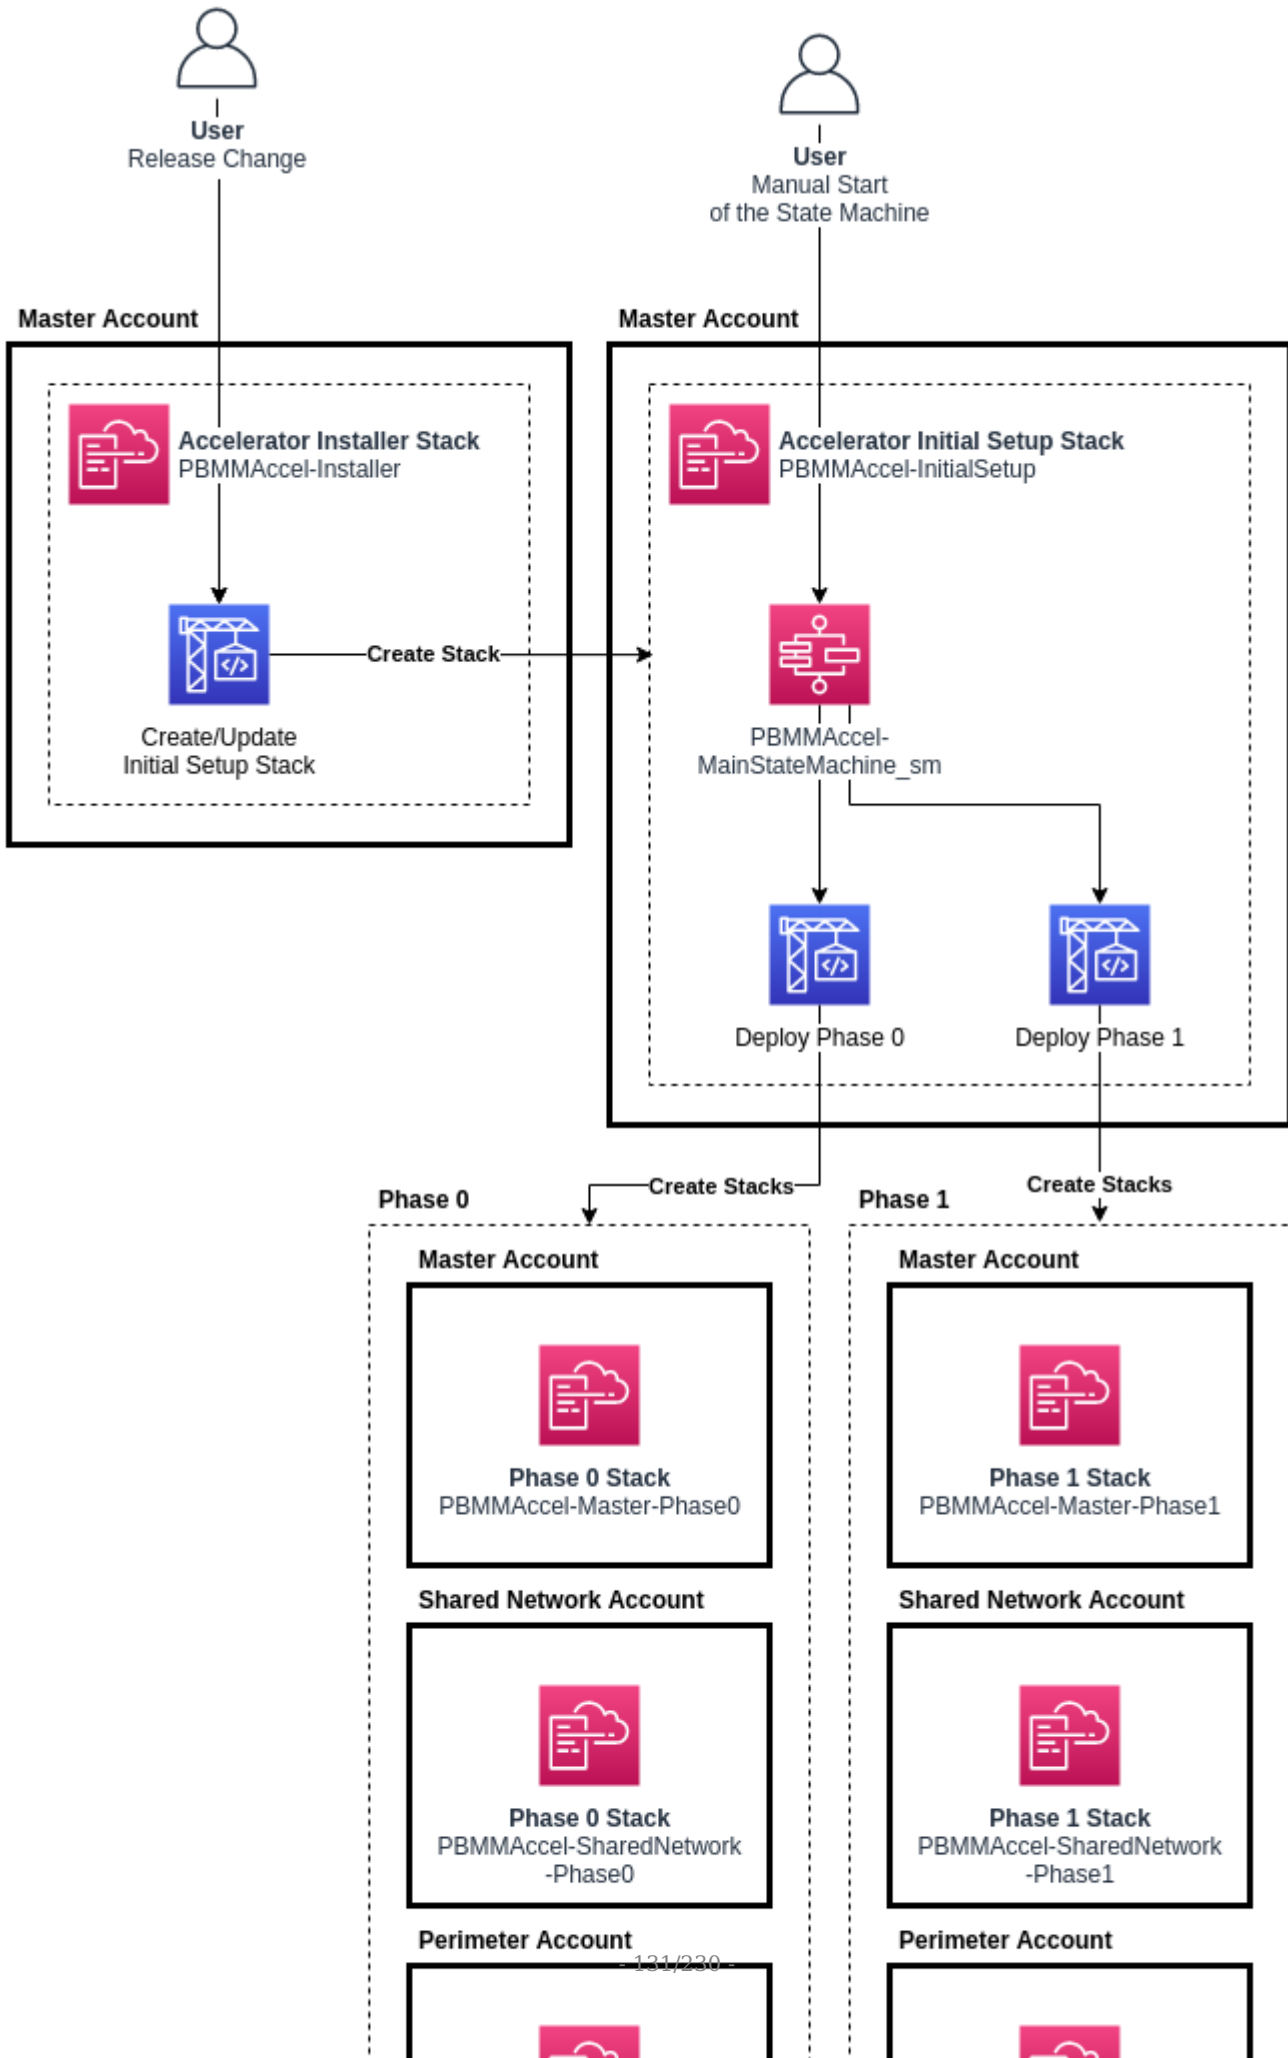
The system can be thought of in two levels. The first level of the system consists of Accelerator stacks and resources. Let's call these the Accelerator-management resource. The second level of the system consists of stacks and resources that are deployed by the Accelerator-management resource. Let's call these the Accelerator-managed resources. The Accelerator-management resources are responsible for deploying the Accelerator-managed resources.

There are two Accelerator-management stacks:

- the `Installer` stack that is responsible for creating the next listed stack;
- the `Initial Setup` stack. This stack is responsible for reading configuration file and creating Accelerator-managed resources in the relevant accounts.

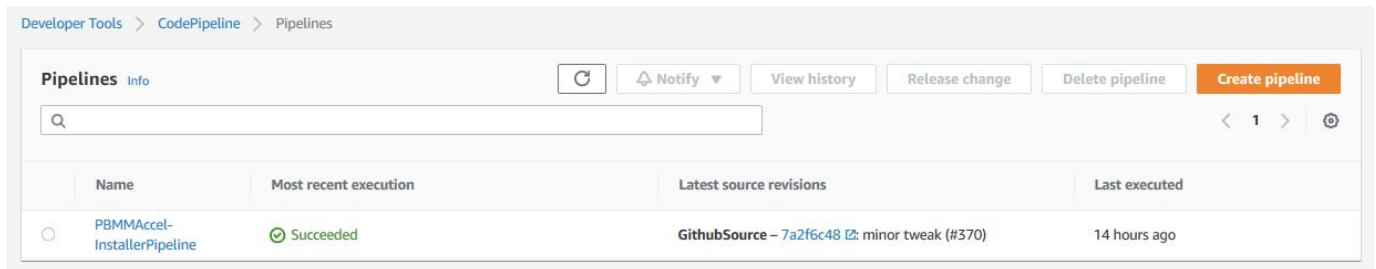
There are multiple Accelerator-managed stacks. Currently there are as many as twelve Accelerator-managed stacks per managed account.

The figure below shows a zoomed-out overview of the Accelerator. The top of the overview shows the Accelerator-management resources, i.e. the `Installer` stack and the `Initial Setup` stack. The bottom of the overview shows the Accelerator-managed resources in the different accounts.




4.2.2 1.2. Installer Stack

The Accelerator-management `Installer` stack contains the necessary resources to deploy the Accelerator-management `Initial Setup` stack in an AWS account. This AWS account will be referred to as the 'root' account in this document.



The Installer stack consists of the following resources:

- `ASEA-InstallerPipeline` : this is a `AWS::CodePipeline::Pipeline` that pulls the latest Accelerator code from GitHub. It launches the CodeBuild project `ASEA-InstallerProject_pl`, executes the `ASEA-Installer-SaveApplicationVersion` Lambda and launches the Accelerator state machine.
- `ASEA-InstallerProject_pl` : this is a `AWS::CodeBuild::Project` that installs the Accelerator in AWS account.
- `ASEA-Installer-SaveApplicationVersion` : this is a `AWS::Lambda::Function` that stores the current Accelerator version into Parameter Store.
- `ASEA-Installer-StartExecution` : this is a `AWS::Lambda::Function` that launches the Accelerator after CodeBuild deploys the Accelerator.
- Creation of `AWS::DynamoDB::Table` - `ASEA-Parameters` and `ASEA-Outputs` which are used for the internal operation of the Accelerator. `ASEA-Outputs` is used to share CloudFormation stack outputs between regions, stacks and phases. `ASEA-Parameters` is used to various configuration items like managed accounts, organizations structure, and limits.



The `ASEA-InstaLLerPipeline` starts when first installed using the CloudFormation template. The administrator can also start the pipeline manually by clicking the `Release Change` button in the AWS Console.

PBMMAccel-InstallerPipeline

This starts the `ASEA-InstaLLerProject_pl` CodeBuild project. The CodeBuild project uses the GitHub source artifact. The CodeBuild projects spins up a new Linux instances and installs the Accelerator dependencies and starts the deployment of the Accelerator using the AWS Cloud Development Kit (CDK¹).

CDK bootstraps its environment and creates the `CDKToolkit` stack in the AWS account. It creates the S3 bucket `cdktoolkit-stagingbucket-*` and the ECR repository `aws-cdk/assets`.

CDK copies assets to the bootstrap bucket and bootstrap repository that are used by the Accelerator. The assets that are stored on S3 include default IAM policies, default SCPs, default firewall configuration. The assets that are pushed to ECR include the Accelerator Docker build image. This Docker image is responsible for deploying Accelerator resources using the CDK.

CDK finally deploys the `Initial Setup` stack. The Accelerator state machine is described in the next section.

This diagram depicts the Accelerator Installer CodePipeline as of v1.2.1:

The screenshot displays a vertical pipeline with four stages, each marked as 'Succeeded'. The stages are: Source, Deploy, UpdateVersion, and Execute. Each stage includes a job name, provider, success status, and time. Between stages, there are 'Disable transition' buttons. A vertical bar on the right side of the pipeline shows four green checkmarks, indicating the overall success of the pipeline.

Source Succeeded
Pipeline execution ID: 73831e89-9116-46ec-ae4a-17537e6b89b1

GithubSource ⓘ
GitHub [🔗](#)
✔ Succeeded - 15 hours ago
7a2f6c48 [🔗](#)

7a2f6c48 [🔗](#) GithubSource: minor tweak (#370)

Disable transition

Deploy Succeeded
Pipeline execution ID: 73831e89-9116-46ec-ae4a-17537e6b89b1

DeployAccelerator ⓘ
AWS CodeBuild
✔ Succeeded - 14 hours ago
[Details](#)

7a2f6c48 [🔗](#) GithubSource: minor tweak (#370)

Disable transition

UpdateVersion Succeeded
Pipeline execution ID: 73831e89-9116-46ec-ae4a-17537e6b89b1

UpdateVersion ⓘ
AWS Lambda [🔗](#)
✔ Succeeded - 14 hours ago
[Details](#)

7a2f6c48 [🔗](#) GithubSource: minor tweak (#370)

Disable transition

Execute Succeeded
Pipeline execution ID: 73831e89-9116-46ec-ae4a-17537e6b89b1

ExecuteAcceleratorStat... ⓘ
AWS Lambda [🔗](#)
✔ Succeeded - 14 hours ago
[Details](#)

7a2f6c48 [🔗](#) GithubSource: minor tweak (#370)

Once the Code Pipeline completes successfully:

- the Accelerator codebase was pulled from GitHub
- the Accelerator codebase was deployed/installed in the Organization Management (root) AWS account
- parameter store `/accelerator/version` was updated with the new version information
- this provides a full history of all Accelerator versions and upgrades
- the newly installed Accelerator state machine is started

At this time the resources deployed by the Installer Stack are no longer required. The Installer stack **could** be removed (which would remove the Code Pipeline) with no impact on Accelerator functionality.

If the Installer Stack was removed, it would need to be re-installed to upgrade the Accelerator. If the stack was not removed, an Accelerator codebase upgrade often only requires updating a single stack parameter to point to the latest Accelerator code branch, and re-releasing the pipeline. No files to manually copy, change or update, an upgrade can be initiated with a simple variable update.

Specify stack details

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

ConfigBranchName

The Code Commit branch name that contains the Accelerator configuration

ConfigRepositoryName

The Code Commit repository name that contains the Accelerator configuration.

ConfigS3Bucket

The S3 bucket name that contains the initial Accelerator configuration.

GithubBranch

The branch of the Github repository containing the Accelerator code.

GithubOwner

The owner of the Github repository containing the Accelerator code.

GithubRepository

The name of the Github repository containing the Accelerator code.

GithubSecretId

The token to use to access the Github repository.

NotificationEmail

The notification email that will get Accelerator State Machine execution notifications.

Cancel

Previous

Next

4.2.3 1.3. Initial Setup Stack

The Accelerator-management Initial Setup stack, named `ASEA-InitialSetup`, consists of a state machine, named `ASEA-MainStateMachine_sm`, that executes various steps to create the Accelerator-managed stacks and resources in the Accelerator-managed accounts. Using a state machine, we can clearly define the deployment process and systematically control branches of execution and handle exceptions.

The Accelerator comprises a primary state machine `ASEA-MainStateMachine_sm`, and nine supporting state machines (as of v1.2.1). Customer will only ever Execute the `ASEA-MainStateMachine_sm`. All troubleshooting will also typically begin with the `ASEA-MainStateMachine_sm`.

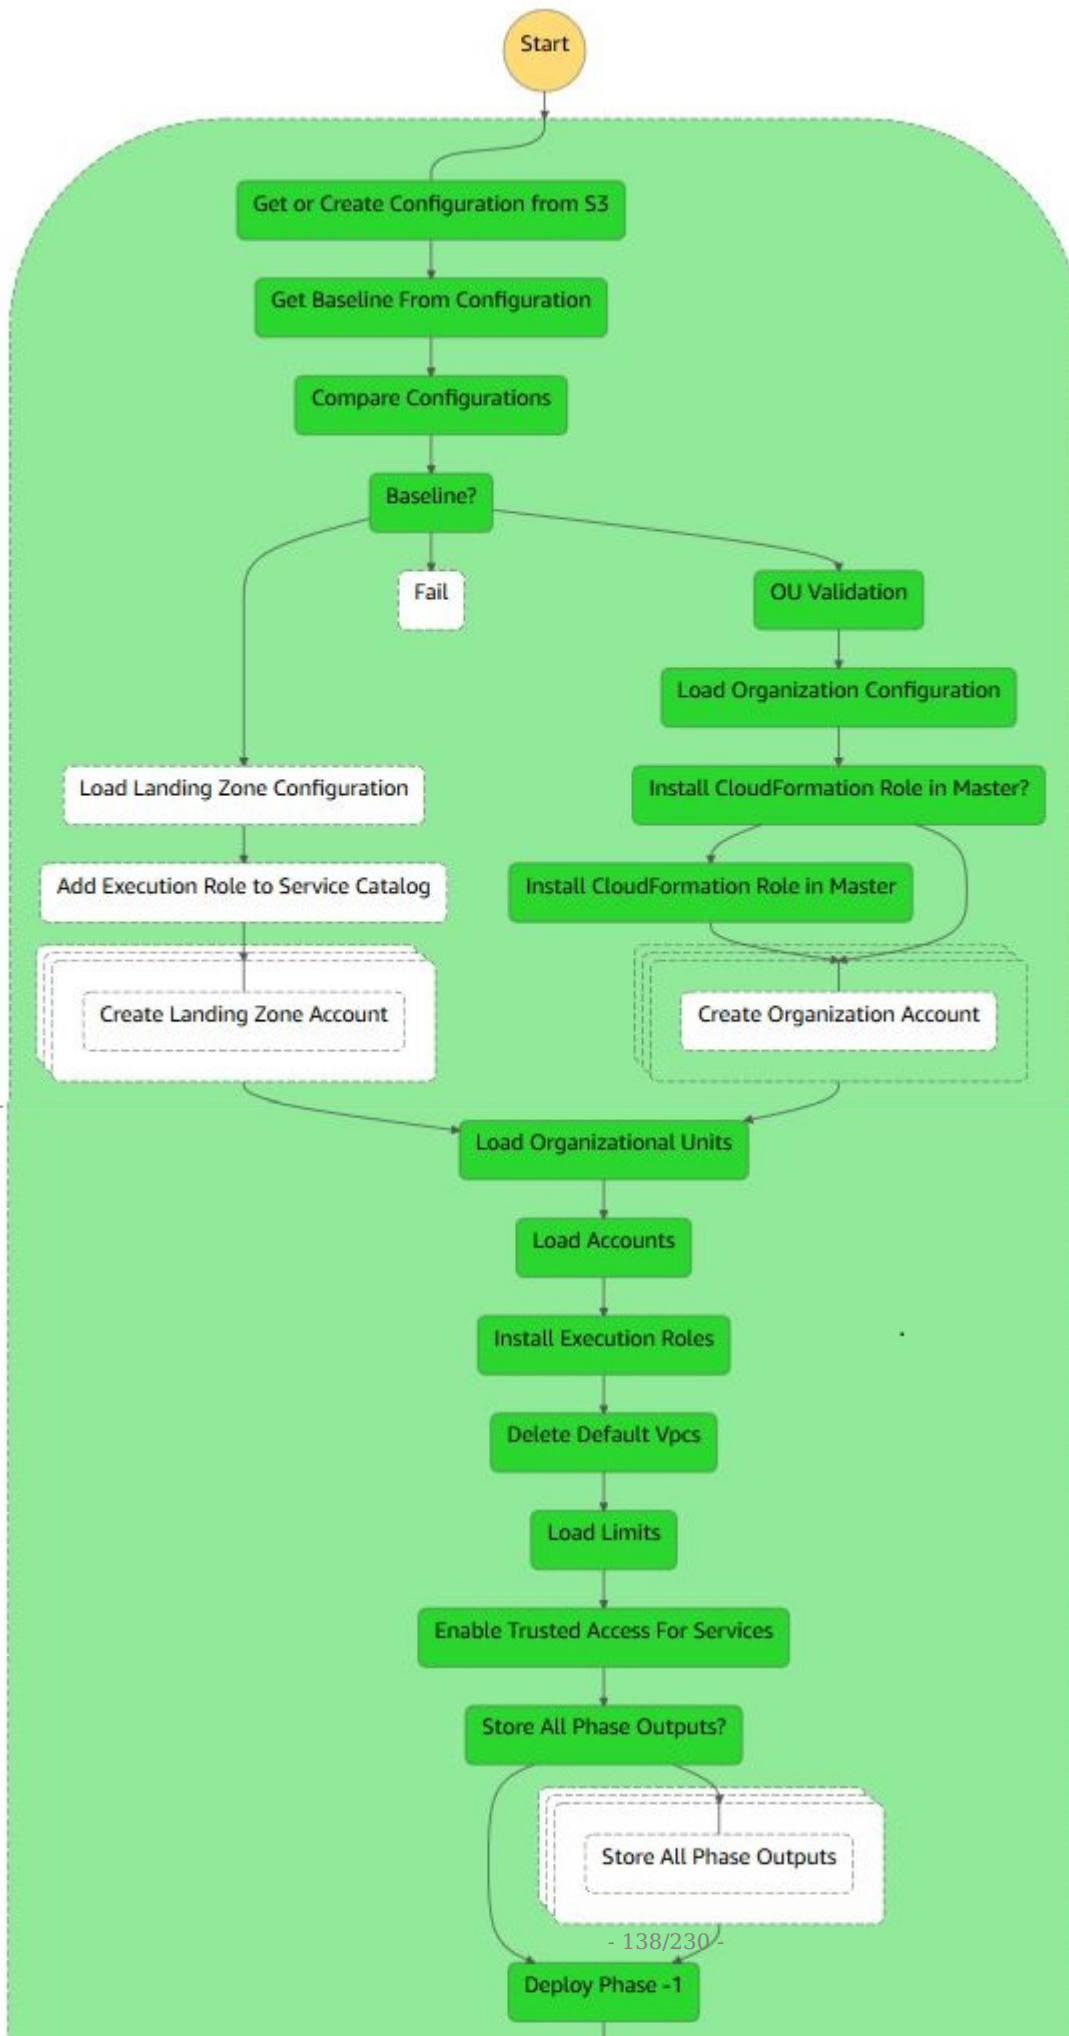
Step Functions > State machines

State machines (9) Refresh View details Edit Copy to new Delete Create state machine

Search for state machines Any type < 1 > Refresh

	Name	Type	Creation date	Status	Logs	Running	Succeeded	Failed	Timed out	Aborted
<input checked="" type="radio"/>	PBMMAccel-MainStateMachine_sm	Standard	Aug 31, 2020 05:50:56.255 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-CodeBuild_sm	Standard	Aug 31, 2020 05:50:30.127 PM	Active	-	0	14	0	0	0
<input type="radio"/>	PBMMAccel-InstallRoles_sm	Standard	Aug 31, 2020 05:50:29.159 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-CreateAdConnector_sm	Standard	Aug 31, 2020 05:50:27.909 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-InstallCfnRoleMaster_sm	Standard	Aug 31, 2020 05:50:27.811 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-OrgCreateAccount_sm	Standard	Aug 31, 2020 05:50:27.749 PM	Active	-	0	16	0	0	0
<input type="radio"/>	PBMMAccel-DeleteDefaultVpcs_sfn	Standard	Aug 31, 2020 05:50:27.364 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-CreateConfigRecorder_sfn	Standard	Aug 31, 2020 05:50:27.312 PM	Active	-	0	2	0	0	0
<input type="radio"/>	PBMMAccel-ALZCreateAccount_sm	Standard	Aug 31, 2020 05:50:26.779 PM	Active	-	0	0	0	0	0

The image below depicts the latest state `ASEA-MainStateMachine_sm` machine. Each green or white square in the image represents a step in the state machine. This all green diagram represents a successful Accelerator state machine execution.



Notice the extremely linear state machine design. This was done to reduce complexity and ease troubleshooting. You may notice a small fork in the state machine. The left path is executed if the Accelerator is deployed on top of an ALZ, The right path is executed if the Accelerator is deployed as a standalone solution. If we eventually add a Control Tower deployment option, it is likely it could also leverage the existing ALZ path.

The state machine contains three different types of steps:

1. steps that execute a Lambda function;
2. steps that start another state machine, e.g. `Create Accounts` step;
3. steps that start another state machine that starts a CodeBuild project, e.g. the `Phase {-1,0,1,2,3,4,5}` steps.

The stack additionally consists of the following resources:


- `AWS::CodeBuild::Project`
- `ASEA-Deploy` or `ASEA-DeployPrebuilt`
- `AWS::CodeCommit::Repository`
- `ASEA-Config-Repo`
- `AWS::IAM::Role`
- `ASEA-L-SFN-MasterRole`
- `ASEA-L-SFN-Execution`
- `AWS::Lambda::Function`
- A Lambda function for every Lambda function step in the state machine.
- `AWS::StepFunctions::StateMachine`
- `ASEA-ALZCreateAccount_sm`: See [Create Landing Zone Account](#);
- `ASEA-OrgCreateAccount_sm`: See [Create Organization Account](#);
- `ASEA-InstallCfnRoleMaster_sm`: See [Install CloudFormation Execution Role](#);
- `ASEA-InstallRoles_sm`: See [Install Execution Roles](#);
- `ASEA-DeleteDefaultVpcs_sfn`: See [Delete Default VPCs](#);
- `ASEA-CodeBuild_sm`: See [Deploy Phase 0](#);
- `ASEA-CreateConfigRecorder_sfn`: See [Create Config Recorders](#);
- `ASEA-CreateAdConnector_sm`: See [Create AD Connector](#);
- `ASEA-StoreOutputs_sm`: See [Share Outputs](#) - new in v1.2.1.

Note: Most resources have a random suffix to their name. This is because we use CDK to deploy the resources. See https://docs.aws.amazon.com/cdk/latest/guide/identifiers.html#identifiers_logical_ids

1.3.1.1. Get or Create Configuration from S3

This step calls a Lambda function that finds or creates the configuration repository. Finds the configuration file(s) in the CodeCommit repository. If the configuration file cannot be found in the repository it is copied from the customer's S3 configuration bucket. If the copy is successful then the configuration file(s) in the S3 bucket will be removed.

The configuration file `config.json` or `config.yaml` is parsed and validated. This step will fail if both file types exist, the configuration file is not valid JSON or YAML or does not adhere to the configuration file specification. Internally the Accelerator always leverages JSON, but accepts JSON or YAML as the source input file and converts it to JSON prior to each execution, storing the converted and fully expanded file in the raw folder.



1.3.2. Get Baseline from Configuration

This step calls a Lambda function that gets the `alz-baseline` of the configuration file to decide which path in the state machine will be taken.

1.3.3. Compare Configurations

This step calls a Lambda function that compares the previous version of the configuration file with the current version of the configuration file. The previous configuration file CodeCommit commit id is stored in the secret `accelerator/config/last-successful-commit` in AWS Secrets Manager in the root account.

The following configuration file changes are not allowed:

- changing ALZ baseline;
- changing root account or region;
- changing central log services account or region;
- changing the organizational unit, name or email address of an account;
- removing an account;
- changing the name, CIDR or region of a VPC;
- disabling a VPC;
- changing the name, availability zone, CIDR of a subnet;
- disabling or removing a subnet;
- changing the name, ASN, region or features of a transit gateway;
- changing the ID, VPC, subnet, region, size, DNS, Netbios of a Managed Active Directory;
- disabling a Managed Active Directory;
- changing the ASN of a virtual private gateway;
- changing the sharing to accounts of a VPC;
- changing the NACLs of a subnet.

It is possible to ignore certain configuration file changes. See [Restart the State Machine](#) how to pass these options to the state machine.

1.3.4. Load Landing Zone Configuration

Executed only when using AWS Landing Zone baseline

This step calls a Lambda function that validates the ALZ deployment and configuration.

This step fails when

- an existing ALZ deployment in the root account can not be found;
- an organizational unit that is used by ALZ or the Accelerator can not be found;
- an account is present in the Accelerator configuration but does not exist in AWS Organizations;
- an account has a name, email address or organizational unit that is different from the ALZ configuration;
- an account has a name, email address or organizational unit that is different from the Accelerator configuration;
- the organizational units in the ALZ configuration are different from the ones in AWS Organizations;
- the account type of ALZ configuration accounts can not be detected;
- the accounts of type `primary`, `security`, `log-archive`, `shared-services` are missing from the ALZ configuration.

1.3.5. Add Execution Role to Service Catalog

Executed only when using AWS Landing Zone baseline

This step calls a Lambda function that adds the state machine's IAM role to the ALZ Account Vending Machine (AVM) service catalog portfolio.

1.3.6. Create Landing Zone Account

Executed only when using AWS Landing Zone baseline

This step starts the `ASEA-ALZCreateAccount_sm` state machine. This state machine is responsible for creating an account using the AVM and waits for the account to be created and configured.

1.3.7. Organizational Unit (OU) Validation

Executed only when using AWS Organizations baseline

This step validates that the OU structure defined in the Accelerator configuration file matches the OU and account structure defined in AWS Organizations. Missing OUs are created. If any AWS Accounts or OUs have been renamed, this step updates the configuration file with the latest information. Accounts located in the incorrect top-level OU cause a state machine failure.

1.3.8. Load Organization Configuration

Executed only when using AWS Organizations baseline

This step calls a Lambda function that validates the Accelerator deployment and configuration.

This step fails when

- an organizational unit that is used by the Accelerator can not be found;
- an account is present in the Accelerator configuration but does not exist in AWS Organizations;
- an account has a name, email address or organizational unit that is different from the Accelerator configuration.
- **returns accounts that have not been created**
- duplicates some validation functionality

1.3.9. Install CloudFormation Role in root

Executed only when using AWS Organizations baseline

This step starts the `ASEA-InstaLLCfnRoleMaster_sm` state machine. This state machine is responsible for creating the IAM role defined in `organization-admin-role` (default: `AWSCloudFormationStackSetAdministrationRole`) in the root account. You can read more about why this role is created [here](#).

1.3.10. Create Organization Account

Executed only when using AWS Organizations baseline

This step starts the `ASEA-OrgCreateAccount_sm` state machine. This state machine is responsible for creating an account using the AWS Organizations and waits for the account to be created and configured. The newly created account will be moved to the correct organizational unit and will have an the `ASEA-Quarantine-New-Object` Service Control Policy (SCP) attached to it. The SCP will only be removed after the Accelerator state machine has completed.

Application of this SCP ensures that new AWS accounts cannot be used until they have been properly configured and had all Accelerator guardrails applied per the configuration file. It is important that guardrails be consistently in place from the time a user first accesses a new AWS account, until the account is decommissioned. Should the state machine fail, still be in-progress, or has not yet processed the account (for whatever reason), users who access it will be prevented from doing *anything*.

1.3.11. Load Organizational Units

This step calls a Lambda function that loads all Organizational Units. The Organizational Units are stored in the DynamoDB table `ASEA-Parameters` in the root account. The Organizational Units will later be used to determine which configuration file options apply to an AWS account.

1.3.12. Load Accounts

This step calls a Lambda function that loads all accounts. The accounts are stored in the DynamoDB table `ASEA-Parameters` in the root account. The accounts will later be used to deploy stacks in the correct accounts.

1.3.13. Install Execution Roles

This step calls a Lambda function that creates stack sets in all Accelerator accounts. This stack sets contains a single resource, i.e. an IAM role `ASEA-PipelineRole` that can be assumed by the `ASEA-L-SFN-MasterRole`. This IAM role allows the root account to administer the Accelerator accounts.

1.3.14. Delete Default VPCs

This step starts the `ASEA-DeleteDefaultVpcs_sfn` state machine. This state machine is responsible for deleting default subnets, internet gateways and VPCs for all regions and accounts in the Accelerator configuration.

This step fails when one or more default VPCs cannot be deleted. This step does not fail on the first error, it executes on all accounts/regions and then fails with a complete list of errors.

1.3.15. Load Limits

This step calls a Lambda function that loads service quotas and requests a service quota increase according to the configuration file. When a service quota increase request has been closed but not increased, then the service quota request will be issued again when the creation of the last request was at least two days ago.

1.3.16. Enable Trusted Access for Services

This step calls a Lambda function that is responsible for

- enabling AWS service access in the organization;
- enabling AWS Resource Access Manager sharing in the organization;
- creating a service-linked role for AWS IAM Access Analyzer;
- setting the security account as delegated administrator for AWS Firewall Manager;
- setting the security account as delegated administrator for AWS IAM Access Analyzer;
- setting the security account as delegated administrator for Amazon GuardDuty.

1.3.17. Store All Phase Outputs

This step only executes on the first run of the state machine after it has been upgraded to v1.2.0 or above. This step exists solely to support upgrades from Accelerator versions prior to v1.2.0 and can be removed when no existing customers are running versions older than v1.2.0. This step populates the DynamoDB Outputs table with the outputs from previous executions which were previously stored in S3 (and at one time even stored in secrets manager).

1.3.18. Deploy Phase -1 (Negative one)

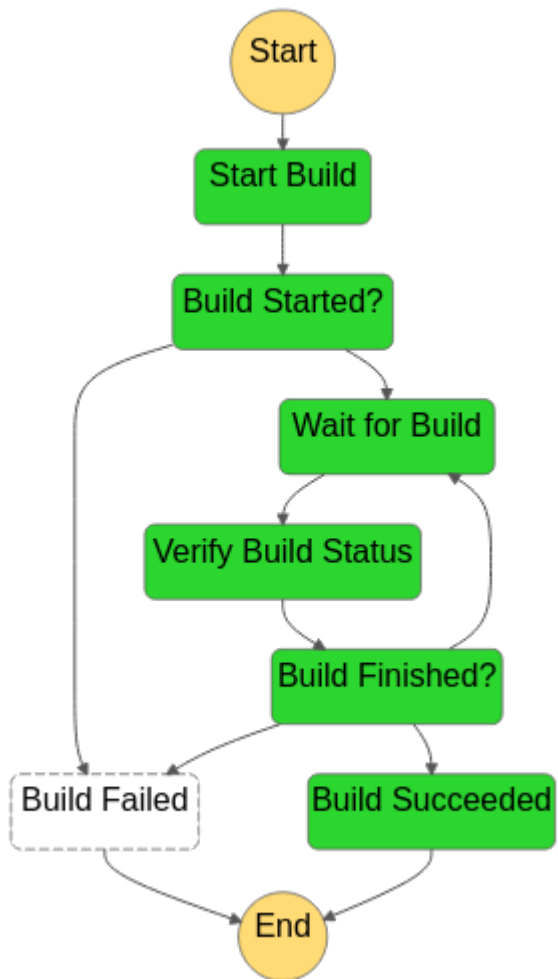
- The following resources are deployed in phase -1:
- Creating required roles for macie custom resources
- Creating required roles for guardDuty custom resources
- Creating required roles for securityHub custom resources
- Creating required roles for iamCreateRole custom resource
- Creating required roles for createSSMDocument custom resource
- Creating required roles for createLogGroup custom resource
- Creating required roles for CWLCentralLoggingSubscriptionFilterRole custom resource
- Creating required roles for TransitGatewayCreatePeeringAttachment custom resource
- Creating required roles for TransitGatewayAcceptPeeringAttachment custom resource
- Creating required roles for createLogsMetricFilter custom resource
- Creating required roles for SnsSubscriberLambda custom resource

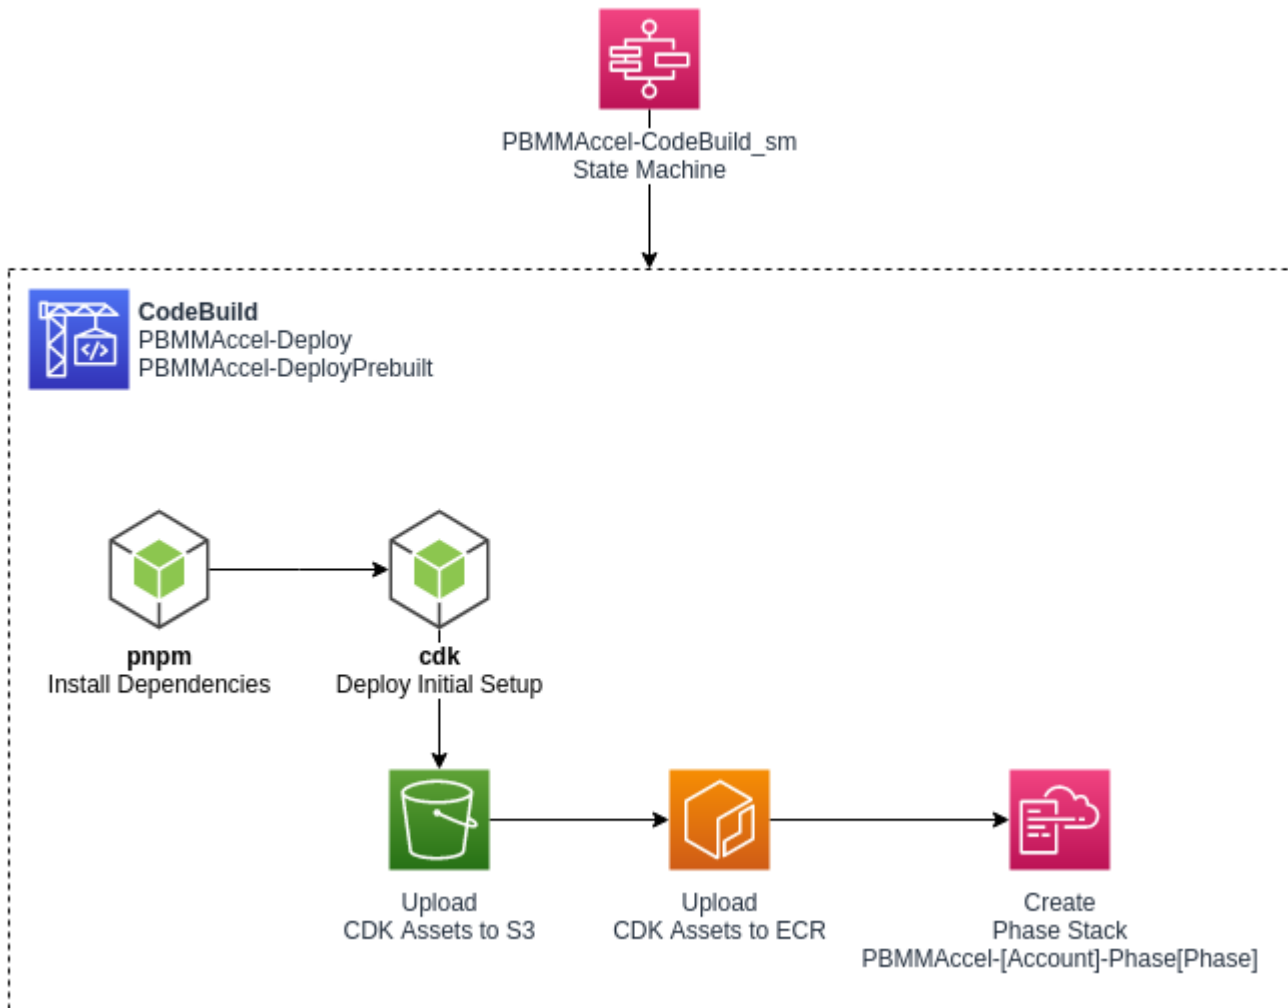
1.3.19. Store Phase -1 Output

This step calls a Lambda function that stores the outputs from the deployed stacks in sub-accounts in the DynamoDB table `ASEA-Outputs` in AWS Secrets Manager in the root account.

1.3.20. Deploy Phase 0

This step starts the `ASEA-CodeBuild_sm` state machine. This state machine is responsible for starting a CodeBuild execution that deploys CloudFormation stacks in Accelerator-managed accounts using CDK.





The same CodeBuild project is used to deploy all phases. The environment variable `ACCELERATOR_PHASE` in the CodeBuild execution decides which phase to deploy.

The following resources are deployed in phase 0:

- create default EBS encryption key;
- create an AWS log bucket with encryption key;
- create the central log services bucket with encryption key;
- create the Accelerator configuration bucket with encryption key;
- copy artifacts to the Accelerator configuration bucket:
- SCPs;
- firewall configuration;
- account warming (step 1);
- set password policy (step 1);
- create IAM users (step 1):
- create passwords and store in Secrets Manager;
- create MAD deployment (step 1):
- create passwords and store in Secrets Manager;
- create service-linked role;
- create `rsyslog` deployment (step 1);
- create firewalls (step 1);
- create budgets (step 1);
- create transit gateways (step 1);
- create Route53 DNS logging log group;
- enable Macie (step 1);
- enable GuardDuty;
- enable Access Analyzer;

1.3.21. Store Phase 0 Output

This step calls a Lambda function that stores the outputs from the deployed stacks in sub-accounts in the secret `ASEA-Outputs` in AWS Secrets Manager in the root account.

1.3.22. Verify Files

This step verifies that all required files to complete the installation have been provided by the user. If any required files have not been provided, this step will fail and provide a list of *all* missing files.

Why do we wait until so late in the state machine execution to perform this task?

1. we do not want to add or delete files from a customer supplied bucket
2. we do not want to force customers to need to copy and supply standard configuration files for 3. prescriptive installation
3. we want to allow customers to override the sample or standard configuration files
4. the config bucket is only created in Phase 0
5. we copy all the relevant sample files from reference-artifacts folders to our config bucket, and then, we copy any customer supplied config files over top of our sample files
6. This allows customers to override the sample configs, with customer custom configs without needing to either fork the repo, or supply the files on a vanilla install
7. Until we do this file copy, we cannot be sure that the combination of customer supplied and reference-artifact sample files provides all the required files.

1.3.23. Create Config Recorders

This step starts the `ASEA-CreateConfigRecorder_sfn` state machine. This state machine is responsible for creating Config recorders in all accounts and regions.

1.3.24. Add SCPs to Organization

This step calls a Lambda function that creates and attaches the SCPs listed in the Accelerator configuration. The SCP policies are loaded from the Accelerator configuration bucket.

This step fails when

- an SCP policy cannot be found in the Accelerator configuration bucket;
- an SCP could not be attached to an organizational unit or account, e.g. when the maximum number of attached SCPs is exceeded

1.3.25. Deploy Phase 1

- Create S3 Bucket in all accounts and replicate to Log Account Bucket
- Deploy VPC:
 - Vpc
 - Subnets
 - Subnet sharing (RAM)
 - Route tables
 - Internet gateways
 - NAT gateways
 - Interface endpoints
 - Gateway endpoints
 - Transit Gateway Attachments
- IAM Role required for VPC Peering Auto accept
- Firewall images subscription check
- Creates the customer gateways for the EIPs of the firewall
- Create IAM Roles, Users in account based on configuration
- Creates the additional budgets for the account stacks.
- Import Certificates
- Setup SSMSessionManagerDocument
- Create Cost and Usage reports
- Enable Macie in root Account
- GuardDuty setup in Security Account
- Setup CWL Central Logging
- Create Roles required for Flow Logs
- Transit Gateway Peering
- Create LogGroup required for DNS Logging

1.3.26. Store Phase 1 Output

See [Deploy Phase 0](#).

1.3.27. Account Default Settings

This step calls a Lambda function that

- enables and sets EBS default encryption for all accounts in the Accelerator configuration;
- enables S3 object level ALZ CloudTrail logging;
- enables Log Insight events;
- enables KMS encryption using the CMK from the central logging account;
- sets AWS Systems Manager Session Manager default configuration in every Accelerator-managed account in every region with a VPC.

1.3.28. Deploy Phase 2

- Create CloudTrail in root account
- Create VPC Peering Connection
- Create Security Groups for shared VPC in sub accounts
- Setup Security Hub in Security Account
- Setup Cross Account CloudWatch logs sharing by creating roles in sub accounts
- Enable VPC FlowLogs
- Create Active Directory (MAD)
- Create Firewall clusters
- Create Firewall Management instance
- Create Transit Gateway Routes, Association and Propagation
- Enable Macie in Security account and Create Members, Update Config
- GuardDuty - Add existing Org accounts as members and allow new accounts to be members and Publish
- Create SNS Topics in Log Account
- TGW Peering Attachments

1.3.29. Store Phase 2 Output

See [Deploy Phase 0](#).

1.3.30. Deploy Phase 3

- create peering connection routes;
- create ALB (step 1);
- create `rsyslog` deployment (step 2);
- create hosted zones, resolver rules and resolver endpoints and Share;
- Enable Security Hub and Invite Sub accounts as members;
- TransitGateway Peering attachment and routes;
- Macie update Session;

1.3.31. Store Phase 3 Output

See [Deploy Phase 0](#).

1.3.32. Deploy Phase 4

- SecurityHub Disable Controls
- Creates CloudWatch Metrics on LogGroups

- Associate Shared Resolver Rules to VPC
- Associate Hosted Zones to VPC

1.3.33. Store Phase 4 Output

See [Deploy Phase 0](#).

1.3.34. Associate Hosted Zones (Step removed in v1.2.1)

This step calls a Lambda function that associates the private zones, all the interface endpoint zones, and the resolver rules with each VPC that leverages endpoint services. This step was removed in v1.2.1 of the Accelerator codebase.

1.3.35. Add Tags to Shared Resources

This step calls a Lambda function that adds tags to shared resources in the share destination account. For example, when a subnet is shared into another account, this step will add the `Name` tag to the subnet in the shared account.

The supported resources are

- VPCs;
- subnets;
- security groups;
- transit gateway attachments.

1.3.36. Enable Directory Sharing

This step calls a Lambda function that shares Managed Active Directory according to the Accelerator configuration. The directory is shared from the source account to the target account. The directory will be accepted in the target account.

1.3.37. Deploy Phase 5

- create Remote Desktop Gateway;
- create launch configuration;
- create autoscaling group;
- enable central logging to S3 (step 2);
- Create CloudWatch Events for `moveAccount`, `policyChanges` and `createAccount`
- Creates CloudWatch Alarms

1.3.38. Create AD Connector

This step starts the `ASEA-DeleteDefaultVpcs_sfn` state machine. This state machine is responsible for creating AD connectors according to the Accelerator configuration.

This step fails when one or more AD connectors failed to be created.

1.3.39. Store Commit ID

This step calls a Lambda function that stores the commit ID of the configuration file for which the state machine ran.

1.3.40. Detach Quarantine SCP

Executed only when using AWS Organizations baseline

This step calls a Lambda function that stores the commit ID for which the state machine just ran.

4.3 1. Troubleshooting

4.3.1 1.1. Overview

Issues could occur in different parts of the Accelerator. We'll guide you through troubleshooting these issues in this section.

4.3.2 1.2. Components

1.2.1. State Machine

Viewing the step function `Graph inspector` (depicted above in 2.2), the majority of the main state machine has a large colored box around which is the functionality to catch state machine failures `Main Try Catch block to Notify users`. This large outer box will be blue while the state machine is still executing, it will be green upon a successful state machine execution and will turn orange/yellow on a state machine failure.

What if my State Machine fails? Why? Previous solutions had complex recovery processes, what's involved?

If your main state machine fails, review the error(s), resolve the problem and simply re-run the state machine. We've put a huge focus on ensuring the solution is idempotent and to ensure recovery is a smooth and easy process.

Ensuring the integrity of deployed guardrails is critical in operating and maintaining an environment hosting protected data. Based on customer feedback and security best practices, we purposely fail the state machine if we cannot successfully deploy guardrails.

Additionally, with millions of active customers each supporting different and diverse use cases and with the rapid rate of evolution of the AWS platform, sometimes we will encounter unexpected circumstances and the state machine might fail.

We've spent a lot of time over the course of the Accelerator development process ensuring the solution can roll forward, roll backward, be stopped, restarted, and rerun without issues. A huge focus was placed on dealing with and writing custom code to manage and deal with non-idempotent resources (like S3 buckets, log groups, KMS keys, etc.). We've spent a lot of time ensuring that any failed artifacts are automatically cleaned up and don't cause subsequent executions to fail. We've put a strong focus on ensuring you do not need to go into your various AWS sub-accounts and manually remove or cleanup resources or deployment failures. We've also tried to provide usable error messages that are easy to understand and troubleshoot. As new scenario's are brought to our attention, we continue to adjust the codebase to better handle these situations.

Will your state machine fail at some point in time, likely. Will you be able to easily recover and move forward without extensive time and effort, YES!

As the state machine executes, each step will turn from white (not started), to blue (executing), to green (Success), or grey/red (failure). To diagnose the problem select the grey/red step that failed. If you miss the step and select the outer box, you will have selected the

`Main Try Catch block to Notify users`. You need to carefully select the failed step.



As stated in section 2.2, the state machine contains 3 different types of states, which are each diagnosed differently.

a. If the step is calling a Lambda function then you will see the following after clicking the failed step.

Code

Step details

Name	Type
Verify Files	Task

Status

✖ Failed

Resource

[arn:aws:lambda:ca-central-1:397069427220:function:PBMMAccel-InitialSetup-PipelineVerifyFilesHandler9-AIKPZN5FIL92](#) [↗](#) | [CloudWatch logs](#) [↗](#)

▶ Input

▶ Output

▼ Exception

Error

Error

Cause

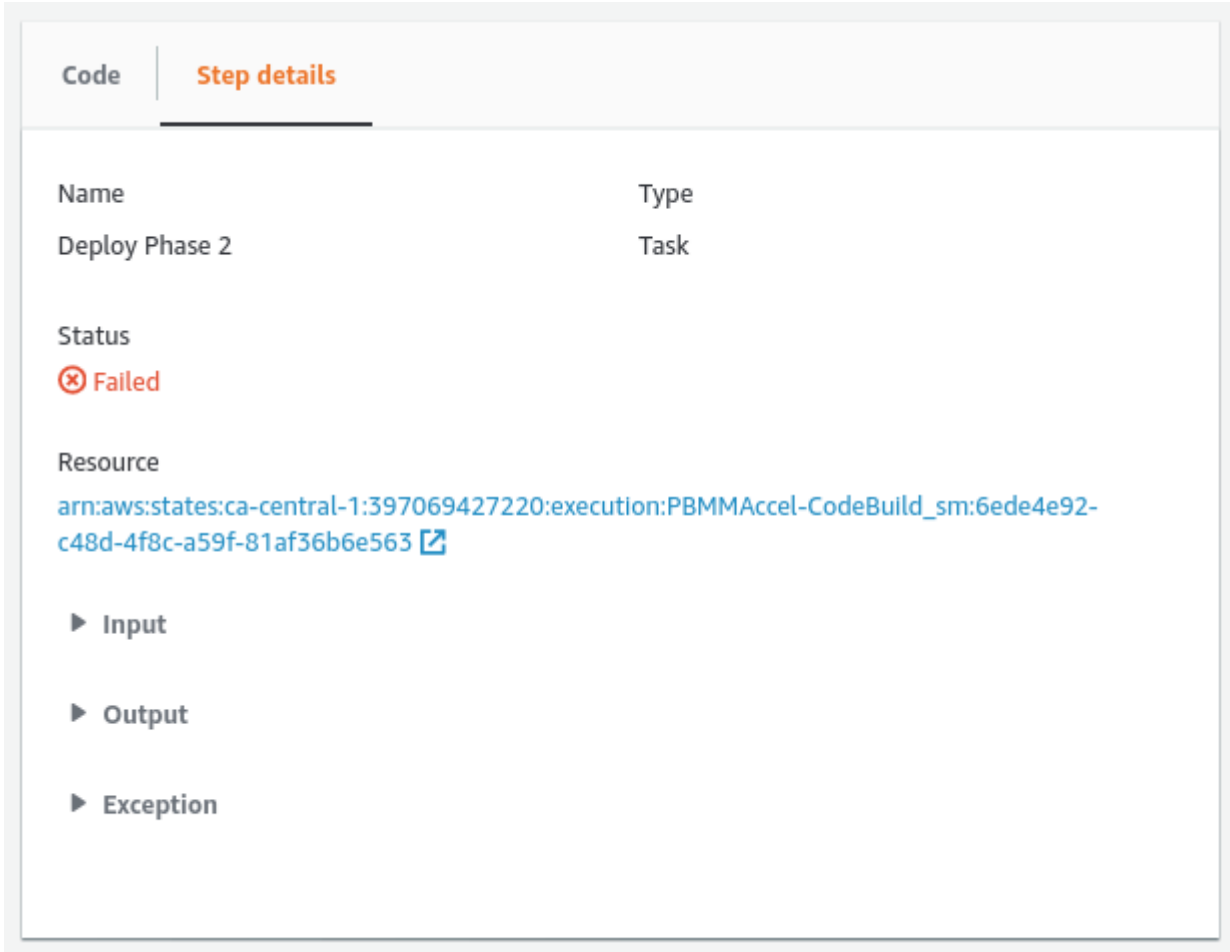
```

{
  "errorType": "Error",
  "errorMessage": "There were errors while loading the
configuration:\nFileCheck: File not found at \"s3://pbmmaccel-master-
phase0-configacental1-9mehzaon40bo/firewall/license2.lic\"\nFileCheck:
File not found at \"s3://pbmmaccel-master-phase0-
configacental1-9mehzaon40bo/firewall/fortigate.txt\"",
  "trace": [
    "Error: There were errors while loading the configuration:",
    "FileCheck: File not found at \"s3://pbmmaccel-master-phase0-
configacental1-9mehzaon40bo/firewall/license2.lic\"",
    "FileCheck: File not found at \"s3://pbmmaccel-master-phase0-
configacental1-9mehzaon40bo/firewall/fortigate.txt\"",
    "    at Runtime.Ni [as handler] (/var/task/index.js:2:4803195)",
    "    at processTicksAndRejections (internal/process
/task_queues.js:97:5)"
  ]
}

```


In this case, you can see that the `Cause` section contains a useful message. This message will differ between Lambda functions. In case this message does not make the issue clear, you can click on the `CloudWatch Logs` link in the `Resource` section to view the output of the Lambda function that was called by the step. See the section [CloudWatch Logs](#). Note: The `Resource` section contains two links that blend together. You need to click the second link (`CloudWatch Logs`), not the first link which will open the actual resource/Lambda.

b. In case the failed step started another state machine, you will see the following after clicking the failed step.



To view the state machine execution that failed you can click the link in the `Resource` section.

In case the failed step started the CodeBuild state machine, `ASEA-CodeBuild_sm`, you will be able to see the CodeBuild project and execution ID that failed by looking at the output of the `Start Build` step in the `ASEA-CodeBuild_sm` state machine.

Code
Step details

Name	Type
Start Build	Task

Status

✔ Succeeded

Resource

[arn:aws:lambda:ca-central-1:397069427220:function:PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-11104YEF7P3JB](#) [🔗](#) | [CloudWatch logs](#) [🔗](#)

▶ **Input**

▼ **Output**

```

{
  "codeBuildProjectName": "PBMMAccel-DeployPrebuilt",
  "environment": {
    "ACCELERATOR_PHASE": "2",
    "CONFIG_COMMIT_ID": "e535270aebac9793f81e1e02777cc327e842ad04",
    "CONFIG_REPOSITORY_NAME": "PBMMAccel-Config-Repo",
    "CONFIG_FILE_PATH": "config.json"
  },
  "startBuildOutput": {
    "status": "SUCCESS",
    "buildArn": "arn:aws:codebuild:ca-central-1:397069427220:build/PBMMAccel-DeployPrebuilt:717584a9-c406-4569-9cc2-0d23e9ff9ef0",
    "buildId": "PBMMAccel-DeployPrebuilt:717584a9-c406-4569-9cc2-0d23e9ff9ef0"
  }
}

```

▶ **Exception**

In the image above the execution of CodeBuild project `ASEA-DeployPrebuilt` with ID `ASEA-DeployPrebuilt:717584a9-c406-4569-9cc2-0d23e9ff9ef0` failed. See the [CodeBuild](#) section to troubleshoot.

1.2.2. CodeBuild

The Accelerator deploys and leverages two CodeBuild projects. The `ASEA-InstallerProject_pl` project is used by the Code Pipeline/Installer stack and `ASEA-DeployPrebuilt` which is used throughout the Accelerator state machine. Both are similar in that they use CDK to deploy stacks. The installer project will not exist, if the installer has been removed.

The screenshot shows the AWS CodeBuild console interface. At the top, there are navigation links for 'Developer Tools', 'CodeBuild', and 'Build projects'. Below this, there's a 'Build projects' section with several action buttons: 'Notify', 'Start build', 'View details', 'Edit', 'Delete build project', and 'Create build project'. A search bar and a 'Your projects' dropdown are also present. The main content is a table with the following columns: Name, Source provider, Repository, Latest build status, and Description. Two projects are listed:

Name	Source provider	Repository	Latest build status	Description
PBMMAccel-DeployPrebuilt	No source	-	✔ Succeeded	-
PBMMAccel-InstallerProject_pl	AWS CodePipeline	-	✔ Succeeded	-

After a successful installation you will see the following in Codebuild, for the `ASEA-DeployPrebuilt` project:

The screenshot shows the configuration and build history for the `PBMMAccel-DeployPrebuilt` project. The configuration section includes:

Configuration			
Source provider: No source	Primary repository: -	Artifacts upload location: -	Build badge: Disabled

Below the configuration, there are tabs for 'Build history', 'Batch history', 'Build details', 'Build triggers', and 'Metrics'. The 'Build history' tab is active, showing a table of build runs:

Build run	Status	Build number	Source version	Submitter	Duration	Completed
PBMMAccel-DeployPrebuilt:1ee98308-950c-4a25-bc42-61e4563038db	✔ Succeeded	7	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	11 minutes 13 seconds	13 hours ago
PBMMAccel-DeployPrebuilt:2a36a971-dd35-4a62-a262-f3cb23711854	✔ Succeeded	6	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	1 minute 13 seconds	13 hours ago
PBMMAccel-DeployPrebuilt:bdbb664e-3a98-4cc7-a930-32fa5f7ee52f	✔ Succeeded	5	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	4 minutes 14 seconds	13 hours ago
PBMMAccel-DeployPrebuilt:b4e6cb38-5fd9-4487-9d2f-5367f3be06aa	✔ Succeeded	4	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	22 minutes 24 seconds	13 hours ago
PBMMAccel-DeployPrebuilt:31451756-1c6f-427a-a26c-a5b3ac5c9253	✔ Succeeded	3	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	12 minutes 14 seconds	14 hours ago
PBMMAccel-DeployPrebuilt:351103e4-033a-45b0-bb3f-3ebe48120da5	✔ Succeeded	2	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	5 minutes 40 seconds	14 hours ago
PBMMAccel-DeployPrebuilt:fc782de1-937c-4060-a3fe-d0b17c4da7e8	✔ Succeeded	1	-	PBMMAccel-L-SFN-MasterRole-DD6508E8/PBMMAccel-InitialSetup-PipelineStartBuildHandlerED-AZEF2501MJOT	3 minutes 0 seconds	14 hours ago

When an error occurs you will see that the CodeBuild project execution fails when looking in the execution overview.

PBMMAccel-
DeployPrebuilt:f77b8
3df-ce3f-4205-
a533-6911e0359094

✘ Failed

206

You can click on the name of the CodeBuild execution and then look inside the logs what caused the failure. These logs can be confusing. We are deploying multiple stacks in parallel and all the messages for all the parallel deployments are interleaved together, so make sure you are correlating the events back to the correct event source. Because we are deploying to 16 regions in parallel, you will also see messages for the same stack deployment interleaved. Even though a task may indicate it is complete and then another seemingly identical task indicates in-progress, the second message is coming from one of the alternate regions.

```

576 1/10 | 10:34:30 PM | UPDATE_COMPLETE | AWS::EC2::Instance | PerimeterPhase2/FirewallManager (FirewallManagerCCB568C3)
577 1/10 | 10:34:30 PM | UPDATE_IN_PROGRESS | AWS::Lambda::Function | PerimeterPhase2/Custom::SecurityHubEnableCustomResourceProvider/Handler
(CustomSecurityHubEnableCustomResourceProviderHandler2B586BC3)
578 1/10 | 10:34:30 PM | UPDATE_COMPLETE | Custom::S3Template | PerimeterPhase2/Firewall/Instance1/License/Resource/Default
(FirewallInstance1LicenseA98B1640)
579 1/10 | 10:34:30 PM | UPDATE_COMPLETE | Custom::S3Template | PerimeterPhase2/Firewall/Instance1/Config/Resource/Default
(FirewallInstance1ConfigF78EF5CB)
580 1/10 | 10:34:30 PM | UPDATE_COMPLETE | AWS::Lambda::Function | PerimeterPhase2/Custom::SecurityHubEnableCustomResourceProvider/Handler
(CustomSecurityHubEnableCustomResourceProviderHandler2B586BC3)
581 Stack PBMMAccel-Perimeter-Phase2 is still not stable (UPDATE_ROLLBACK_IN_PROGRESS)
582 0/10 | 10:34:32 PM | UPDATE_IN_PROGRESS | Custom::S3Template | PerimeterPhase2/Firewall/Instance0/Config/Resource/Default
(FirewallInstance0Config984094C3) Requested update required the provider to create a new physical resource
583 0/10 | 10:34:32 PM | UPDATE_COMPLETE | Custom::S3Template | PerimeterPhase2/Firewall/Instance0/Config/Resource/Default
(FirewallInstance0Config984094C3)
584 0/10 | 10:34:33 PM | UPDATE_ROLLBACK_COMP | AWS::CloudFormation::Stack | PBMMAccel-Perimeter-Phase2
585 0/10 | 10:34:34 PM | DELETE_IN_PROGRESS | AWS::CloudFormation::CustomResource | PerimeterPhase2/Firewall/Instance0/Config/Resource/Default
(FirewallInstance0Config984094C3)
586 0/10 | 10:34:34 PM | DELETE_IN_PROGRESS | AWS::CloudFormation::CustomResource | PerimeterPhase2/Firewall/Instance1/Config/Resource/Default
(FirewallInstance1ConfigF78EF5CB)
587 0/10 | 10:34:34 PM | DELETE_IN_PROGRESS | AWS::CloudFormation::CustomResource | PerimeterPhase2/Firewall/Instance1/License/Resource/Default
(FirewallInstance1LicenseA98B1640)
588 0/10 | 10:34:34 PM | DELETE_COMPLETE | AWS::EC2::Instance | PerimeterPhase2/FirewallManager (FirewallManagerCCB568C3)
589 0/10 | 10:34:35 PM | DELETE_IN_PROGRESS | AWS::CloudFormation::CustomResource | PerimeterPhase2/Firewall/Instance0/License/Resource/Default
(FirewallInstance0License05FC980)
590 0/10 | 10:34:35 PM | DELETE_COMPLETE | AWS::CloudFormation::CustomResource | PerimeterPhase2/Firewall/Instance0/Config/Resource/Default
(FirewallInstance0Config984094C3)
591 Error: Error: The stack named PBMMAccel-Perimeter-Phase2 is in a failed state: UPDATE_ROLLBACK_COMPLETE
592 at Object.fulfillAll (/app/initial-setup/templates/promise.ts:8:11)
593 at processTicksAndRejections (internal/process/task_queues.js:97:5)
594 at cdkToolkit.deployAllStacks (/app/initial-setup/templates/toolkit.ts:143:27)
595 at main (/app/initial-setup/templates/cdk.ts:53:21)
596
597 [Container] 2020/07/04 22:34:38 Command did not exit successfully sh docker-entrypoint.sh exit status 1
598 [Container] 2020/07/04 22:34:38 Phase complete: BUILD State: FAILED
599 [Container] 2020/07/04 22:34:38 Phase context status code: COMMAND_EXECUTION_ERROR Message: Error while executing command: sh docker-entrypoint.sh. Reason: exit
status 1
600 [Container] 2020/07/04 22:34:38 Entering phase POST_BUILD
601 [Container] 2020/07/04 22:34:38 Phase complete: POST_BUILD State: SUCCEEDED
602 [Container] 2020/07/04 22:34:38 Phase context status code: Message:

```

You can for example see the error message `The stack named ASEA-Perimeter-Phase2 is in a failed state: UPDATE_ROLLBACK_COMPLETE`. This means the stack `ASEA-Perimeter-Phase2` failed to update and it had to rollback. The error indicated at the bottom of the Codebuild screen is typically NOT the cause of the failure, just the end result. You need to scroll up and find the FIRST occurrence of an error in the log file. Often starting at the top of the log file and searching for the text `FAIL` (case sensitive), will allow you to find the relevant error message(s) quickly. The failure is typically listed in the CloudFormation update logs.

```

550 4/10 | 10:34:22 PM | UPDATE_FAILED | AWS::EC2::Instance | PerimeterPhase2/FirewallManager (FirewallManagerCCB568C3) Interface:
[eni-0dd94b35aa8be7b3d] in use. (Service: AmazonEC2; Status Code: 400; Error Code: InvalidNetworkInterface.InUse; Request ID: df3728f6-04cb-4fad-97a4-0fd929a27381)
551 new FirewallManager (/app/common/constructs/lib/firewall/manager.ts:22:21)
552   \ createFirewallManager (/app/initial-setup/templates/src/deployments/firewall/manager/step-1.ts:93:19)
553   \ Object.step1 (/app/initial-setup/templates/src/deployments/firewall/manager/step-1.ts:55:11)
554   \ deploy (/app/initial-setup/templates/src/apps/phase-2.ts:239:28)
555   \ processTicksAndRejections (internal/process/task_queues.js:97:5)
556   \ Object.deploy (/app/initial-setup/templates/src/app.ts:62:3)
557   \ main (/app/initial-setup/templates/cdk.ts:36:16)

```

In this example we can see that the resource `FirewallManager` failed to create through CloudFormation. One way to solve this issue is to deprovision the firewall manager in the configuration file and then run the state machine. Next, provision the firewall manager and run the state machine again.

If the error message is not clear, or the error occurred in a nested stack, then a more detailed error will be available in the CloudFormation stack events. See the [CloudFormation](#) section below.


```

584 1/3 | 10:42:52 PM | CREATE_FAILED | AWS::CloudFormation::Stack | FunAcctPhase1/VpcStackBrianFunVpc.NestedStack/VpcStackBrianFunVpc.NestedStackResource
(VpcStackBrianFunVpcNestedStackVpcStackBrianFunVpcNestedStackResource72499135) Embedded stack arn:aws:cloudformation:ca-central-1:144226684814:stack/PBMMAccel-
FunAcct-Phase1-VpcStackBrianFunVpcNestedStackVpcStackBrianFunVpcNestedStackR-1JUUYAR0LBT5J/F49e4700-bb22-11ea-8ced-066237fd0b7e was not successfully created: The
following resource(s) failed to create: [BrianFunVPCAppFunBrianFunVPCaza339451C7].
585 new NestedStack (/app/node_modules/.pnpm/@aws-cdk/core@1.46.0/node_modules/@aws-cdk/core/lib/nested-stack.ts:117:21)
586   \ new NestedStack (/app/node_modules/.pnpm/@aws-cdk/aws-cloudformation@1.46.0/node_modules/@aws-cdk/aws-cloudformation/lib/nested-stack.ts:67:5)
587   \ new VpcStack (/app/initial-setup/templates/src/common/vpc.ts:101:5)
588   \ createVpc (/app/initial-setup/templates/src/apps/phase-1.ts:160:22)
589   \ deploy (/app/initial-setup/templates/src/apps/phase-1.ts:262:17)
590   \ processTicksAndRejections (internal/process/task_queues.js:97:5)
591   \ Object.deploy (/app/initial-setup/templates/src/app.ts:62:3)
592   \ main (/app/initial-setup/templates/cdk.ts:36:16)

```

1.2.3. CloudFormation

In case you want to troubleshoot errors that occurred in CloudFormation, the best way is to look in the CloudFormation stack's events. This requires you to assume a role into the relevant sub-account, and to locate the relevant failed, rolled-back, or deleted stack. Unfortunately, we are unable to log the region of the error message, so depending on what's being deployed, you may need to search all 16 regions for the failed stack.



Events (300+)				
Timestamp	Logical ID	Status	Status reason	
2020-06-06 14:35:58 UTC+0200	NotebookepnotebookF1DF839C	DELETE_IN_PROGRESS	-	
2020-06-06 14:35:58 UTC+0200	NotebookEndpoint52F5778D	DELETE_COMPLETE	-	
2020-06-06 14:35:51 UTC+0200	PBMMAccel-SharedNetwork-Phase1-Endpoint1NestedStackEndpoint1NestedStackResourceBBB2BAD-1NI7TLNB83A1N	UPDATE_ROLLBACK_COMPLETE_CLEANUP_IN_PROGRESS	-	
2020-06-06 14:34:38 UTC+0200	PBMMAccel-SharedNetwork-Phase1-Endpoint1NestedStackEndpoint1NestedStackResourceBBB2BAD-1NI7TLNB83A1N	UPDATE_ROLLBACK_IN_PROGRESS	The following resource(s) failed to create: [NotebookEndpoint52F5778D].	
2020-06-06 14:34:37 UTC+0200	NotebookEndpoint52F5778D	CREATE_FAILED	Limit of 50 VPC endpoints per VPC exceeded. (Service: AmazonEC2; Status Code: 400; Error Code: VpcEndpointLimitExceeded; Request ID: 76f60902-2ba8-4681-99d0-46fa7d586100)	
2020-06-06 14:34:37 UTC+0200	NotebookEndpoint52F5778D	CREATE_IN_PROGRESS	-	
2020-06-06 14:34:33 UTC+0200	NotebookepnotebookF1DF839C	CREATE_COMPLETE	-	
2020-06-06 14:34:32 UTC+0200	NotebookepnotebookF1DF839C	CREATE_IN_PROGRESS	Resource creation Initiated	
2020-06-06 14:34:27 UTC+0200	NotebookepnotebookF1DF839C	CREATE_IN_PROGRESS	-	

[Load more](#)

When a native resource fails to create or update there are no additional logs available except what is displayed in the `Status reason` column. When a custom resource fails to create or update -- i.e. not a native CloudFormation resource but a resource backed by a custom Lambda function -- then we can find additional logs in CloudWatch.

Often the stack failure occurs in a managed account instead of the root account. See [Switch To a Managed Account](#) to switch to the CloudFormation console in the managed account.

1.2.4. Custom Resource

Custom resources are backed by a Lambda function that implements the creation, modification or deletion of the resource. Every Lambda function has a CloudWatch log group that contains logs about the custom resource creation. To troubleshoot errors in custom resource, you need to check the custom resource's log group.

Example custom resource log group names:

```
/aws/Lambda/ASEA-Master-Phase1-CustomCurReportDefinitionL-14IHLQCC1LY8L  
/aws/Lambda/ASEA-Master-Phase2-AWS679f53fac002430cb0da5b7-Z75Q4GG9LIV5  
/aws/Lambda/ASEA-Operations-Phas-AWS679f53fac002430cb0da5-HMV2YF60KJET  
/aws/Lambda/ASEA-Operations-Phas-CustomGetDetectorIdLambd-HEM07DR0D00J
```

1.2.5. CloudWatch

When you arrived in CloudWatch logs by clicking on the state machine's step `CloudWatch Logs` link you will immediately see the list of log streams. Every log stream represents an instance of the Lambda function.

You can find errors in multiple log groups using CloudWatch Log Insights.

CloudWatch > CloudWatch Logs > Logs Insights [Switch to the original interface.](#)

5m 30m 1h 3h 12h Custom (20w)

Clear /aws/lambda/PBMMAccel-Master-Phase0-CustomS3CopyFilesLambdaAF2-7Y1XBVZD29LL

```
fields @timestamp, @message
| sort @timestamp desc
| limit 100
| filter strcontains(@message, 'ERROR')
```

Run query
Save
History

Logs
Visualization
Export results ▼
Add to dashboard
⚙️

Showing 2 of 2 records matched ⓘ [Hide histogram](#)

27 records (6.2 kB) scanned in 2.7s @ 9 records/s (2.3 kB/s)

#	@timestamp	@message
▶ 1	2020-06-06T13:55:01...	2020-06-06T11:55:01.567Z c1c10696-20a9-4a46-b493-d8ae8660b726 ERROR Er
▶ 2	2020-06-06T13:55:01...	2020-06-06T11:55:01.546Z c1c10696-20a9-4a46-b493-d8ae8660b726 ERROR Se

```
fields @timestamp, @message
| sort @timestamp desc
| filter strcontains(@message, 'ERROR')
| limit 100
```

1.2.6. CodePipeline

- "Internal Failure" incorrect Github token, repo or branch

4.3.3 1.3. Examples

Lets walk through a couple of example:

1.3.1. Example 1

State Machine failed (Lambda), click on the grey box, then click on the Resource object:

The screenshot shows the AWS Step Functions console. On the left, a state machine graph is visible with a yellow arrow pointing to a failed step. On the right, the 'Details' tab is active, displaying the following information:

- Name:** Delete Default Vpcs
- Type:** Task
- Status:** Cancelled
- Resource:** `arn:aws:states:ca-central-1:991961644336:execution:PBMMAccel-DeleteDefaultVpcs_sfme751dc26-85f0-4476-a41e-83c27fec8b81`

Click on the red failed box, click on Step Input. The error is clearly indicated, we could not delete a Default VPC because the default VPC had dependencies, in a specified account and region. In this case several dependencies exist and need to be cleaned up to proceed (EIP's and something less obvious like security groups).

The screenshot shows the AWS Step Functions console. On the left, a state machine graph is visible with a yellow arrow pointing to a failed step. On the right, the 'Step Input' tab is active, displaying the following error message:

```

1 {
2   "status": "FAILED",
3   "errors": [
4     "991961644336:us-east-1: DependencyViolation: The subnet 'subnet-86187c08' has dependencies and cannot be deleted.",
5     "991961644336:us-east-1: DependencyViolation: Network vpc-2377885e has some mapped public address(es). Please unmap those public address(es) before detaching the gateway.",
6     "991961644336:us-east-1: DependencyViolation: The vpc 'vpc-2377885e' has dependencies and cannot be deleted."
7   ]
8 }

```

1.3.2. Example 2

In the next example the state machine failed (sub-state machine) on the create accounts step. In this case rather than clicking on the Graph inspector we are going to scroll down through the Execution event history underneath the Graph inspector. We are going to find the FIRST failed task from the top of the list and then select the state machine from the prior task:

Task ID	Task Name	Task Type	Task Status	Task ID	Task Time
50	TaskStateEntered	Create Organization Account	-	184422	Sep 4, 2020 09:39:08.517 AM
51	TaskScheduled	Create Organization Account	-	184422	Sep 4, 2020 09:39:08.517 AM
52	TaskStarted	Create Organization Account	-	184433	Sep 4, 2020 09:39:08.528 AM
53	TaskSubmitted	Create Organization Account	-	184507	Sep 4, 2020 09:39:08.602 AM
54	TaskFailed	Create Organization Account	-	188490	Sep 4, 2020 09:39:12.585 AM
55	MapIterationFailed	Create Organization Accounts	-	188490	Sep 4, 2020 09:39:12.585 AM
56	TaskStateAborted	Create Organization Account	-	188490	Sep 4, 2020 09:39:12.585 AM
57	MapStateFailed	Create Organization Account	-	188490	Sep 4, 2020 09:39:12.585 AM

We will then click on the red failed box, select Exception and we can see a clear error message - we have exceeded the maximum number of AWS accounts allowed in your organization:

The screenshot shows a state machine diagram on the left and the 'Exception' details on the right. The state machine starts with 'Start Account Creation', followed by 'Account Creation Started?', 'Wait for Org Account Creation', 'Verify Account Creation', 'Account Creation Done?', 'Account Creation Failed', 'Account Creation Succeeded', 'Move Account to Organizational Unit', 'Attach Quarantine SCP To Account', and 'End'. The 'Exception' tab shows the following error details:

```

Error
ConstraintViolationException
Cause
{
  "errorType": "ConstraintViolationException",
  "errorMessage": "You have exceeded the allowed number of AWS accounts.",
  "trace": [
    "ConstraintViolationException: You have exceeded the allowed number of AWS accounts.",
    "    at constructor.extractError (/var/task/index.js:125:6861)",
    "    at constructor.callListeners (/var/task/index.js:261:31984)",
    "    at constructor.emit (/var/task/index.js:261:31694)",
    "    at constructor.emitEvent (/var/task/index.js:314:556300)",
    "    at constructor.e (/var/task/index.js:314:551839)",
    "    at r.runTo (/var/task/index.js:314:558142)",
    "    at /var/task/index.js:314:558348".
  ]
}
    
```

Alternatively, in case the Exception error is not clear, we can select Details and then select CloudWatch logs for the end of the Resource section:

The screenshot shows the same state machine diagram on the left and the 'Details' tab on the right. The 'Resource' section contains the following information:

Name	Type
Start Account Creation	Task

Status: ⊗ Failed

Resource: `arn:aws:lambda:ca-central-1:991961644336:function:PBMMAccel-InitialSetup-StartAccountCreationHandler-1I2N4EP29D72`

[CloudWatch logs](#)

If you open the latest log stream in the opened log group (`/aws/Lambda/ASEA-InitialSetup-StartAccountCreationHandler-1I2N4EP29D72`) and review the last several messages in the stream, the following clear message also appears:

```

2020-09-04T13:39:11.411Z      a66ca383-77a1-4cdf-a431-732e10727cbd      ERROR      Invoke Error
{
  "errorType": "ConstraintViolationException",
  "errorMessage": "You have exceeded the allowed number of AWS accounts.",
  "code": "ConstraintViolationException",
  "message": "You have exceeded the allowed number of AWS accounts.",
  "time": "2020-09-04T13:39:11.410Z",
  "requestId": "2e6e879c-0968-4227-8b7d-61fbdf8abc84",
  "statusCode": 400,
  "retryable": false,
  "retryDelay": 22.311573790276995,
  "stack": [
    "ConstraintViolationException: You have exceeded the allowed number of AWS accounts.",
    "    at constructor.extractError (/var/task/index.js:125:6861)",
    "    at constructor.callListeners (/var/task/index.js:261:31984)",
    "    at constructor.emit (/var/task/index.js:261:31694)",
  ]
}
    
```

1.3.3. Example 3

In the next example the state machine failed in one of the CodeBuild state machine steps, based on the Resource name of the failed step.

Rather than tracing this failure through the sub-state machine and then into the failed CodeBuild task, we are simply going to open AWS CodeBuild, and open the ASEA-DeployPrebuilt task. The failed task should be on the top of the Codebuild build run list. Open the build job.

Using your browser, from the top of the page, search for "FAIL", and we are immediately brought to the error. In this particular case we had an issue with the creation of VPC endpoints. We defined something not supported by the current configuration file. The solution was to simply remove the offending endpoints from the config file and re-run the state machine.

```

792 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
793 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
794 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
795 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
796 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
797 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
798 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
799 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
800 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
801 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
802 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
803 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
804 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
805 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
806 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_IN_PROGRESS)
807 | 3:49:16 AM | CREATE_FAILED | AWS::CloudFormation::Stack | SharedNetworkPhaseUsEast_1/Endpoint1.NestedStack/Endpoint1.NestedStackResource
  (Endpoint1NestedStackEndpoint1NestedStackResourcee882BAD0) Embedded stack arn:aws:cloudformation:us-east-1:629131167886:stack/PBMMAccel-SharedNetwork-Phase1-
  Endpoint1NestedStackEndpoint1NestedStackResourcee882BAD0-CXP4HEBTHQ1Q/26ae0f20-fc86-11ea-99a7-0a3d3edde975 was not successfully created: The following resource(s) failed to create:
  [ElasticbeanstalkHealthPhz4237B56A, LogsEndpoint9655188C, EmailSmtptEndpoint4862062B, WorkspacesEndpoint22A2C83F, SageMakerRuntimeEndpointF8108889, CassandraEndpoint6F8BF867, GitCodecommitEndpoint21D91288,
  TransferPhz2088217B, StatesEndpoint23924C67, SageMakerApiPhzCB00297, KmsPhzC708808B, SnsEndpoint220C709B, SqsEndpoint338099E, AwsconnectorEndpointAD8802A, ClueEndpoint3M4416, Ec2ApiEndpointCD96854,
  KinesisStreamsEndpoint9E8AD98, ElasticbeanstalkEndpointD1FA8625, SnsPhz694BA985, LicenseManagerPhzAD68798E, ServicecatalogEndpoint3FFA91B, NorebookEndpoint3f5778D, GitCodecommitFipsEndpoint14ECF7F9,
  Macie2EndpointD07FD17, KinesisFirehosePhzAD38C94B, StoragegatewayEndpoint892D996A, AcmPcaEndpoint8D728D3C].
808   new NestedStack (/app/node_modules/.pnpm/@aws-cdk/core@1.46.0/node_modules/@aws-cdk/core/lib/nested-stack.ts:117:21)
809   \ new NestedStack (/app/node_modules/.pnpm/@aws-cdk/aws-cloudformation@1.46.0/node_modules/@aws-cdk/aws-cloudformation/lib/nested-stack.ts:67:5)
810   \ createVpc (/app/src/deployments/cdk/src/apps/phase-1.ts:168:27)
811   \ deploy (/app/src/deployments/cdk/src/apps/phase-1.ts:233:17)
812   \ processTicksAndRejections (internal/process/task_queues.js:97:5)
813   \ Object.deploy (/app/src/deployments/cdk/src/app.ts:78:3)
814   \ main (/app/src/deployments/cdk/cdk.ts:36:16)
815 Stack PBMMAccel-SharedNetwork-Phase1 is still not stable (UPDATE_ROLLBACK_IN_PROGRESS (The following resource(s) failed to create: [Endpoint0NestedStackEndpoint0NestedStackResource7E23D9FF,
  Endpoint1NestedStackEndpoint1NestedStackResourcee882BAD0, Endpoint2NestedStackEndpoint2NestedStackResource5E89388D].))
816 | 3:49:17 AM | CREATE_FAILED | AWS::CloudFormation::Stack | SharedNetworkPhaseUsEast_1/Endpoint2.NestedStack/Endpoint2.NestedStackResource
  (Endpoint2NestedStackEndpoint2NestedStackResource5E89388D) Resource creation cancelled
817   new NestedStack (/app/node_modules/.pnpm/@aws-cdk/core@1.46.0/node_modules/@aws-cdk/core/lib/nested-stack.ts:117:21)
818   \ new NestedStack (/app/node_modules/.pnpm/@aws-cdk/aws-cloudformation@1.46.0/node_modules/@aws-cdk/aws-cloudformation/lib/nested-stack.ts:67:5)
819   \ createVpc (/app/src/deployments/cdk/src/apps/phase-1.ts:168:27)
820   \ deploy (/app/src/deployments/cdk/src/apps/phase-1.ts:233:17)
821   \ processTicksAndRejections (internal/process/task_queues.js:97:5)
822   \ Object.deploy (/app/src/deployments/cdk/src/app.ts:78:3)
823   \ main (/app/src/deployments/cdk/cdk.ts:36:16)
  
```

4.4 1. Common Tasks

4.4.1 1.1. Restart the State Machine

The state machine can be stopped and restarted at any time. The Accelerator has been design to be able to rollback to a stable state, such that should the state machine be stopped or fail for any reason, subsequent state machine executions can simply proceed through the failed step without manual cleanup or issues (assuming the failure scenario has been resolved). An extensive amount of effort was placed on ensuring seamless customer recovery in failure situations. The Accelerator is idempotent - it can be run as many or as few times as desired with no negative effect. On each state machine execution, the state machine, primarily leveraging the capabilities of CDK, will evaluate the delta's between the old previously deployed configuration and the new configuration and update the environment as appropriate.

The state machine will execute:

- automatically after each execution of the Code Pipeline (new installs, code upgrades, or manual pipeline executions)
- automatically when new AWS accounts are moved into any Accelerator controller OU in AWS Organizations
- when someone manual starts it: `Step Functions`, `ASEA-MainStateMachine_sm`, `Start Execution`, `Start Execution` (leave default values in name and json box)

The state machine prevents users from accidentally performing certain major breaking changes, specifically unsupported AWS platform changes, changes that will fail to deploy, or changes that could be catastrophic to users. If someone knows exactly what they are doing and the full implications of these changes, we provide the option to override these checks. Customers should expect that items we have blocked CANNOT be changed after the Accelerator installation.

These flags should be used with extreme caution. Specifying any of these flags without proper guidance will likely leave your Accelerator in a state of disrepair. These flags were added for internal purposes only - we do NOT support customers providing these flags.

Providing this parameter to the state machine overrides *all* checks:

```
{
  "overrideComparison": true
}
```

Providing any one or more of the following flags will only override the specified check(s):


```
{
  "configOverrides": {
    "ov-global-options": true,
    "ov-del-accts": true,
    "ov-ren-accts": true,
    "ov-acct-email": true,
    "ov-acct-ou": true,
    "ov-acct-vpc": true,
    "ov-acct-subnet": true,
    "ov-tgw": true,
    "ov-mad": true,
    "ov-ou-vpc": true,
    "ov-ou-subnet": true,
    "ov-share-to-ou": true,
    "ov-share-to-accounts": true,
    "ov-nac": true,
    "ov-nfw": true
  }
}
```

Providing this value allows for the forced rebuilding of the DynamoDB Outputs table:

```
{
  "storeAllOutputs": true
}
```

4.4.2 1.2. Switch To a Managed Account

To switch from the root account to a managed account you can click on your account name in the AWS Console. Then choose `Switch Role` in the menu.



In the page that appears next you need to fill out the account ID of the managed account you want to switch to. Next, you need to enter the role name defined in `organization-admin-role` (default: `AWSCloudFormationStackSetAdministrationRole`). And lastly, you need to enter a relevant name so you can later switch roles by using this name.

TBD: This role may be locked down starting in v1.2.5 - Update process once direction finalized

Caution: This mechanism is ONLY to be used for troubleshooting Accelerator problems. This role is outside the Accelerator governance process and bypasses **all** the preventative guardrails that protect the Accelerator contracts and prevent users from performing activities in violation of the security guardrails. This role should NOT be used outside this context, all users should be authenticating and logging into the environment through AWS SSO.

Switch Role

Allows management of resources across AWS accounts using a single user ID and password. You can switch roles after an AWS administrator has configured a role and given you the account and role details. [Learn more.](#)

Account* ⓘ

Role* ⓘ

Display Name ⓘ

Color

*Required



Cancel

Switch Role

After switching to the managed account, the AWS Console header will look like the following image.



You can switch to the same account again quickly by clicking the name you entered previously in the menu.

<p>Logged in as: Admin</p> <p>Account: 3420-0203-3641</p> <hr/> <p>Role History:</p> <p> DevAcct </p> <p>Switch Role</p> <hr/> <p>Sign Out</p>	<p>Currently active as: PBMMAccel-PipelineRole</p> <p>Account: 6694-2100-9277</p> <hr/> <p>My Account</p> <p>My Organization</p> <p>My Service Quotas</p> <p>My Billing Dashboard</p> <p>Orders and Invoices</p> <p>Back to Admin</p> <hr/>
---	---

[1]: <https://docs.aws.amazon.com/cdk/latest/guide/home.html>

5. Developer Guide

5.1 Accelerator Developer Guide

It is important to read the [Operations Guide](#) before reading this document.

- [Development Guide](#)
- [Tech Stack](#)
- [Best Practices](#)
- [How to Contribute](#)
- [Release Process](#)

5.2 1. Development Guide

This document is a reference document. Instead of reading through it in linear order, you can use it to look up specific issues as needed.

It is important to read the [Operations Guide](#) before reading this document. If you're interested in actively contributing to the project, you should also review the [Governance and Contributing Guide](#).

5.2.1 1.1. Overview

There are different types of projects in this monorepo.

1. Projects containing CDK code that compiles to CloudFormation templates and deploy to AWS using the CDK toolkit;
2. Projects containing runtime code that is used by the CDK code to deploy Lambda functions;
3. Projects containing reusable code; both for use by the CDK code and/or runtime code.

The CDK code either deploys Accelerator-management resources or Accelerator-managed resources. See the [Operations Guide](#) for the distinction between Accelerator-management and Accelerator-managed resources.

The only language used in the project is TypeScript and exceptionally JavaScript. We do not write CloudFormation templates, only CDK code.

When we want to enable functionality in a managed account we try to

1. use native CloudFormation/CDK resource to enable the functionality;
2. create a custom resource to enable the functionality; or
3. lastly create a new step in the `Initial Setup` state machine to enable the functionality.

5.2.2 1.2. Project Structure

The folder structure of the project is as follows:

- `src/installer/cdk`: See [Installer Stack](#);
- `src/core/cdk`: See [Initial Setup Stack](#);
- `src/core/runtime`: See [Initial Setup Stack](#) and [Phase Steps and Phase Stacks](#);
- `src/deployments/runtime`: See [Phase Steps and Phase Stacks](#);
- `src/deployments/cdk`: See [Phase Steps and Phase Stacks](#);
- `src/lib/accelerator-cdk`: See [Libraries & Tools](#);
- `src/lib/cdk-constructs`: See [Libraries & Tools](#);
- `src/lib/cdk-plugin-assume-role`: See [CDK Assume Role Plugin](#).
- `src/lib/common-config`: See [Libraries & Tools](#);
- `src/lib/common-outputs`: See [Libraries & Tools](#);
- `src/lib/common-types`: See [Libraries & Tools](#);
- `src/lib/common`: See [Libraries & Tools](#);
- `src/lib/custom-resources/**/cdk`: See [Custom Resources](#);
- `src/lib/custom-resources/**/runtime`: See [Custom Resources](#);

5.2.3 1.3. Installer Stack

Read the [Operations Guide](#) first before reading this section. This section is a technical addition to the section in the Operations Guide.

As stated in the Operations Guide, the `Installer` stack is responsible for installing the `Initial Setup` stack. It is an Accelerator-management resource. The main resource in the `Installer` stack is the `ASEA-Installer` CodePipeline. The CodePipeline uses this GitHub repository as source action and runs CDK in a CodeBuild step to deploy the `Initial Setup` stack.

```
new codebuild.PipelineProject(stack, 'InstallerProject', {
  buildSpec: codebuild.BuildSpec.fromObject({
    version: '0.2',
    phases: {
      install: {
        'runtime-versions': {
          nodejs: 14,
        },
        // The flag '--unsafe-perm' is necessary to run npm scripts in Docker
        commands: ['npm install --global pnpm@6.2.3', 'pnpm install --unsafe-perm --frozen-lockfile'],
      },
      pre_build: {
        // The flag '--unsafe-perm' is necessary to run npm scripts in Docker
        commands: ['pnpm recursive run build --unsafe-perm'],
      },
      build: {
        commands: [
          'cd src/core/cdk',
          // Bootstrap the environment for use by CDK
          'pnpx cdk bootstrap --require-approval never',
          // Deploy the Initial Setup stack
          'pnpx cdk deploy --require-approval never',
        ],
      },
    },
  }),
});
```

When the CodePipeline finishes deploying the `Initial Setup` stack, it starts a Lambda function that starts the execution of the `Initial Setup` stack's main state machine.

The `Initial Setup` stack deployment receives environment variables from the CodePipeline's CodeBuild step. The most notable environment variables are:

- `ACCELERATOR_STATE_MACHINE_NAME`: The `Initial Setup` will use this name for the main state machine. So it is the `Installer` stack that decides the name of the main state machine. This way we can confidently start the main state machine of the `Initial Setup` stack from the CodePipeline;
- `ENABLE_PREBUILT_PROJECT`: See [Prebuilt Docker Image](#).

5.2.4 1.4. Initial Setup Stack

Read [Operations Guide](#) first before reading this section. This section is a technical addition to the section in the Operations Guide.

As stated in the Operations Guide, the `Initial Setup` stack consists of a state machine, named `ASEA-MainStateMachine_sm`, which executes steps to create the Accelerator-managed stacks and resources in the managed accounts. It is an Accelerator-management resource.

The `Initial Setup` stack is defined in the `src/core/cdk` folder.

The `Initial Setup` stack is similar to the `Installer` stack, as in that it runs a CodeBuild project to deploy others stacks using CDK. In case of the `Initial Setup` stack

- we use a AWS Step Functions State Machine to run steps instead of using a CodePipeline;
- we deploy multiple stacks, called `Phase` stacks, in Accelerator-managed accounts. These `Phase` stacks contain Accelerator-managed resources.

In order to install these `Phase` stacks in Accelerator-managed accounts, we need access to those accounts. We create a stack set in the Organization Management (root) account that has instances in all Accelerator-managed accounts. This stack set contains what we call the `PipelineRole`.

The code for the steps in the state machine is in `src/core/runtime`. All the steps are in different files but are compiled into a single file. We used to compile all the steps separately but we would hit a limit in the amount of parameters in the generated CloudFormation template. Each step would have its own CDK asset that would introduce three new parameters. We quickly reached the limit of 60 parameters in a CloudFormation template and decided to compile the steps into a single file and use it across all different Lambda functions.

1.4.1.1. CodeBuild and Prebuilt Docker Image

The CodeBuild project that deploys the different `Phase` stacks is constructed using the `CdkDeployProject` or `PrebuiltCdkDeployProject` based on the value of the environment variable `ENABLE_PREBUILT_PROJECT`.

The first, `CdkDeployProject` constructs a CodeBuild project that copies this whole Github repository as a ZIP file to S3 using [CDK S3 assets](#). This ZIP file is then used as source for the CodeBuild project. When the CodeBuild project executes, it runs `pnpm recursive install` which in turn will run all `prepare` scripts in all `package.json` files in the project -- as described in section [CDK Code Dependency on Lambda Function Code](#).

After installing the dependencies, the CodeBuild project deploys the `Phase` stacks.

```
cd src/deployments/cdk
sh codebuild-deploy.sh
```

We have more than 50 workspace projects in the monorepo with a `prepare` script, so the `pnpm recursive install` step can take some time. Also, the CodeBuild project will run for each deployed `Phase` stack in each Accelerator-managed account.

This is where the `PrebuiltCdkDeployProject` CodeBuild project comes in. The `PrebuiltCdkDeployProject` contains a Docker image that contains the whole project in the `/app` directory and has all the dependencies already installed.

```
FROM node:12-alpine3.11
# Install the package manager
RUN npm install --global pnpm
RUN mkdir /app
WORKDIR /app
# Copy over the project root to the /app directory
ADD . /app/
# Install the dependencies
RUN pnpm install --unsafe-perm --frozen-lockfile
# Build all Lambda function runtime code
RUN pnpm recursive run build --unsafe-perm
```

When this CodeBuild project executes, it uses the Docker image as base -- the dependencies are already installed -- and runs the same commands as the `CdkDeployProject` to deploy the `Phase` stacks.

1.4.2. Passing Data to Phase Steps and Phase Stacks

Some steps in the state machine write data to Amazon DynamoDB. This data is necessary to deploy the `Phase` stacks later on. At one time this data was written to Secrets Manager and/or S3, these mechanisms were deemed ineffective due to object size limitations or consistency challenges and were all eventually migrated to DynamoDB.

- `Load Accounts` step: This step finds the Accelerator-managed accounts in AWS Organizations and stores the account key -- the key of the account in `mandatory-account-configs` or `workload-account-configs` object in the Accelerator config -- and account ID and other useful information in the `ASEA-Parameters` table, `accounts/#` key and `accounts-items-count` key;
- `Load Organizations` step: More or less the same as the `Load Accounts` step but for organizational units in AWS Organizations and stores the values in the `ASEA-Parameters` table, `organizations` key;
- `Load Limits` step: This step requests limit increases for Accelerator-managed accounts and stores the current limits in the the `ASEA-Parameters` table, `Limits` key.
- `Store Phase X Output`: This step loads stack outputs from all existing `Phase` stacks and stores the outputs in the DynamoDB table `ASEA-Outputs`.

Other data is passed through environment variables:

- `ACCELERATOR_NAME`: The name of the Accelerator;
- `ACCELERATOR_PREFIX`: The prefix for all named Accelerator-managed resources;
- `ACCELERATOR_EXECUTION_ROLE_NAME`: The name of the execution role in the Accelerator-managed accounts. This is the `PipelineRole` we created using stack sets.

5.2.5 1.5. Phase Steps and Phase Stacks

Read [Operations Guide](#) first before reading this section. This section is a technical addition to the `Deploy Phase X` sections in the Operations Guide.

The `Phase` stacks contain the Accelerator-managed resources. The reason the deployment of Accelerator-managed resources is split into different phases is because there cannot be cross account/region references between CloudFormation stacks. See [Cross-Account/Region References](#).

The `Phase` stacks are deployed by a CodeBuild project in the `Initial Setup` stack as stated in the previous paragraphs. The CodeBuild project executes the `codebuild-deploy.sh` script. See [initial-setup.ts](#).

The `codebuild-deploy.sh` script executes the `cdk.ts` file.

The `cdk.ts` file is meant as a replacement for the `cdk` CLI command. To deploy a phase stack you would **not** run `pnpx cdk deploy` but `cdk.sh --phase 1`. See [CDK API](#) for more information why we use the CDK API instead of using the CDK CLI.

The `cdk.ts` command parses command line arguments and creates all the `cdk.App` for all accounts and regions for the given `--phase`. When you pass the `--region` or `--account-key` command, all the `cdk.App` for all accounts and regions will still be created, except that only the `cdk.App` s matching the parameters will be deployed. This behavior could be optimized in the future. See [Stacks with Same Name in Different Regions](#) for more information why we're creating multiple `cdk.App` s.

5.2.6 1.6. Store outputs to SSM Parameter Store

Customers need the CloudFormation outputs of resources that are created by the accelerator in order to deploy their own resources in AWS. eg. `vpcId` in shared-network account to create an `ec2` instance, etc.

This step loads the stack outputs from our DynamoDB Table `ASEA-Outputs` and stores as key value pairs in SSM Parameter Store in each account.

Example values are

- `/ASEA/network/vpc/1/name => Endpoint`
- `/ASEA/network/vpc/1/id => vpc-XXXXXXXXXX`

`ASEA-Outputs-Util` s DynamoDB Table is used extensively to maintain same index irrespective of configuration changes.

This allows customers to reliably build Infrastructure as Code (IaC) which depends on accelerator deployed objects like VPC's, security groups, subnets, ELB's, KMS keys, IAM users and policies. Rather than making the parameters dependent on object names, we used an indexing scheme, which we maintain and don't update as a customers configuration changes. We have attempted to keep the index values consistent across accounts (based on the config file), such that when code is promoted through the SDLC cycle from Dev to Test to Prod, the input parameters to the IaC scripts do not need to be updated, the `App` subnet, for example, will have the same index value in all accounts.

1.6.1. Phases and Deployments

The `cdk.ts` file calls the `deploy` method in the `apps/app.ts`. This `deploy` method loads the Accelerator configuration, accounts, organizations from DynamoDB; loads the stack outputs from Amazon DynamoDB; and loads required environment variables.

```
/**
 * Input to the `deploy` method of a phase.
 */
export interface PhaseInput {
  // The config.json file
  acceleratorConfig: AcceleratorConfig;
  // Auxiliary class to construct stacks
  accountStacks: AccountStacks;
  // The list of accounts, their key in the configuration file and their ID
  accounts: Account[];
  // The parsed environment variables
  context: Context;
  // The list of stack outputs from previous phases
  outputs: StackOutput[];
  // Auxiliary class to manage limits
  limiter: Limiter;
}
```

It is important to note that no configuration is hard-coded. The CloudFormation templates are generated by CDK and the CDK constructs are created according to the configuration file. Changes to the configuration will change the CDK construct tree and that will result in a different CloudFormation template that is deployed.

The different phases are defined in `apps/phase-x.ts`. Historically we created all CDK constructs in the `phase-x.ts` files. After a while the `phase-x.ts` files started to get too big and we moved to separating the logic into separate deployments. Every logical component has a separate folder in the `deployments` folder. Every `deployment` consists of so-called steps. Separate steps are put in loaded in phases.

For example, take the `deployments/defaults` deployment. The deployment consists of two steps, i.e. `step-1.ts` and `step-2.ts`. `deployments/defaults/step-1.ts` is created in `apps/phase-0.ts` and `deployments/defaults/step-2.ts` is created in `apps/phase-1.ts`. You can find more details about what happens in each phase in the [Operations Guide](#).

`apps/phase-0.ts`

```
export async function deploy({ acceleratorConfig, accountStacks, accounts, context }: PhaseInput) {
  // Create defaults, e.g. S3 buckets, EBS encryption keys
  const defaultsResult = await defaults.step1({
    acceleratorPrefix: context.acceleratorPrefix,
    accountStacks,
    accounts,
    config: acceleratorConfig,
  });
}
```

`apps/phase-1.ts`

```
export async function deploy({ acceleratorConfig, accountStacks, accounts, outputs }: PhaseInput) {
  // Find the central bucket in the outputs
  const centralBucket = CentralBucketOutput.getBucket({
    accountStacks,
    config: acceleratorConfig,
    outputs,
  });

  // Find the log bucket in the outputs
  const logBucket = LogBucketOutput.getBucket({
    accountStacks,
    config: acceleratorConfig,
    outputs,
  });

  // Find the account buckets in the outputs
  const accountBuckets = await defaults.step2({
    accounts,
    accountStacks,
    centralLogBucket: logBucket,
    config: acceleratorConfig,
  });
}
```

1.6.2. Passing Outputs between Phases

The CodeBuild step that is responsible for deploying a `Phase` stack runs in the Organization Management (root) account. We wrote a CDK plugin that allows the CDK deploy step to assume a role in the Accelerator-managed account and create the CloudFormation `Phase` stack in the managed account. See [CDK Assume Role Plugin](#).

After a `Phase-X` is deployed in all Accelerator-managed accounts, a step in the `Initial Setup` state machine collects all the `Phase-X` stack outputs in all Accelerator-managed accounts and regions and stores these outputs in DynamoDB.

Then the next `Phase-X+1` deploys using the outputs from the previous `Phase-X` stacks.

See [Creating Stack Outputs](#) for helper constructs to create outputs.

1.6.3. Decoupling Configuration from Constructs

At the start of the project we created constructs that had tight coupling to the Accelerator config structure. The properties to instantiate a construct would sometimes have a reference to an Accelerator-specific interface. An example of this is the `Vpc` construct in `src/deployments/cdk/common/vpc.ts`.

Later on in the project we started decoupling the Accelerator config from the construct properties. Good examples are in `src/lib/cdk-constructs/`.

Decoupling the configuration from the constructs improves reusability and robustness of the codebase.

5.2.7 1.7. Libraries & Tools

1.7.1. CDK Assume Role Plugin

At the time of writing, CDK does not support cross-account deployments of stacks. It is possible however to write a CDK plugin and implement your own credential loader for cross-account deployment.

We wrote a CDK plugin that can assume a role into another account. In our case, the Organization Management (root) account will assume the `PipelineRole` in an Accelerator-managed account to deploy stacks.

1.7.2. CDK API

We use the internal CDK API to deploy the `Phase` stacks instead of the CDK CLI for the following reasons:

- It allows us to deploy multiple stacks in parallel;
- Disable stack termination before destroying a stack;
- Delete a stack after it initially failed to create;
- Deploy multiple apps at the same time -- see [Stacks with Same Name in Different Regions](#).

The helper class `CdkToolkit` in `toolkit.ts` wraps around the CDK API.

The risk of using the CDK API directly is that the CDK API can change at any time. There is no stable API yet. When upgrading the CDK version, the `CdkToolkit` wrapper might need to be adapted.

1.7.3. AWS SDK Wrappers

You can find `aws-sdk` wrappers in the `src/lib/common/src/aws` folder. Most of the classes and functions just wrap around `aws-sdk` classes and implement promises and exponential backoff to retryable errors. Other classes, like `Organizations` have additional functionality such as listing all the organizational units in an organization in the function `ListOrganizationalUnits`.

Please use the `aws-sdk` wrappers throughout the project or write an additional wrapper when necessary.

1.7.4. Configuration File Parsing

The configuration file is defined and validated using the [io-ts](#) library. See `src/lib/common-config/src/index.ts`. In case any changes need to be made to the configuration file parsing, this is the place to be.

We wrap a class around the `AcceleratorConfig` type that contains additional helper functions. You can add your own additional helper functions.

1.7.4.1. ACCELERATORNAME_TAGGER

`AcceleratorNameTagger` is a [CDK aspect](#) that sets the name tag on specific resources based on the construct ID of the resource.

The following example illustrates its purpose.

```
const stack = new cdk.Stack();
new ec2.CfnVpc(stack, 'SharedNetwork', {});
Aspects.of(stack).add(new AcceleratorNameTagger());
```

The example above synthesizes to the following CloudFormation template.

```
Resources:
  SharedNetworkAB7JKF7:
    Properties:
      Tags:
        - Key: Name
          Value: SharedNetwork_vpc
```

1.7.4.2. ACCELERATORSTACK

`AcceleratorStack` is a class that extends `cdk.Stack` and adds the `Accelerator` tag to all resources in the stack. It also applies the aspect `AcceleratorNameTagger`.

It is also used by the `accelerator-name-generator.ts` functions to find the name of the `Accelerator`.

1.7.4.3. NAME GENERATOR

The `accelerator-name-generator.ts` file contains methods that create names for resources that are optionally prefixed with the Accelerator name, and optionally suffixed with a hash based on the path of the resource, the account ID and region of the stack.

The functions should be used to create pseudo-random names for IAM roles, KMS keys, key pairs and log groups.

1.7.4.4. ACCOUNTSTACKS

`AccountStacks` is a class that manages the creation of an `AcceleratorStack` based on a given account key and region. If an account with the given account key cannot be found in the accounts object -- which is loaded by `apps/app.ts` then no stack will be created. This class is used extensively throughout the phases and deployment steps.

```
export async function step1(props: CertificatesStep1Props) {
  const { accountStacks, centralBucket: centralBucket, config } = props;

  for (const { accountKey, certificates } of config.getCertificateConfigs()) {
    if (certificates.length === 0) {
      continue;
    }

    const accountStack = accountStacks.tryGetOrCreateAccountStack(accountKey);
    if (!accountStack) {
      console.warn("Cannot find account stack ${accountKey}");
      continue;
    }

    for (const certificate of certificates) {
      createCertificate({
        centralBucket,
        certificate,
        scope: accountStack,
      });
    }
  }
}
```

1.7.4.5. VPC AND IMPORTEDVPC

`Vpc` is an interface in the `src/lib/cdk-constructs/src/vpc/vpc.ts` file that attempts to define an interface for a VPC. The goal of the interface is to be implemented by a `cdk.Construct` that implements the interface. This CDK issue provides more background [<https://github.com/aws/aws-cdk/issues/5927>].

Another goal of the interface is to provide an interface on top of imported VPC outputs. This is what the `ImportedVpc` class implements. The class loads outputs from VPC in a previous phase and implements the `Vpc` interface on top of those outputs.

1.7.4.6. LIMITER

So far we haven't talked about limits yet. There is a step in the `Initial Setup` state machine that requests limit increases according to the desired limits in the configuration file. The step saves the current limits to the `Limits` key in the DynamoDB table `ASEA-Parameters`. The `apps/app.ts` file loads the limits and passes them as an input to the phase deployment.

The `Limiter` class helps keeps track of resource we create and prevents exceeding these limits.

```
for (const { ouKey, accountKey, vpcConfig, deployments } of acceleratorConfig.getVpcConfigs()) {
  if (!limiter.create(accountKey, Limit.VpcPerRegion, region)) {
    console.log("Skipping VPC "${vpcConfig.name}" deployment.");
    console.log("Reached maximum VPCs per region for account "${accountKey}" and region "${region}");
    continue;
  }

  createVpc({ ouKey, accountKey, vpcConfig });
}
```

Action Item: This functionality could be redesigned to scan all the constructs in a `cdk.App` and remove resource that are exceeding any limits.

1.7.5. Creating Stack Outputs

Initially we would create stack outputs like this:

```
new cdk.CfnOutput(stack, 'BucketOutput', {
  value: bucket.bucketArn,
});
```

But then we'd get a lot of outputs in a stack. We started some outputs together using JSON. This allowed us to store structured data inside the stack outputs.

```
new JsonOutputValue(stack, 'Output', {
  type: 'FirewallInstanceOutput',
  value: {
    instanceId: instance.instanceId,
    name: firewallConfig.name,
    az,
  },
});
```

Using the solution above, we'd not have type checking when reading or writing outputs. That's what the class `StructuredOutputValue` has a solution for. It uses the `io-ts` library to serialize and deserialize structured types.

```
export const FirewallInstanceOutput = t.interface(
  {
    id: t.string,
    name: t.string,
    az: t.string,
  },
  'FirewallInstanceOutput',
);

export type FirewallInstanceOutput = t.TypeOf<typeof FirewallInstanceOutput>;

new StructuredOutputValue<FirewallInstanceOutput>(stack, 'Output', {
  type: FirewallInstanceOutput,
  value: {
    instanceId: instance.instanceId,
    name: firewallConfig.name,
    az,
  },
});
```

And we can even improve on this a bit more.

```
export const CfnFirewallInstanceOutput = createCfnStructuredOutput(FirewallInstanceOutput);

new CfnFirewallInstanceOutput(stack, 'Output', {
  vpcId: vpc.ref,
  vpcName: vpcConfig.name,
});
```

```
export const FirewallInstanceOutputFinder = createStructuredOutputFinder(FirewallInstanceOutput, () => ({}));

// Create an OutputFinder
const firewallInstances = FirewallInstanceOutputFinder.findAll({
  outputs,
  accountKey,
});

// Example usage of the OutputFinder
const firewallInstance = firewallInstances.find(i => i.name === target.name && i.az === target.az);
```

Generally you would place the output type definition inside `src/lib/common-outputs` along with the output finder. Then in the deployment folder in `src/deployments/cdk/deployments` you would create an `output.ts` file where you would define the CDK output type with `createCfnStructuredOutput`. You would not define the CDK output type in `src/lib/common-outputs` since that project is also used by runtime code that does not need to know about CDK and CloudFormation.

1.7.5.1. ADDING TAGS TO SHARED RESOURCES IN DESTINATION ACCOUNT

There is another special type of output, `AddTagsToResourcesOutput`. It can be used to attach tags to resources that are shared into another account.

```
new AddTagsToResourcesOutput(this, 'OutputSharedResourcesSubnets', {
  dependencies: sharedSubnets.map(o => o.subnet),
  produceResources: () =>
    sharedSubnets.map(o => ({
      resourceId: o.subnet.ref,
      resourceType: 'subnet',
      sourceAccountId: o.sourceAccountId,
      targetAccountIds: o.targetAccountIds,
      tags: o.subnet.tags.renderTags(),
    })),
});
```

This will add the outputs to the stack in the account that is initiating the resource share.

Next, the state machine step `Add Tags to Shared Resources` looks for all those outputs. The step will assume the `PipelineRole` in the `targetAccountIds` and attach the given tags to the shared resource.

1.7.6. Custom Resources

There are different ways to create a custom resource using CDK. See the [Custom Resource](#) section for more information.

All custom resources have a `README.md` that demonstrates their usage.

1.7.6.1. EXTERNALIZING `aws-sdk`

Some custom resources set the `aws-sdk` as external dependency and some do not.

Example of setting `aws-sdk` as external dependency.

`src/lib/custom-resources/cdk-kms-grant/runtime/package.json`

```
{
  "externals": ["aws-lambda", "aws-sdk"],
  "dependencies": {
    "aws-lambda": "1.0.6",
    "aws-sdk": "2.631.0"
  }
}
```

Example of setting `aws-sdk` as embedded dependency.

`src/lib/custom-resources/cdk-guardduty-enable-admin/runtime/package.json`

```
{
  "externals": ["aws-lambda"],
  "dependencies": {
    "aws-lambda": "1.0.6",
    "aws-sdk": "2.711.0"
  }
}
```

Setting the `aws-sdk` library as external is sometimes necessary when a newer `aws-sdk` version is necessary for the Lambda runtime code. At the time of writing the NodeJS 12 runtime uses `aws-sdk` version `2.631.0`

For example the method `AWS.GuardDuty.enableOrganizationAdminAccount` was only introduced in `aws-sdk` version `2.660`. That means that Webpack has to embed the `aws-sdk` version specified in `package.json` into the compiled JavaScript file. This can be achieved by removing `aws-sdk` from the `external` array.

`src/lib/custom-resources/cdk-kms-grant/runtime/package.json`

1.7.6.2. CFN-RESPONSE

This library helps you send a custom resource response to CloudFormation.

`src/lib/custom-resources/cdk-kms-grant/runtime/src/index.ts`

```
export const handler = errorHandler(onEvent);

async function onEvent(event: CloudFormationCustomResourceEvent) {
  console.log('Creating KMS grant...');
  console.log(JSON.stringify(event, null, 2));

  // eslint-disable-next-line default-case
  switch (event.RequestType) {
    case 'Create':
      return onCreate(event);
    case 'Update':
      return onUpdate(event);
    case 'Delete':
      return onDelete(event);
  }
}
```

1.7.6.3. CFN-TAGS

This library helps you send attaching tags to resource created in a custom resource.

1.7.6.4. WEBPACK-BASE

This library defines the base Webpack template to compile custom resource runtime code.

```
src/lib/custom-resources/cdk-kms-grant/runtime/package.json
```

```
{
  "name": "@aws-accelerator/custom-resource-kms-grant-runtime",
  "version": "0.0.1",
  "private": true,
  "scripts": {
    "prepare": "webpack-cli --config webpack.config.ts"
  },
  "source": "src/index.ts",
  "main": "dist/index.js",
  "types": "dist/index.d.ts",
  "externals": ["aws-lambda", "aws-sdk"],
  "devDependencies": {
    "@aws-accelerator/custom-resource-runtime-webpack-base": "workspace:*0.0.1",
    "@types/aws-lambda": "8.10.46",
    "@types/node": "14.14.31",
    "ts-loader": "7.0.5",
    "typescript": "3.8.3",
    "webpack": "4.42.1",
    "webpack-cli": "3.3.11"
  },
  "dependencies": {
    "@aws-accelerator/custom-resource-runtime-cfn-response": "workspace:*0.0.1",
    "aws-lambda": "1.0.6",
    "aws-sdk": "2.668.0"
  }
}
```

```
src/lib/custom-resources/cdk-ec2-image-finder/runtime/webpack.config.ts
```

```
import { webpackConfigurationForPackage } from '@aws-accelerator/custom-resource-runtime-webpack-base';
import pkg from './package.json';

export default webpackConfigurationForPackage(pkg);
```

5.2.8 1.8. Workarounds

1.8.1. Stacks with Same Name in Different Regions

The reason we're creating a `cdk.App` per account and per region and per phase is because stack names across environments might overlap, and at the time of writing, the CDK CLI does not handle stacks with the same name well. For example, when there is a stack `Phase1` in `us-east-1` and another stack `Phase1` in `ca-central-1`, the stacks will both be synthesized by CDK to the `cdk.out/Phase1.template.json` file and one stack will overwrite another's output. Using multiple `cdk.App`s overcomes this issues as a different `outdir` can be set on each `cdk.App`. These `cdk.App`s are managed by the `AccountStacks` abstraction.

5.2.9 1.9. Local Development

1.9.1. Local Installer Stack

Use CDK to synthesize the CloudFormation template.

```
cd src/installer/cdk
pnpx cdk synth
```

The installer template file is now in `cdk.out/AcceleratorInstaller.template.json`. This file can be used to install the installer stack.

You can also deploy the installer stack directly from the command line but then you'd have to pass some stack parameters. See [CDK documentation: Deploying with parameters](#).

```
cd accelerator/installer
pnpx cdk deploy --parameters GithubBranch=main --parameters ConfigS3Bucket=ASEA-myconfigbucket
```

1.9.2. Local Initial Setup Stack

There is a script called `cdk.sh` in `src/core/cdk` that allows you to deploy the Initial Setup stack.

The script sets the required environment variables and makes sure all workspace projects are built before deploying the CDK stack.

1.9.3. Phase Stacks

There is a script called `cdk.sh` in `src/deployments/cdk` that allows you to deploy a phase stack straight from the command-line without having to deploy the Initial Setup stack first.

The script enables development mode which means that accounts, organizations, configuration, limits and outputs will be loaded from the local environment instead of loading the values from DynamoDB. The local files that need to be available in the `src/deployments/cdk` folder are the following.

1. `accounts.json` based on `accelerator/accounts` (-Parameters table)

```
[
  {
    "key": "shared-network",
    "id": "000000000001",
    "arn": "arn:aws:organizations::000000000000:account/o-0123456789/000000000001",
    "name": "myacct-ASEA-shared-network",
    "email": "myacct+ASEA-mandatory-shared-network@example.com",
    "ou": "core"
  },
  {
    "key": "operations",
    "id": "000000000002",
    "arn": "arn:aws:organizations::000000000000:account/o-0123456789/000000000002",
    "name": "myacct-ASEA-operations",
    "email": "myacct+ASEA-mandatory-operations@example.com",
    "ou": "core"
  }
]
```

1. `organizations.json` based on `accelerator/organizations` (-Parameters table)

```
[
  {
    "ouId": "ou-0000-00000000",
    "ouArn": "arn:aws:organizations::000000000000:ou/o-0123456789/ou-0000-00000000",
    "ouName": "core",
    "ouPath": "core"
  },
  {
    "ouId": "ou-0000-00000001",
    "ouArn": "arn:aws:organizations::000000000000:ou/o-0123456789/ou-0000-00000001",
    "ouName": "prod",
    "ouPath": "prod"
  }
]
```

1. `limits.json` based on `accelerator/limits` (-Parameters table)

```
[
  {
    "accountKey": "shared-network",
    "limitKey": "Amazon VPC/VPCs per Region",
    "serviceCode": "vpc",
    "quotaCode": "L-F678F1CE",
    "value": 15,
    "region": "ca-central-1"
  },
  {
    "accountKey": "shared-network",
    "limitKey": "Amazon VPC/Interface VPC endpoints per VPC",
    "serviceCode": "vpc",
    "quotaCode": "L-29B6F2EB",
    "value": 50,
    "region": "ca-central-1"
  }
]
```

1. `outputs.json` based on the -Outputs table

```
[
  {
    "accountKey": "shared-network",
    "outputKey": "DefaultBucketOutputC7CE5936",
    "outputValue": "{\n  \"type\": \"AccountBucket\",\n  \"value\": {\n    \"bucketArn\": \"arn:aws:s3:::ASEA-sharednetwork-phase1-ca-central-1-18vq0emthri3h\",\n    \"bucketName\": \"ASEA-sharednetwork-phase1-ca-central-1-18vq0emthri3h\",\n    \"encryptionKeyArn\": \"arn:aws:kms:ca-central-1:0000000000001:key/d54a8acb-694c-4fc5-9afe-ca2b263cd0b3\",\n    \"region\": \"ca-central-1\"\n  }\n}"
  }
]
```

1. `context.json` that contains the default values for values that are otherwise passed as environment variables.

```
{
  "acceleratorName": "ASEA",
  "acceleratorPrefix": "ASEA-",
  "acceleratorExecutionRoleName": "ASEA-PipelineRole",
  "defaultRegion": "ca-central-1"
}
```

1. `config.json` that contains the Accelerator configuration.

The script also sets the default execution role to allow CDK to assume a role in subaccounts to deploy the phase stacks.

Now that you have all the required local files you can deploy the phase stacks using `cdk.sh`.

```
cd src/deployments/cdk
./cdk.sh deploy --phase 1 # deploy all phase 1 stacks
./cdk.sh deploy --phase 1 --parallel # deploy all phase 1 stacks in parallel
./cdk.sh deploy --phase 1 --account shared-network # deploy phase 1 stacks for account shared-network in all regions
./cdk.sh deploy --phase 1 --region ca-central-1 # deploy phase 1 stacks for region ca-central-1 for all accounts
./cdk.sh deploy --phase 1 --account shared-network --region ca-central-1 # deploy phase 1 stacks for account shared-network and region ca-central
```

Other CDK commands are also available.

```
cd src/deployments/cdk
./cdk.sh bootstrap --phase 1
./cdk.sh synth --phase 1
```

5.2.10 1.10. Testing

We use `jest` for unit testing. There are no integration tests but this could be set-up by configuring the `Installer CodePipeline` to have a webhook on the repository and deploying changes automatically.

To run unit tests locally you can run the following command in the monorepo.

```
pnpx recursive run test -- --pass-with-no-tests --silent
```

See CDK's documentation on [Testing constructs](#) for more information on how to tests CDK constructs.

1.10.1. Validating Immutable Property Changes and Logical ID Changes

The most important unit test in this project is one that validates that logical IDs and immutable properties do not change unexpectedly. To avoid the issues described in section [Resource Names and Logical IDs](#), [Changing Logical IDs](#) and [Changing \(Immutable\) Properties](#).

This test can be found in the `src/deployments/cdk/test/apps/unsupported-changes.spec.ts` file. It synthesizes the `Phase` stacks using mocked outputs and uses [jest snapshots](#) to compare against future changes.

The test will fail when changing immutable properties or changing logical IDs of existing resources. In case the changes are expected then the snapshots will need to be updated. You can update the snapshots by running the following command.

```
pnpx run test -- -u
```

See [Accept Unit Test Snapshot Changes](#).

1.10.2. Upgrade CDK

There's a test in the file `src/deployments/cdk/test/apps/unsupported-changes.spec.ts` that is currently commented out. The test takes a snapshot of the whole `Phase` stack and compares the snapshot to changes in the code.

```
test('templates should stay exactly the same', () => {
  for (const [stackName, resources] of Object.entries(stackResources)) {
    // Compare the relevant properties to the snapshot
    expect(resources).toMatchSnapshot(stackName);
  }
});
```

Before upgrading CDK we uncomment this test. We run the test to update all the snapshots. Then we update all CDK versions and run the test again to compare the snapshots with the code using the new CDK version. If the test passes, then the upgrade should be stable.

Action Item: Automate this process.

5.3 1. Technology Stack

5.3.1 1.1. Overview

We use TypeScript, NodeJS, CDK and CloudFormation. You can find some more information in the sections below.

5.3.2 1.2. TypeScript and NodeJS

In the following sections we describe the tools and libraries used along with TypeScript.

1.2.1. pnpm

We use the `pnpm` package manager along with `pnpm workspaces` to manage all the packages in this monorepo.

<https://pnpm.js.org>

<https://pnpm.js.org/en/workspaces>

The binary `pnpx` runs binaries that belong to `pnpm` packages in the workspace.

<https://pnpm.js.org/en/pnpx-cli>

1.2.2. prettier

We use [prettier](#) to format code in this repository. A GitHub action makes sure that all the code in a pull requests adheres to the configured `prettier` rules. See [Github Actions](#).

1.2.3. eslint

We use [eslint](#) as a static analysis tool that checks our TypeScript code. A GitHub action makes sure that all the code in a pull requests adheres to the configured `eslint` rules. See [Github Actions](#).

5.3.3 1.3. CloudFormation

CloudFormation deploys both the Accelerator stacks and resources and the deployed stacks and resources. See [Operations Guide: System Overview](#) for the distinction between Accelerator resources and deployed resources.

5.3.4 1.4. CDK

AWS CDK defines the cloud resources in a familiar programming language. While AWS CDK supports TypeScript, JavaScript, Python, Java, and C#.Net, the contributions should be made in Typescript, as outlined in the [Accelerator Development First Principles](#).

Developers can use programming languages to define reusable cloud components known as Constructs. You compose these together into Stacks and Apps. Learn more at <https://docs.aws.amazon.com/cdk/latest/guide/home.html>

5.4 1. Best Practices

5.4.1 1.1. TypeScript and NodeJS

1.1.1.1. Handle Unhandled Promises

Entry point TypeScript files -- files that start execution instead of just defining methods and classes -- should have the following code snippet at the start of the file.

```
process.on('unhandledRejection', (reason, _) => {
  console.error(reason);
  process.exit(1);
});
```

This prevents unhandled promise rejection errors by NodeJS. Please read <https://medium.com/dailyjs/how-to-prevent-your-node-js-process-from-crashing-5d40247b8ab2> for more information.

5.4.2 1.2. CloudFormation

1.2.1. Cross-Account/Region References

When managing multiple AWS accounts, the Accelerator may need permissions to modify resources in the managed accounts. For example, a transit gateway could be created in a shared network account and it need to be shared to the perimeter account to create a VPN connection.

In a single-account environment we would could just:

1. create a single stack and use `!Ref` to refer to the transit gateway;
2. or deploy two stacks
 - one stack that contains the transit gateway and creates a CloudFormation exported output that contains the transit gateway ID;
 - another stack that imports the exported output value from the previous stack and uses it to create a VPN connection.

In a multi-account environment this is not possible and we had to find a way to share outputs across accounts and regions.

See [Passing Outputs Between Phases](#).

1.2.2. Resource Names and Logical IDs

Some resources, like `AWS::S3::Bucket`, can have an explicit name. Setting an explicit name can introduce some possible issues.

The first issue that could occur goes as follows:

- the named resource has a retention policy to retain the resource after deleting;
- then the named resource is created through a CloudFormation stack;
- next, an error happens while creating or updating the stack and the stack rolls back;
- and finally the named resource is deleted from the stack but has a retention policy to retain, so the resource not be deleted;

Suppose then that the stack creation issue is resolved and we retry to create the named resource through the CloudFormation stack:

- the named resource is created through a CloudFormation stack;
- the named resource will fail to create because a resource with the given name already exists.

The best way to prevent this issue from happening is to not explicitly set a name for the resource and let CloudFormation generate the name.

Another issue could occur when changing the logical ID of the named resource. This is documented in the following section.

1.2.3. Changing Logical IDs

When changing the logical ID of a resource CloudFormation assumes the resource is a new resource since it has a logical ID it does not know yet. When updating a stack, CloudFormation will always prioritize resource creation before deletion.

The following issue could occur when the resource has an explicit name. CloudFormation will try to create the resource anew and will fail since a resource with the given name already exists. Example of resources where this could happen are `AWS::S3::Bucket`, `AWS::SecretManager::Secret`.

1.2.4. Changing (Immutable) Properties

Not only changing logical IDs could cause CloudFormation to replace resources. Changing immutable properties also cause replacement of resources. See [Update behaviors of stack resources](#).

Be especially careful when:

- changing immutable properties for a named resource. Example of a resource is `AWS::Budgets::Budget`, `AWS::ElasticLoadBalancingV2::LoadBalancer`.
- updating network interfaces for an `AWS::EC2::Instance`. Not only will this cause the instance to re-create, it will also fail to attach the network interfaces to the new EC2 instance. CloudFormation creates the new EC2 instance first before deleting the old one. It will try to attach the network interfaces to the new instance, but the network interfaces are still attached to the old instance and CloudFormation will fail.

For some named resources, like `AWS::AutoScaling::LaunchConfiguration` and `AWS::Budgets::Budget`, we append a hash to the name of the resource that is based on its properties. This way when an immutable property is changed, the name will also change, and the resource will be replaced successfully. See for example `src/lib/cdk-constructs/src/autoscaling/launch-configuration.ts` and `src/lib/cdk-constructs/src/billing/budget.ts`.

```
export type LaunchConfigurationProps = autoscaling.CfnLaunchConfigurationProps;

/**
 * Wrapper around CfnLaunchConfiguration. The construct adds a hash to the launch configuration name that is based on
 * the launch configuration properties. The hash makes sure the launch configuration gets replaced correctly by
 * CloudFormation.
 */
export class LaunchConfiguration extends autoscaling.CfnLaunchConfiguration {
  constructor(scope: cdk.Construct, id: string, props: LaunchConfigurationProps) {
    super(scope, id, props);

    if (props.launchConfigurationName) {
      const hash = hashSum({ ...props, path: this.node.path });
      this.launchConfigurationName = `${props.launchConfigurationName}-${hash}`;
    }
  }
}
```

5.4.3 1.3. CDK

CDK makes heavy use of CloudFormation so all best practices that apply to CloudFormation also apply to CDK.

1.3.1. Logical IDs

The logical ID of a CDK component is calculated based on its path in the construct tree. Be careful moving around constructs in the construct tree -- e.g. changing the parent of a construct or nesting a construct in another construct -- as this will change the logical ID of the construct. Then you could end up with the issues described in section [Changing Logical IDs](#) and section [Changing \(Immutable\) Properties](#).

See [Logical ID Stability](#) for more information.

1.3.2. Moving Resources between Nested Stacks

In some cases we use nested stacks to overcome [the limit of 200 CloudFormation resources per stack](#).

In the code snippet below you can see how we generate a dynamic amount of nested stack based on the amount of interface endpoints we construct. The `InterfaceEndpoint` construct contains CloudFormation resources so we have to be careful to not exceed the limit of 200 CloudFormation resources per nested stack. That is why we limit the amount of interface endpoints to 30 per nested stack.

```
let endpointCount = 0;
let endpointStackIndex = 0;
let endpointStack;
for (const endpoint of endpointConfig.endpoints) {
  if (endpointStack || endpointCount >= 30) {
    endpointStack = new NestedStack(accountStack, `Endpoint${endpointStackIndex++}`);
    endpointCount = 0;
  }
  new InterfaceEndpoint(endpointStack, pascalCase(endpoint), {
    serviceName: endpoint,
  });
  endpointCount++;
}
```

We have to be careful here though. Suppose the configuration file contains 40 interface endpoints. The first 30 interface endpoints will be created in the first nested stack; the next 10 interface endpoints will be created in the second nested stack. Suppose now that we remove the first nested endpoint from the configuration file. This will cause the 31st interface endpoint to become the 30th interface endpoint in the list and it will cause the interface endpoint to be moved from the second nested stack to the first nested stack. This will cause the stack updates to fail since CloudFormation will first try to create the interface endpoint in the first nested stack before removing it from the second nested stack. We do currently not support changes to the interface endpoint configuration because of this behavior.

1.3.3. L1 vs. L2 Constructs

See [AWS Construct library](#) for an explanation on L1 and L2 constructs.

The L2 constructs for EC2 and VPC do not map well onto the Accelerator-managed resources. For this reason we mostly use L1 CDK constructs -- such as `ec2.CfnVPC`, `ec2.CfnSubnet` -- instead of using L2 CDK constructs -- such as `ec2.Vpc` and `ec2.Subnet`.

1.3.4. CDK Code Dependency on Lambda Function Code

You can read about the distinction between CDK code and runtime code in the introduction of the [Development](#) section.

CDK code can depend on runtime code. For example when we want to create a Lambda function using CDK, we need the runtime code to define the Lambda function. We use `npm scripts`, `npm dependencies` and the `NodeJS modules` API to define this dependency between CDK code and runtime code.

First of all, we create a separate folder that contains the workspace and runtime code for our Lambda function. Throughout the project we've called these workspaces `...-lambda` but it could also be named `...-runtime`. See `src/lib/custom-resources/cdk-acm-import-certificate/runtime/package.json`.

This workspace's `package.json` file needs a `prepare` script that compiles the runtime code. See [npm-scripts](#).

The `package.json` file also needs a `name` and a `main` entry that points to the compiled code.

`runtime/package.json`

```
{
  "name": "lambda-fn-runtime",
  "main": "dist/index.js",
  "scripts": {
    "prepare": "webpack-cli --config webpack.config.ts"
  }
}
```

Now when another workspace depends on our Lambda function runtime code workspace, the `prepare` script will run and it will compile the Lambda function runtime code.

Next, we add the dependency to the new workspace to the workspace that contains the CDK code using `pnpm` or by adding it to `package.json`.

cdk/package.json

```
{
  "devDependencies": {
    "lambda-fn-runtime": "workspace:*0.0.1"
  }
}
```

In the CDK code we can now resolve the path to the compiled code using the `NodeJS modules API`. See [NodeJS modules API](#).

cdk/src/index.ts

```
class LambdaFun extends cdk.Construct {
  constructor(scope: cdk.Construct, id: string) {
    super(scope, id);

    // Find the runtime package folder and resolves the `main` entry of `package.json`.
    // In our case this is `node_modules/lambda-fn-runtime/dist/index.js`.
    const runtimeMain = resolve.require("lambda-fn-runtime");

    // Find the directory containing our `index.js` file.
    // In our case this is `node_modules/lambda-fn-runtime/dist`.
    const runtimeDir = path.dirname(lambdaPath);

    new lambda.Function(this, 'Resource', {
      runtime: lambda.Runtime.NODEJS_14_X,
      code: lambda.Code.fromAsset(runtimeDir),
      handler: 'index.handler', // The `handler` function in `index.js`
    });
  }
}
```

You now have a CDK Lambda function that uses the compiled Lambda function runtime code.

Note: The runtime code needs to recompile every time it changes since the `prepare` script only runs when the runtime workspace is installed.

1.3.5. Custom Resource

We create custom resources for functionality that is not supported natively by CloudFormation. We have two types of custom resources in this project:

1. Custom resource that calls an SDK method;
2. Custom resource that needs additional functionality and is backed by a custom Lambda function.

CDK has a helper construct for the first type of custom resources. See [CDK `AwsCustomResource` documentation](#). This helper construct is for example used in the custom resource [ds-log-subscription](#).

The second type of custom resources requires a custom Lambda function runtime as described in the previous section. For example [acm-import-certificate](#) is backed by a custom Lambda function.

Only a single Lambda function is created per custom resource, account and region. This is achieved by creating only a single Lambda function in the construct tree.

src/lib/custom-resources/custom-resource/cdk/index.ts

```
class CustomResource extends cdk.Construct {
  constructor(scope: cdk.Construct, id: string, props: CustomResourceProps) {
    super(scope, id);

    new cdk.CustomResource(this, 'Resource', {
      resourceType: 'Custom::CustomResource',
      serviceToken: this.lambdaFunction.functionArn,
    });
  }

  private get lambdaFunction() {
    const constructName = 'CustomResourceLambda';

    const stack = cdk.Stack.of(this);
    const existing = stack.node.tryFindChild(constructName);
    if (existing) {
      return existing as lambda.Function;
    }

    // The package '@aws-accelerator/custom-resources/cdk-custom-resource-runtime' contains the runtime code for the custom resource
    const lambdaPath = require.resolve('@aws-accelerator/custom-resources/cdk-custom-resource-runtime');
    const lambdaDir = path.dirname(lambdaPath);

    return new lambda.Function(stack, constructName, {
      code: lambda.Code.fromAsset(lambdaDir),
    });
  }
}
```

1.3.6. Escape Hatches

Sometimes CDK does not support a property on a resource that CloudFormation does support. You can then override the property using the `addOverride` or `addPropertyOverride` methods on CDK CloudFormation resources. See [CDK escape hatches](#).

1.3.6.1. AUTOSCALING GROUP METADATA

An example where we override metadata is when we create a launch configuration.S

```
const launchConfig = new autoscaling.CfnLaunchConfiguration(this, 'LaunchConfig', { ... });

launchConfig.addOverride('Metadata.AWS::CloudFormation::Authentication', {
  S3AccessCreds: {
    type: 'S3',
    roleName,
    buckets: [bucketName],
  },
});

launchConfig.addOverride('Metadata.AWS::CloudFormation::Init', {
  configSets: {
    config: ['setup'],
  },
  setup: {
    files: {
      // Add files here
    },
    services: {
      // Add services here
    },
    commands: {
      // Add commands here
    },
  },
});
```

1.3.6.2. SECRET SECRETVALUE

Another example is when we want to use `secretsmanager.Secret` and set the secret value.

```
function setSecretValue(secret: secrets.Secret, value: string) {
  const cfnSecret = secret.node.defaultChild as secrets.CfnSecret; // Get the L1 resource that backs this L2 resource
  cfnSecret.addPropertyOverride('SecretString', value); // Override the property 'SecretString' on the L1 resource
  cfnSecret.addPropertyDeletionOverride('GenerateSecretString'); // Delete the property 'GenerateSecretString' from the L1 resource
}
```

5.5 1. How to Contribute

5.5.1 1.1. General

Please first refer to and comply with the Contributing and Governance document found [here](#)

5.5.2 1.2. Adding New Functionality?

Before making a change or adding new functionality you have to verify what kind of functionality is being added.

- Is it an Accelerator-management change?
- Is the change related to the `Installer` stack?
- Is the change CDK related?
- Make the change in `src/installer/cdk`.
- Is the change runtime related?
- Make the change in `src/installer/cdk/assets`.
- Is the change related to the `Initial Setup` stack?
- Is the change CDK related?
- Make the change in `src/core/cdk`
- Is the change runtime related?
- Make the change in `src/core/runtime`
- Is it an Accelerator-managed change?
- Is the change related to the `Phase` stacks?
- Is the change CDK related?
- Make the change in `src/deployments/cdk`
- Is the change runtime related?
- Make the change in `src/deployments/runtime`

5.5.3 1.3. Create a CDK Lambda Function with Lambda Runtime Code

See [CDK Code Dependency on Lambda Function Code](#) for a short introduction.

5.5.4 1.4. Create a Custom Resource

See [Custom Resource](#) and [Custom Resources](#) for a short introduction.

1. Create a separate folder that contains the CDK and Lambda function runtime code, e.g. `src/lib/custom-resources/my-custom-resource`;
2. Create a folder `my-custom-resource` that contains the CDK code;
 - a. Create a `package.json` file with a dependency to the `my-custom-resource/runtime` package;
 - b. Create a `cdk` folder that contains the source of the CDK code;
3. Create a folder `my-custom-resource/runtime` that contains the runtime code;
 - a. Create a `runtime/package.json` file with a `"name"`, `"prepare"` script and a `"main"`;
 - b. Create a `runtime/webpack.config.ts` file that compiles TypeScript code to a single JavaScript file;
 - c. Create a `runtime/src` folder that contains the source of the Lambda function runtime code;

You can look at the `src/lib/custom-resources/cdk-acm-import-certificate` custom resource as an example.

It is best practice to add tags to any resources that the custom resource creates using the `cfn-tags` library.

5.5.5 1.5. Run All Unit Tests

Run in the root of the project.

```
pnpm recursive run test --no-bail --stream --silent
```

5.5.6 1.6. Accept Unit Test Snapshot Changes

Run in `src/deployments/cdk`.

```
pnpm run test -- -u
```

5.5.7 1.7. Validate Code with Prettier

Run in the root of the project.

```
pnpx prettier --check **/*.ts
```

5.5.8 1.8. Format Code with Prettier

Run in the root of the project.

```
pnpx prettier --write **/*.ts
```

5.5.9 1.9. Validate Code with `tslint`

Run in the root of the project.

```
pnpm recursive run lint --stream --no-bail
```

5.6 1. AWS Internal - Accelerator Release Process

5.6.1 1.1. Creating a new Accelerator Code Release

1. Ensure `main` branch is in a suitable state
 2. Disable branch protection for both the `main` branch and for the `release/` branches
 3. Create a version branch with [SemVer](#) semantics and a `release/` prefix: e.g. `release/v1.0.5` or `release/v1.0.5-b` using github UI or using the commands below
 - On latest `main`, run: `git checkout -b release/vX.Y.Z`
 - **Important:** Certain git operations are ambiguous if tags and branches have the same name. Using the `release/` prefix reserves the actual version name for the tag itself; i.e. every `release/vX.Y.Z` branch will have a corresponding `vX.Y.Z` tag.
 - Push that branch to GitHub (if created locally)
 - `git push origin release/vX.Y.Z`
 4. The release workflow will run, and create a **DRAFT** release if successful with all commits since the last tagged release.
 5. Prune the commits that have been added to the release notes (e.g. remove any low-information commits)
 6. Publish the release - this creates the git tag in the repo and marks the release as latest. It also bumps the `version` key in several project `package.json` files.
 7. Re-enable branch protection for both the `main` branch and for the `release/` branches
 - Note: The `Publish` operation will run [the following GitHub Action](#), which merges the `release/vX.Y.Z` branch to `main`. **Branch Protection in GitHub will cause this to fail**, and why we are momentarily disabling branch protection.
-

6. Sample Sensitive Architecture

6.1 Accelerator Sample Sensitive Architecture

- [Overview](#)
- [Account Structure](#)
- [Authentication & Authorization](#)
- [Logging and Monitoring](#)
- [Networking](#)
- [Architecture Diagrams](#)

6.2 1. AWS Secure Environment Accelerator Reference Architecture

6.2.1 1.1. Overview

The *AWS Secure Environment Accelerator (ASEA) Reference Architecture* is a comprehensive, multi-account AWS cloud architecture, which was designed working backwards from AWS customers with high compliance requirements like federal, provincial and municipal governments. The *ASEA Reference Architecture* has been designed to address central identity and access management, governance, data security, comprehensive logging, and network design/segmentation per security frameworks like NIST 800-53, ITSG-33, FEDRAMP Moderate, IRAP and other [Sensitive](#) or Medium level security profiles.

This document is solely focused on the deployed or resulting reference architecture and does NOT talk about the tooling, mechanisms, or automation engine used to deploy the architecture. The [AWS Secure Environment Accelerator \(ASEA\)](#) is one tool capable of deploying this architecture (along with many other architectures), but customers are free to choose whichever mechanism they deem appropriate to deploy it. Readers should refer to the [ASEA documentation](#) for references to the ASEA architecture, design, operation, and troubleshooting. If using the ASEA automation engine, this document reflects the resulting architecture deployed using one of the Medium ASEA sample configuration files. This architecture document should stand on its own in depicting the `deployed` architecture.

The *ASEA Architecture* is a standalone architecture, irrespective of how it was delivered into a customer AWS environment. It is nonetheless anticipated that most customers will choose to realize their *ASEA Architecture* via the delivery mechanism of the ASEA automation engine. Except where absolutely necessary, this document will refrain from referencing the ASEA automation engine further.

6.2.2 1.2. Introduction

The *AWS Secure Environment Accelerator (ASEA) Reference Architecture* has been built with the following design principles:

1. Deliver security outcomes aligned with a medium level security control profile;
2. Maximize agility, scalability, and availability, while minimizing cost;
3. Enable the full capabilities of the AWS cloud;
4. Be adaptable to evolving technological capabilities in the underlying platform being used in the *AWS Secure Environment Accelerator Architecture*;
5. Allow for seamless auto-scaling and provide unbounded bandwidth as bandwidth requirements increase (or decrease) based on actual customer load (a key aspect of the value proposition of cloud computing);
6. Architect for high availability: the design stretches across two physical AWS Availability Zones (AZ), such that the loss of any one AZ does not impact application availability. The design can be easily extended to a third availability zone;
7. Operate as least privilege: all principals in the accounts are intended to operate with the lowest-feasible permission set.

1.2.1. Purpose of Document

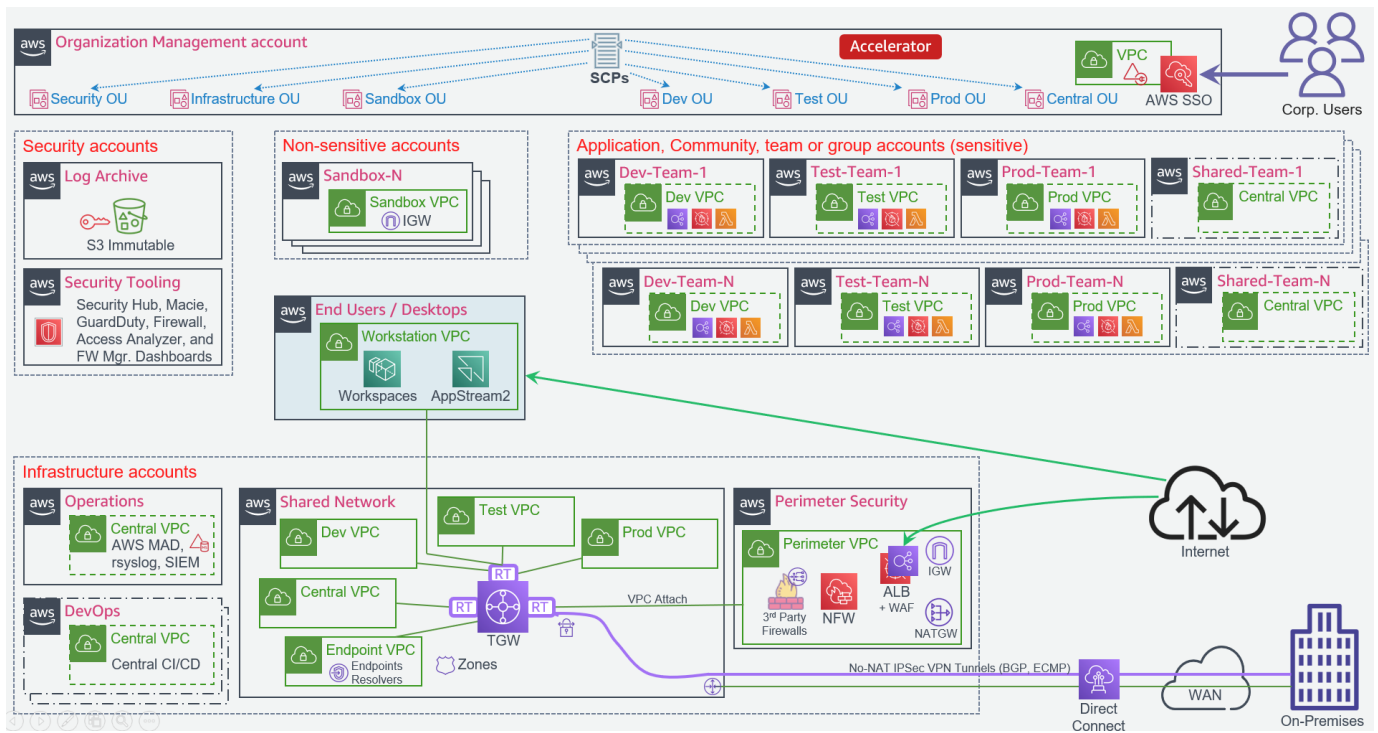
This document is intended to outline the technical measures that are delivered by the *AWS Secure Environment Accelerator Reference Architecture*. This includes a summary of the major architectural decisions.

1.2.2. Architecture Summary

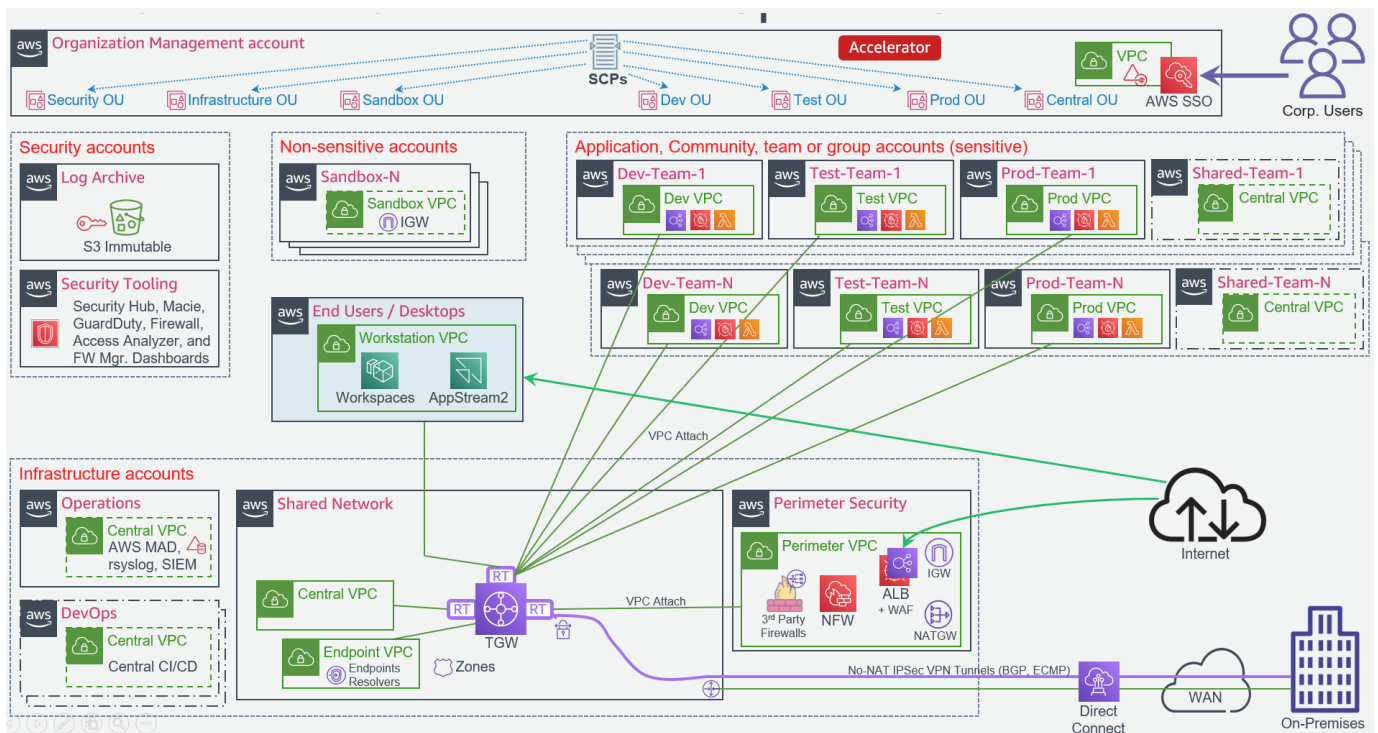
The central features of the *AWS Secure Environment Accelerator Architecture* are as follows:

- **AWS Organization with multiple-accounts:** An [AWS Organization](#) is a grouping construct for a number of separate AWS accounts that are controlled by a single customer entity. This provides consolidated billing, account grouping using organizational units, and facilitates the deployment of pan-organizational guardrails such as API logging with CloudTrail and preventative controls using AWS Service Control Policies (SCPs). Separate AWS accounts provide strong control-plane and data-plane isolation between workloads and/or environments, as if they were owned by different AWS customers. The solution provides a prescriptive AWS account structure, giving different accounts different security personas based on its grouping.
- **Preventative security controls:** Protect the architecture, prevent the disablement of guardrails, and block undesirable user behavior. These are implemented using AWS Service Control Policies ([SCPs](#)). SCPs provide a guardrail mechanism principally used to deny specific or entire categories of API operations at an AWS account, OU, or organization level. These can be used to ensure workloads are deployed only in prescribed regions, or deny access to specific AWS services. The solution provides prescriptive SCPs.
- **Encryption:** AWS Key Management Service (KMS) with customer-managed keys is used to encrypt data stored at rest using FIPS 140-2 validated encryption, whether in S3 buckets, EBS volumes, RDS databases, or other AWS services storing data. Data in-transit is protected using TLS 1.2 or higher encryption.
- **Centralized, isolated networking:** [AWS Virtual Private Clouds](#) (VPCs) are used to create data-plane isolation between workloads, centralized in a shared-network account. Centralization enables strong segregation of duties and cost optimization. Connectivity to on-premises environments, internet egress, shared resources and AWS APIs are mediated at a central point of ingress/egress via the use of [Transit Gateway](#), [Site-to-Site VPN](#), Next-Gen Firewalls, and [AWS Direct Connect](#) (where applicable). The centralized VPC architecture is not for all customers; for customers less concerned with cost optimization, an option exists for local account based VPCs interconnected via the Transit Gateway in the central shared-network account. Under both options, the architecture prescribes moving AWS public API endpoints into the customer's private VPC address space, using centralized endpoints for cost efficiency.
- **Segmentation and Separation:** Not only does the architecture provide strong segmentation and separation between workloads belonging to different stages of the SDLC cycle, or between different IT administrative roles (like between networking, ingress/egress firewalls, and workloads), it offers a strong network zoning architecture, micro-segmenting the environment by wrapping every instance or component in a stateful firewall which is enforced in the hardware of the AWS Nitro System. All flows are tightly enforced, with lateral movement prevented between applications, tiers within an application, and nodes in a tier of an application unless explicitly allowed. Further, routing is prevented between Dev, Test, and Prod with recommendations on a CI/CD architecture to enable developer agility and ease code promotion between environments with appropriate approvals.
- **Centralized DNS management:** [Amazon Route 53](#) is used to provide unified public and private hosted zones across the cloud environment. Inbound and Outbound Route 53 Resolvers extend this unified view of DNS to on-premises networks.
- **Centralized logging:** CloudTrail logs are enabled organization-wide to provide full control plane auditability across the cloud environment. CloudWatch Logs, AWS' cloud native logging service is used to capture a wide variety of logs including OS and application logs, VPC flow logs, and DNS logs which are then centralized and deletion is prevented using SCPs. The architecture prescribes comprehensive log collection and centralization across AWS services and accounts.
- **Centralized security monitoring:** Compliance drift and security threats are surfaced across the customer's AWS organization via the automatic deployment of a multitude of different types of detective security controls. This includes enabling the multitude of AWS security services in every account in the customer's AWS organization including Amazon GuardDuty, AWS Security Hub, AWS Config, AWS Firewall Manager, Amazon Macie, Access Analyzer, CloudWatch Alarms with control and visibility delegated across the multi-account environment to a single central security tooling account for easy organization-wide visibility to all security findings and compliance drift. In addition, the security account has been provided View-Only access across the organization (including access to each account's CloudWatch console) to enable investigation during an incident. View-Only access is different from Read-Only access in that it does not provide any access to any data. Finally, an optional add-on is available to consume the comprehensive set of centralized logs making them searchable, providing correlation and basic dashboards.
- **Single Sign-On:** AWS SSO is used to provide centralized IAM role assumption into AWS accounts across the organization for authorized principals. An organization's existing identities can be sourced from a customer's existing Active Directory identity store or other 3rd party identity provider (IdP). AWS enables MFA enforcement using Authenticator apps, Security keys and built-in authenticators, supporting WebAuthn, FIDO2, and U2F based authentication and devices.
- **Centralized ingress/egress IaaS inspection:** It is common to see centralized ingress/egress requirements for IaaS based workloads. The architecture provides said functionality, enabling customers to decide if native AWS ingress/egress firewall inspection services meet their requirements, or to augment those capabilities with 3rd party firewall appliances. The architecture supports starting with an AWS firewall, switching to a 3rd party firewall, or a combination of ingress/egress firewall technologies.
- **Automation:** Automation is a critical component of the architecture. It ensures guardrails are consistently applied as new AWS accounts are added to the organization as new teams and workloads are brought onboard. It remediates compliance drift and provides guardrails in the root organization account.

The following diagram illustrates the AWS ASEA reference architecture:



The following diagram illustrates an alternative AWS ASEA reference architecture which uses spoke VPCs (instead of centralized VPCs):



1.2.3. Relationship to other AWS reference architectures

The AWS Secure Environment Accelerator Architecture builds upon AWS standardized design patterns and best practices. The architecture aligns with AWS multi-account guidance, the foundation provided by AWS Control Tower, the AWS Secure Environment Accelerator sample architecture, and the AWS Security reference architecture.

1.2.4. Document Conventions

The following conventions are used throughout this document.

1.2.4.1. AWS ACCOUNT NUMBERS

AWS account numbers are decimal-digit pseudorandom identifiers with 12 digits (e.g. 651278770121). This document will use the convention that an AWS Organization Management (root) account has the account ID 123456789012 , and child accounts are represented by 111111111111 , 222222222222 , etc.

For example the following ARN would refer to a VPC subnet in the ca-central-1 region in the Organization Management (root) account:

```
arn:aws:ec2:ca-central-1:123456789012:subnet/subnet-024759b61fc305ea3
```

1.2.4.2. JSON ANNOTATION

Throughout the document, JSON snippets may be annotated with comments (starting with //). The JSON language itself does not define comments as part of the specification; these must be removed prior to use in most situations, including the AWS Console and APIs.

For example:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::123456789012:root" // Trust the Organization Management (root) account.
  },
  "Action": "sts:AssumeRole"
}
```

The above is not valid JSON without first removing the comment on the fourth line.

1.2.4.3. IP ADDRESSES

The design makes use of [RFC1918](#) addresses (e.g. 10.1.0.0/16) and [RFC6598](#) (e.g. 100.96.250.0/23) for various networks; these will be labeled accordingly. Any specific range or IP shown is purely for illustration purposes only.

1.2.5. Customer Naming

This document will make no reference to specific AWS customers. Where naming is required (e.g. in domain names), this document will use a placeholder name as needed; e.g. example.ca .

6.3 1. Account Structure

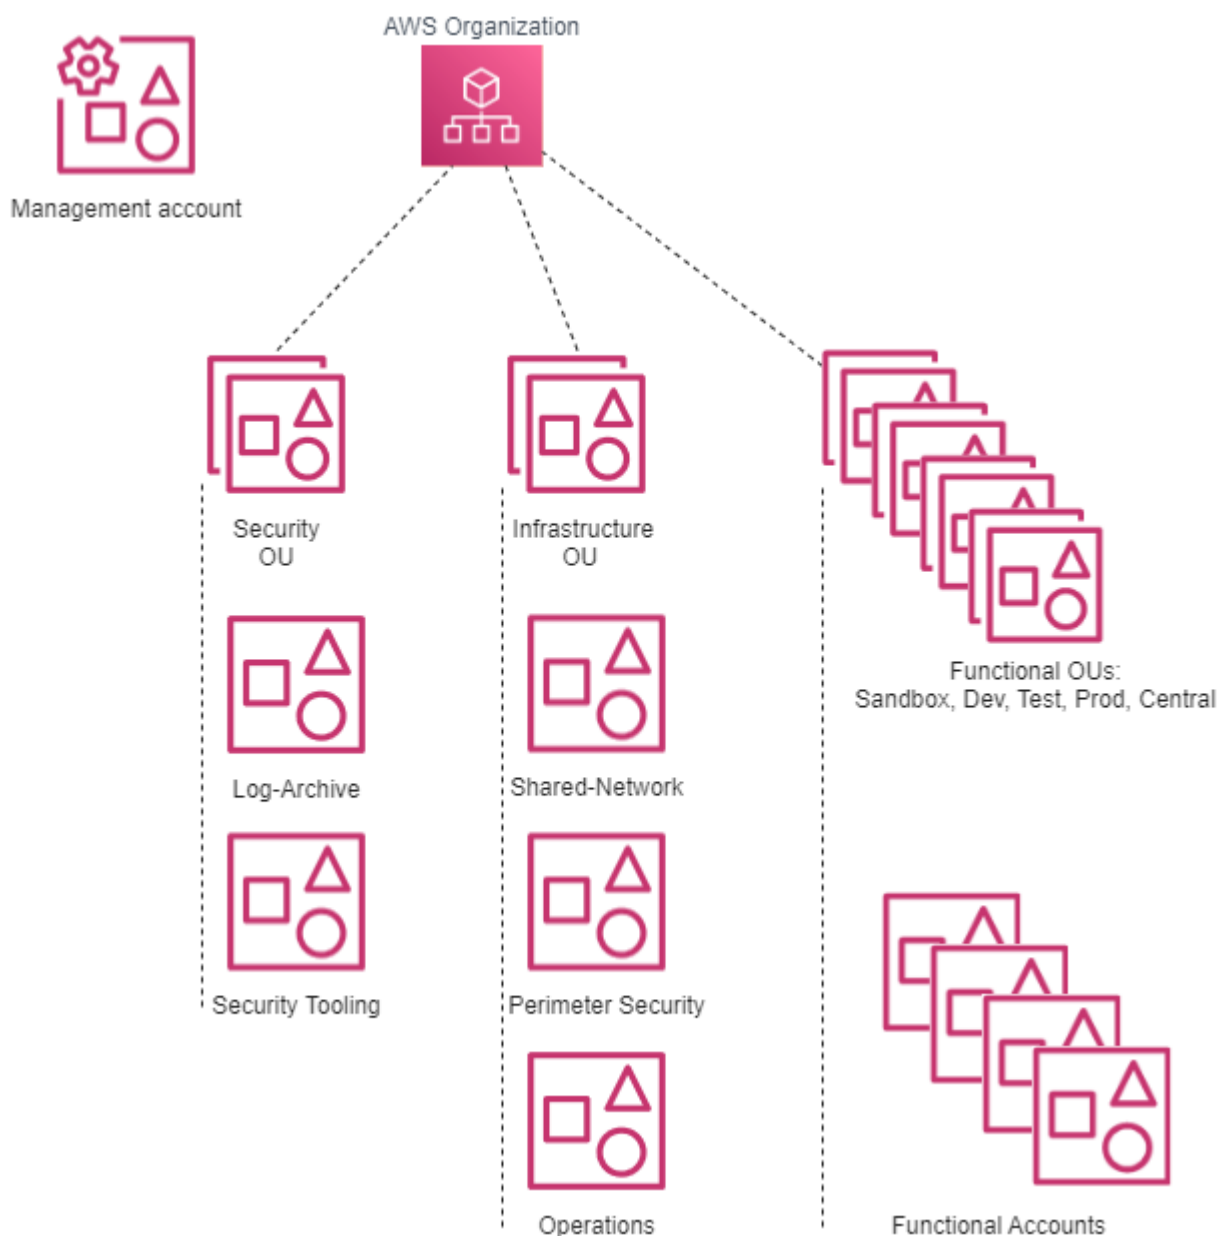
6.3.1 1.1. Overview

AWS accounts are a strong isolation boundary; by default there is no control plane or data plane access from one AWS account to another. AWS accounts provide different AWS customers an isolated private cloud tenancy inside the AWS commercial cloud. It is worth noting that users and roles reside within AWS accounts, and are the constructs used to grant permissions within an AWS account to people, services and applications. AWS Organizations is a service that provides centralized billing across a fleet of accounts, and optionally, some integration-points for cross-account guardrails and cross-account resource sharing. The *AWS Secure Environment Accelerator Architecture* uses these features of AWS Organizations to realize its outcomes.

6.3.2 1.2. Organization structure

The *AWS Secure Environment Accelerator Architecture* includes the following default AWS organization and account structure.

Note that the AWS account structure is strictly a control plane concept - nothing about this structure implies anything about the network architecture or network flows.



1.2.1. Organization Management (root) AWS Account

The AWS Organization resides in the Organization Management (root) AWS account and is traditionally an organization's first AWS account. This account is not used for workloads - it functions primarily as a billing aggregator, and a gateway to the entire cloud footprint for high-trust principals. Access to the Management account must be strictly controlled to a small set of highly trusted individuals from the organization. Additionally, the Organization Management account is where the automation engine or tooling is installed to automate the deployment of the ASEA architecture and its security guardrails. There exists a trust relationship which is used by the automation engine between child AWS accounts in the organization and the Organization Management (root) account; from the Management account, users and roles can assume a role of the following form in child accounts:

```
{
  "Role": {
    "Path": "/",
    "RoleName": "OrganizationAccountAccessRole",
    "Arn": "arn:aws:iam::111111111111:role/OrganizationAccountAccessRole", // Child account
    "AssumeRolePolicyDocument": {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Effect": "Allow",
          "Principal": {
            "AWS": "arn:aws:iam::123456789012:root" // Organization Management (root) account may assume this role
          },
          "Action": "sts:AssumeRole"
        }
      ]
    }
  }
}
```

Note: this is the default role installed by AWS Organizations (`OrganizationAccountAccessRole`) when new AWS accounts are created using AWS organizations. This role changes to `AWSControlTowerExecution` when Control Tower is being leveraged.

1.2.2. AWS SSO

AWS Single Sign-On (SSO) resides in the Organization Management account. Once deployed from the Organization Management account it is recommended that AWS SSO administration is delegated to the Operations account (sometimes referred to as the Shared Services account). AWS SSO is where you create, or connect, your workforce identities in AWS, once, and manage access centrally across your AWS organization. You can create user identities directly in AWS SSO, or you can bring them from your Microsoft Active Directory or a standards-based identity provider, such as Okta Universal Directory or Azure AD. AWS SSO provides a unified administration experience to define, customize, and assign fine-grained access. Your workforce users get a user portal to access all of their assigned AWS accounts. The AWS SSO service deploys IAM roles into accounts in the organization and associates them with the designated workforce identities . More details on SSO are available in the **Authentication and Authorization** section of this document.

6.3.3 1.3. Organizational Units

Underneath the root of the organization, Organizational Units (OUs) provide a mechanism for grouping accounts into logical collections. Aside from the benefit of the grouping itself, these collections serve as the attachment points for SCPs (preventative API-blocking controls), and Resource Access Manager sharing (cross-account resource sharing). Additionally, the ASEA leverages OUs to assign AWS accounts a persona which includes a consistent security personality.

The OU an AWS account is placed in determines the account's purpose, its security posture and the applicable guardrails. An account placed in the Sandbox OU would have the least restrictive, most agile, and most cloud native functionality, whereas an account placed in the Prod OU would have the most restrictive set of guardrails applied.

OUs are NOT designed to reflect an organization's structure, and should instead reflect major shifts in permissions. OUs should not be created for every stage in the SDLC cycle, but those that represent a major permissions shift. For example, organizations that have multiple test stages would often locate the Test and QA Test instances of a workload within the same AWS test account. Customers with a Pre-Prod requirement would often either place their Pre-Prod workloads into their Prod account (alongside the Prod workloads), or in cases requiring more extensive isolation, in a second AWS account located in the Prod OU.

The screenshot shows the AWS Organizations console interface. On the left, there is a navigation menu with 'AWS Organizations' at the top, followed by 'AWS accounts' (expanded), and sub-items: 'Invitations', 'Services', 'Policies', 'Settings', and 'Get started'. Below the menu, the 'Organization ID' is shown as 'o-jjpbdt1ob38'. The main content area is titled 'AWS accounts' and includes a breadcrumb 'AWS Organizations > AWS accounts'. A prominent orange button 'Add an AWS account' is in the top right. Below the title, a descriptive paragraph states: 'The accounts listed below are members of your organization. The organization's management account is responsible for paying the bills for all accounts in the organization. You can use the tools provided by AWS Organizations to centrally manage these accounts. [Learn more](#)'. A search bar is provided with the placeholder text: 'Find AWS accounts by name, email, or account ID. Find an OU by the exact OU ID.' To the right of the search bar are two buttons: 'Hierarchy' (selected) and 'List'. Below this is a section titled 'Organization structure' with a sub-header 'Account created/joined date'. The structure is shown as a tree view starting with 'Root' (ID: r-o3uc). Under 'Root', there are eight child OUs: 'Central' (ou-o3uc-qtat15tu6), 'Dev' (ou-o3uc-8q4ukfoe), 'Infrastructure' (ou-o3uc-thkttgkh), 'Prod' (ou-o3uc-qzfpfw5), 'Sandbox' (ou-o3uc-k1753rmk), 'Security' (ou-o3uc-s9w2a1dy), 'Suspended' (ou-o3uc-tl6ct8iw), and 'Test' (ou-o3uc-h1bn2293). Each OU entry includes a folder icon, a right-pointing arrow, and a checkbox.

Example use cases are as follows:

- An SCP is attached to the Infrastructure OU to prevent the deletion of Transit Gateway resources in the associated accounts.
- The Shared Network account uses RAM sharing to share the development line-of-business VPC with accounts in the development OU. This makes the VPC available to a functional account in that OU used by developers, despite residing logically in the shared network account.

OUs may be nested (to a total depth of five), with SCPs and RAM sharing being controlled at the top level by the automation tooling. A typical AWS *Secure Environment Accelerator Architecture* environment will have the following OUs:

1.3.1. Security OU

The accounts in this OU are considered administrative in nature with access often restricted to *IT security personnel*.

The Security OU is used to hold AWS accounts containing AWS security resources shared or utilized by the rest of the organization. The accounts in the security OU (Log-Archive and Security Tooling) often represent the core or minimum viable set of accounts for organizations wishing to slim the architecture down. No application accounts or application workloads are intended to exist within this OU.

1.3.2. Infrastructure OU

The accounts in this OU are also considered administrative in nature with access often restricted to *IT operations personnel*.

The Infrastructure OU is used to hold AWS accounts containing AWS infrastructure resources shared or utilized by the rest of the organization. The accounts in the Infrastructure OU are also considered core accounts, including accounts like the centralized Shared Network account, the Perimeter Security centralized ingress/egress account, and the Operations account. No application accounts or application workloads are intended to exist within this OU.

These accounts can be optionally removed depending on the outcomes a customer desires. For example, some small customers merge the Operations account into the Security Tooling account, while others may deploy local account-based networking instead of using the Shared Network account and may at the same time drop the central ingress/egress requirements supported by the Perimeter Security account. Eliminating these three accounts should only be considered in the smallest of organizations.

1.3.3. Functional OUs

It is envisioned that most major applications, groups of workloads or teams will work across a set of three or four dedicated AWS accounts provisioned across different functional OUs (Dev, Test, Prod and optionally Central). As new teams, major applications or groups of workloads are onboarded to AWS, they will be provided with this group of new AWS accounts. New teams, groups of applications and major workloads do not share AWS accounts, they each get their own group of unique AWS accounts providing strong segregation and isolation both between stages in the SDLC cycle and between other workloads or teams.

1.3.4. Functional OU: Sandbox

Accounts in this OU are used by development and IT teams for proof of concept / prototyping work. The Sandbox OU offers the most cloud native, agile experience and is used for experimentation. It is not to be used to hold production workloads or sensitive data as it offers the fewest security controls.

These accounts are isolated at a network level and are not connected to the VPCs hosting development, test, production, or shared workloads, nor do they have on-premises network connectivity. These accounts have direct internet access via an internet gateway (IGW). They do not route through the Perimeter Security services account or VPC for internet access.

1.3.5. Functional OU: Dev

The Dev OU is used to hold accounts at the Development or similarly permissioned stage of the SDLC cycle, often containing sensitive data and workloads.

This is the primary account type that an organization's developers would typically work within and hosts development tools and line of business application solutions that are in active development. These accounts are provided internet access for IaaS based workloads via the Perimeter Security account.

1.3.6. Functional OU: Test

The Test OU is used to hold accounts at the test or similarly permissioned (i.e. QA) stage of the SDLC cycle and typically contain sensitive data and workloads. Accounts in this OU host testing tools and line of business application solutions that are being tested prior to promotion to production. These accounts are provided internet access for IaaS based workloads via the Perimeter Security account. As test workloads can easily be destroyed and recreated between test cycles, and temporarily scaled on-demand during performance testing, accounts in the Test OU are generally small in comparison to their Dev and Prod counterparts.

1.3.7. Functional OU: Prod

The Prod OU is used to hold accounts at the Production or similarly permissioned (i.e. Pre-Prod) stage of the SDLC cycle containing sensitive unclassified data or workloads.

Accounts in this OU host production tools and production line of business application solutions. These accounts are provided internet access for IaaS based workloads via the Perimeter Security account. Accounts in this OU are ideally locked down with only specific Operations and Security personnel having access.

1.3.8. Central OU

The Central OU is used to optionally hold AWS accounts which contain group or team resources used and shared across functional OUs or SDLC boundaries like CI/CD or code promotion tools and software development tooling (source control, testing infrastructure, asset repositories). The architecture supports creating a single central CI/CD or DevOps account, a DevOps account per set of team or application accounts, or combinations in-between. The account structure can be customized to meet each organization's own code promotion and shared tooling requirements. These accounts are provided internet access for IaaS based workloads via the Perimeter Security account.

1.3.9. Functional OU: UnClass (Optional)

Non-sensitive workloads should generally be placed with sensitive workloads (Dev/Test/Prod/Central OUs), gaining the extra security benefits of these environments. This OU is used when an organization needs to provide AWS console access to users (internal or external) without appropriate security clearance, to enable deploying and testing AWS services not approved for use with sensitive data, or when services are not available in the local AWS regions which support data locality and sovereignty requirements. These accounts are provided internet access for IaaS based workloads via the Perimeter Security account. Unless this specific use case applies, we generally discourage the use of this OU.

1.3.10. Suspended OU

A suspended OU is created to act as a container for end-of-life accounts which have been closed or suspended, as suspended accounts continue to appear in AWS organizations even after they have been closed and suspended. The `DenyAll` SCP is applied, which prevents all control-plane API operations from taking place by any account principal. Should a suspended account be unintentionally re-activated, no API operations can be performed without intervention of the cloud team.

6.3.4 1.4. Mandatory Accounts

The *AWS Secure Environment Accelerator Architecture* is an opinionated architecture, which partly manifests in the accounts that are deemed mandatory within the organization. The following accounts are assumed to exist, and each has an important function with respect to the goals of the overall Architecture.

AWS Organizations > AWS accounts

AWS accounts Add an AWS account

The accounts listed below are members of your organization. The organization's management account is responsible for paying the bills for all accounts in the organization. You can use the tools provided by AWS Organizations to centrally manage these accounts. [Learn more](#)

Organization Actions ▾

Organizational units (OUs) enable you to group several accounts together and administer them as a single unit instead of one at a time.

Hierarchy List

	AWS accounts	Account created/joined date
<input type="checkbox"/>	SharedNetwork [redacted] [redacted]20-network@amazon.com	Created 2021/08/04
<input type="checkbox"/>	TheFunAccount [redacted] [redacted]20-funacct@amazon.com	Created 2021/08/04
<input type="checkbox"/>	MyDev1 [redacted] [redacted]20-dev1@amazon.com	Created 2021/08/04
<input type="checkbox"/>	log-archive [redacted] [redacted]20-log@amazon.com	Created 2021/08/04
<input type="checkbox"/>	security [redacted] [redacted]20-sec@amazon.com	Created 2021/08/04
<input type="checkbox"/>	Operations [redacted] [redacted]20-operations@amazon.com	Created 2021/08/04
<input type="checkbox"/>	Perimeter [redacted] [redacted]20-perimeter@amazon.com	Created 2021/08/04
<input type="checkbox"/>	ASEA-Main20 management account [redacted] [redacted]20-management@amazon.com	Joined 2021/08/04

1.4.1. Organization Management (root)

This is the Organization Management or root AWS account. Access to this account must be highly restricted, and should not contain customer resources.

As discussed above, the Organization Management (root) account functions as the root of the AWS Organization, the billing aggregator, and attachment point for SCPs. Workloads are not intended to run in this account.

Note: Customers deploying the *AWS Secure Environment Accelerator Architecture* via the [ASEA automation engine][accel_tool] will deploy into this account.

1.4.2. Perimeter Security

This account is used for internet facing IaaS based ingress/egress security services. The perimeter account, and in particular the perimeter VPC therein, functions as the single point of IaaS based ingress/egress from the sensitive cloud environment to the public internet and, in some cases, on-premises networks. This provides a central point of network control through which all workload-generated IaaS traffic, both ingress and egress, must transit. The perimeter VPC can host AWS and/or 3rd party next-generation firewalls that provide security services such as virus scanning, malware protection, intrusion protection, TLS inspection, and web application firewall functionality. More details can be found in the Networking section of this document.

1.4.3. Shared Network

This account is used for centralized or shared networking resources. The shared network account hosts the vast majority of the AWS-side of the networking resources throughout the *AWS Secure Environment Accelerator Architecture*. Workload-scoped VPCs (`Dev` , `Test` , `Prod` , etc.) are deployed in the Shared Network account, and shared via RAM sharing to the respective workload accounts based on their associated OUs. A Transit Gateway provides connectivity from the workloads to the internet or on-premises, without permitting cross-environment (AKA "East/West traffic") traffic (e.g. the Transit Gateway provides VRF like separation between the `Dev` VPC, the `Test` VPC, and the `Prod` VPC). More details can be found in the Networking section of this document.

1.4.4. Operations (or alternatively called Shared Services)

This account is used for centralized IT Operational resources (Active Directory, traditional syslog tooling, ITSM, etc.). The operations account provides a central location for the cloud team to provide cloud operation services to other AWS accounts across the organization and is where an organizations cloud operations team "hangs out" or delivers tooling applicable across all accounts in the organization. The Operations account has View-Only access to CloudWatch logs and metrics across the organization. It is where centralized Systems Manager Session Manager Automation (remediation) documents are located. It is where organizations centralize backup automation (if automated), SSM inventory and patch jobs are automated, and where AWS Managed Active Directory would typically be deployed and accessible to all workloads in the organization. In some AWS documentation this is referred to as the Shared Services account.

1.4.5. Log Archive

The Log archive account is used to centralize and store immutable logs for the organization. The Log Archive account provides a central aggregation and secure storage point for all audit logs created within the AWS Organization. This account contains a centralized storage location for copies of every account's audit, configuration compliance, and operational logs. It also provides a storage location for any other audit/compliance logs, as well as application/OS logs.

As this account is used to provide long-term retention and immutable storage of logs, it is generally recommended nobody have access to this account. Logs should generally be made directly available for local use by teams working in any account on a shorter-term retention basis. Logs will be auto-ingested into a SIEM-like solution as they are centralized into the Log Archive account for analysis and correlation by auditors and security teams. Access to the Log Archive account would be restricted to deep forensic analysis of logs by auditors and security teams during a forensic investigation.

1.4.6. Security Tooling

This account is used to centralize access to AWS security tooling and consoles, as well as provide View-Only access for investigative purposes into all accounts in the organization. The security account is restricted to authorized security and compliance personnel, and related security or audit tools. This is an aggregation point for security services, including AWS Security Hub, GuardDuty, Macie, Config, Firewall Manager, Detective, Inspector, and IAM Access Analyzer.

1.4.7. DevOps account and/or Shared Team/Application accounts

These accounts are used to deliver CI/CD capabilities or shared tooling - two patterns are depicted in the architecture diagram in section 1.2.2 of the Overview - The first has a single organization wide central CI/CD or shared tooling account (named the DevOps account), the other has a CI/CD and shared tooling account per major application team/grouping of teams/applications (named Shared-Team-N). Which pattern is used will depend entirely on the organization's size, maturity model, delegation model of the organization and their team structures. We still generally recommend CI/CD tooling in each developer account (i.e. using Code Commit). When designated branch names are leveraged, the branch/PR will automatically be pulled into the centralized CI/CD tooling account. After approval(s), the code will again be automatically pushed through the SDLC cycle to the Test and/or Prod accounts, as appropriate. Refer to [this](#) blog for more details on automating this pattern.

1.4.8. End User or Desktop account

When an organization configures hybrid on-premises to cloud networking, end-users can directly access the cloud environment using their on-premises organization provided desktops or laptops. When hybrid access is not available, or to enable access from locations and devices whose posture may be not be fully assessed, users can access the organization's cloud environment via virtualized desktops and/or virtual applications provisioned through the End User or Desktop account. The End User or Desktop account is used to provide your organization's end-user community access to the applications running in the environment.

A dedicated Workstation VPC would be created within the End User account to hold the virtual desktops. Virtual desktops would be created within a security group that prevents and blocks all lateral movement or communications between virtual desktops.

Different pools of desktops would exist, with most users only being provisioned desktops which could access the front-end of production applications (or the web tier). Additional pools would be created that provide developers full access to their development environments.

6.3.5 1.5. Functional Accounts

Functional accounts are created on demand, and placed into an appropriate OU in the organization structure. The purpose of functional accounts is to provide a secure and managed environment where project teams can use AWS resources. They provide an isolated control plane so that the actions of one team in one account cannot inadvertently affect the work of other teams in other accounts.

Functional accounts will gain access to the RAM shared resources based on their respective parent OU. Accounts created for `systemA` and `systemB` in the `Dev` OU would have control plane isolation from each other; however these would both have access to the `Dev` VPC (shared from the `Shared Network` account).

Data plane isolation within the same VPC is achieved by default, by using appropriate security groups whenever ingress is warranted. For example, the app tier of `systemA` should only permit ingress from the `systemA-web` security group, not an overly broad range such as `0.0.0.0/0`, or even the entire VPCs address range.

6.3.6 1.6. Account Level Security Settings

The *AWS Secure Environment Accelerator Architecture* enables certain account-wide features on account creation. Namely, these include:

1. [S3 Public Access Block](#)
2. [Default encryption of EBS volumes](#) using a customer managed local account KMS key

6.3.7 1.7. Private Marketplace

The *AWS Secure Environment Accelerator Architecture* enables the AWS Private Marketplace for the organization. Private Marketplace helps administrators govern which products they want their users to run on AWS by making it possible to see only products that have been allow-listed by the organization based on compliance with an organization's security and procurement policies. When Private Marketplace is enabled, it will replace the standard AWS Marketplace for all users, with the new custom branded and curated Private Marketplace.

The screenshot shows the AWS Marketplace search results page for Private Marketplace (PMP). The browser address bar shows the URL: https://aws.amazon.com/marketplace/search/results?page=1&filters=procurement&procurement=approved&ref=pmp_home. The page header includes "AWS-PMP Private Marketplace" and a search bar. The user is logged in as "Hello, assumed-role/Admin/...".

The left sidebar contains filters for Categories, Procurement, Vendors, Operating System, Pricing Plan, Delivery Method, Average Rating, Architecture, and Region. The main content area displays "All Categories (2 results) showing 1 - 2".

The first result is "Fortinet FortiGate (BYOL) Next-Generation Firewall", which is "Approved for procurement". It is Version 6.4.1, sold by Fortinet, Inc., with a 5-star rating from 2 AWS reviews and 68 external reviews. The description states: "FortiGate Next-Generation Firewall technology delivers complete content and network protection by combining stateful inspection with a comprehensive suite of powerful security features. Application control, firewall, antivirus, IPS, Web filtering and VPN along with advanced features such as an...". The delivery method is "Linux/Unix, Other 6.4.1 - 64-bit Amazon Machine Image (AMI)".

The second result is "Fortinet FortiManager (BYOL) Centralized Security Management", also "Approved for procurement". It is Version 6.4.1, sold by Fortinet, Inc. The description states: "Fortinet FortiManager Security Management appliances allow you to centrally manage any number of Fortinet Network Security devices, from several to thousands, including FortiGate, FortiWiFi, and FortiCarrier. Network administrators can better control their network by logically grouping devices into...". The delivery method is "Linux/Unix, Other 6.4.1 - 64-bit Amazon Machine Image (AMI)".

At the bottom of the main content area, it says "showing 1 - 2".

6.4 1. Authorization and Authentication

6.4.1 1.1. Overview

The *AWS Secure Environment Accelerator Architecture* makes extensive use of AWS authorization and authentication primitives from the Identity and Access Management (IAM) service as a means to enforce the guardrail objectives of the *AWS Secure Environment Accelerator Architecture*, and govern access to the set of accounts that makes up the organization.

6.4.2 1.2. Relationship to the Organization Management (root) AWS Account

AWS accounts, as a default position, are entirely self-contained with respect to IAM principals - their Users, Roles, Groups are independent and scoped only to themselves. Accounts created by AWS Organizations deploy a default role with a trust policy back to the Organization Management account. While it can be customized, by default this role is named the `OrganizationAccountAccessRole` (or `AWSControlTowerExecution` when Control Tower is deployed).

As discussed, the AWS Organization resides in the Organization Management (root) account. This account is not used for workloads and is primarily a gateway to the entire cloud footprint for a high-trust principal. It is therefore crucial that all Organization Management account credentials be handled with extreme diligence, and with a U2F hardware key enabled as a second-factor (and stored in a secure location such as a safe) for all users created within this account, including the root user, regardless of privilege level assigned directly within the Management account.

6.4.3 1.3. Break Glass Accounts

The Organization Management account is used to provide break glass access to AWS accounts within the organization. Break glass (which draws its name from breaking the glass to pull a fire alarm) refers to a quick means for a person who does not have access privileges to certain AWS accounts to gain access in exceptional circumstances, using an approved process. Access to AWS accounts within the organization is provided through AWS SSO. The use and creation of IAM users is highly discouraged, with one exception, break glass users. It is generally recommended that organizations create between 2 to 4 IAM break glass users within the Organization Management account. These users would have hardware based MFA enabled and would be leveraged in exceptional circumstances to gain access to the Organization Management account or sub-accounts within the organization by assuming a role. Use cases for break glass access include failure of the organizations IdP, an incident involving the organizations IdP, a failure of AWS SSO, or a disaster involving the loss of an organization's entire cloud or IdP teams.

To re-iterate, access to the Organization Management account grants 'super admin' status, given the organizational-wide trust relationship to the management account. Therefore access to the 2 to 4 break glass IAM users must be tightly controlled, yet accessible via a predefined and strict process. This process often involves one trusted individual having access to a safe containing the password and a different trusted individual having access to a safe with the hardware MFA key – requiring 2 people to access the break glass credentials.

It is worth noting that AWS SCPs are not applicable to the Organization Management account. It is also worth noting that from within the Organization Management account, roles can be assumed in any account within the organization which include broad exclusions from the SCPs (discussed below). These roles are needed to allow the automation tooling to apply and update the guardrails as required, to troubleshoot and resolve issues with the automation tooling, and to bypass the guardrails under approved exception scenarios.

Two primary roles are available for access across the organization from the Management account: the `{AcceleratorPrefix}PipelineRole` which is excluded from the majority of the SCPs to enable the automation tooling to deploy, manage and update the guardrails and provide access to troubleshoot and resolve issues with the automation tooling; and the standard `OrganizationAccountAccessRole` which has been only been excluded from SCPs which strictly deliver preventative security controls. The `OrganizationAccountAccessRole` is within the bounds of the SCPs which protect automation tooling deployed guardrails and functionality. Access to either of these roles is available to any IAM users or roles in the Organization Management account.

6.4.4 1.4. Multi-Factor Authentication

The following are commonly used MFA mechanisms, supported by AWS:

- RSA tokens are a strong form of hardware based MFA authentication but can only be assigned on a 1:1 basis. A unique token is required for every user in every account. You cannot utilize the same token for multiple users or across AWS accounts.
- Yubikeys are U2F compliant devices and also a strong form of hardware based MFA authentication. Yubikeys have the advantage of allowing many:1 assignment, with multiple users and accounts able to use a single Yubikey.
- Virtual MFA like Google Authenticator on a mobile device is generally considered a good hardware based MFA mechanism, but is not considered as strong as tokens or Yubikeys. Virtual MFA also adds considerations around device charge and is not suitable for break glass type scenarios.
- SMS text messages and email based one time tokens are generally considered a weak form of MFA based authentication, but still highly desirable over no MFA.

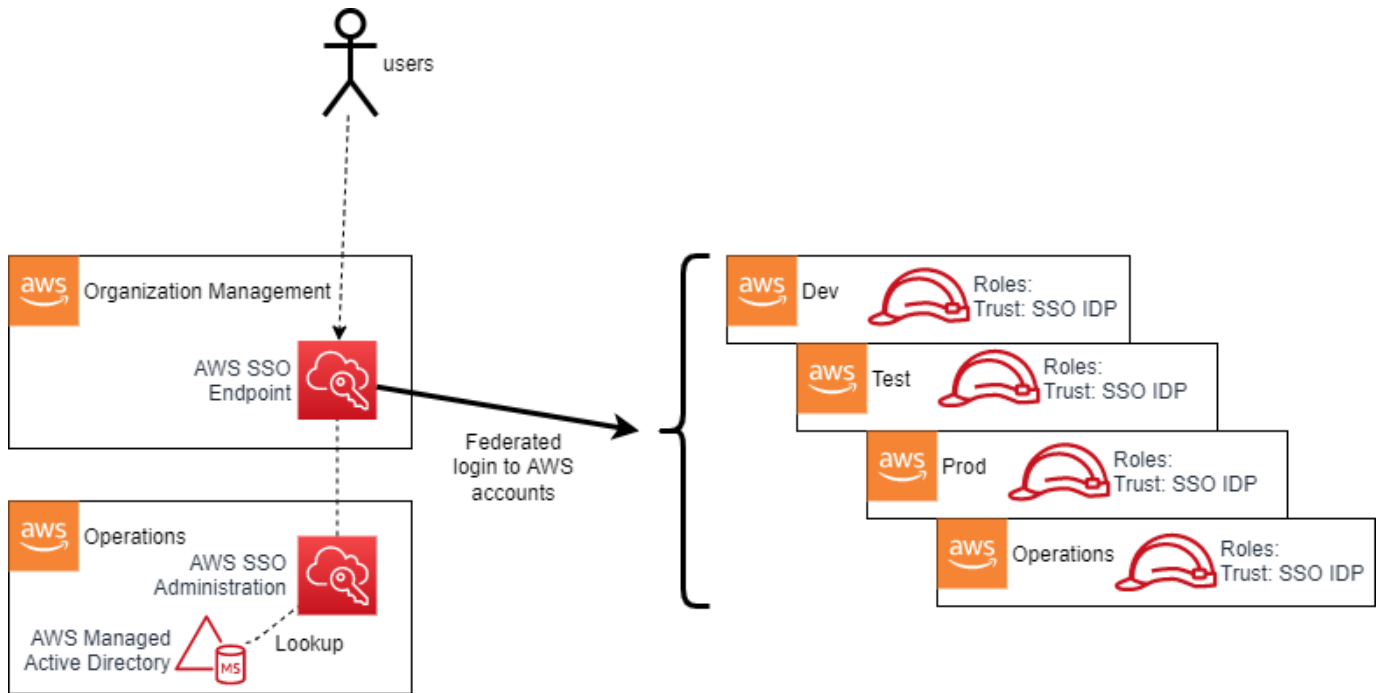
MFA should be used by all users regardless of privilege level with some general guidelines:

- Yubikeys provide the strongest form of MFA protection and are strongly encouraged for all account root users and all IAM users in the Organization Management (root) account;
- the Organization Management (root) account requires a dedicated Yubikey, such that when access is required to a sub-account root user, you do not expose the Organization Management account's Yubikey;
- every ~50 sub-accounts requires a dedicated Yubikey to protect the root user, minimizing the required number of Yubikeys and the scope of impact should a Yubikey be lost or compromised;
- each IAM break glass user requires a dedicated Yubikey, as do any additional IAM users in the Organization Management (root) account. While some CSPs do not recommend MFA on the break glass users, it is strongly encouraged in AWS;
- the MFA devices for all account root users including the management account and the IAM break glass users should be securely stored, with well defined access policies and procedures;
- all other AWS users (AWS SSO, IAM in sub-accounts, etc.) regardless of privilege level should leverage virtual MFA devices (like Google Authenticator on a mobile device).

6.4.5 1.5. Control Plane Access via AWS SSO

The vast majority of end-users of the AWS cloud within the organization will never use or interact with the Organization Management account, or the root users of any child account in the organization. The *AWS Secure Environment Accelerator Architecture* recommends instead that AWS SSO be provisioned in the Organization Management account (a rare case where Organization Management account deployment is mandated). SSO administration is then delegated to the Operations account, to further minimize access to the highly restricted management account. Once delegation is in place, the location of the AWS SSO identity source is also delegated, enabling AWS SSO to directly connect to a Managed Active Directory (AD) or other IdP in the Operations account (this previously required an AWS Directory Connector deployed in the Organization Management account).

Users will login to AWS via the web-based endpoint for the AWS SSO service:



AWS SSO then authenticates the user based on the connected Managed Microsoft AD installation (in the Operations account). Based on group membership, the user will be presented with a set of roles to assume into assigned accounts. For example, a developer may be placed into groups that permit Administrative access in a specific developer account and Read-Only access in a test account; meanwhile an IT Cloud Administrator may have high-privilege access to most, or all, accounts. In effect, AWS SSO adds SAML IdP capabilities to the AWS Managed Microsoft AD, with the AWS Console acting as a service-provider (SP) in SAML parlance. Other SAML-aware SPs may also be used with AWS SSO.

1.5.1.1. SSO User Roles

AWS SSO automatically creates an identity provider (IdP) and associated roles in each account in the organization. The roles used by end users have a trust policy to this IdP. When a user authenticates to AWS SSO (via the underlying Managed AD) and selects a role to assume based on their group membership, the SSO service provides the user with temporary security credentials unique to the role session. In such a scenario, the user has no long-term credentials (e.g. password, or access keys) and instead uses their temporary security credentials.

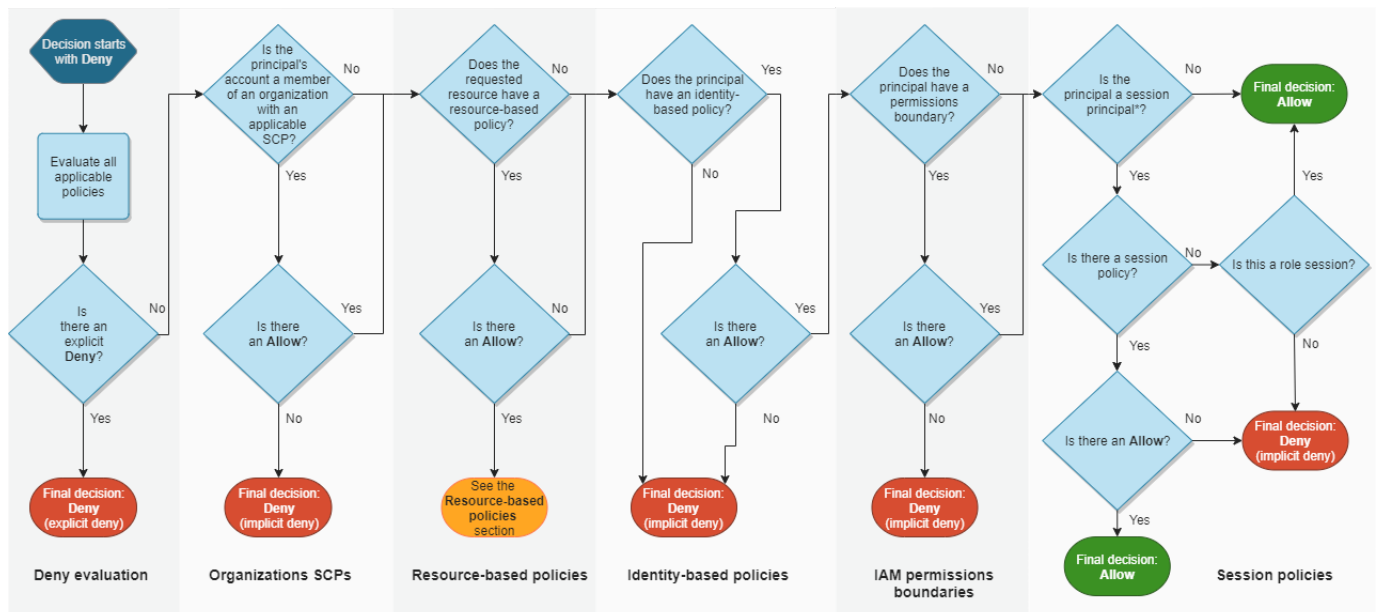
Users, via their AD group membership, are ultimately assigned to SSO user roles via the use of AWS SSO permission sets. A permission set is an assignment of a particular permission policy to an AWS account. For example:

An organization might decide to use **AWS Managed Policies for Job Functions** that are located within the SSO service as the baseline for role-based-access-control (RBAC) separation within an AWS account. This enables job function policies such as:

- **Administrator** - This policy provides full access to all AWS services and resources in the account;
- **Power User** - Provides full access to AWS services and resources, but does not allow management of users, groups and policies;
- **Database Administrator** - Grants full access permissions to AWS services and actions required to set up and configure AWS database services;
- **View-Only User** - This policy grants permissions to view resources and basic metadata across all AWS services. It does not provide access to get or read workload data.

1.5.2. Principal Authorization

Having assumed a role, a user's permission level within an AWS account with respect to any API operation is governed by the IAM policy evaluation logic flow ([detailed here](#)):



Having an `allow` to a particular API operation on the role (i.e. session policy) does not necessarily imply that API operation will succeed. As depicted above, a `deny` at any level in the evaluation logic will block access to the API call; for example a restrictive permission boundary or an explicit `deny` at the resource or SCP level will block the call. SCPs are used extensively as a guardrail mechanism in the *AWS Secure Environment Accelerator Architecture*, and are discussed in a later section.

6.4.6 1.6. Root Authorization

Every AWS account has a set of root credentials. These root credentials are generated on account creation with a random 64-character password. It is important that the root credentials for each account be recovered and MFA enabled via the AWS root credential password reset process using the account's unique email address. To further protect these credentials, the *AWS Secure Environment Accelerator Architecture* specifically denies the use of the root user via SCP. Root credentials authorize all actions for all AWS services and for all resources in the account (except anything denied by SCPs). There are some actions which only root has the capability to perform which are found within the [AWS documentation](#). These are typically rare operations (e.g. creation of X.509 keys), and should not be required in the normal course of business. Root credentials should be handled with extreme diligence, with MFA enabled per the guidance in the previous section.

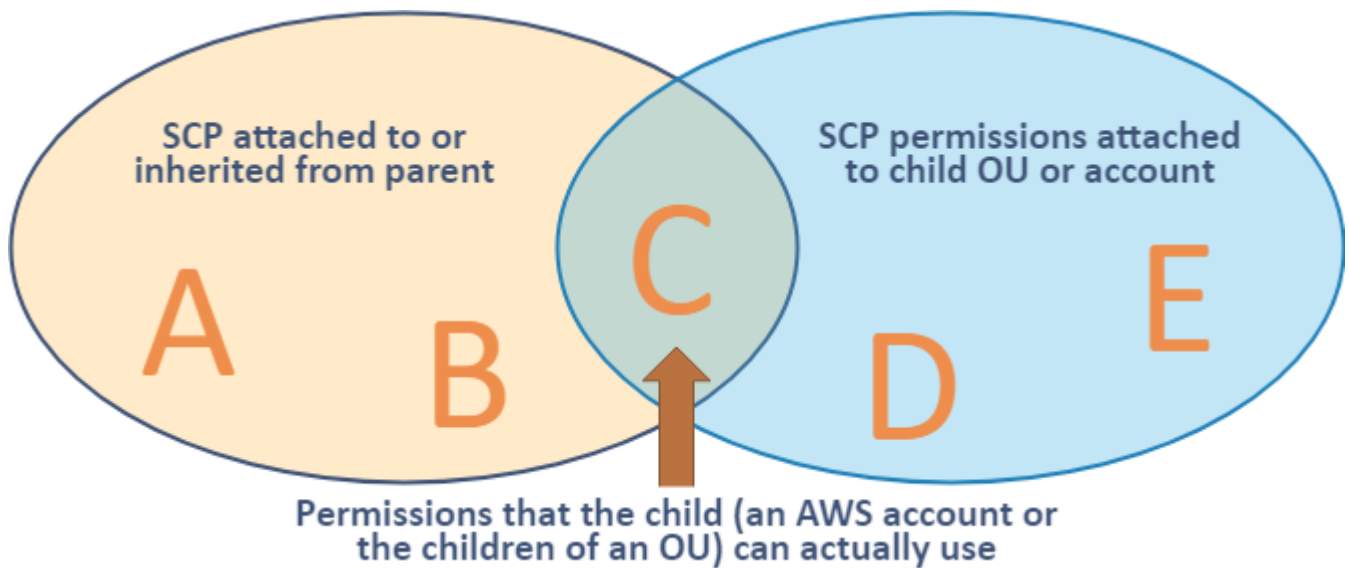
6.4.7 1.7. Service Roles

A service role is an IAM role that a service assumes to perform actions in an account on the user's behalf. When a user sets up an AWS service, the user must define an IAM role for the service to assume. This service role must include all the permissions that are required for the service to access the AWS resources that it needs. Service roles provide access only within a single account and cannot be used to grant access to services in other accounts. Users can create, modify, and delete a service role from within the IAM service. For example, a user can create a role that allows Amazon Redshift to access an Amazon S3 bucket on the user's behalf and then load data from that bucket into an Amazon Redshift cluster. In the case of SSO, during the process in which AWS SSO is enabled, the AWS Organizations service grants AWS SSO the necessary permissions to create subsequent IAM roles.

6.4.8 1.8. Service Control Policies

Service Control Policies are a key preventative control used by the *AWS Secure Environment Accelerator Architecture*. It is crucial to note that SCPs, by themselves, never *grant* permissions. They are most often used to *Deny* certain actions at an OU, or account level within an AWS Organization. Since *deny* always overrides *allow* in the IAM policy evaluation logic, SCPs can have a powerful effect on all principals in any account, and can wholesale deny entire categories of actions irrespective of the permission policy attached to the principal itself - even the root user of the account.

SCPs follow an inheritance pattern from all levels of the hierarchy down to the account of the organization:



In order for any principal to be able to perform an action A, it is necessary (but not sufficient) that there is an `Allow` on action A from all levels of the hierarchy down to the account, and no explicit `Deny` anywhere. This is discussed in further detail in [How SCPs Work](#).

The *AWS Secure Environment Accelerator Architecture* leverages the following SCPs in the organization:

1.8.1. Sensitive policy

This is a comprehensive policy whose main goal is to provide a compliant cloud environment for medium sensitivity workloads, namely prohibiting any non-centralized networking, data-at-rest encryption and mandating data residency in the home region. It should be attached to all top-level OUs with the exception of Unclassified and Sandbox.

Policy Statement ID (SID)	Description
PMP	Prevents the modification or creation of AWS Private Marketplace
ROOT	Prevents the use of the root user
EBS1	Prevents starting EC2 instances without volume level encryption
EBS2	Prevents the creation of an unencrypted EBS volume
EFS1	Prevents the creation of an unencrypted EFS volume
RDS	Prevents the creation of an unencrypted RDS instance
AUR	Prevents the creation of an unencrypted RDS Aurora cluster
OTHS	Blocks miscellaneous operations and services including Leave Organization, Modify Payment Methods, Object Sharing, Disabling SSO, and Lightsail, Sumerian, Gamelift, AppFlow, and IQ.
NET2	Prevents the creation of any networking infrastructure such as VPCs, gateways, peering, VPN, etc. Additionally blocks making objects public (RDS, EMR, EC2, etc.) and the creation of IAM users, groups and access keys
GBL2	Within services that are exempted from GBL1, scope the use of those services to the us-east-1 region (ACM, KMS, SNS)
GBL1	Prevents the use of any service in any non-approved AWS region with the exception of services that are considered global; e.g. CloudFront, IAM, STS, etc.

1.8.1.1. ENCRYPTION AT REST

Note that the Encryption SCP statements above, taken together, mandate encryption at rest for block storage volumes used in EC2 and RDS instances.

1.8.2. Unclassified policy

This policy is only used on the Unclassified OU. This policy is broadly similar to Sensitive; however it relaxes the requirement for non-approved region usage to include additional approved regions. The `GLBL2` services are moved into the `GLBL1` policy.

1.8.3. Sandbox policy

This policy is only used on the Sandbox OU and is only appropriate for accounts in which AWS service experimentation is taking place. This policy is broadly similar to Unclassified; however it does not prohibit network infrastructure creation (e.g. VPCs, IGWs), dropping the `NET2` section.

1.8.4. Guardrail Protection (Parts 0 and 1)

These guardrails apply across the organization and protect the guardrails and infrastructure deployed by the automation tooling. Note that this policy is split into two parts due to a current limitation of SCP sizing, but logically it should be considered a single policy.

Policy Statement ID (SID)	Description
TAG1	Prevents creation, deletion and modification of a protected security group
TAG2	Prevents creation, deletion and modification and use of any protected IAM resource
S3	Prevents deletion and modification of any S3 bucket used by the automation tooling incl. centralized logs
CFN	Prevents creation, deletion or modification of any CloudFormation stacks deployed by the automation tooling
ALM	Prevents deletion and modification of protected cloudwatch alarms which alert on significant control plane events
ROL	Prevents any IAM operation on protected IAM resources
SSM	Prevents creation, deletion or modification of any SSM resource deployed by the automation tooling
LOG	Prevents the deletion and modification of any CloudWatch Log groups and VPC flow logs
LOG2	Additional CloudWatch Log group protections
LAM	Prevents the creation, deletion and modification of any Lambda functions deployed by the automation tooling
NET1	Prevents deletion of any protected networking (VPC) constructs like subnets and route tables
NFW	Prevents destructive operations on protected AWS Network Firewalls
CT	Prevents deletion and modification of protected Cloud Trails
CON	Protects AWS Config configuration from modification or deletion
CWE	Prevents deletion and modification of protected CloudWatch events
RUL	Protects AWS Config rules from modification or deletion
Deny	Protects deletion and modification of protected KMS keys
IAM	Protects creation, deletion, and modification of protected IAM policies
ACM	Prevents deletion of a protected certificates and Load Balancers
SEC	Prevents the deletion and modification to AWS security services like GuardDuty, Security Hub, Macie, Firewall Manager, Access Analyzer, password policies, and resource shares
SNS	Prevents creation, deletion and modification of a protected SNS topics
RDGW	Prevents the modification of the role used for Remote Desktop Gateway
KIN	Prevents creation, deletion and modification of a protected Kinesis streams

Note: Two versions of the Part-0 policy exist (`Core0Us` and `Wkld0Us`), with the `Wkld` OUs version of the policy removing `SNS`, `RDGW`, and `KIN` sections as they are not relevant outside the Security and Infrastructure OUs.

1.8.5. Quarantine New Object

This policy is attached to an account to 'quarantine' it - to prevent any AWS operation from taking place. This is useful in the case of an account with credentials which are believed to have been compromised. This policy is also applied to new accounts upon creation. After the installation of guardrails, it is removed. In the meantime, it prevents all AWS control plane operations except by principals required to deploy guardrails.

Policy Statement ID (SID)	Description
DenyAllAWSservicesExceptBreakglassRoles	Blanket denial on all AWS control plane operations for all non-break-glass roles

6.5 1. Logging and Monitoring

6.5.1 1.1. Overview

The *AWS Secure Environment Accelerator Architecture* requires the following security services be deployed across the organization. These services, taken together, provide a comprehensive picture of the security posture of the organization.

6.5.2 1.2. CloudTrail

The AWS CloudTrail service provides a comprehensive log of control plane and data plane operations (audit history) of all actions taken against most AWS services, including users logging into accounts. A CloudTrail Organizational trail should be deployed into the organization. For each account, this captures management events and optionally S3 data plane events taking place by every principal in every account in the organization. These records are sent to both a CloudWatch log group in the Organization Management account and an S3 bucket in the Log Archive account. The trail itself cannot be modified or deleted by any principal in any child account. This provides an audit trail for detective purposes in the event of the need for forensic analysis into account usage. The logs themselves provide an integrity guarantee: every hour, CloudTrail produces a digest of that hour's log files, with a hash, and signs it with its own private key. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. This process is [detailed here](#). The Log Archive bucket is protected with SCPs and has versioning enabled ensuring deleted or overwritten files are retained.

6.5.3 1.3. VPC Flow Logs

VPC Flow Logs capture information about the IP traffic going to and from network interfaces in an AWS Account VPC such as source and destination IPs, protocol, ports, and success/failure of the flow. The *AWS Secure Environment Accelerator Architecture* enables ALL (i.e. both accepted and rejected traffic) logs for all VPCs in all accounts in both a local CloudWatch log group and an S3 bucket in the log-archive account. It is important to use custom flow log formats to ensure all fields are captured as important fields are not part of the basic format. More details about VPC Flow Logs are [available here](#).

It should be noted that certain categories of network flows are not captured, including traffic to and from the instance metadata service (169.254.169.254), and DNS traffic with an Amazon VPC resolver (available in DNS resolver logs).

6.5.4 1.4. GuardDuty

Amazon GuardDuty is a cloud native threat detection and Intrusion Detection Service (IDS) that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. The service uses machine learning, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats. GuardDuty uses a number of data sources including VPC Flow Logs, DNS logs, CloudTrail logs and several threat feeds.

The *AWS Secure Environment Accelerator Architecture* requires GuardDuty be enabled [at the Organization level](#), and delegating the Security account as the GuardDuty Administrative account. The GuardDuty Administrative account should be auto-enabled to add new accounts as they come online. Note that this should be done in every region as a defense in depth measure, with the understanding that the SCPs will prevent service usage in all other regions.

6.5.5 1.5. Config

[AWS Config](#) provides a detailed view of the resources associated with each account in the AWS Organization, including how they are configured, how they are related to one another, and how the configurations have changed on a recurring basis. Resources can be evaluated on the basis of their compliance with Config Rules - for example, a Config Rule might continually examine EBS volumes and check that they are encrypted.

Config may be [enabled at the Organization](#) level - this provides an overall view of the compliance status of all resources across the organization. The AWS Config multi-account multi-region data aggregation capability has been located in both the Organization Management account and the Security account.

6.5.6 1.6. CloudWatch Logs

CloudWatch Logs is AWS' log aggregator service, used to monitor, store, and access log files from EC2 instances, AWS CloudTrail, Route 53, and other sources. The *AWS Secure Environment Accelerator Architecture* requires that log subscriptions are created for all log groups in all workload accounts, and streamed into S3 in the log-archive account (via Kinesis) for analysis and long-term audit purposes. The CloudWatch log group retention period on all log groups should be set to a medium retention period (such as 2 years) to enable easy online access.

6.5.7 1.7. SecurityHub

The primary dashboard for Operators to assess the security posture of the AWS footprint is the centralized AWS Security Hub service. Security Hub needs to be configured to aggregate findings from Amazon GuardDuty, Amazon Macie, AWS Config, Systems Manager, Firewall Manager, Amazon Detective, Amazon Inspector and IAM Access Analyzers. Events from security integrations are correlated and displayed on the Security Hub dashboard as 'findings' with a severity level (informational, low, medium, high, critical).

The AWS Secure Environment Accelerator Architecture recommends that the current 3 Security Hub frameworks be enabled, specifically:

- [AWS Foundational Security Best Practices v1.0.0](#)
- [PCI DSS v3.2.1](#)
- [CIS AWS Foundations Benchmark v1.2.0](#)

These frameworks will perform checks against the accounts via Config Rules that are evaluated against the AWS Config resources in scope. See the above links for a definition of the associated controls.

6.5.8 1.8. Systems Manager Session Manager

Session Manager is a fully managed AWS Systems Manager capability that lets you manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, on-premises instances, and virtual machines (VMs) through an interactive one-click browser-based shell or through the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable instance management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also makes it easy to comply with corporate policies that require controlled access to instances, strict security practices, and fully auditable logs with instance access details, while still providing end users with simple one-click cross-platform access to your managed instances.¹

The *AWS Secure Environment Accelerator Architecture* stores encrypted session log data in both CloudWatch logs and in the centralized S3 bucket for auditing purposes.

6.5.9 1.9. Systems Manager Inventory

AWS Systems Manager Inventory provides visibility into your AWS computing environment. AWS ASEA architecture uses SSM Inventory to collect metadata from your managed nodes and stores this metadata in the central Log Archive S3 bucket. These logs enable customers to quickly determine which nodes are running the software and configurations required by your software policy, and which nodes need to be updated.

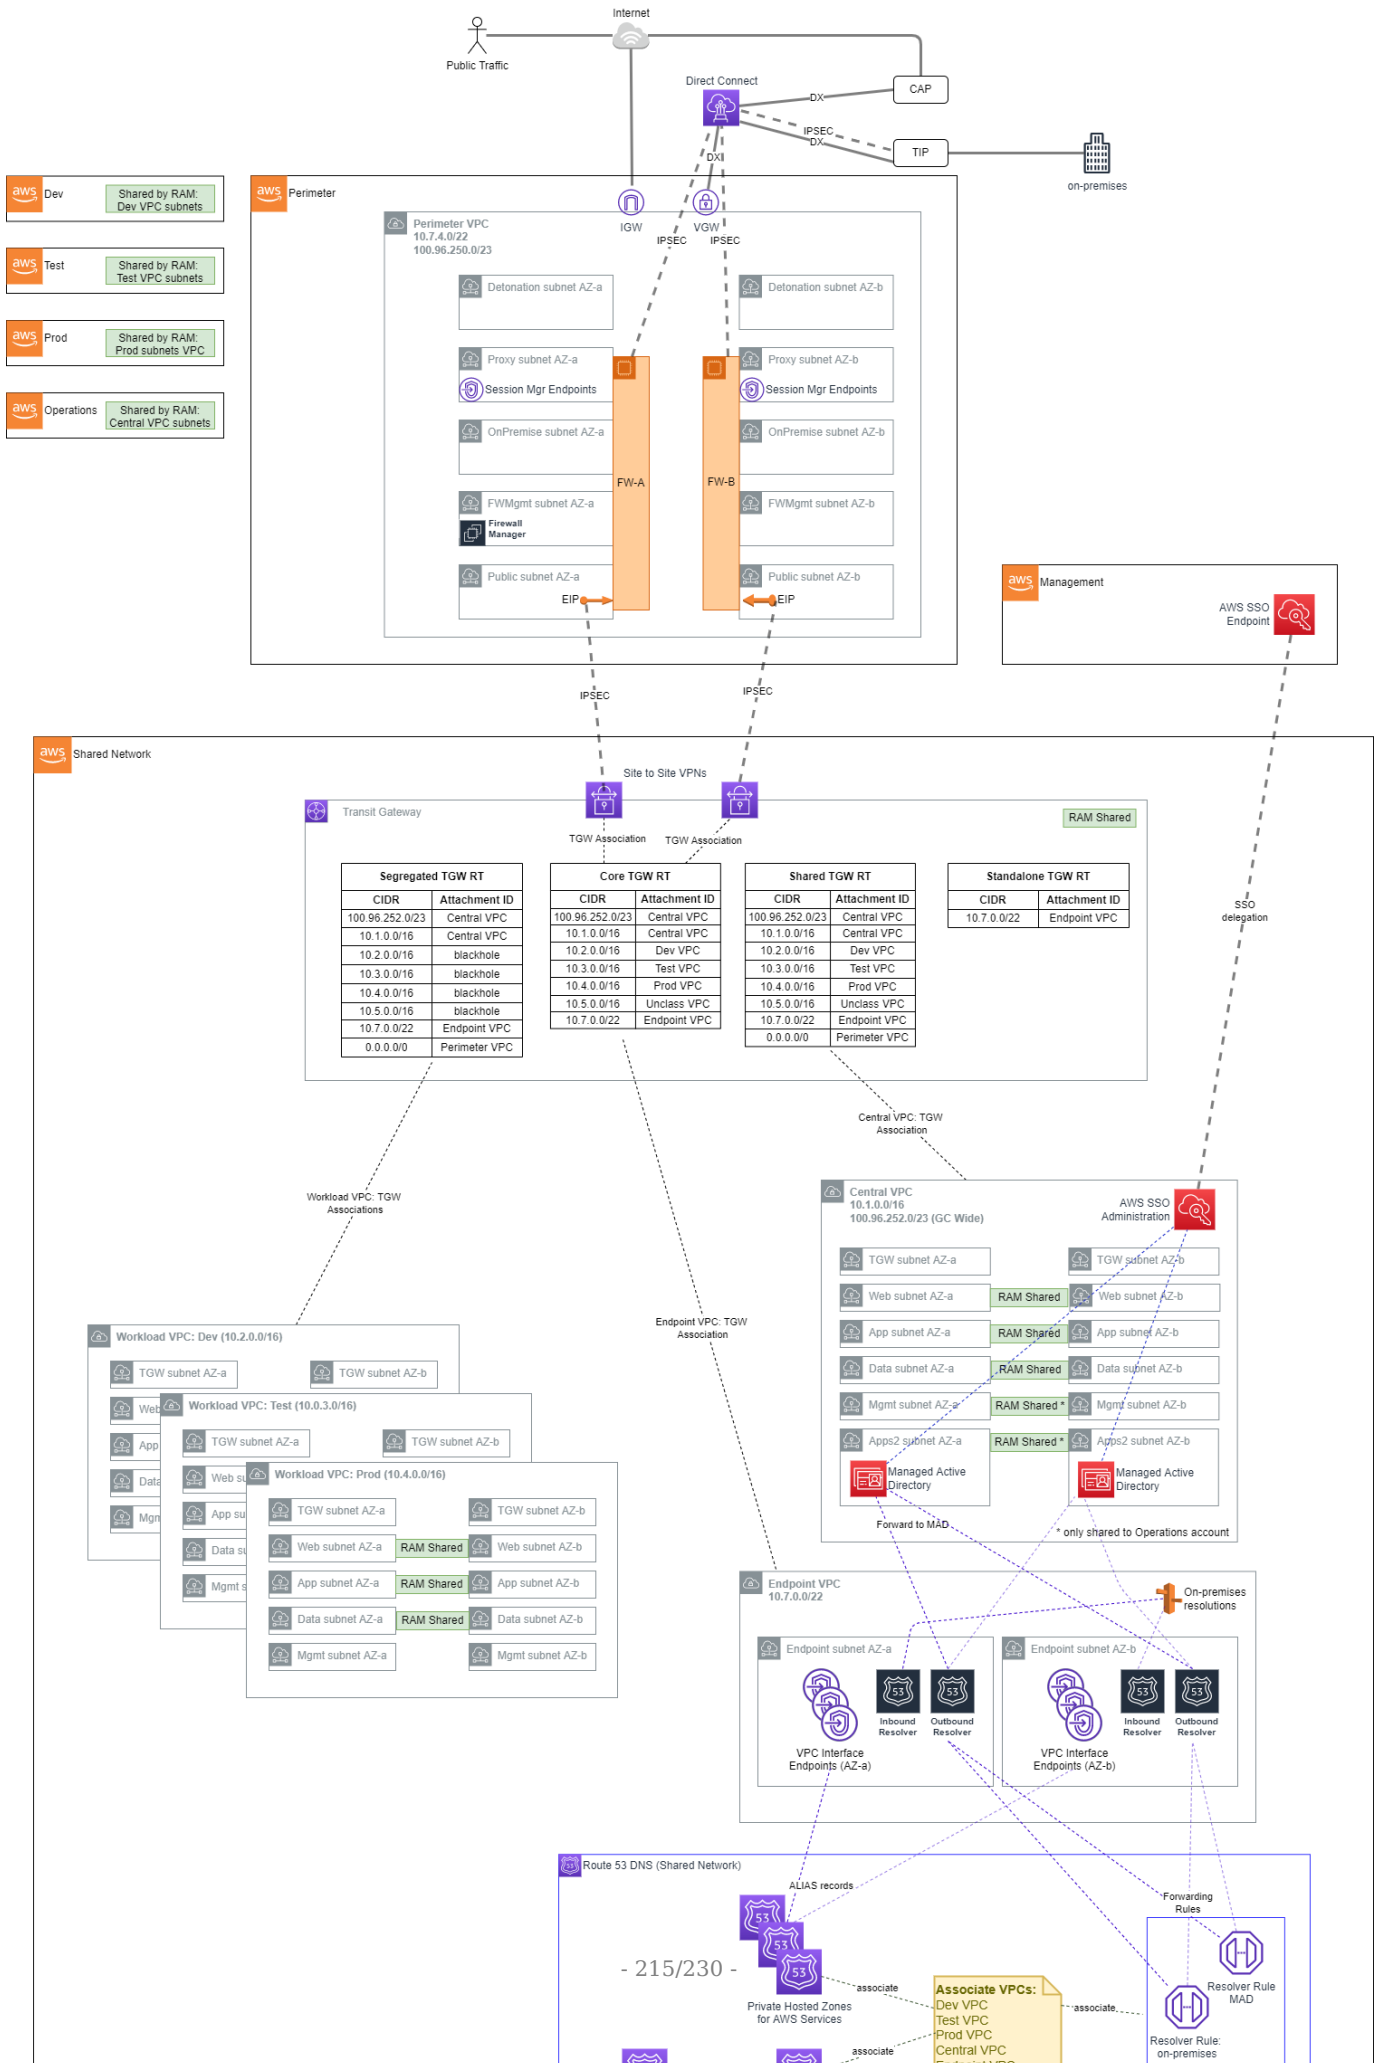
6.5.10 1.10. Other Services

The following additional services are configured with their organization-wide administrative and visibility capabilities centralized to the Security account: Macie, Firewall Manager, Access Analyzer. The following additional logging and reporting services are configured: CloudWatch Alarms, Cost and Usage Reports, rsyslog, MAD, R53 logs, OS/App, ELB, OpenSearch SIEM.

6.6 1. Networking

6.6.1 1.1. Overview

The *AWS Secure Environment Accelerator Architecture* networking is built on a principle of centralized on-premises and internet ingress/egress, while enforcing data plane isolation between workloads in different environments. Connectivity to on-premises environments, internet egress, shared resources and AWS APIs are mediated at a central point of ingress/egress via the use of a [Transit Gateway](#). Consider the following overall network diagram:



All functional accounts use RAM-shared networking infrastructure as depicted above. The workload VPCs (Dev, Test, Prod, etc) are hosted in the Shared Network account and made available to the appropriate accounts based on their OU in the organization.

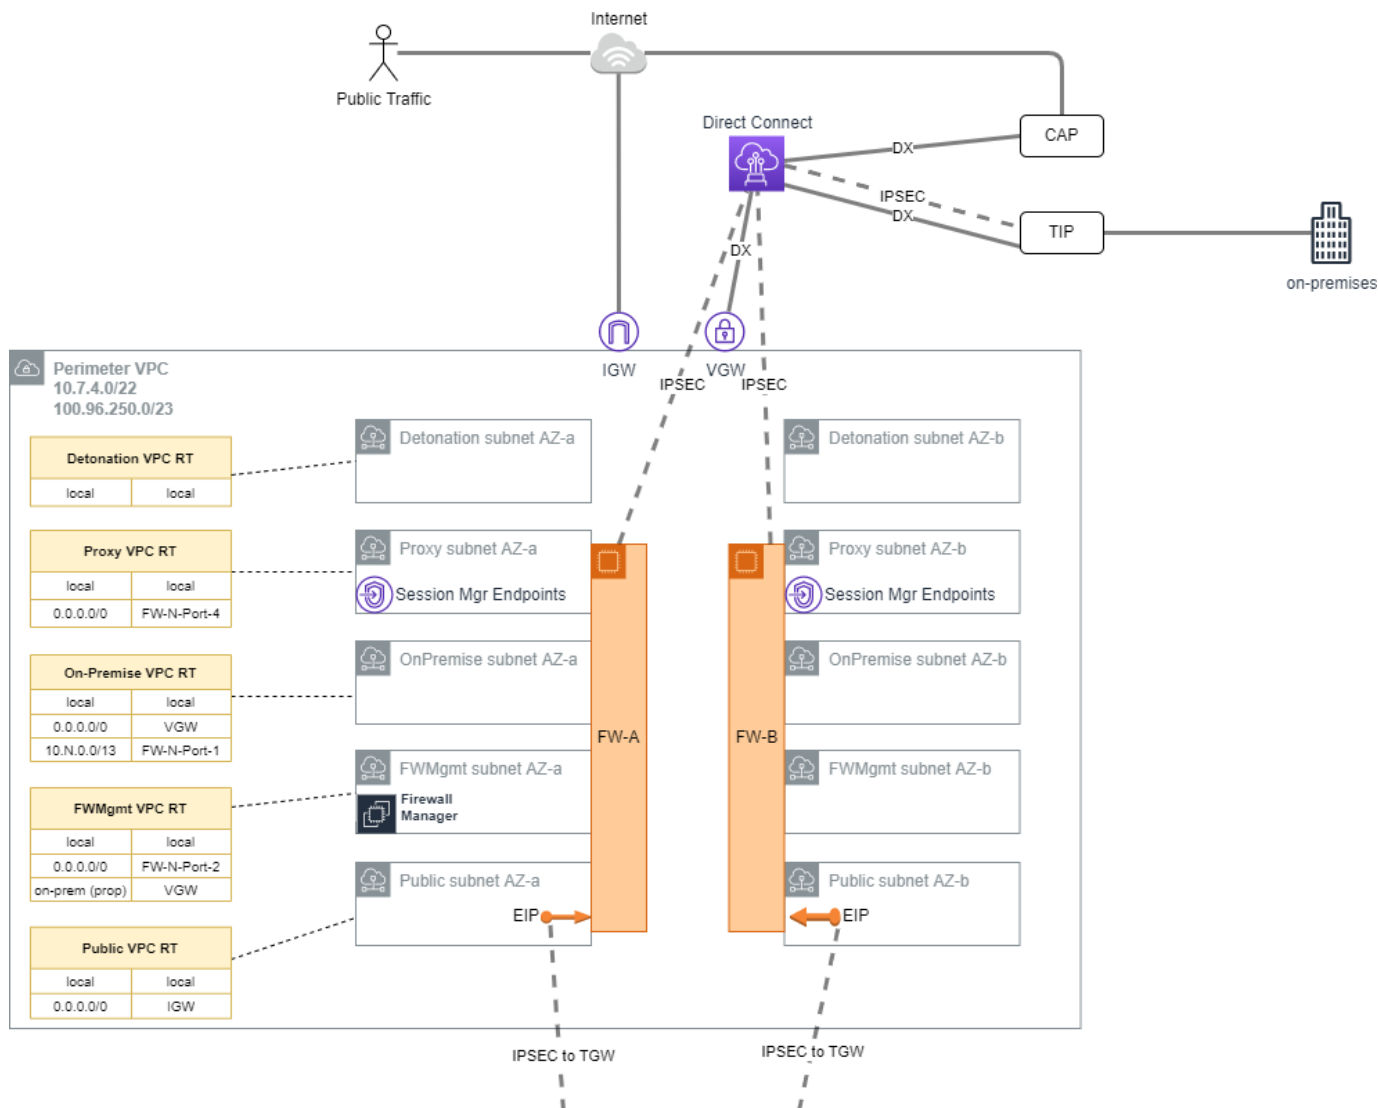
6.6.2 1.2. Perimeter

The perimeter VPC hosts the organization's perimeter security services. The Perimeter VPC is used to control the flow of traffic between AWS Accounts and external networks for IaaS workloads: both public (internet) and in some cases private (access to on-premises datacenters). This VPC hosts AWS Network Firewall and/or 3rd party Next Generation Firewalls (NGFW) that provide perimeter security services including virus scanning / malware protection, Intrusion Protection services, TLS Inspection and Web Application Firewall protection. If applicable, this VPC also hosts reverse proxy servers.

Note that this VPC is in its own isolated account, separate from Shared Network, in order to facilitate networking and security 'separation of duties'. Internal networking teams may administer the cloud networks in Shared Network without being granted permission to administer the security perimeter itself.

1.2.1. IP Ranges

- **Primary Range:** The *AWS Secure Environment Accelerator Architecture* recommends that the perimeter VPC have a primary range in the [RFC1918](#) block (e.g. 10.7.4.0/22), used only for subnets dedicated to 'detonation' purposes. This primary range, in an otherwise-unused [RFC1918](#) range, is not intended to be routable outside of the VPC, and is reserved for future use with malware detonation capabilities of NGFW devices.
- **Secondary Range:** This VPC should also have a secondary range in the [RFC6598](#) block (e.g. 100.96.250.0/23) used for the overlay network (NGFW devices inside VPN tunnel) for all other subnets. This secondary range is assigned by an external entity, and should be carefully selected in order to co-exist with *AWS Secure Environment Accelerator Architecture* deployments that exist at peer organizations; for instance other government departments that maintain a relationship with the same shared entity in a carrier-grade NAT topology. Although this is a 'secondary' range in VPC parlance, this VPC CIDR should be interpreted as the more 'significant' of the two with respect to Transit Gateway routing; the Transit Gateway will only ever interact with this 'secondary' range.



This VPC has four subnets per AZ, each of which hosts a port used by the NGFW devices, which are deployed in an HA pair. The purpose of these subnets is as follows.

- **Detonation:** This is an unused subnet reserved for future use with malware detonation capabilities of the NGFW devices (e.g. 10.7.4.0/24 - not routable except local);
- **Proxy:** This subnet hosts reverse proxy services for web and other protocols. It also contains the [three interface endpoints](#) necessary for AWS Systems Manager Session Manager, which enables SSH-less CLI access to authorized and authenticated principals in the perimeter account (e.g. 100.96.251.64/26);
- **On-Premises:** This subnet hosts the private interfaces of the firewalls, corresponding to connections from the on-premises network (e.g. 100.96.250.192/26);
- **FW-Management:** This subnet is used to host management tools and the management of the Firewalls itself (e.g. 100.96.251.160/27 - a smaller subnet is permissible due to modest IP requirements for management instances);
- **Public:** This subnet is the public-access zone for the perimeter VPC. It hosts the public interface of the firewalls, as well as application load balancers that are used to balance traffic across the firewall pair. There is one Elastic IPv4 address per public subnet that corresponds to the IPsec Customer Gateway (CGW) for the VPN connections into the Transit Gateway in Shared Networking (e.g. 100.96.250.0/26).

Outbound internet connections (for software updates, etc.) can be initiated from within the workload VPCs, and use the transparent proxy feature of the next-gen Firewalls.

Note on VPN Tunnel Redundancy: Each NGFW device manifests as a unique CGW on the AWS side (shared network account) of the IPSec VPNs. Moreover, there are **two Site-to-Site VPNs** in this architecture, each with two active tunnels; this configuration is highly redundant as only one tunnel is required per firewall to provide complete redundancy.

6.6.3 1.3. Shared Network

The Shared Network account, and the AWS networking resources therein, form the core of the cloud networking infrastructure across the account structure. Rather than the individual accounts defining their own networks, these are instead centralized here and shared out to the relevant OUs. Principals in a Dev OU will have access to a Dev VPC, Test OU will have access to a Test VPC and so on - all of which are owned by this account.

You can share AWS Transit Gateways, Subnets, AWS License Manager configurations, and Amazon Route 53 Resolver rules resources with AWS Resource Access Manager (RAM). The RAM service eliminates the need to create duplicate resources in multiple accounts, reducing the operational overhead of managing those resources in every single account.

1.3.1. Transit Gateway

The Transit Gateway is a central hub that performs several core functions within the Shared Network account.

1. Routing of permitted flows; for example a Workload to On-premises via the Perimeter VPC.
 - All routing tables in Shared Network VPCs send `0.0.0.0/0` traffic to the TGW, where its handling will be determined by the TGW Route Table (TGW-RT) that its attachment is associated with. For example:
 - an HTTP request to `registry.hub.docker.com` from the Test VPC will go to the TGW
 - The Segregated TGW RT will direct that traffic to the Perimeter VPC via the IPSec VPNs
 - The request will be proxied to the internet, via GC-CAP if appropriate
 - The return traffic will again transit the IPSec VPNs
 - The `10.3.0.0/16` bound response traffic will arrive at the Core TGW RT, where a propagation in that TGW RT will direct the response back to the Test VPC.
2. Defining separate routing domains that prohibit undesired east-west flows at the network level; for example, by prohibiting Dev to Prod traffic. For example:
 - All routing tables in Shared Network VPCs send `0.0.0.0/0` traffic to the TGW, which defines where the next permissible hop is. For example, `10.2.0.0/16` Dev traffic destined for the `10.0.4.0/16` Prod VPC will be blocked by the blackhole route in the Segregated TGW RT.
3. Enabling centralization of shared resources; namely a shared Microsoft AD installation in the Central VPC, and access to shared VPC Endpoints in the Endpoint VPC.
 - The Central VPC, and the Endpoint VPC are routable from Workload VPCs. This provides an economical way to share organization-wide resources that are nonetheless isolated into their own VPCs. For example:
 - a `git` request in the Dev VPC to `git.private-domain.ca` resolves to a `10.1.0.0/16` address in the Central VPC.
 - The request from the Dev VPC will go to the TGW due to the VPC routing table associated with that subnet
 - The TGW will send the request to the Central VPC via an entry in the Segregated TGW RT
 - The `git` response will go to the TGW due to the VPC routing table associated with that subnet
 - The Shared TGW RT will direct the response back to the Dev VPC

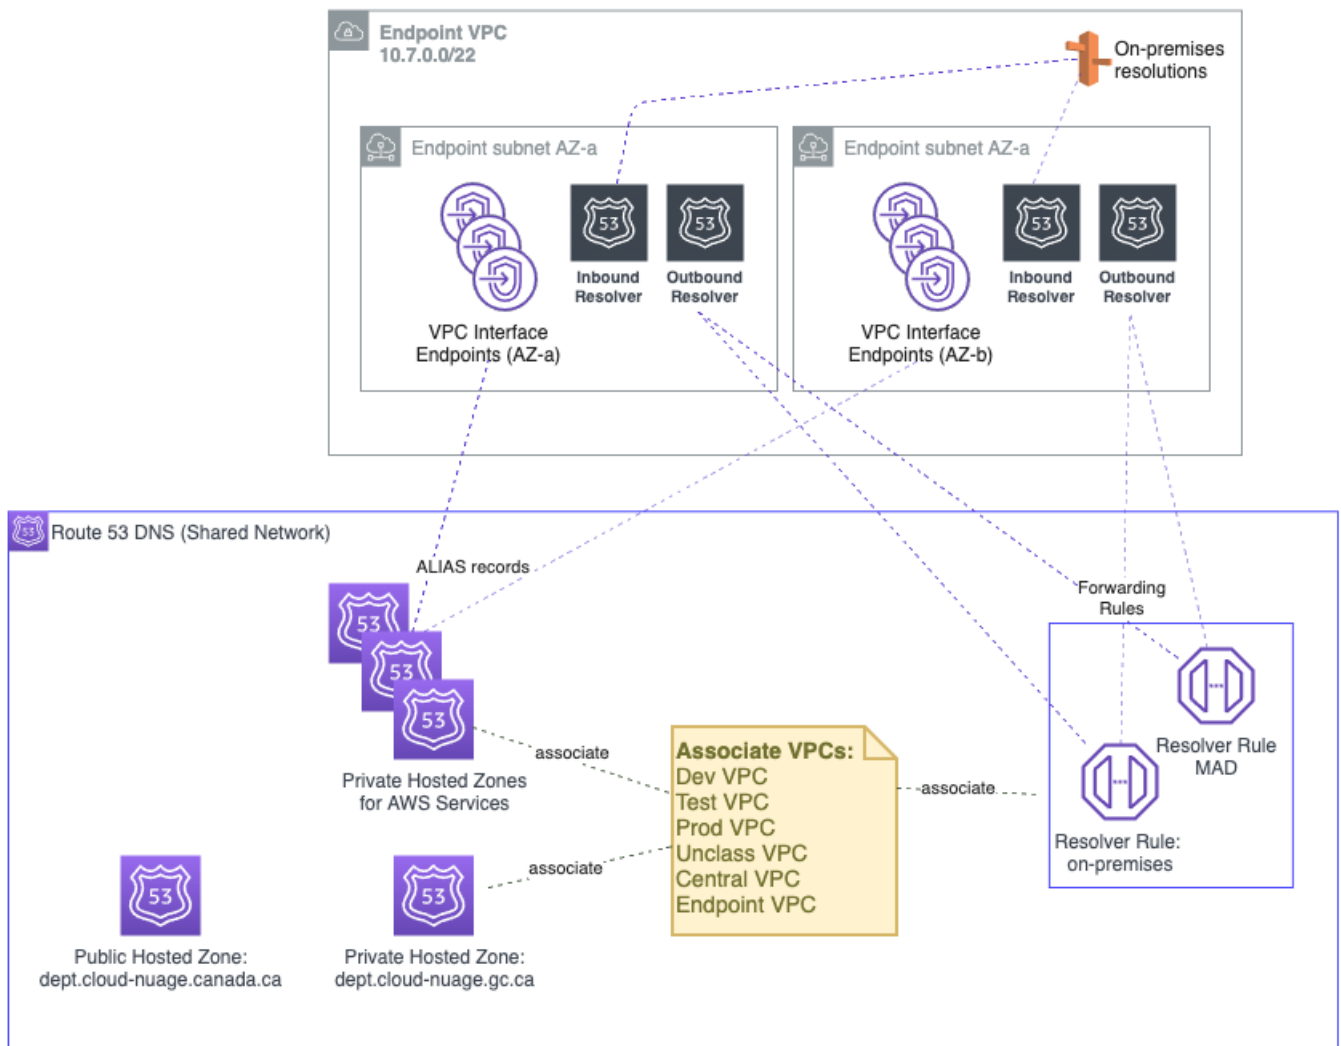
The four TGW RTs exist to serve the following main functions:

- **Segregated TGW RT:** Used as the association table for the workload VPCs; prevents east-west traffic, except to shared resources.
- **Core TGW RT:** Used for internet/on-premises response traffic, and Endpoint VPC egress.
- **Shared TGW RT:** Used to provide Central VPC access east-west for the purposes of response traffic to shared workloads
- **Standalone TGW RT:** Reserved for future use. Prevents TGW routing except to the Endpoint VPC.

Note that a unique BGP ASN will need to be used for the TGW.

1.3.2. Endpoint VPC

DNS functionality for the network architecture is centralized in the Endpoint VPC. It is recommended that the Endpoint VPC use a [RFC1918](#) range - e.g. 10.7.0.0/22 with sufficient capacity to support 60+ AWS services and future endpoint expansion, and inbound and outbound resolvers (all figures per AZ).



1.3.3. Endpoint VPC: Interface Endpoints

The endpoint VPC hosts VPC Interface Endpoints (VPCs) and associated Route 53 private hosted zones for all applicable services in the designated home region. This permits traffic destined for an eligible AWS service; for example SQS, to remain entirely within the Shared Network account rather than transiting via the IPv4 public endpoint for the service:

The screenshot shows the AWS Route 53 Management Console. On the left is a navigation menu with categories like Dashboard, Hosted zones, Health checks, Traffic flow, Domains, Resolver, and VPCs. The main area displays a list of 64 hosted zones, with 'sqs.ca-central-1.amazonaws.com' selected. The details panel on the right shows the domain name, type (Private Hosted Zone for Amazon VPC), hosted zone ID, record set count (3), and a list of associated VPCs including Prod_vpc, Test_vpc, Central_vpc, UnClass_vpc, Dev_vpc, and Endpoint_vpc. An 'Important' note states that to use private hosted zones, certain Amazon VPC settings must be set to true: 'enabledDnsHostnames' and 'enabledDnsSupport'.

From within an associated workload VPC such as `Dev`, the service endpoint (e.g. `sqs.ca-central-1.amazonaws.com`) will resolve to an IP in the `Endpoint` VPC:

```
sh-4.2$ nslookup sqs.ca-central-1.amazonaws.com
Server:      10.2.0.2      # Dev VPC's .2 resolver.
Address:    10.2.0.2#53

Non-authoritative answer:
Name:   sqs.ca-central-1.amazonaws.com
Address: 10.7.1.190      # IP in Endpoint VPC - AZ-a.
Name:   sqs.ca-central-1.amazonaws.com
Address: 10.7.0.135     # IP in Endpoint VPC - AZ-b.
```

This cross-VPC resolution of the service-specific private hosted zone functions via the association of each VPC to each private hosted zone, as depicted above.

1.3.4. Endpoint VPC: Hybrid DNS

The Endpoint VPC also hosts the common DNS infrastructure used to resolve DNS queries:

- within the cloud
- from the cloud to on-premises
- from on-premises to the cloud

1.3.4.1. WITHIN THE CLOUD

In-cloud DNS resolution applies beyond the DNS infrastructure that is put in place to support the Interface Endpoints for the AWS services in-region. Other DNS zones, associated with the Endpoint VPC, are resolvable the same way via an association to workload VPCs.

1.3.4.2. FROM CLOUD TO ON-PREMISES

DNS Resolution from the cloud to on-premises is handled via the use of a Route 53 Outbound Endpoint, deployed in the Endpoint VPC, with an associated Resolver rule that forwards DNS traffic to the outbound endpoint. Each VPC is associated to this rule.

The screenshot displays the AWS Route 53 Resolver console for a specific Resolver rule. The rule is named "private-domain7-example-ca-phz-rule" and is in a "Complete" status. The configuration details are as follows:

ID	Name	Sharing status
rslvr-rr-d908f2538a614f718	private-domain7-example-ca-phz-rule	Not shared

Status	Type	Owner
Complete	Forward	372303786158

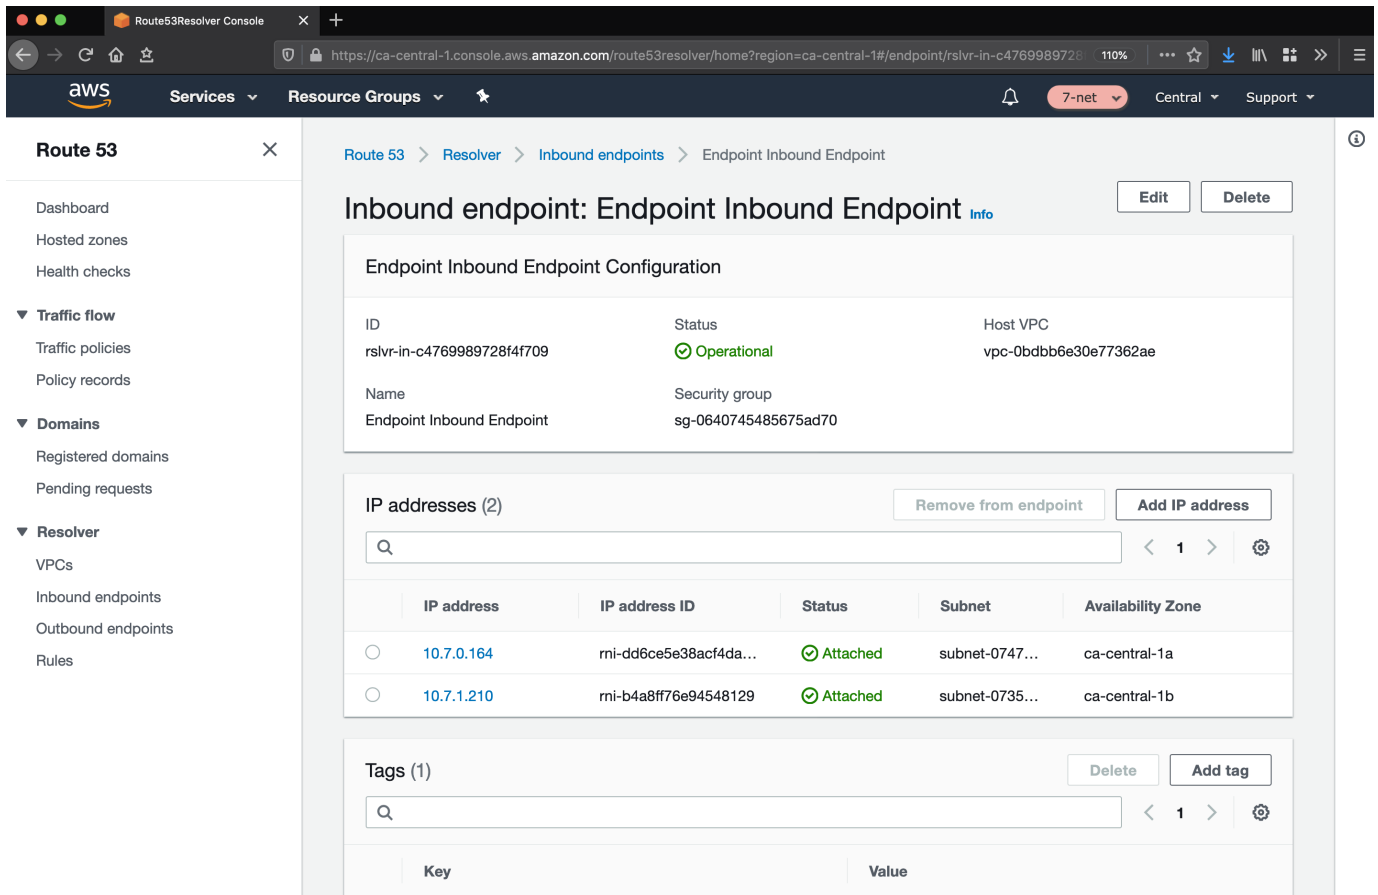
Domain name	Outbound endpoint
private-domain7.example.ca.	rslvr-out-3871bb4cc2b54e508

Below the configuration, there is a section for "VPCs (6)" with a search bar and "Disassociate" and "Associate VPC" buttons. A table lists the associated VPCs:

VPC ID	Name	Association ID	Status
vpc-0bdbb6e30e77362ae		rslvr-rrassoc-7341e948e25c482f9	Complete
vpc-05cf294070fc4437		rslvr-rrassoc-3cdf18977ecd4c5da	Complete
vpc-0a16024113a653f00		rslvr-rrassoc-e884860fa7874ab7a	Complete
vpc-05531dc7c5134850a		rslvr-rrassoc-9c142067197d4dcd8	Complete
vpc-0276b1f739be64487		rslvr-rrassoc-ece73f2aa3d14e8a8	Complete
vpc-089b46016fa3f9020		rslvr-rrassoc-484e2098da1a4058b	Complete

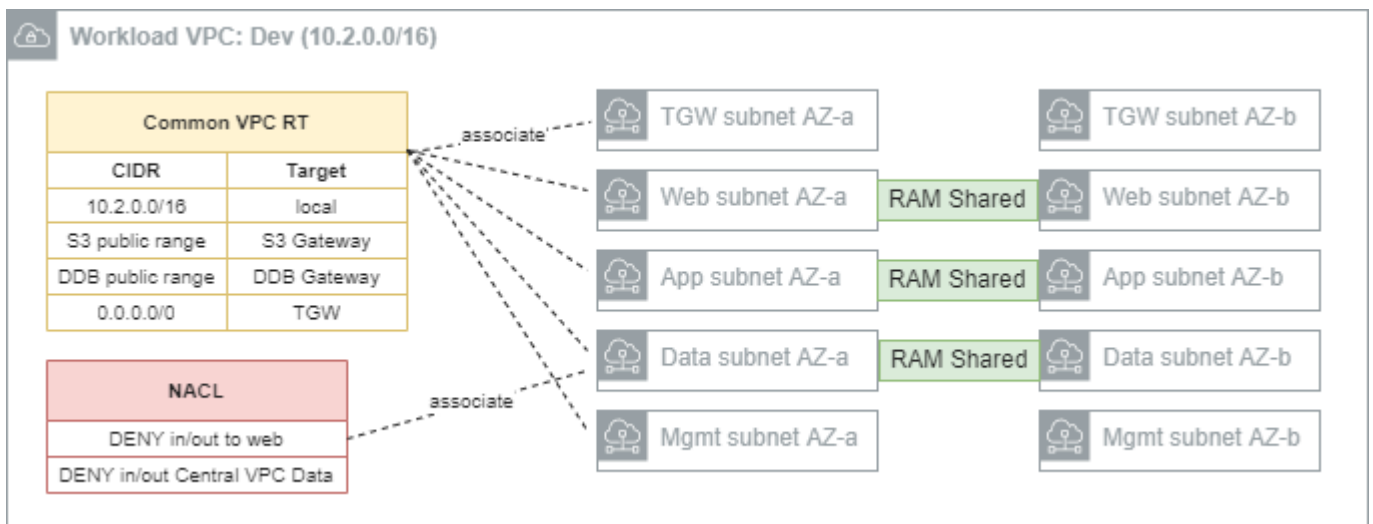
1.3.4.3. FROM ON-PREMISES TO CLOUD

Conditional forwarding from on-premises networks is made possible via the use of a Route 53 Inbound Endpoint. On-premises networks send resolution requests for relevant domains to the endpoints deployed in the Endpoint VPC:



1.3.5. Workload VPCs

The workload VPCs are where line of business applications ultimately reside, segmented by environment (Dev, Test, Prod, etc). It is recommended that the Workload VPC use a [RFC1918](#) range (e.g. 10.2.0.0/16 for Dev, 10.3.0.0/16 for Test, etc).



Note that security groups are recommended as the primary data-plane isolation mechanism between applications that may coexist in the same VPC. It is anticipated that unrelated applications would coexist in their respective tiers without ever permitting east-west traffic flows.

The following subnets are defined by the *AWS Secure Environment Accelerator Architecture*:

- **TGW subnet:** This subnet hosts the elastic-network interfaces for the TGW attachment. A /27 subnet is sufficient.
- **Web subnet:** This subnet hosts front-end or otherwise 'client' facing infrastructure. A /20 or larger subnet is recommended to facilitate auto-scaling.
- **App subnet:** This subnet hosts app-tier code (EC2, containers, etc). A /19 or larger subnet is recommended to facilitate auto-scaling.
- **Data subnet:** This subnet hosts data-tier code (RDS instances, ElastiCache instances). A /21 or larger subnet is recommended.
- **Mgmt subnet:** This subnet hosts bastion or other management instances. A /21 or larger subnet is recommended.

Each subnet is associated with a Common VPC Route Table, as depicted above. Gateway Endpoints for relevant services (Amazon S3, Amazon DynamoDB) are installed in the Common route tables of all Workload VPCs. Aside from local traffic or gateway-bound traffic, 0.0.0.0/0 is always destined for the TGW.

1.3.5.1. SECURITY GROUPS

Security Groups are instance level stateful firewalls, and represent a foundational unit of network segmentation across AWS networking. Security groups are stateful, and support ingress/egress rules based on protocols and source/destinations. While CIDR ranges are supported by the latter, it is preferable to instead use other security groups as sources/destinations. This permits a higher level of expressiveness that is not coupled to particular CIDR choices and works well with autoscaling; e.g.

"permit port 3306 traffic from the App tier to the Data tier"

versus

"permit port 3306 traffic from 10.0.1.0/24 to 10.0.2.0/24 ."

Security group egress rules are often used in 'allow all' mode (0.0.0.0/0), with the focus primarily being on consistently allow listing required ingress traffic. This ensures day to day activities like patching, access to Windows DNS, and to directory services can function without friction. The provided sample security groups in the workload accounts offers a good balance that considers both security, ease of operations, and frictionless development. They allow developers to focus on developing, enabling them to simply use the pre-created security constructs for their workloads, and avoid the creation of wide-open security groups. Developers can equally choose to create more appropriate least-privilege security groups more suitable for their application, if they are skilled in this area. It is expected as an application is promoted through the SDLC cycle from Dev through Test to Prod, these security groups will be further refined by the extended customer teams to further reduce privilege, as appropriate. It is expected that each customer will review and tailor their Security Groups based on their own security requirements.

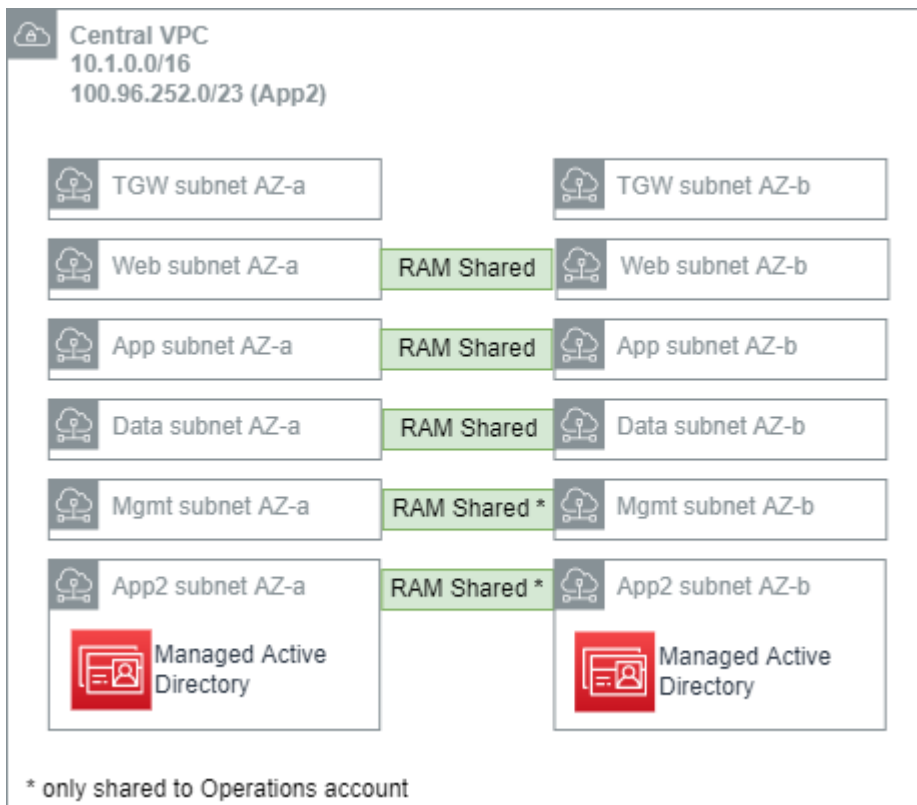
1.3.5.2. NACLs

Network Access-Control Lists (NACLs) are stateless constructs used sparingly as a defense-in-depth measure in this architecture. AWS generally discourages the use of NACLs given the added complexity and management burden, given the availability and ease of use provided by security groups. Each network flow often requires four NACL entries (egress from ephemeral, ingress to destination, egress from destination, ingress to ephemeral). The architecture recommends NACLs as a segmentation mechanism for Data subnets; i.e. deny all inbound traffic to such a subnet except that which originates in the App subnet for the same VPC. As with security groups, we encourage customers to review and tailor their NACLs based on their own security requirements.

1.3.6. Central VPC

The Central VPC is a network for localizing operational infrastructure that may be needed across the organization, such as code repositories, artifact repositories, and notably, the managed Directory Service (Microsoft AD). Instances that are domain joined will connect to this AD domain - a network flow that is made possible from anywhere in the network structure due to the inclusion of the Central VPC in all relevant association TGW RTs.

It is recommended that the Central VPC use a [RFC1918](#) range (e.g. 10.1.0.0/16) for the purposes of routing from the workload VPCs, and a secondary range from the [RFC6598](#) block (e.g. 100.96.252.0/23) to support the Microsoft AD workload.



1.3.6.1. DOMAIN JOINING

An EC2 instance deployed in the Workload VPCs can join the domain corresponding to the Microsoft AD in Central, provided the following conditions are all true:

1. The instance needs a network path to the Central VPC (given by the Segregated TGW RT), and appropriate security group assignment;
2. The Microsoft AD should be 'shared' with the account the EC2 instance resides in (The *AWS Secure Environment Accelerator Architecture* recommends these directories are shared to workload accounts);
3. The instance has the AWS managed policies `AmazonSSMManagedInstanceCore` and `AmazonSSMDirectoryServiceAccess` attached to its IAM role, or runs under a role with at least the permission policies given by the combination of these two managed policies; and
4. The EC2's VPC has an associated resolver rule that directs DNS queries for the AD domain to the Central VPC.

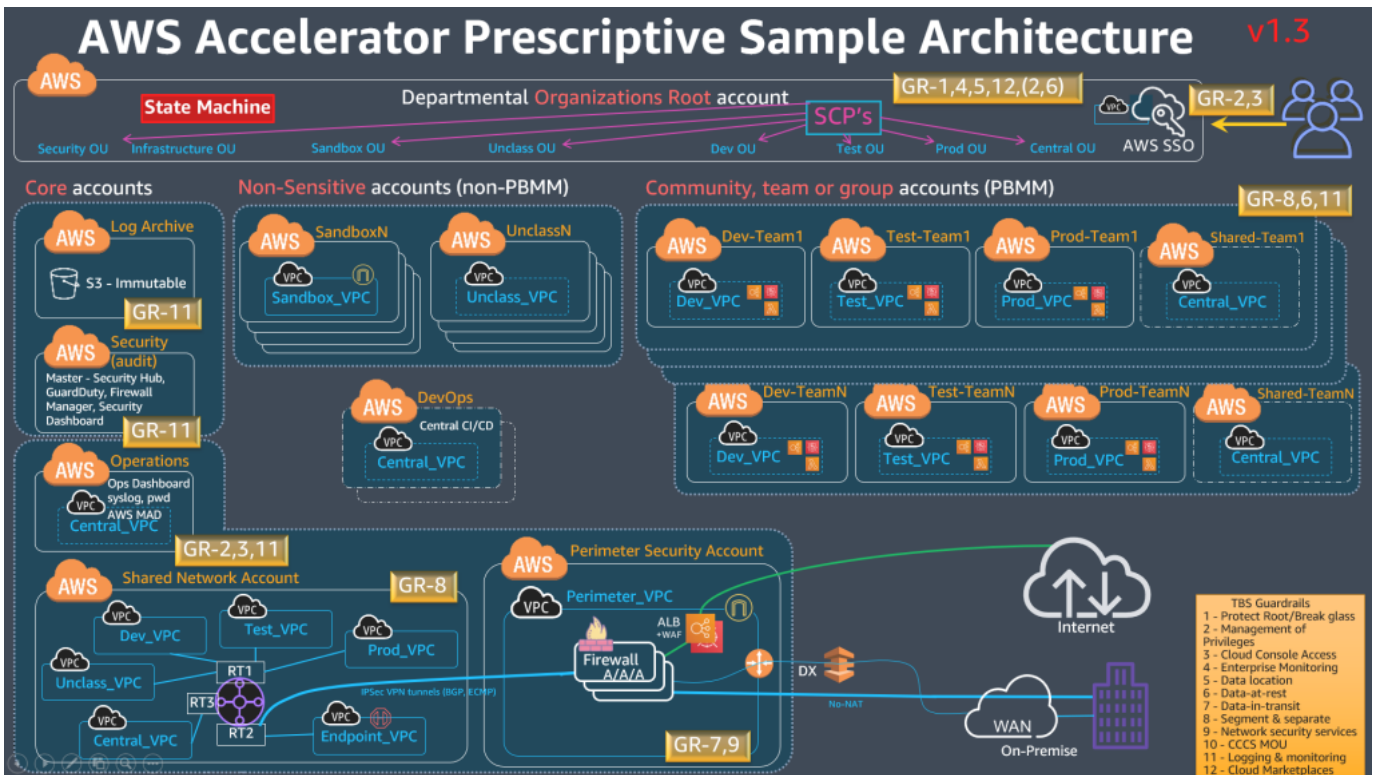
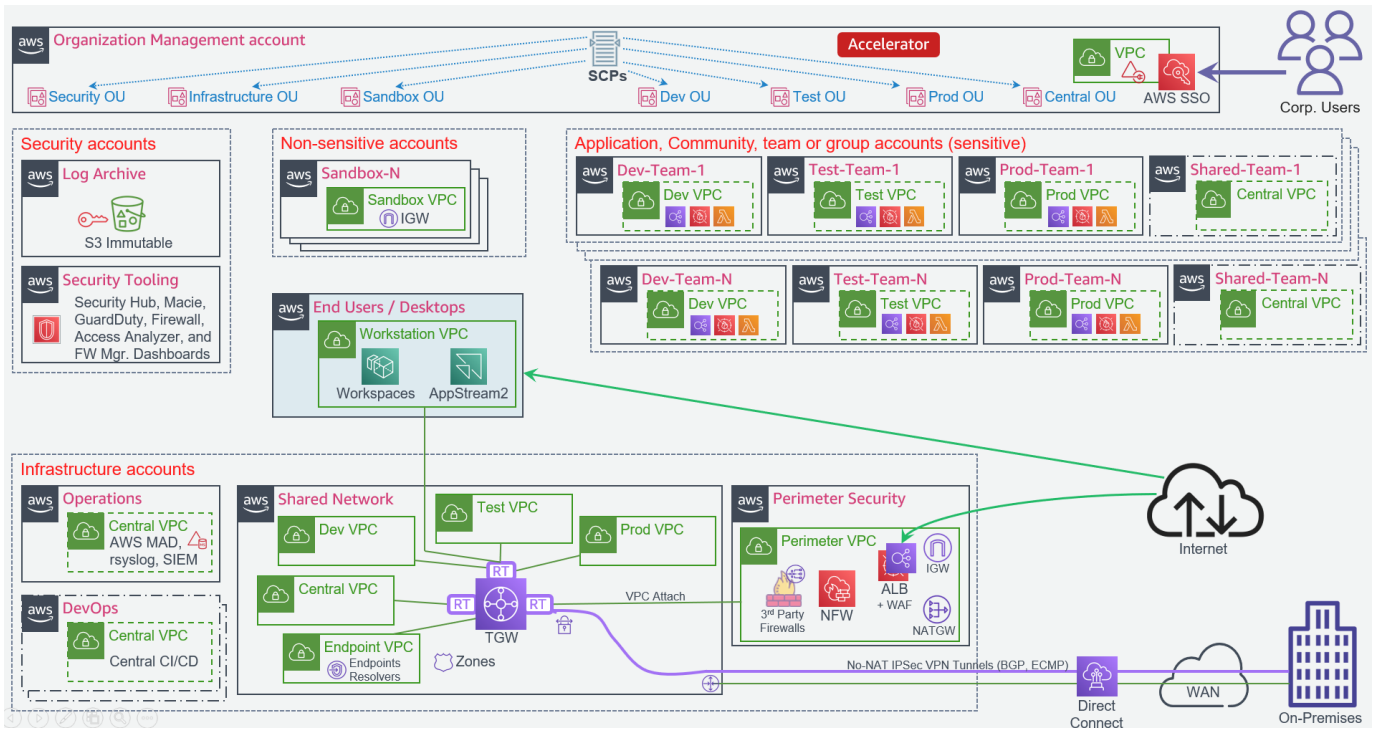
1.3.7. Sandbox VPC

A sandbox VPC, not depicted, may be included in the *AWS Secure Environment Accelerator Architecture*. This is **not** connected to the Transit Gateway, Perimeter VPC, on-premises network, or other common infrastructure. It contains its own Internet Gateway, and is an entirely separate VPC with respect to the rest of the *AWS Secure Environment Accelerator Architecture*.

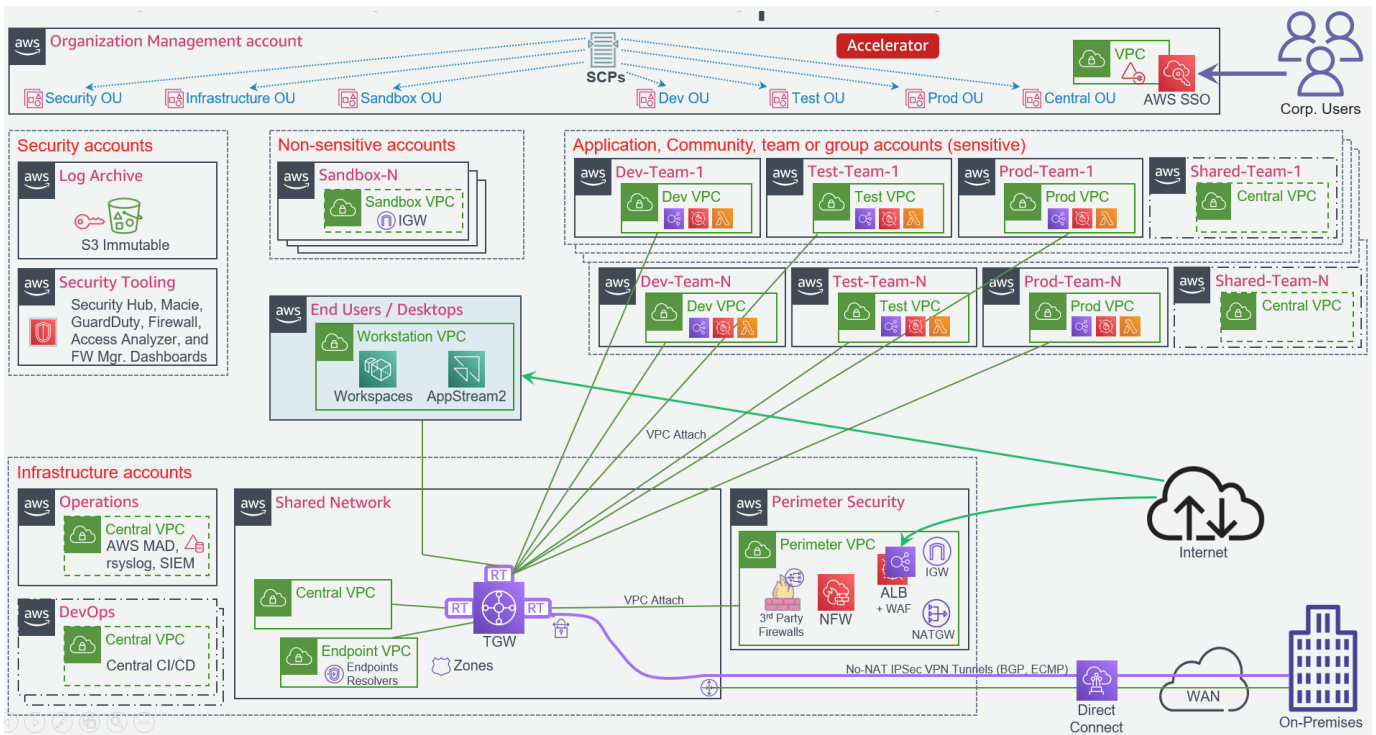
The sandbox VPC should be used exclusively for time-limited experimentation, particularly with out-of-region services, and never used for any line of business workload or data.

6.7 1. Prescriptive Sensitive Sample Architecture Diagrams

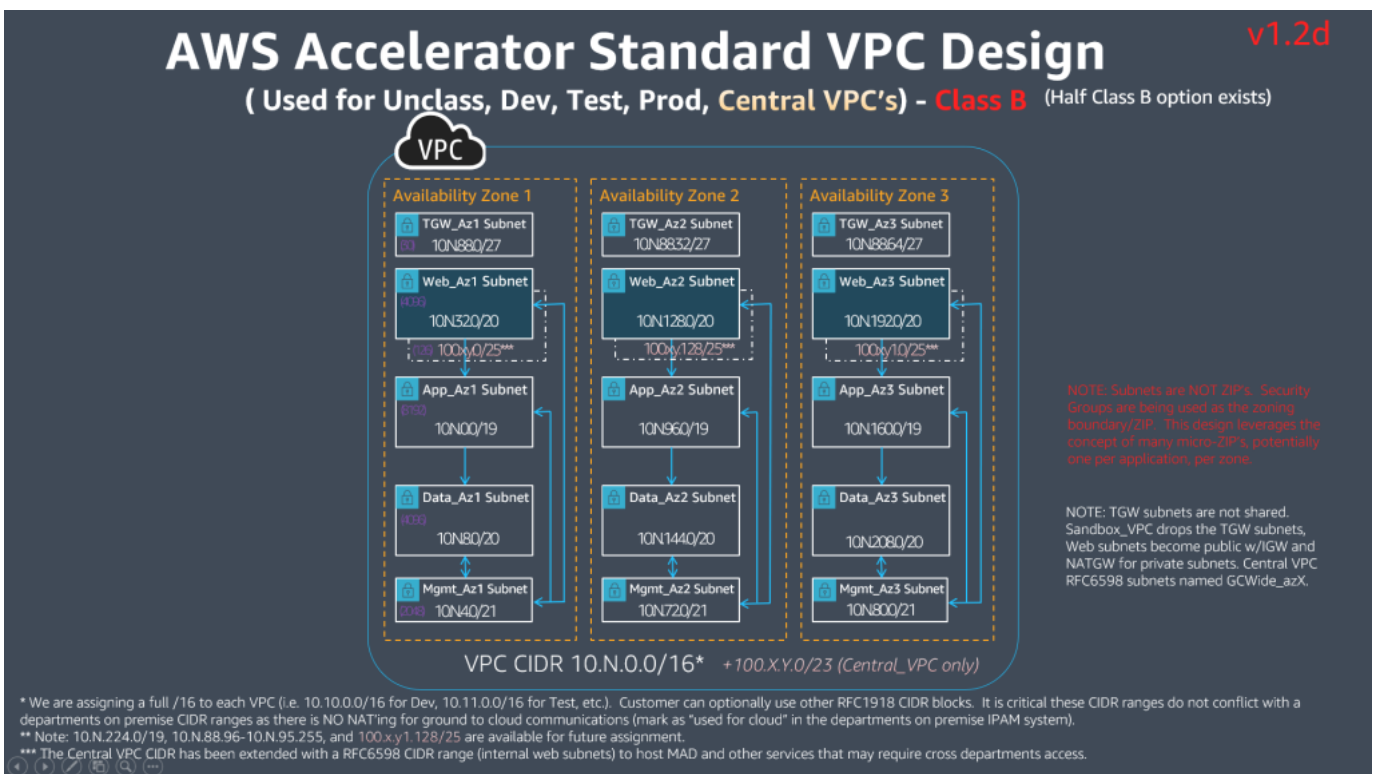
6.7.1 1.1. Shared VPC Architecture

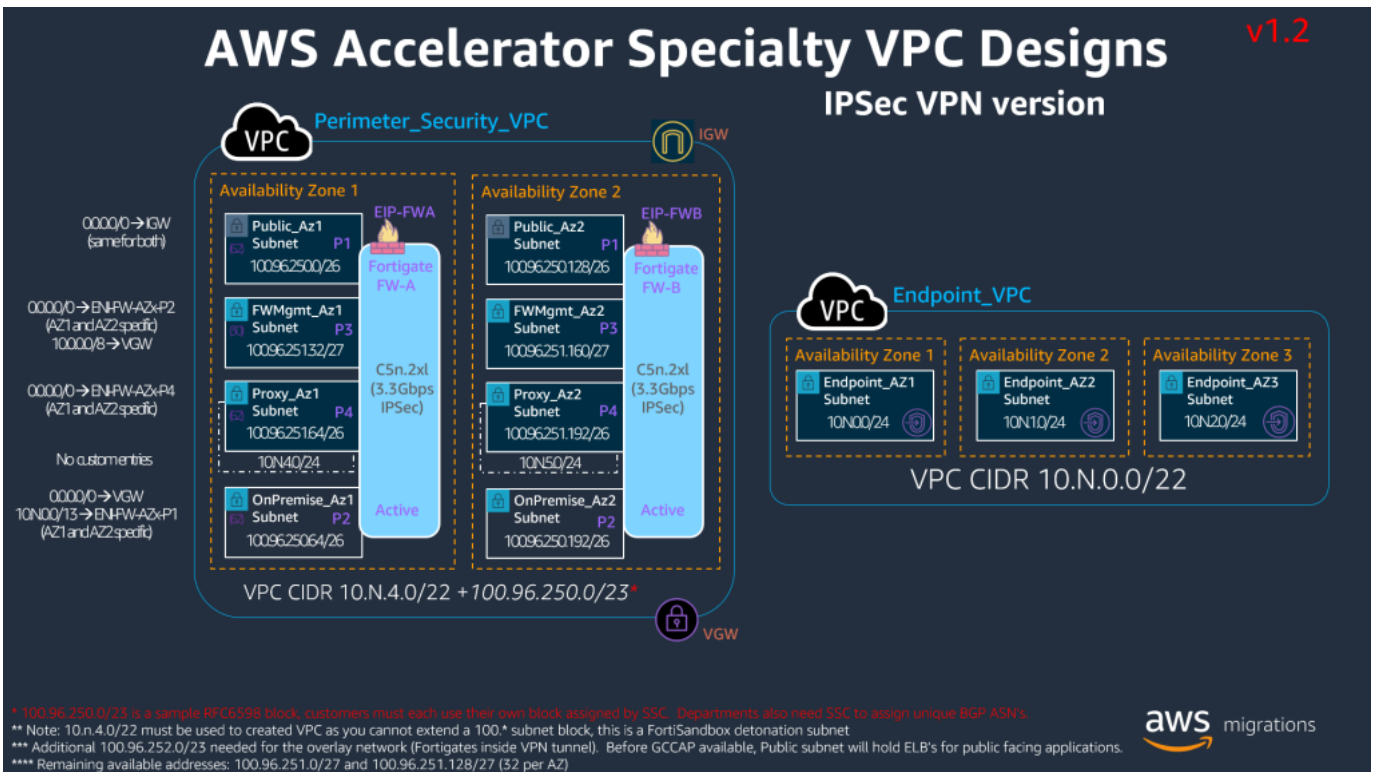
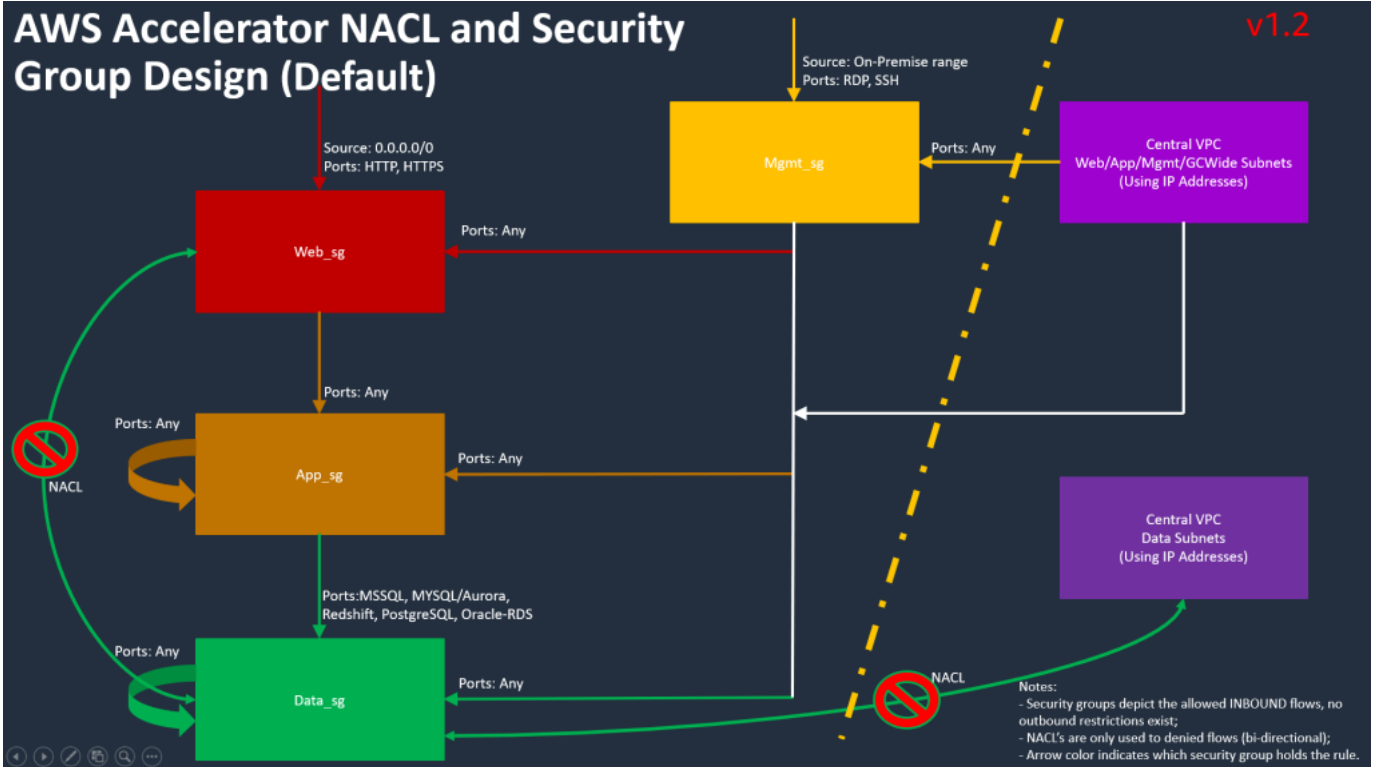


6.7.2 1.2. Spoke VPC Architecture




6.7.3 1.3. VPC and Security Group Patterns





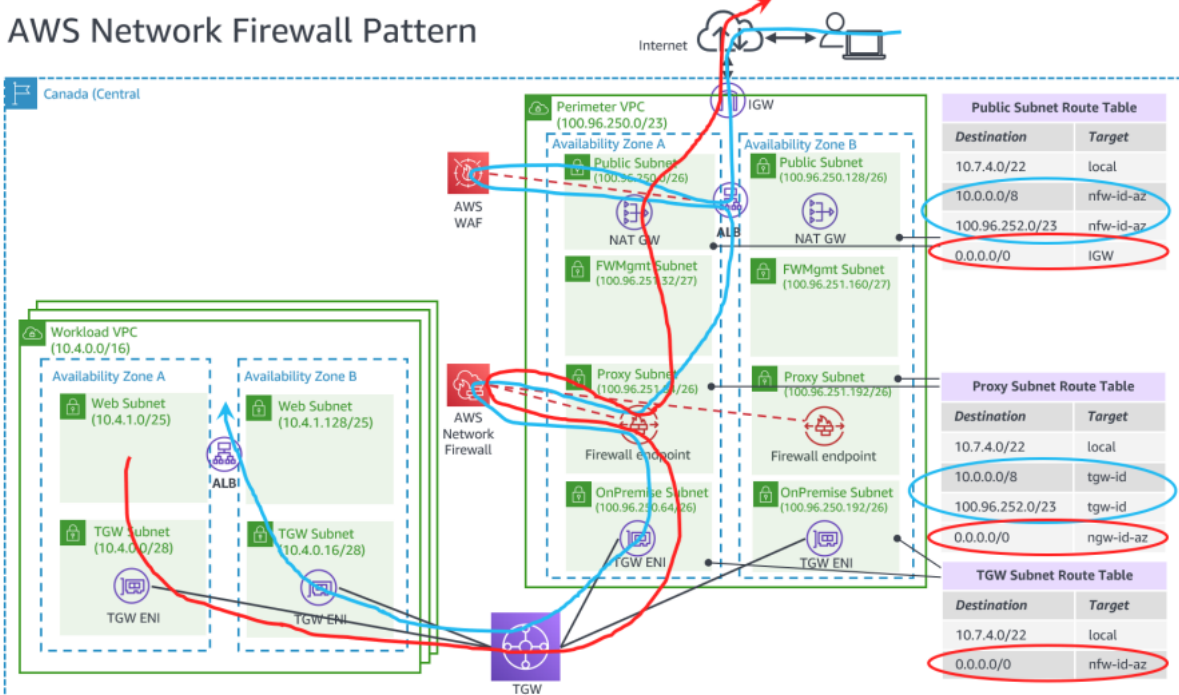
6.7.4 1.4. Additional Perimeter Patterns

AWS Network Firewall Pattern



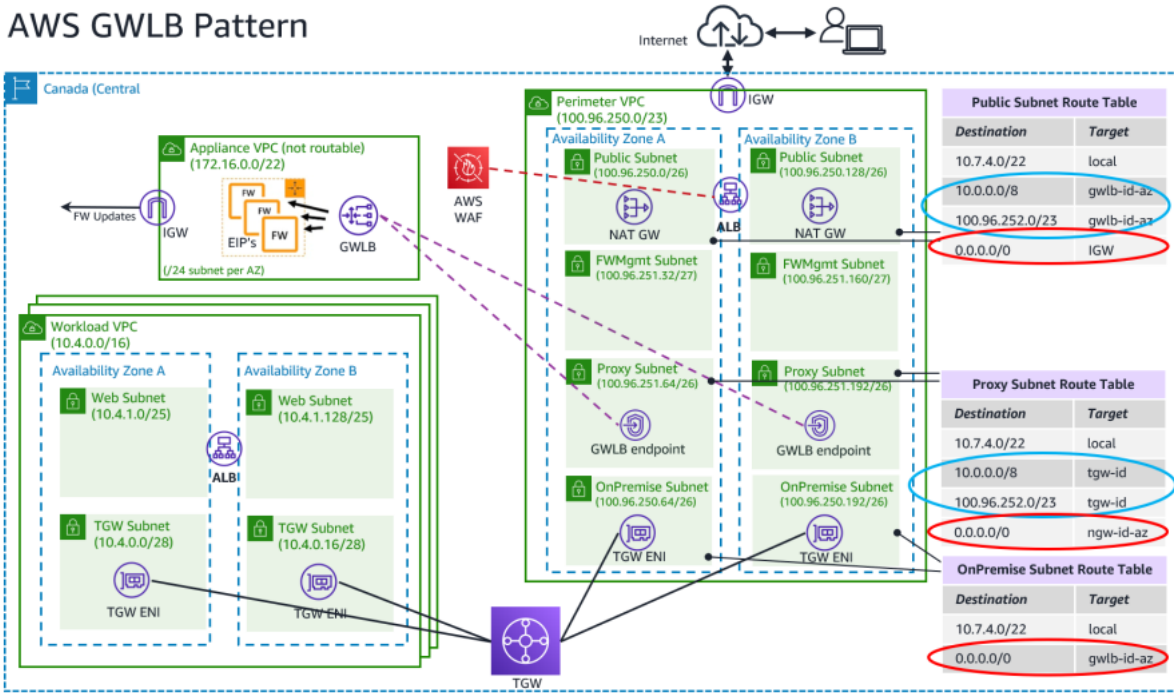
NOTE1: Distinct route tables required per AZ, which targets the local AZ's nfw, gwlb or ngw endpoint

AWS Network Firewall Pattern



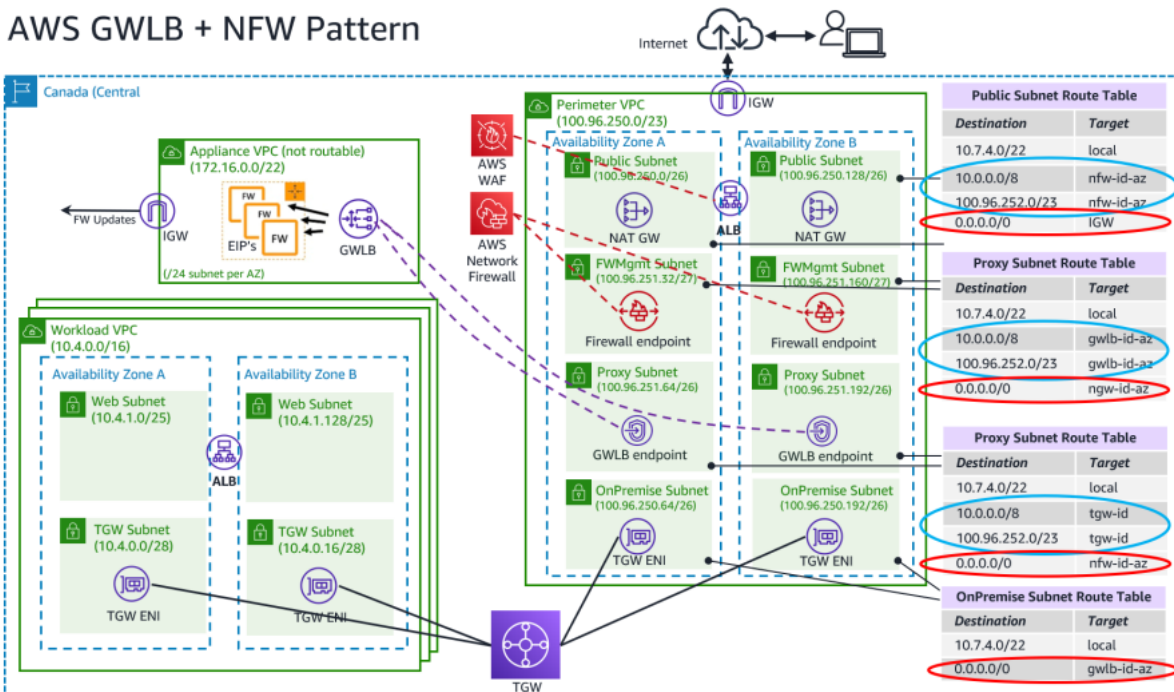
NOTE1: Distinct route tables required per AZ, which targets the local AZ's nfw, gwlb or ngw endpoint

AWS GWLB Pattern



NOTE1: Distinct route tables required per AZ, which targets the local AZ's nfw, gwlb or ngw endpoint

AWS GWLB + NFW Pattern



NOTE1: Distinct route tables required per AZ, which targets the local AZ's nfw, gwlb or ngw endpoint

7. Workshops

7.1 Accelerator Workshops

7.1.1 Accelerator Administrator Immersion Day

The Accelerator Administrator Immersion Day is focused on helping administrators who will be administering the landing zone understand how they can design, build and operate the components in ASEA. Click [here](#) for an overview of the topics covered.

7.1.2 Accelerator Workload/Application Team Immersion Day

The Accelerator Workload/Application Team Immersion Day is focused on helping project teams understand what it means to operate within an ASEA managed environment. Click [here](#) for an overview of the topics covered.