Monitoring Amazon Web Services

Amazon Web Services PowerPack version 127

AWS Extended Services PowerPack version 102

Table of Contents

Introduction ... **7**

What is AWS? ... 8

What is an AWS Region? ... 8

What is an AWS Zone? ... 9

What Does the Amazon Web Services PowerPack Monitor? ... 9

Installing the Amazon Web Services PowerPack ... 12

PowerPacks ... 13

Controlling What is Discovered by the PowerPack ... **14**

Configuring Region Discovery ... 14

Configuring AWS for Monitoring Regions Using CloudWatch Namespaces ... 15

Configuring AWS Monitoring ... **17**

Configuring AWS to Report Billing Metrics ... 18

Filtering EC2 Instances By Tag ... 18

Filtering EC2 Instances by Tag in the Skylar One Classic User Interface ... 20

Automatic Skylar One Organization Creation ... 21

Monitoring Consolidated Billing Accounts ... 22

ScienceLogic Events and AWS Alarms ... 22

Using a Proxy Server ... 22

Configuring "AWS: Lambda Service Discovery" ... 23

Configuring "AWS: Lambda Function Qualified Discovery" ... 24

Configuring AWS Integration with Docker ... 25

Configuring AWS Integration with Kubernetes ... 26

Enabling the Prometheus Metrics Server ... 28

Define the Cluster Role ... 28

Define the ClusterRoleBinding ... 30

Map the IAM User or Role to the Kubernetes RBAC Role ... 31

Example 1 ... 31

Example 2 ... 33

Amazon API Throttling Events ... 34

Support for AWS GovCloud Regions ... 34

Migrating from Using an IAM Key Per Account to Using AssumeRole ... 35

Minimum Permissions ... **37**

Minimum Permissions Needed to Monitor Your AWS Accounts ... 37

Creating a Minimum Permissions Policy for GovCloud Accounts ... 51

AWS Discovery **57**
Discovering Amazon Web Services 57
Preparing your AWS Organization or AWS Account for Discovery 58
Organization Account 58
IAM User 59
EC2 "Credential-less" 67
Individual Account 76
Discovering your AWS Organization or AWS Account 83
AWS Universal EC2 Guided Discovery (Organization Credential-less) 84
AWS Universal IAM Guided Discovery (Individual Account) 86
AWS Universal Assume Role Guided Discovery (Organization IAM) 88
Discovering your AWS Organization or AWS Account in the Classic Skylar One User Interface 90
Creating Credential in the Classic Skylar One User Interface 90
AWS Credential (Individual Account) 90
AWS Credential - Master Account (Organization IAM) 92
AWS Credential - EC2 Instance (Organization "Credential-less") 94
Launching the Discovery Session in the Classic Skylar One User Interface 96
AWS Organizations 96
Individual Accounts 97
Appendix 98
The AWS Credential Test and Viewing Component Devices **104**
Testing the AWS Credential 104
Testing the AWS Credential in the Skylar One Classic User Interface 107
Viewing AWS Component Devices 108
Relationships Between Component Devices 110
Vanishing Component Devices 111
Configuring Inbound CloudWatch Alarms **113**
CloudWatch Alarm Event Policies 113
Creating Custom CloudWatch Metrics 115
Configuring CloudWatch to Send Alarms for a Metric 118
Enabling Custom Metrics Collection in Skylar One 119
Configuring the "AWS: CloudWatch Alarms Performance" Dynamic Application 119
Enabling CloudWatch Alarm Events in Skylar One 121
Preserving CloudWatch Alarm Event Changes 122
AWS Reports **123**

AWS Billing Report 124
AWS Inventory Report 126
AWS Running Config Report 129
AWS Dashboards 131
Installing the Amazon Web Services: Dashboards PowerPack 131
AWS Account Billing Dashboard 132
AWS Health Status Dashboard 132
Configuring the AWS Dashboards 133
AWS Service Instance Performance Dashboards 133
AWS Run Book Actions and Automations 135
About the Run Book Actions and Automations 135
Disabling EC2 and EBS Instances by EC2 Tag 137
Modifying the Parameters of the Automation Actions 138
Enabling the Component Device Record Created Event Policy 138
Enabling the Automation Policies 139
Preserving Automation Changes 139
Discovering EC2 Instances by Public or Private IP Address 140
Modifying the Parameters of the Automation Actions 141
Enabling the Component Device Record Created Event Policy 143
Enabling the Device Record Created Event Policy 143
Enabling the Automation Policies 144
Preserving Automation Changes 144
Aligning AWS Regions to the AWS Region Device Class 145
Vanishing Terminated or Terminating EC2 Instances 145
Enabling the Automation Policies 146
Preserving Automation Changes 146
Key Metrics Collected by the PowerPack 148
Amazon MQ Service 150
API Gateway Service 156
Application ELB Service 159
AppRunner Service 161
AppStream Service 162
AppSync Service 162
Athena Service 165
Auto Scale Service 166

CloudFront Service 168
CloudHSM Service 173
CloudSearch Service 174
CloudTrail Service 174
CloudWatch Service 175
CloudWatch Synthetics Service 176
CodeBuild Service 177
Cognito Service 179
Connect Service 180
DAX Service 183
DDB Service 187
Direct Connect Service 189
DMS Service 193
DocumentDB Service 195
DynamoDB Service 201
EBS Service 202
EC2 Service 204
ECS Service 212
EFS Service 219
EKS Service 221
Elastic Beanstalk Service 223
Elastic Transcoder Service 227
ElastiCache Service 228
ELB Service 234
EMR Service 238
FSx Service 242
Glacier Service 247
Glue Service 248
IoT Service 250
Kafka/MSK Service 251
Kinesis Service 252
KMS Service 259
Lambda Service 260
Lex V2 Service 267
LightSail Service 269

MediaConnect Service 271
MediaConvert Service 280
MediaPackage Service 281
MediaTailor Service 282
Neptune Service 285
Network ELB Service 290
OpenSearch Service 292
OpsWorks Service 303
RDS Service 305
Redshift Service 317
Route 53 Service 320
S3 Service 322
SES Service 325
Shield Standard Service 325
SNS Service 327
SQS Service 329
Storage Gateway Service 331
STS Service 332
SWF Service 333
Transfer Server Service 333
Transit Gateway Service 334
VPC Service 336
WAF Service 343
Workspaces Service 345
WorkMail Service 348
Other Supported Services 348

Chapter 1

Introduction

Overview

This manual describes how to monitor Amazon Web Services (AWS) in Skylar One using the "Amazon Web Services" and "AWS Extended Services" PowerPacks. It also describes the reports you can generate and the dashboards you can view after you collect data from AWS, as well as the Run Book Action and Automation policies you can use to automate certain aspects of monitoring AWS.

The following sections provide an overview of Amazon Web Services and the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

What is AWS? 8

What is an AWS Region? 8

What is an AWS Zone? 9

What Does the Amazon Web Services PowerPack Monitor? 9

Installing the Amazon Web Services PowerPack 12

PowerPacks 13

NOTE: For more information about setting up a Skylar One appliance on an Amazon Web Services EC2 instance, see the ***Installation and Initial Configuration*** manual.

NOTE: For more information about setting up an AWS Elasticsearch, Logstash, and Kibana (ELK) stack, see the ***Monitoring AWS ELK Stacks*** manual.

NOTE: ScienceLogic provides this documentation for the convenience of ScienceLogic customers. Some of the configuration information contained herein pertains to third-party vendor software that is subject to change without notice to ScienceLogic. ScienceLogic makes every attempt to maintain accurate technical information and cannot be held responsible for defects or changes in third-party vendor software. There is no written or implied guarantee that information contained herein will work for all third-party variants. See the End User License Agreement (EULA) for more information.

What is AWS?

Amazon Web Services is Amazon's "Infrastructure as a Service" offering. AWS includes multiple products (called ***Services***) including compute, DNS, networking, content delivery, analytics, storage, and database services, among many others.

What is an AWS Region?

An AWS region is a geographical area made up of availability zones located within that region. Each zone may have multiple data centers. Regions have a canonical naming scheme of:

country/continent-direction-number

For example, the 'us-east-1' region is located in the United States, on the east coast, and it is the #1 data center in that region.

AWS regions are also commonly referred to by the city or state in which the data center is located. For example, us-west-2 is commonly referred to as "Oregon", ap-northeast-1 is commonly referred to as "Tokyo", etc.

The Dynamic Applications in the "Amazon Web Services" PowerPack create a "region" component device for each discovered region. The component devices for regions include both the region name and city/state description. For example, the Dynamic Applications might discover a component device called "US West (Oregon): us-west-2". Component devices that represent region-specific AWS services reside under the appropriate "region" component device and appropriate "zone" component device.

NOTE: For more information about AWS regions, see https://docs.aws.amazon.com/general/latest/gr/rande.html.

What is an AWS Zone?

All instances of an AWS service reside in one or more Zones. A zone is a physical network and power partition (air-gap firewall) within a regional data center. Some AWS instances, like EC2 instances, are in a single zone. Other AWS instances, like an SNS queue, exist in all zones simultaneously.

The AWS naming convention for a zone is:

region[a-z]

For example, zone 'a' for the region 'us-east-1' is named 'us-east-1a'.

When a user deploys a service instance, the user can specify a "zone preference", but the final zone for that service instance is decided by AWS, not the user.

The Dynamic Applications in the "Amazon Web Services" PowerPack create a "zone" component device for each discovered zone.

AWS services with a specific zone affinity reside under the appropriate zone component device. For example, the Dynamic Applications in the PowerPack might discover the zone "us-west-1b" and create a component device called "us-west-1b".

AWS services that are specific to a zone reside under the appropriate "region" component device and appropriate "zone" component device. The Dynamic Applications in the PowerPack create a "multi-zoned" component device for services that are inherently zone agnostic such as the Simple Queue Service (SQS).

Component devices that represent Zones are a named container with no associated performance metrics.

What Does the Amazon Web Services PowerPack Monitor?

To collect data from Amazon Web Services, the ScienceLogic Data Collector or All-In-One Appliance connects via HTTPS to the URLs listed in the following AWS document: http://docs.aws.amazon.com/general/latest/gr/rande.html.

The "Amazon Web Services" PowerPack includes Dynamic Applications that can monitor performance metrics and collect configuration data for numerous AWS Services and components, which can be viewed in the *Key Metrics Collected by the PowerPack* section.

NOTE: The following services are not monitored for GovCloud accounts:

- API Gateway private integrations
- AppStream
- AppSync
- Cassandra (Keyspaces)
- Chime
- CloudFront
- CloudSearch
- DAX
- Elastic Transcoder
- GameLift
- HealthLake
- Inspector2
- IoT Analytics
- IoT TwinMaker
- Kinesis Video
- Lex
- Lightsail
- MediaConnect
- MediaConvert
- MediaLive
- MediaPackage
- MediaPackage VOD
- MediaTailor
- OpsWorks
- Replica Lambda functions
- Shield
- TimeStream
- Web Application Firewall
- WorkMail

NOTE: Not all AWS services are supported by all AWS regions. For more information about which AWS services are supported by which AWS regions, see https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services.

NOTE: To monitor performance metrics for an AutoScale group, you must activate detailed instance monitoring for that group. For instructions on how to perform this task, see http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-instance-monitoring.html.

NOTE: When monitoring EC2-backed ECS clusters, you can optionally use the "Docker" PowerPack to collect container information in addition to what the AWS API provides for the ECS service. For more information, see the section on *Configuring AWS Integration with Docker*.

NOTE: To monitor Lambda services, you must first configure some of the Dynamic Applications in the "Amazon Web Services" PowerPack prior to discovery. For more information, see the *Configuring "AWS Lambda Service Discovery"* and *Configuring "AWS Lambda Function Qualified Discovery"* sections.

The Dynamic Applications in the PowerPack also monitor:

- The general health of each AWS service
- Current billing metrics for each service aligned with the account
- Custom, application-specific performance metrics configured on the account
- The state of any AWS Alarms set on metrics in Cloudwatch

In addition to Dynamic Applications, the PowerPack includes the following features:

- Event Policies and corresponding alerts that are triggered when AWS component devices meet certain status criteria
- Device Classes for each of the AWS component devices monitored
- Sample Credentials for discovering AWS component devices
- Reports and dashboards that display information about AWS instances and component devices
- Run Book Action and Automation policies that can automate certain AWS monitoring processes

NOTE: To view Amazon Web Services dashboards, you must first install the "Amazon Web Services: Dashboards" PowerPack. For more information, see the *AWS Dashboards* chapter.

Installing the Amazon Web Services PowerPack

Before completing the steps in this manual, you must import and install the latest version of the "Amazon Web Services" PowerPack.

NOTE: If you are upgrading from an earlier version of the PowerPack, see the Release Notes for the version you are installing for upgrade instructions.

TIP: By default, when you install a new version of a PowerPack, it overwrites all content from that PowerPack that is already installed on the target system. You can use the ***Enable Selective PowerPack Field Protection*** setting on the **Behavior Settings** page (System > Settings > Behavior) to prevent the new version of the PowerPack from overwriting local changes for some commonly customized fields.

NOTE: For details on upgrading Skylar One, see the relevant *Skylar One Platform Release Notes*.

To download and install the PowerPack:

1. Search for and download the PowerPack from the **PowerPacks** page at the ScienceLogic Support Center (Skylar One > PowerPacks, login required).
2. In Skylar One, go to the **PowerPacks** page (System > Manage > PowerPacks).
3. Click the **[Actions]** menu and select *Import PowerPack*. The **Import PowerPack** modal appears.
4. Click **[Browse]** and go to the PowerPack file from step 1.
5. Select the PowerPack file and click **[Import]**. The **PowerPack Installer** modal displays a list of the PowerPack contents.
6. Click **[Install]**. The PowerPack appears on the **PowerPacks** page.

NOTE: If you exit the **PowerPack Installer** modal without installing the imported PowerPack, the imported PowerPack does not appear on the **PowerPacks** page. However, the imported PowerPack appears in the **Imported PowerPacks** modal. This modal appears when you click the **[Actions]** menu and select *Install PowerPack*.

CAUTION: If you need to upgrade to a new version of the "Amazon Web Services" PowerPack, do **not** uninstall the PowerPack on Skylar One versions 11.2.0 and above first, as guided discovery workflows will be lost. Follow the upgrade instructions in the "Amazon Web Services" PowerPack release notes.

PowerPacks

This manual describes content from the following PowerPack versions:

- "Amazon Web Services" version 127
- "AWS Extended Services" version 102

Chapter 2

Controlling What is Discovered by the PowerPack

Overview

The following sections describe the different methods to control what you can discover and monitor with the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

Configuring Region Discovery *14*

Configuring AWS for Monitoring Regions Using CloudWatch Namespaces *15*

Configuring Region Discovery

The ***Regions*** header is an optional header that can be inserted into the AWS SOAP/XML credential you will create to restrict which regions are discovered. This header supports a comma-separated list of regions that will be discovered and monitored. For example, the credential below shows the header with two specific regions. In this case, only those two regions would be discovered and monitored. To restrict discovery of a region, type a "!" in front of the region, for example, `Regions: !us-east-1`.

NOTE: When using guided discovery, ignore the ***Configuration*** field in your credential as the PowerPack will automatically determine which regions will be utilized.

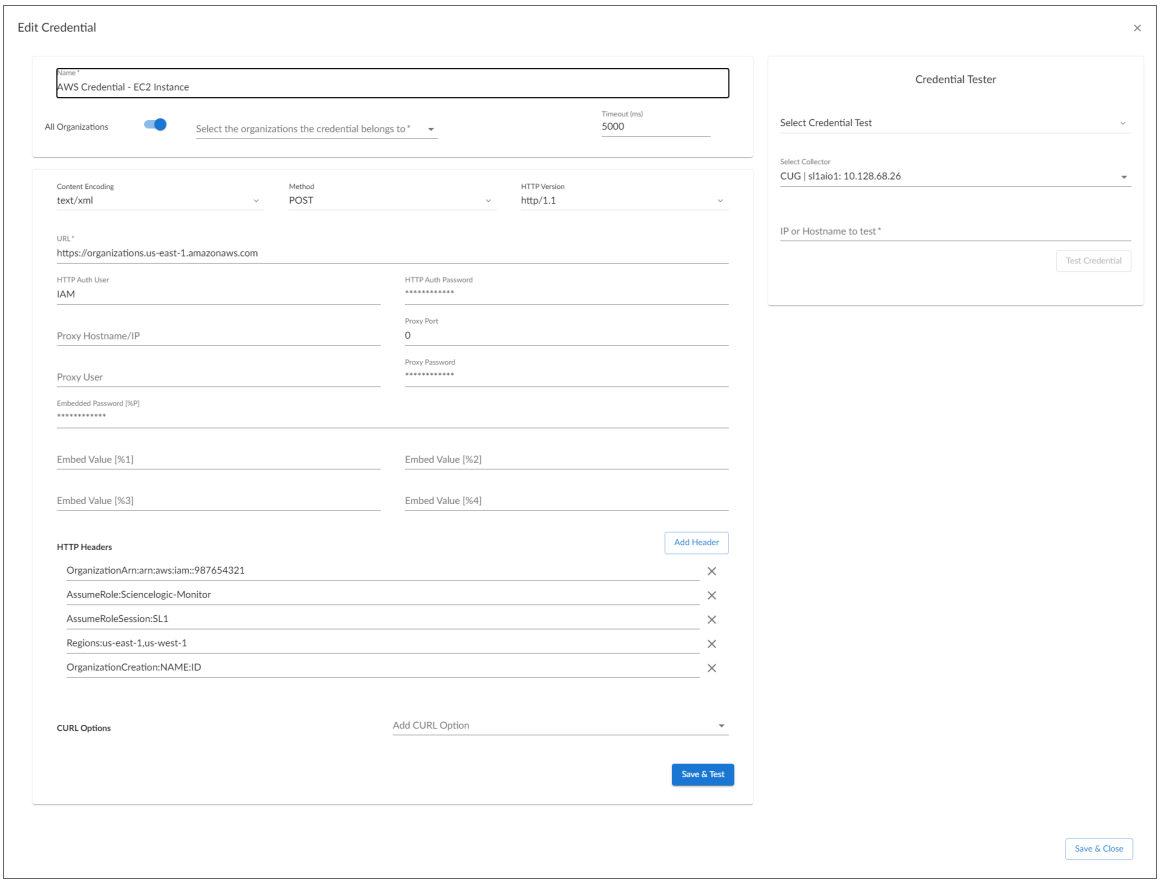

Configuring AWS for Monitoring Regions Using CloudWatch Namespaces

NOTE: These steps will be applied to all discovered AWS accounts on your Skylar One system.

Users discovering with CloudWatch metrics can also discover regions where a specific namespace is available by editing the NAMESPACES field in the aws_region_discovery snippet in the "AWS: Region Discovery" Dynamic Application.

To edit the snippet:

1. Find the "AWS: Region Discovery" Dynamic Application in the **Dynamic Applications Manager** page (System > Manage > Applications) and click its wrench icon (🔧).
2. Click the **[Snippets]** tab and then click the wrench icon (🔧) for the aws_regon_discovery snippet.
3. Edit the NAMESPACES field to include the namespace for your region. For example:

```
NAMESPACES = [ 'AWS/SQS' ]
```

4. Click **[Save]**.

Only regions that have services grouped in the specified namespace will be discovered. Global services will also be discovered.

> **NOTE:** For more information about namespaces, see https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/viewing_metrics_with_cloudwatch.html.

Chapter 3

Configuring AWS Monitoring

Overview

This section describes the options that are available for configuring the "Amazon Web Services" PowerPack.

This chapter covers the following topics:

Configuring AWS to Report Billing Metrics 18

Filtering EC2 Instances By Tag 18

Automatic Skylar One Organization Creation 21

Monitoring Consolidated Billing Accounts 22

ScienceLogic Events and AWS Alarms 22

Using a Proxy Server 22

Configuring "AWS: Lambda Service Discovery" 23

Configuring "AWS: Lambda Function Qualified Discovery" 24

Configuring AWS Integration with Docker 25

Configuring AWS Integration with Kubernetes 26

Amazon API Throttling Events 34

Support for AWS GovCloud Regions 34

Migrating from Using an IAM Key Per Account to Using AssumeRole 35

Configuring AWS to Report Billing Metrics

To use the "AWS: Billing Performance Percent" Dynamic Application, your AWS account must meet the following requirements:

- The user account you supplied in the AWS credential must have permission to view the us-east-1 zone.
- Your AWS account must be configured to export billing metrics to the CloudWatch service.

If your AWS account is not configured to export billing metrics to the CloudWatch service, the "AWS: Billing Performance Percent" Dynamic Application will generate the following event:

```
No billing metrics can be retrieved. Your AWS account is not configured to export billing metrics into CloudWatch.
```

To configure your AWS account to export billing metrics to the CloudWatch service, perform the following steps:

1. Open a browser session and go to aws.amazon.com.
2. Click **[My Account]** and then select *Billing & Cost Management*. If you are not currently logged in to the AWS site, you will be prompted to log in.
3. After logging in, the **Billing & Cost Management Dashboard** page appears. In the left navigation bar, click **[Preferences]**. The **Preferences** page appears.
4. Select the ***Receive Billing Alerts*** checkbox.

CAUTION: If you enable this option, this option cannot be disabled.

5. Click the **[Save Preferences]** button.

Filtering EC2 Instances By Tag

To discover EC2 instances and filter them by tag, you can use the "AWS Credential - Tag Filter" sample credential to enter EC2 tag keys and values.

NOTE: Filtering EC2 instance by tag will apply to all accounts discovered.

NOTE: Any EC2 instances that have already been discovered, but do not match the tag filter, will be set to "Unavailable."

To define an AWS credential:

NOTE: If you are using a Skylar One system prior to version 11.1.0, the new user interface does not include the ***Duplicate*** option for sample credential(s). ScienceLogic recommends that you use *the classic user interface and the Save As button* to create new credentials from sample credentials. This will prevent you from overwriting the sample credential(s).

1. Go to the **Credentials** page (Manage > Credentials).
2. Locate the **AWS Credential - Tag Filter** sample credential, click its **[Actions]** icon (⋮) and select ***Duplicate***. A copy of the credential, called **AWS Credential - Tag Filter copy** appears.
3. Click the **[Actions]** icon (⋮) for the **AWS Credential - Tag Filter copy** credential and select ***Edit***. The **Edit Credential** modal page appears:

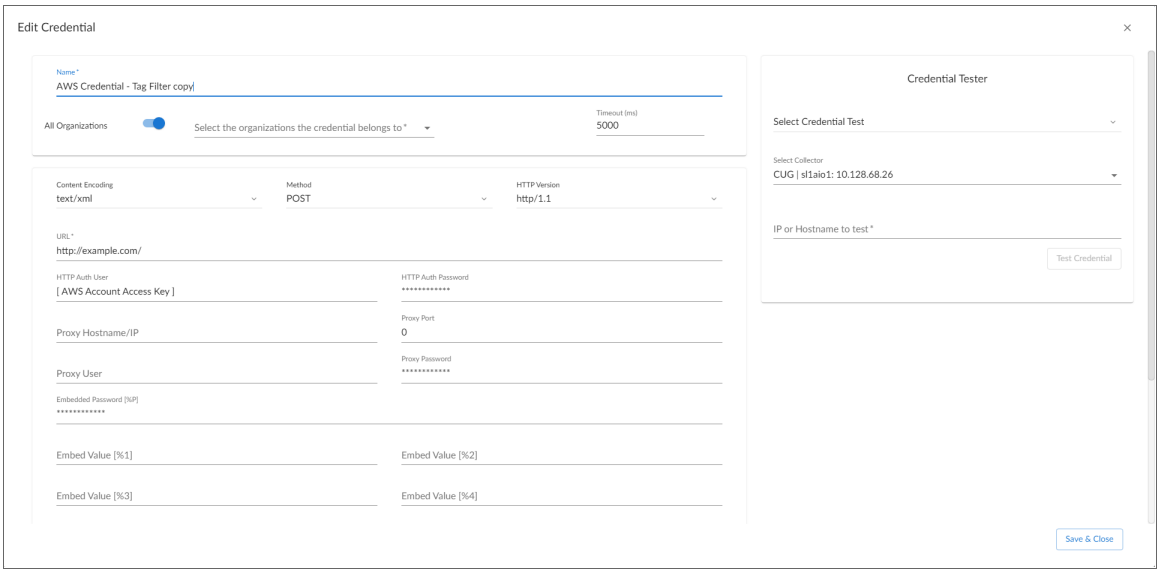

4. Supply values in the following fields:
    - ***Name***. Type a new name for your AWS credential.
    - ***All Organizations***. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the ***What organization manages this service?*** drop-down field to align the credential with those specific organizations.
    - ***Timeout (ms)***. Keep the default value.
    - ***URL***. Enter a valid URL. This field is not used for this discovery method but must be populated with a valid URL for discovery to complete.
    - ***HTTP Auth User***. Type your AWS access key ID.
    - ***HTTP Auth Password***. Type your AWS secret access key.
    - Under **HTTP Headers**, edit the header provided:

- *Tags:<operation>#<EC2-Tag-Key>#<EC2-Tag-Value>*. Type the tag, followed by its operation, tag key, or tag value. For example, if you want to filter by Tag Name, you would type the following:

  ```
  Tags:equals#Name#Example
  ```

  Valid operations include:

  - equals
  - notEquals
  - contains
  - notContains

  You can chain together multiple filters separating them by a comma. For example:

  ```
  Tags:equals#Name#Example,contains#Owner#Someone
  ```

5. Click the **[Save& Close]** button.

Filtering EC2 Instances by Tag in the Skylar One Classic User Interface

To discover EC2 instances and filter them by tag, you can use the "AWS Credential - Tag Filter" sample credential to enter EC2 tag keys and values.

NOTE: Filtering EC2 instance by tag will apply to all accounts discovered.

NOTE: Any EC2 instances that have already been discovered, but do not match the tag filter, will be set to "Unavailable."

To define an AWS credential to discover EC2 instances and filter them by tag:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Locate the **AWS Credential - Tag Filter** sample credential and click its wrench icon (🔧). The **Credential Editor** modal page appears.
3. Enter values in the following fields:

   Basic Settings

   - ***Profile Name***. Type a new name for your AWS credential.
   - ***HTTP Auth User***. Type your AWS access key ID.
   - ***HTTP Auth Password***. Type your AWS secret access key.

HTTP Headers

- Edit the HTTP header provided:
  - *Tags:<operation>#<EC2-Tag-Key>#<EC2-Tag-Value>*. Type the tag, followed by its operation, tag key, or tag value. For example, if you want to filter by Tag Name, you would type the following:

    ```
    Tags:equals#Name#Example
    ```

  Valid operations include:

  - equals
  - notEquals
  - contains
  - notContains

  You can chain together multiple filters separating them by a comma. For example:

  ```
  Tags:equals#Name#Example,contains#Owner#Someone
  ```

4. Click the **[Save As]** button, and then click **[OK]**.

Automatic Skylar One Organization Creation

This feature is only applicable to the two discovery methods that use the Assume Role and automatically discover multiple accounts.

When multiple accounts are discovered, this feature places each account in its own Skylar One organization. This feature requires an optional header in the SOAP/XML credential you will create. When this header is present, it will place each account into a new Skylar One organization. When this header is not present, each account will be placed in the Skylar One organization selected in the discovery session. The name of the organization can be controlled depending on what is provided in the header as follows:

- ***OrganizationCreation:NAME:ID***. Autocreates a Skylar One organization for accounts using AssumeRole. You can enter one of the following options:
  - ***OrganizationCreation:NAME***. The name of the organization will contain the name of the user.
  - ***OrganizationCreation:ID***. The name of the organization will contain the ID of the user.
  - ***OrganizationCreation:ID:NAME***. The name of the organization will contain both the ID and name of the user, in that order.
  - ***OrganizationCreation:NAME:ID***. The name of the organization will contain both the name and ID of the user, in that order.

Monitoring Consolidated Billing Accounts

Consolidated billing is an option provided by Amazon that allows multiple AWS accounts to be billed under a single account. For more information about consolidated billing, see http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html.

If a consolidated billing account is monitored by Skylar One, the billing metrics associated with that account include only the consolidated amounts, per service. If you use consolidated billing and want to collect billing metrics per-account, you must discover each account separately. To monitor only the billing metrics for an AWS account, you can create credentials that include only billing permissions.

ScienceLogic Events and AWS Alarms

In addition to Skylar One collecting metrics for AWS instances, you can configure CloudWatch to send alarm information to Skylar One via API. Skylar One can then generate an event for each alarm.

For instructions on how configure CloudWatch and Skylar One to generate events based on CloudWatch alarms, see the *Configuring Inbound CloudWatch Alarms* section.

Using a Proxy Server

You can use a proxy server with the *Manual Discovery* and the *Automated Discovery Using AssumeRole with a Single IAM Key from the AWS Master Account* discovery methods.

To use a proxy server in both cases, you must fill in the proxy settings in the SOAP/XML credential.

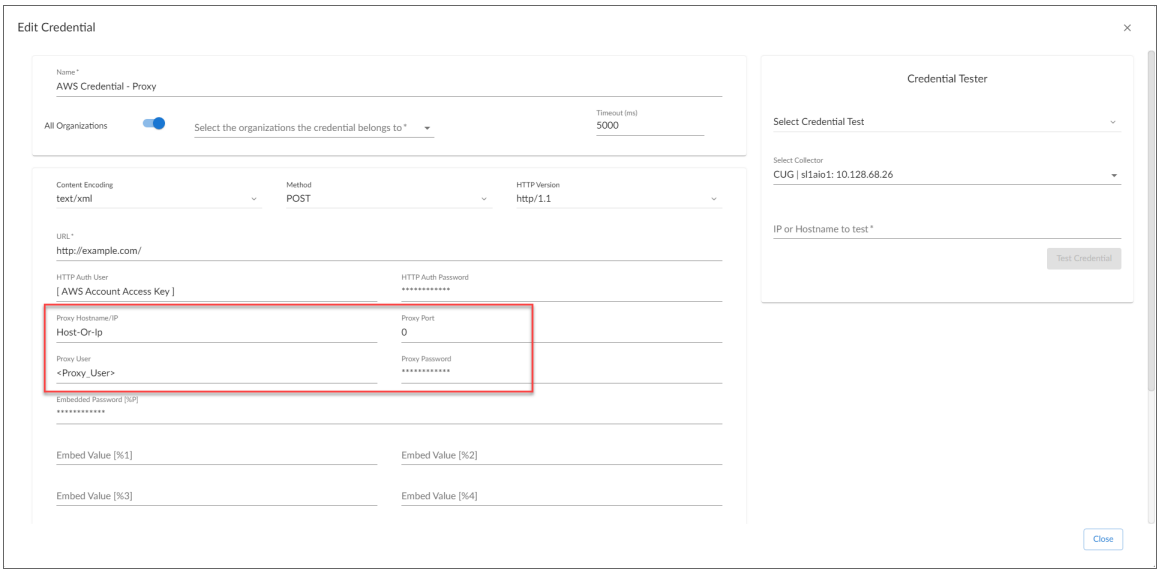

For the *Automated Discovery Using AssumeRole with a Single IAM Key from the AWS Master Account* discovery method, if the proxy does not support ping passthrough you will also need to follow the steps in the ***Automated Discovery Using AssumeRole with a Single IAM Key from the AWS Master Account*** section without ping support.

Configuring "AWS: Lambda Service Discovery"

By default, the "AWS: Lambda Service Discovery" Dynamic Application is configured to discover only regular Lambda functions, not replica functions. If you want to discover both regular and replica Lambda functions, then you must configure the "AWS: Lambda Service Discovery" Dynamic Application to do so ***prior*** to discovering your Lambda service.

To configure the "AWS: Lambda Service Discovery" Dynamic Application to discover both regular and replica Lambda functions:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Locate the "AWS: Lambda Service Discovery" Dynamic Application and click its wrench icon (). The **Dynamic Applications Properties Editor** page appears.
3. In the ***Operational State*** field, select *Disabled*, and then click **[Save]**. This disables the Dynamic Application from collecting data.
4. Click the **[Snippets]** tab. The **Dynamic Applications Snippet Editor & Registry** page appears.
5. In the **Snippet Registry** pane, click the wrench icon () for the "aws_lambda_service_discovery" snippet.
6. In the ***Active State*** field, select *Disabled*, and then click **[Save]**. This disables the "aws_lambda_service_discovery" snippet.
7. In the **Snippet Registry** pane, click the wrench icon () for the "aws_lambda_service_discovery_show_replicas" snippet.
8. In the ***Active State*** field, select *Enabled*, and then click **[Save]**. This enables the "aws_lambda_service_discovery_show_replicas" snippet.
9. Click the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page appears.
10. Click the wrench icon () for the first Collection Object listed in the **Collection Object Registry** pane, select *aws_lambda_service_discovery_show_replicas* in the ***Snippet*** field for that Collection Object, and then click **[Save]**.
11. Repeat step 10 for all of the remaining Collection Objects listed in the **Collection Object Registry** pane.
12. Click the **[Properties]** tab.
13. In the ***Operational State*** field, select *Enabled*, and then click **[Save]**. This re-enables data collection for the Dynamic Application.

NOTE: If you configure the "AWS: Lambda Service Discovery" Dynamic Application to discover both regular and replica Lambda functions, then when you run discovery, the Dynamic Applications in the "Amazon Web Services" PowerPack will create *parent/child relationships* between replica Lambda functions and their corresponding master Lambda functions. In this scenario, the *Device View and other device component maps* will display the relationship in this order: Lambda Function Service > Lambda Replica Function > Master Lambda Function. The replica appears as the parent to the master Lambda function because the replica could be in the same or a different region than the master Lambda function.

Configuring "AWS: Lambda Function Qualified Discovery"

By default, the "AWS: Lambda Function Qualified Discovery" Dynamic Application is configured to discover and model all Lambda alias components. An ***alias*** is a qualifier inside an AWS Lambda function that enables the user to control which versions of the Lambda function are executable–for instance, a production version and a test version.

When the "AWS: Lambda Function Qualified Discovery" Dynamic Application is configured to discover alias components, Skylar One collects data only for the Lambda function versions specified in the alias.

Depending on your needs, you can optionally configure the Dynamic Application to instead do one of the following:

- Discover and model all Lambda version components. If you select this configuration, Skylar One collects data for all existing versions of the Lambda function.
- Discover and model only Lambda version components with AWS configurations filtered by a trigger. If you select this configuration, Skylar One collects data only for versions of the Lambda function that have triggers or are specified in an alias.

NOTE: If you have *configured the "AWS: Lambda Service Discovery" Dynamic Application* to discover both regular and replica Lambda functions and you want Skylar One to *create dynamic component map relationships* between replica Lambda functions and their parent Lambda function versions, you must follow these instructions to configure the "AWS: Lambda Function Qualified Discovery" Dynamic Application to discover and model all Lambda version components.

To configure the "AWS: Lambda Function Qualified Discovery" Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Locate the "AWS: Lambda Function Qualified Discovery" Dynamic Application and click its wrench icon (🔧). The **Dynamic Applications Properties Editor** page appears.
3. In the ***Operational State*** field, select *Disabled*, and then click **[Save]**. This disables the Dynamic Application from collecting data.

4. Click the **[Snippets]** tab. The **Dynamic Applications Snippet Editor & Registry** page appears. The **Snippet Registry** pane includes the following snippets:
    - *aws_lambda_function_aliases_discovery*. When this snippet is enabled, the Dynamic Application discovers all Lambda alias components.
    - *aws_lambda_function_all_versions_discovery*. When this snippet is enabled, the Dynamic Application discovers all Lambda version components.
    - *aws_lambda_function_versions_by_triggers_discovery*. When this snippet is enabled, the Dynamic Application discovers Lambda version components with AWS configurations containing a trigger or those with an alias.
5. One at a time, click the wrench icon (🔧) for each of the snippets, select *Enabled* or *Disabled* in the ***Active State*** field, and then click **[Save]** to enable the appropriate snippet and disable the others.

NOTE: You can enable only one of these snippets at a time.

6. Click the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page appears.
7. Click the wrench icon (🔧) for the first Collection Object listed in the **Collection Object Registry** pane, select the snippet you enabled in step 5 in the ***Snippet*** field for that Collection Object, and then click **[Save]**.
8. Repeat step 7 for all of the remaining Collection Objects listed in the **Collection Object Registry** pane.
9. Click the **[Properties]** tab.
10. In the ***Operational State*** field, select *Enabled*, and then click **[Save]**. This re-enables data collection for the Dynamic Application. The next time discovery is run, new component devices might be discovered and some previously discovered components might become unavailable, depending on how you configured the Dynamic Application.

NOTE: If you configure the "AWS: Lambda Function Qualified Discovery" Dynamic Application to discover Lambda alias or version components and your AWS service includes an API Gateway that triggers a Lambda Function, then the Dynamic Applications in the "Amazon Web Services" PowerPack will create *a device relationship* between that Lambda Function and its corresponding Lambda alias or version component device.

Configuring AWS Integration with Docker

If you have discovered EC2-backed ECS clusters using the "Amazon Web Services" PowerPack, you can optionally use the "Docker" PowerPack to collect container information in addition to what the AWS API provides for the ECS service.

NOTE: This integration does not work with Fargate-backed ECS clusters.

To configure this integration, cURL version 7.40 or later must be installed on the ECS AMI image. For example, the 2018.03 ECS AMI image is compatible is compatible because it includes cURL 7.43.1.

Additionally, you must install the most recent version of the "Docker" PowerPack on your Skylar One system and run a discovery session using an SSH credential that will work on the EC2 host(s). This discovery session will discover the EC2 instances that comprise the ECS cluster and align the Docker host Dynamic Applications with those EC2 instances. Optionally, you can merge the EC2 host with the Docker host if you so choose.

NOTE: For more information about the "Docker" PowerPack, including instructions about creating the SSH credential and running discovery, see the ***Monitoring Docker*** manual.

NOTE: ScienceLogic does not recommend enabling and securing the Docker HTTP API when aligning EC2 instances with Docker hosts. Doing so requires you to complete manual steps on each EC2 host. Furthermore, if you use this method and then merge the EC2 host with the Docker host, data collection will fail for all containers that are children of the merged host.

Configuring AWS Integration with Kubernetes

If you are using the AWS EKS service you can optionally use the "Kubernetes" PowerPack to provide visibility into your Kubernetes worker nodes and their associated workloads.

To use the "Kubernetes" PowerPack with the "Amazon Web Services" PowerPack, you must have the following versions of these PowerPacks installed:

- "Amazon Web Services" version 118 or later
- "Kubernetes" version 104 or later

If you are using AWS EKS but do ***not*** want to use this feature, then it is recommended to disable the "AWS EKS Cluster Virtual Discovery" Dynamic Application. To do this:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Dynamic Applications).
2. Search for "AWS EKS" in the ***Dynamic Application Name*** column.
3. Click on the wrench icon (🔧) for the "AWS EKS Cluster Virtual Device Discovery" Dynamic Application and set the ***Operational State*** dropdown to *Disabled*.
4. Click the **[Save]** button.

Using the "Kubernetes" PowerPack is completely automated on Skylar One. If the proper credentials have been assigned on AWS and the AWS EKS Cluster, then Skylar One will automatically discover the Kubernetes worker nodes and the associated workloads. The following additional components will be automatically created:

1. A new DCM tree root device to represent the Kubernetes cluster. This will be a virtual device of the type "Kubernetes Cluster".
2. A child component of the cluster will be created for each worker node in the cluster. This will be a component device of the type "Kubernetes Node".

3. A child component of the cluster will be created that represents the Namespaces. This will be a component device of the type "Kubernetes Namespace Folder".
4. A child component of the Namespace Folder will be created for each Namespace discovered. This will be a component device of the type "Kubernetes Namespace".
5. A child component of the Namespace will be created for each controller discovered as follows:
    - Kubernetes Daemon Set
    - Kubernetes Deployment

    > **NOTE**: At most only a single component is created to represent a controller. If a deployment and replica set exists, Skylar One models only the deployment and replica set info as provided by the deployment component.

    - Kubernetes Job
    - Kubernetes Cronjob
    - Kubernetes Replication Controller
    - Kubernetes Replication Set
    - Kubernetes Stateful Set
6. A child component of the cluster will be created for each ingress defined. This will be a component device of the type "Kubernetes: Ingress".

For Skylar One to automatically discovery the EKS cluster, you must perform the following steps:

> **NOTE**: When logging into the Kubernetes cluster, ensure that the AWS credentials that `kubectl` is using are already authorized for your cluster. The IAM user that created the cluster has these permissions by default.

1. *Enable the Prometheus Metrics Server*. AWS EKS does not have the metrics server enabled by default. This is highly recommended as it will provide CPU and memory utilization metrics for both the worker nodes as well as the pods.

> **NOTE**: Skylar One automatically aggregates the CPU and memory utilization for pods and presents data at the controller level.

2. *Define the cluster role* needed by Skylar One so that it can access the necessary APIs. This is done on the EKS Cluster.
3. *Define the ClusterRoleBinding*. This is done on the EKS Cluster.
4. *Map the IAM user or role to the RBAC role and groups* using the aws-auth ConfigMap. This is done on the EKS Cluster.

Enabling the Prometheus Metrics Server

The Prometheus Metrics Server is required to provide CPU and memory utilization for pods and for nodes. The metrics server can be easily installed on Kubernetes clusters with the following:

```
kubectl apply -f https://github.com/kubernetes-sigs/metrics-server/releases/latest/download/components.yaml
```

To verify that the server is running, execute the command:

```
kubectl get deployment metrics-server -n kube-system
```

The following output will show that the metrics server is running:

```
            NAME             READY   UP-TO-DATE   AVAILABLE   AGE
metrics-server   1/1    1           14h
```

Define the Cluster Role

The cluster role defines the minimum permissions that Skylar One needs to monitor the Kubernetes cluster. ClusterRole is used as it provides access to all namespaces. Since Skylar One is directly monitoring the Kubernetes cluster via the Kubernetes API, this role's permissions need to be defined on the cluster itself.

To define the cluster role in Kubernetes:

1. Log in to the EKS cluster with the same user or role that created the cluster.
2. Create a new file called `SL1_cluster_role.yaml` and cut and paste the following text into that file:

CAUTION: YAML requires specific spacing. Please double-check the spacing after cutting-and-pasting code into YAML files.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
   name: eks-readonly-clusterrole
rules:
- apiGroups:
  - ""
  resources:
```

```
  - nodes
  - namespaces
  - pods
  - replicationscontroller
  - events
  - persistentvolumes
  - persistentvolumeclaims
  - componentstatuses
  - services
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - apps
  resources:
  - deployments
  - daemonsets
  - statefulsets
  - replicasets
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - batch
  resources:
  - jobs
  - cronjobs
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - metrics.k8s.io
  resources:
  - nodes
  - pods
  verbs:
  - get
```

```
  - list
  - watch
- apiGroups:
  - networking.k8s.io
  resources:
  - ingresses
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - autoscaling
  resources:
  - horizontalpodautoscalers
  verbs:
  - get
  - list
  - watch
```

The above file defines the minimum read-only permissions needed for Skylar One to monitor Kubernetes.

3. Once the file is defined, execute the following command to apply the file:

```
kubectl apply -f cluster_role.yaml
```

Define the ClusterRoleBinding

Once the role is defined, it must be bound to users, groups, or services. This is done by defining a ClusterRoleBinding:

1. Log in to the EKS cluster with the same user or role that created the cluster.
2. Create a new file called `SL1_ClusterRoleBinding.yaml` and cut and paste the following text into that file:

CAUTION: YAML requires specific spacing. Please double-check the spacing after cutting-and-pasting code into YAML files.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
```

```
metadata:
 name: eks-cluster-role-binding
subjects:
 - kind: User
   name: Sciencelogic-Monitor
   apiGroup: rbac.authorization.k8s.io
roleRef:
 kind: ClusterRole
 name: eks-readonly-clusterrole
 apiGroup: rbac.authorization.k8s.io
```

3. Once the file is created, apply the ClusterRoleBinding by executing the following command:

```
kubectl apply -f SL1_ClusterRoleBinding.yaml
```

NOTE: Under subjects, "name: Sciencelogic-Monitor" defines the Kubernetes user and it must match the username field in the config map shown below.

NOTE: Under roleRef, "name: eks-readonly-clusterrole" must match the name defined in the cluster role.

Map the IAM User or Role to the Kubernetes RBAC Role

After defining the ClusterRoleBinding, you must map the AWS credentials that Skylar One is using to the username created above in the `SL1_ClusterRoleBinding.yaml` file. To do this, perform the following steps:

1. Enter the `kubectl edit -n kube-system configmap/aws-auth` command. This will bring up the `configmap`. How the `configmap` is updated depends on what type of IAM was used to discover Skylar One.

NOTE: If the `configmap/aws-auth` does not exist, follow the procedures defined in https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html

Example 1

If Skylar One has discovered your AWS organization using assume role, add the following text to the `mapRoles:` section in the `configmap`:

CAUTION: YAML requires specific spacing. Please double-check the spacing after cutting-and-pasting code into YAML files.

```
- groups:
   - eks-cluster-role-binding
   rolearn:arn:aws:iam::<Account number that hosts the Kubernetes cluster-
>:role/Sciencelogic-Monitor
   username: Sciencelogic-Monitor
```

NOTE: If `mapRoles` does not exist, then you can add the `mapRoles` section to the `configmap`.

The text should appear in the `configmap` as the highlighted text below:

```
# Please edit the object below. Lines beginning with a '#' will be
ignored,
 # and an empty file will abort the edit. If an error occurs while saving,
this field will be
 # reopened with the relevant failures
 #
apiVersion: v1
data:
 mapRoles: |
  - groups:
    - system:bootstrappers
    - system:nodes
    rolearn: arn:aws-us-gov:iam::<account number>:role/eksctl-eks-cluster-
testfriday-nod-NodeInstanceRole-6VCMS669U9NA
    username: system:node:{{EC2PrivateDNSName}}
   - groups:
     - eks-cluster-role-binding
     rolearn: arn:aws:iam::<account number>:role/Sciencelogic-Monitor
     username: Sciencelogic-Monitor
kind: ConfigMap
metadata:
 creationTimestamp: "2021-07-30T20:43:55Z"
 name: aws-auth
```

```
 namespace: kube-system
 resourceVersion: "173718"
 selfLink: /api/v1/namespaces/kube-system/configmaps/aws-auth
 uid: d1bcdafd-fc40-44e6-96d4-9a079b407d06
```

Example 2

If Skylar One has been discovered with a single IAM key for the account, add the following text to the `mapUsers:` section of the `configmap`:

> **CAUTION:** YAML requires specific spacing. Please double-check the spacing after cutting-and-pasting code into YAML files.

```
- groups:
   - eks-cluster-role-binding
   userarn:arn:aws:iam::<Account number that hosts the Kubernetes cluster-
>:user/<Name of the user associated with the IAM key
   username: Sciencelogic-Monitor
```

The text should appear in the `configmap` as the highlighted text below:

```
 # Please edit the object below. Lines beginning with a '#' will be
ignored,
 # and an empty file will abort the edit. If an error occurs while saving,
this field will be
 # reopened with the relevant failures
 #
apiVersion: v1
data:
 mapRoles: |
  - groups:
    - system:bootstrappers
    - system:nodes
    rolearn: arn:aws-us-gov:iam::<account number>:role/eksctl-eks-cluster-
testfriday-nod-NodeInstanceRole-6VCMS669U9NA
    username: system:node:{{EC2PrivateDNSName}}
 mapUsers: |
    - groups:
```

```
    - eks-cluster-role-binding
    userarn: arn:aws:iam::<account number>:user/<username>
    username: Sciencelogic-Monitor
kind: ConfigMap
metadata:
 creationTimestamp: "2021-07-30T20:43:55Z"
 name: aws-auth
 namespace: kube-system
 resourceVersion: "173718"
 selfLink: /api/v1/namespaces/kube-system/configmaps/aws-auth
 uid: d1bcdafd-fc40-44e6-96d4-9a079b407d06
```

NOTE: In `userarn: arn:aws:iam::<account number>:user/<username>`, the username is the userarn that Skylar One is using to monitor the Kubernetes cluster.

NOTE: Under `mapUsers`, the `username:` is the name used in the ClusterRoleBinding.

Amazon API Throttling Events

By default, Skylar One will use the Collector Group aligned with the root AWS virtual device to retrieve data from AWS devices and services.

If Skylar One must collect data from a large set of AWS devices and services, Skylar One might generate Notify events with a message ending in the text "Retry #1-10 Sleeping: ... seconds". Skylar One generates these events when the Amazon API throttles collection in response to a large number of requests to the API. Even though Skylar One is generating Notify "Retry" events, Skylar One is still collecting data from AWS. This issue commonly occurs when a specific Amazon data center edge is close to capacity.

If Skylar One generates the Minor event "Collection missed on <device> on 5 minute poll", this indicates that Skylar One was unable to retrieve that specific datum from the Amazon cloud during the most recent five-minute polling cycle. If you frequently see the "Collection missed" event across your cloud, you must contact Amazon support to whitelist the IP address of your Data Collector. This will prevent further throttling from occurring.

Support for AWS GovCloud Regions

AWS GovCloud Regions can be discovered using all discovery methods as defined below:

- For an individual account using the *Manual Discovery* method, type the name of the AWS GovCloud region in the ***Embed Value %1*** field in the *SOAP/XML credential*.

- For those using one of the discovery methods with AssumeRole, enter one of the following URLs in the ***URL*** field of the *SOAP/XML credential* to specify the specific government region:
    - https://organizations.us-gov-west-1.amazonaws.com
    - https://organizations.us-gov-east-1.amazonaws.com

NOTE: All examples shown are for commercial AWS accounts. When AWS Gov is being monitored, the JSON data that refers to ARN will need to be modified from "aws" to "aws-us-gov". For example: `Resource": "arn:aws:iam::<account number>:role/Sciencelogic-Monitor` would need to be `Resource": "arn:aws:iam-us-gov::<account number>:role/Sciencelogic-Monitor`

Migrating from Using an IAM Key Per Account to Using AssumeRole

Skylar One supports the ability to migrate accounts that were originally discovered using an IAM key per account to start using AssumeRole. To upgrade, perform the following steps:

1. Run steps 1-4 in the *Automated Discovery when the Data Collector Runs as an EC2 Instance* section

   or

   Run steps 1-3 of the *Automated Discovery Using AssumeRole with a Single IAM Key from the AWS Master Account* section.
2. Next you must disable collection for every account that is being migrated. To do this, go to the **Device Components** page (Devices > Device Components, or Registry > Devices > Device Components in the Skylar One classic user interface) and enter "AWS | Service" in the ***Device Class | Sub-class*** column.
3. Select the checkbox for each account being migrated.
4. In the ***Select Action*** menu, select *_Disabled* under **Change Collection State**. Click **[Go]**.

WARNING: Failing to disable collection for accounts that will be migrated will result in a loss of data.

If you have a large number of accounts that will be migrated, it is recommended to start with a single account to ensure that all settings migrate correctly. To limit the accounts that are migrated, put only a single account in the ec2-collector policy so Skylar One will assume the role for that single account. Once you have ensured that the one account has been migrated successfully, you can add the other accounts back into the ec2-collector policy.

NOTE: If you had previously changed the account name in Skylar One, the upgrade process will overwrite any changes made to the name of the account component.

NOTE: If you had previously placed your accounts into different Skylar One organizations, these organization names will be preserved upon upgrading. However, if you add the headers for Skylar One organizations to your SOAP/XML credential, then the Skylar One organizations will be set according to the SOAP header.

Once discovery has completed successfully, the Dynamic Applications aligned to the root device and the component device will display the new credential used in the discovery process, while the Dynamic Applications for the child device(s) of the account device will still display the old credential. While you cannot delete the credentials, you can remove the IAM keys from those credentials as they are no longer being used by the Dynamic Applications.

Chapter 4

Minimum Permissions

Overview

The following sections describe the minimum permissions that must be set before you can run discovery with the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

Minimum Permissions Needed to Monitor Your AWS Accounts 37

Minimum Permissions Needed to Monitor Your AWS Accounts

The following table displays the minimum permissions required for Dynamic Applications in the "Amazon Web Services" PowerPack to collect data.

Service	Actions	
Application Certificate Manager - Private Certificate Authority	List	ListCertificateAuthorities
Airflow	List	ListEnvironments
Amplify	List	ListApps
API Gateway	Read	GET

Service	Actions	
AppFlow	List	ListFlows
AppRunner	List	ListServices
AppStream	Read	DescribeFleets
AppSync	List	ListGraphqlApis
Athena	List	ListDatabases ListDataCatalogs ListWorkGroups
Backup	List	ListBackupPlans
Cassandra	Read	Select
Chime	List	ListAccounts ListVoiceConnectors
CloudFront	List	ListDistributions ListInvalidations ListStreamingDistributions
	Read	GetDistribution GetStreamingDistribution
CloudHSM	Read	DescribeClusters
CloudSearch	List	DescribeDomains
CloudTrail	List	DescribeTrails
	Read	GetTrailStatus
CloudWatch	List	ListMetrics ListMetricStreams
	Read	DescribeAlarmHistory DescribeAlarms GetMetricData GetMetricStatistics
CloudWatch Logs	List	DescribeLogGroups
CloudWatch RUM	List	ListAppMonitors
CloudWatch Synthetics	Read	DescribeCanaries
CodeBuild	List	ListProjects
CodeGuru Profiler	List	ListProfilingGroups
CodeGuru Reviewer	List	ListRepositoryAssociations
Cognito Identity	List	ListIdentityPools
Config	Read	GetDiscoveredResourceCounts

Service	Actions	
Connect	Read	ListInstances
Data Exchange	List	ListDataSets
Data Lifecycle Manager	List	GetLifecyclePolicies
Data Pipeline	List	ListPipelines
Database Migration Service	Read	DescribeReplicationInstances
DataSync	List	ListAgents
DAX	List	DescribeClusters
Direct Connect	Read	DescribeConnections DescribeTags DescribeVirtualInterfaces
DMS	Read	DescribeReplicationInstances
DynamoDB	List	ListTables
	Read	DescribeTable
EC2	List	DescribeAvailabilityZones DescribeClientVpnEndpoints DescribeFleets DescribeElasticGpus DescribeImages DescribeInstances DescribeNatGateways DescribeRegions DescribeRouteTables DescribeSecurityGroups DescribeSnapshots DescribeSpotFleetRequests DescribeSubnets DescribeTransitGatewayRouteTables DescribeTransitGateways DescribeTransitGatewayAttachments DescribeVolumes DescribeVpcEndpointConnections DescribeVpcPeeringConnections DescribeVpcs DescribeVpnGateways

Service	Actions	
	Read	DescribeVpnConnections
EC2 Auto Scaling	List	DescribeAutoScalingGroups DescribeAutoScalingInstances DescribeLaunchConfigurations
EFS	List	DescribeFileSystems
Elastic Beanstalk	List	DescribeEnvironments
	Read	DescribeConfigurationSettings DescribeEnvironmentResources DescribeEnvironmentHealth DescribeInstancesHealth
Elastic Container Services (ECS)	List	ListClusters ListContainerInstances ListServices ListTagsForResource ListTasks
	Read	DescribeClusters DescribeContainerInstances DescribeServices DescribeTaskDefinition DescribeTasks
ElasticCache	List	DescribeCacheClusters ListTagsForResource
Elastic Inference	List	DescribeAccelerators
Elastic Kubernetes Service (EKS)	List	ListClusters
	Read	DescribeCluster
Elastic Transcoder	List	ListPipelines
ELB	List	DescribeLoadBalancers
	Read	DescribeTags
ELB v2	Read	DescribeListeners DescribeLoadBalancers DescribeTags DescribeTargetGroups DescribeTargetHealth

Service	Actions	
EMR	List	ListClusters
	Read	ListInstances
ES	List	ListDomainNames
EventBridge	List	ListRules
FinSpace	List	ListEnvironments
Firehose	List	ListDeliveryStreams
Forecast	List	ListDatasetGroups
Fraud Detector	List	GetDetectors
FSx	Read	DescribeFileSystems
Gamelift	List	ListFleets
Glacier	List	ListTagsForVault ListVaults
	Read	GetVaultNotifications
Glue	List	ListJobs
Glue Databrew	List	ListProjects
Ground Station	List	ListSatellites
GuardDuty	List	ListDetectors
HealthLake	List	ListFHIRDatastores
IAM	Read	GetUser GetAccountAuthorizationDetails
Inspector2	List	ListCoverage
Interactive Video Service	List	ListChannels
IoT	List	ListScheduledAudits ListSecurityProfiles ListThings ListTagsForResource
	Read	DescribeThing
IoT Analytics	List	ListDatastores
IoT Events	List	ListDetectorModels
IoT GreenGrass V2	List	ListCoreDevices
IoT Sitewise	List	ListAssetModels
IoT Twinmaker	List	ListWorkspaces
Kafka	List	ListClusters
Kendra	List	ListIndices
Key Management Service	List	ListKeys

Service	Actions	
(KMS)		ListAliases
	Read	DescribeKey
		ListResourceTags
Kinesis	List	ListStreams
Kinesis Firehose	List	ListDeliveryStreams
Kinesis Video	List	ListStreams
Lambda	List	ListFunctions
		ListAliases
		ListEventSourceMappings
	Read	GetAccountSettings
		GetPolicy
		ListTags
Lex V2		ListBots
Lightsail	List	GetBundles
		GetRegions
	Read	GetInstanceMetricData
		GetInstances
Location	List	ListMaps
Lookout for Equipment	List	ListDatasets
Lookout for Metrics	List	ListAnomalyDetectors
Lookout for Vision	List	ListProjects
MediaConnect	List	ListFlows
MediaConvert	List	DescribeEndpoints
		ListQueues
MediaLive	List	ListChannels
MediaPackage - VOD	List	ListPackagingConfigurations
		ListPackagingGroups
MediaPackage	Read	ListChannels
MediaTailor	List	ListPlaybackConfigurations
MQ	List	ListBrokers
Network Firewall	List	ListFirewalls
Network Manager	List	DescribeGlobalNetworks
Nimble Studio	Read	ListStudios
OpsWorks	List	DescribeInstances

Service	Actions	
		DescribeStacks
Personalize	List	ListDatasets
Polly	List	ListLexicons
QLDB	List	ListLedgers
QuickSight	List	ListDashboards
RDS	List	DescribeDBClusters DescribeDBInstances DescribeDBSubnetGroups
	Read	ListTagsForResource
Redshift	List	DescribeClusters
	Read	DescribeLoggingStatus
RoboMaker	List	ListSimulationJobs
Route 53	List	GetHostedZone ListHealthChecks ListHostedZones ListResourceRecordSets
S3	List	ListAllMyBuckets ListBucket
	Read	GetBucketLocation GetBucketLogging GetBucketTagging GetBucketWebsite GetObject (Restrict access to specific resources of Elastic Beanstalk. For instance, Bucket name: elasticbeanstalk-*, Any Object name.)
SageMaker	List	ListDomains
Secrets Manager	List	ListSecrets
Service Catalog	List	SearchProductsAsAdmin
Shield	List	ListAttacks ListProtections
	Read	DescribeEmergencyContactSettings GetSubscriptionState
Simple Email Service (SES)	List	ListIdentities
Simple Notification Service	List	ListTopics

Service	Actions	
(SNS)		ListSubscriptions
SQS	List	ListQueues
	Read	GetQueueAttributes
SSM	Read	GetParameters
Storage Gateway	List	ListGateways
		ListVolumes
States	List	ListStateMachines
STS	Read	GetCallerIdentity
SWF	List	ListDomains
Tag Editor	Read	GetResources
		GetTagKeys
		GetTagValues
TimeStream	List	ListDatabases
Transfer	List	ListServers
WAF	List	ListWebACLs
	Read	GetRateBasedRule
		GetRule
		GetRuleGroup
		GetWebACL
WAF Regional	List	ListResourcesForWebACL
		ListWebACLs
	Read	GetRateBasedRule
		GetRule
		GetRuleGroup
		GetWebACL
WorkMail	List	ListOrganizations
WorkSpaces	List	DescribeWorkspaces
		DescribeWorkspaceDirectories

To create the Minimum Permission policy:

NOTE: If you are monitoring a GovCloud account, see the section on creating the minimum permission policy for those accounts *here*.

1. Go to the AWS console and select **IAM > Policies > Create Policy**. Select **JSON** and cut and paste the following JSON document:

NOTE: If you receive an error message that the policy exceeds the character limit, split the following JSON into two policies.

```
{
        "Version": "2012-10-17",
        "Statement": [
        {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": [
        "acm-pca:ListCertificateAuthorities",
        "airflow:ListEnvironments",
        "amplify:ListApps",
        "apigateway:GET",
        "appflow:ListFlows",
        "apprunner:ListServices",
        "appstream:DescribeFleets",
        "appsync:ListGraphqlApis",
        "athena:ListDatabases",
        "athena:ListDataCatalogs",
        "athena:ListWorkGroups",
        "autoscaling:DescribeAutoScalingGroups",
        "autoscaling:DescribeAutoScalingInstances",
        "autoscaling:DescribeLaunchConfigurations",
        "backup:ListBackupPlans",
        "cassandra:Select",
        "chime:ListAccounts",
        "chime:ListVoiceConnectors",
        "cloudfront:GetDistribution",
        "cloudfront:GetStreamingDistribution",
        "cloudfront:ListDistributions",
        "cloudfront:ListInvalidations",
        "cloudfront:ListStreamingDistributions",
        "cloudhsm:DescribeClusters",
        "cloudsearch:DescribeDomains",
        "cloudtrail:DescribeTrails",
        "cloudtrail:GetTrailStatus",
```

```
        "cloudwatch:DescribeAlarmHistory",
        "cloudwatch:DescribeAlarms",
        "cloudwatch:GetMetricData",
        "cloudwatch:GetMetricStatistics",
        "cloudwatch:ListMetrics",
        "cloudwatch:ListMetricStreams",
        "codebuild:ListProjects",
        "codeguru-profiler:ListProfilingGroups",
        "codeguru-reviewer:ListRepositoryAssociations",
        "cognito-identity:ListIdentityPools",
        "config:GetDiscoveredResourceCounts",
        "connect:ListInstances",
        "databrew:ListProjects",
        "dataexchange:ListDataSets",
        "datasync:ListAgents",
        "datapipeline:ListPipelines",
        "dax:DescribeClusters",
        "devops-guru:ListMonitoredResources",
        "directconnect:DescribeConnections",
        "directconnect:DescribeTags",
        "directconnect:DescribeVirtualInterfaces",
        "dlm:GetLifecyclePolicies",
        "dms:DescribeReplicationInstances",
        "dynamodb:DescribeTable",
        "dynamodb:ListTables",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeClientVpnEndpoints",
        "ec2:DescribeFleets",
        "ec2:DescribeElasticGpus",
        "ec2:DescribeImages",
        "ec2:DescribeInstances",
        "ec2:DescribeNatGateways",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSpotFleetRequests",
        "ec2:DescribeSubnets",
        "ec2:DescribeTransitGatewayAttachments",
        "ec2:DescribeTransitGatewayRouteTables",
```

```
        "ec2:DescribeTransitGateways",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcPeeringConnections",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpcEndpointConnections",
        "ec2:DescribeVpnConnections",
        "ec2:DescribeVpnGateways",
        "ecs:DescribeClusters",
        "ecs:DescribeContainerInstances",
        "ecs:DescribeServices",
        "ecs:DescribeTaskDefinition",
        "ecs:DescribeTasks",
        "ecs:ListClusters",
        "ecs:ListContainerInstances",
        "ecs:ListServices",
        "ecs:ListTagsForResource",
        "ecs:ListTasks",
        "eks:DescribeCluster",
        "eks:ListClusters",
        "elasticache:DescribeCacheClusters",
        "elasticache:ListTagsForResource",
        "elasticbeanstalk:DescribeConfigurationSettings",
        "elasticbeanstalk:DescribeEnvironmentHealth",
        "elasticbeanstalk:DescribeEnvironmentResources",
        "elasticbeanstalk:DescribeEnvironments",
        "elasticbeanstalk:DescribeInstancesHealth",
        "elasticfilesystem:DescribeFileSystems",
        "elastic-inference:DescribeAccelerators",
        "elasticloadbalancing:DescribeListeners",
        "elasticloadbalancing:DescribeLoadBalancers",
        "elasticloadbalancing:DescribeTags",
        "elasticloadbalancing:DescribeTargetGroups",
        "elasticloadbalancing:DescribeTargetHealth",
        "elasticmapreduce:ListClusters",
        "elasticmapreduce:ListInstances",
        "elastictranscoder:ListPipelines",
        "es:ListDomainNames",
        "events:ListRules",
        "finspace:ListEnvironments",
        "firehose:ListDeliveryStreams",
```

```
        "forecast:ListDatasetGroups",
        "frauddetector:GetDetectors",
        "fsx:DescribeFileSystems",
        "gamelift:ListFleets",
        "geo:ListMaps",
        "glacier:GetVaultNotifications",
        "glacier:ListTagsForVault",
        "glacier:ListVaults",
        "glue:ListJobs",
        "greengrass:ListCoreDevices",
        "groundstation:ListSatellites",
        "guardduty:ListDetectors",
        "healthlake:ListFHIRDatastores",
        "iam:GetAccountAuthorizationDetails",
        "iam:GetUser",
        "inspector2:ListUsageTotals",
        "iot:DescribeThing",
        "iot:ListScheduledAudits",
        "iot:ListSecurityProfiles",
        "iot:ListTagsForResource",
        "iot:ListThings",
        "iotanalytics:ListDatastores",
        "iotevents:ListDetectorModels",
        "iotsitewise:ListAssetModels",
        "iottwinmaker:ListWorkspaces",
        "ivs:ListChannels",
        "ivschat:ListRooms",
        "kafka:ListClusters",
        "kendra:ListIndices",
        "kinesis:ListStreams",
        "kinesisvideo:ListStreams",
        "kms:DescribeKey",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ListResourceTags",
        "lakeformation:ListResources",
        "lambda:GetAccountSettings",
        "lambda:GetPolicy",
        "lambda:ListAliases",
        "lambda:ListEventSourceMappings",
```

```
        "lambda:ListFunctions",
        "lambda:ListTags",
        "lex:ListBots",
        "lightsail:GetBundles",
        "lightsail:GetInstanceMetricData",
        "lightsail:GetInstances",
        "lightsail:GetRegions",
        "logs:DescribeLogGroups",
        "lookoutequipment:ListDatasets",
        "lookoutmetrics:ListAnomalyDetectors",
        "lookoutvision:ListProjects",
        "mediaconnect:ListFlows",
        "mediaconvert:DescribeEndpoints",
        "mediaconvert:ListQueues",
        "medialive:ListChannels",
        "mediapackage-vod:ListPackagingConfigurations",
        "mediapackage-vod:ListPackagingGroups",
        "mediapackage:ListChannels",
        "mediatailor:ListPlaybackConfigurations",
        "mq:ListBrokers",
        "network-firewall:ListFirewalls",
        "nimble:ListStudios",
        "opsworks:DescribeInstances",
        "opsworks:DescribeStacks",
        "panorama:ListDevices",
        "personalize:ListDatasets",
        "polly:ListLexicons",
        "qldb:ListLedgers",
        "rds:DescribeDBClusters",
        "rds:DescribeDBInstances",
        "rds:DescribeDBSubnetGroups",
        "rds:ListTagsForResource",
        "redshift:DescribeClusters",
        "redshift:DescribeLoggingStatus",
        "rekognition:DescribeProjects",
        "robomaker:ListSimulationJobs",
        "route53:GetHostedZone",
        "route53:ListHealthChecks",
        "route53:ListHostedZones",
        "route53:ListResourceRecordSets",
```

```
        "rum:ListAppMonitors",
        "s3:GetBucketLocation",
        "s3:GetBucketLogging",
        "s3:GetBucketTagging",
        "s3:GetBucketWebsite",
        "s3:GetObject",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "sagemaker:ListDomains",
        "states:ListStateMachines",
        "secretsmanager:ListSecrets",
        "ses:ListIdentities",
        "servicecatalog:SearchProductsAsAdmin",
        "shield:DescribeEmergencyContactSettings",
        "shield:GetSubscriptionState",
        "shield:ListAttacks",
        "shield:ListProtections",
        "sns:ListSubscriptions",
        "sns:ListTopics",
        "sqs:GetQueueAttributes",
        "sqs:ListQueues",
        "ssm:GetInventory",
        "ssm:GetParameters",
        "storagegateway:ListGateways",
        "storagegateway:ListVolumes",
        "sts:GetCallerIdentity",
        "swf:ListDomains",
        "synthetics:DescribeCanaries",
        "tag:GetResources",
        "tag:GetTagKeys",
        "tag:GetTagValues",
        "timestream:ListDatabases",
        "transcribe:ListTranscriptionJobs",
        "transfer:ListServers",
        "translate:ListTextTranslationJobs",
        "waf-regional:GetRateBasedRule",
        "waf-regional:GetRule",
        "waf-regional:GetRuleGroup",
        "waf-regional:GetWebACL",
        "waf-regional:ListResourcesForWebACL",
```

```
        "waf-regional:ListWebACLs",
        "waf:GetRateBasedRule",
        "waf:GetRule",
        "waf:GetRuleGroup",
        "waf:GetWebACL",
        "waf:ListWebACLs",
        "workmail:ListOrganizations",
        "workspaces-web:ListPortals",
        "workspaces:DescribeWorkspaceDirectories",
        "workspaces:DescribeWorkspaces"
        ],
        "Resource": "*"
        }
        ]
        }
```

2. Click **[Next: Tags]**. If applicable, enter your Tags.
3. Click **[Next: Review]**. Name the policy "Skylar OneMinimumPermissions" and click **[Create Policy]**.

This policy needs to be available in each account that is to be monitored and will be referenced in the following sections.

Creating a Minimum Permissions Policy for GovCloud Accounts

If you are on an AWS GovCloud account, perform the following steps to create the Minimum Permission policy:

1. Go to the AWS console and select **IAM > Policies > Create Policy**. Select **JSON** and cut and paste the following JSON document:

> **NOTE:** If you receive an error message that the policy exceeds the character limit, split the following JSON into two policies.

```
{
        "Version": "2012-10-17",
        "Statement": [
        {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": [
        "acm-pca:ListCertificateAuthorities",
```

```
        "apigateway:GET",
        "appstream:DescribeFleets",
        "athena:ListWorkGroups",
        "athena:ListDatabases",
        "athena:ListDataCatalogs",
        "autoscaling:DescribeAutoScalingGroups",
        "autoscaling:DescribeAutoScalingInstances",
        "autoscaling:DescribeLaunchConfigurations",
        "backup:ListBackupPlans",
        "cloudhsm:DescribeClusters",
        "cloudtrail:DescribeTrails",
        "cloudtrail:GetTrailStatus",
        "cloudwatch:DescribeAlarmHistory",
        "cloudwatch:DescribeAlarms",
        "cloudwatch:GetMetricData",
        "cloudwatch:GetMetricStatistics",
        "cloudwatch:ListMetrics",
        "cloudwatch:ListMetricStreams",
        "codebuild:ListProjects",
        "cognito-identity:ListIdentityPools",
        "config:GetDiscoveredResourceCounts",
        "connect:ListInstances",
        "databrew:ListProjects",
        "datasync:ListAgents",
        "directconnect:DescribeConnections",
        "directconnect:DescribeTags",
        "directconnect:DescribeVirtualInterfaces",
        "dms:DescribeReplicationInstances",
        "dynamodb:DescribeTable",
        "dynamodb:ListTables",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeClientVpnEndpoints",
        "ec2:DescribeFleets",
        "ec2:DescribeImages",
        "ec2:DescribeInstances",
        "ec2:DescribeNatGateways",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
```

```
        "ec2:DescribeSpotFleetRequests",
        "ec2:DescribeSubnets",
        "ec2:DescribeTransitGatewayAttachments",
        "ec2:DescribeTransitGatewayRouteTables",
        "ec2:DescribeTransitGateways",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcEndpointConnections",
        "ec2:DescribeVpcPeeringConnections",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpnConnections",
        "ec2:DescribeVpnGateways",
        "ecs:DescribeClusters",
        "ecs:DescribeContainerInstances",
        "ecs:DescribeServices",
        "ecs:DescribeTaskDefinition",
        "ecs:DescribeTasks",
        "ecs:ListClusters",
        "ecs:ListContainerInstances",
        "ecs:ListServices",
        "ecs:ListTasks",
        "ecs:ListTagsForResource",
        "eks:DescribeCluster",
        "eks:ListClusters",
        "elasticache:DescribeCacheClusters",
        "elasticbeanstalk:DescribeConfigurationSettings",
        "elasticbeanstalk:DescribeEnvironmentHealth",
        "elasticbeanstalk:DescribeEnvironmentResources",
        "elasticbeanstalk:DescribeEnvironments",
        "elasticbeanstalk:DescribeInstancesHealth",
        "elasticfilesystem:DescribeFileSystems",
        "elasticloadbalancing:DescribeListeners",
        "elasticloadbalancing:DescribeLoadBalancers",
        "elasticloadbalancing:DescribeTags",
        "elasticloadbalancing:DescribeTargetGroups",
        "elasticloadbalancing:DescribeTargetHealth",
        "elasticmapreduce:ListClusters",
        "elasticmapreduce:ListInstances",
        "es:ListDomainNames",
        "events:ListRules",
        "firehose:ListDeliveryStreams",
```

```
        "fsx:DescribeFileSystems",
        "glacier:GetVaultNotifications",
        "glacier:ListTagsForVault",
        "glacier:ListVaults",
        "glue:ListJobs",
        "greengrass:ListCoreDevices",
        "guardduty:ListDetectors",
        "iam:GetAccountAuthorizationDetails",
        "iam:GetUser",
        "inspector2:ListUsageTotals",
        "iot:DescribeThing",
        "iot:ListScheduledAudits",
        "iot:ListSecurityProfiles",
        "iot:ListTagsForResource",
        "iot:ListThings",
        "iotevents:ListDetectorModels",
        "iotsitewise:ListAssetModels",
        "kafka:ListClusters",
        "kendra:ListIndices",
        "kinesis:ListStreams",
        "kms:DescribeKey",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ListResourceTags",
        "lakeformation:ListResources",
        "lambda:GetAccountSettings",
        "lambda:GetPolicy",
        "lambda:ListAliases",
        "lambda:ListEventSourceMappings",
        "lambda:ListFunctions",
        "lambda:ListTags",
        "logs:DescribeLogGroups",
        "mq:ListBrokers",
        "network-firewall:ListFirewalls",
        "polly:ListLexicons",
        "rds:DescribeDBClusters",
        "rds:DescribeDBInstances",
        "rds:DescribeDBSubnetGroups",
        "rds:ListTagsForResource",
        "redshift:DescribeClusters",
```

```
        "redshift:DescribeLoggingStatus",
        "rekognition:DescribeProjects",
        "route53:GetHostedZone",
        "route53:ListHealthChecks",
        "route53:ListHostedZones",
        "route53:ListResourceRecordSets",
        "s3:GetBucketLocation",
        "s3:GetBucketLogging",
        "s3:GetBucketTagging",
        "s3:GetBucketWebsite",
        "s3:GetObject",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "sagemaker:ListDomains",
        "states:ListStateMachines",
        "secretsmanager:ListSecrets",
        "servicecatalog:SearchProductsAsAdmin",
        "ses:ListIdentities",
        "sns:ListSubscriptions",
        "sns:ListTopics",
        "sqs:GetQueueAttributes",
        "sqs:ListQueues",
        "ssm:GetInventory",
        "ssm:GetParameters",
        "storagegateway:ListGateways",
        "storagegateway:ListVolumes",
        "sts:GetCallerIdentity",
        "swf:ListDomains",
        "synthetics:DescribeCanaries",
        "tag:GetResources",
        "tag:GetTagKeys",
        "tag:GetTagValues",
        "transcribe:ListTranscriptionJobs",
        "transfer:ListServers",
        "translate:ListTextTranslationJobs",
        "workspaces:DescribeWorkspaceDirectories",
        "workspaces:DescribeWorkspaces"
        ],
        "Resource": "*"
        }
```

```
        ]
        }
```

2. Click **[Next: Tags]**. If applicable, enter your Tags.
3. Click **[Next: Review]**. Name the policy "Skylar OneMinimumPermissions" and click **[Create Policy]**.

This policy needs to be available in each account that is to be monitored and will be referenced in the following sections.

Chapter 5

AWS Discovery

Overview

The following sections describe the different methods of discovery that can be used with the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

Discovering Amazon Web Services 57

Preparing your AWS Organization or AWS Account for Discovery 58

Discovering your AWS Organization or AWS Account 83

Discovering your AWS Organization or AWS Account in the Classic Skylar One User Interface 90

Appendix 98

Discovering Amazon Web Services

Skylar One currently supports the following methods to discover your AWS accounts:

- ***Manual Discovery***. Requires the creation of a virtual device, manual alignment of Dynamic Applications, and an IAM key. This process needs to be repeated for each AWS account.
- ***Automated Discovery using Assume Role with single IAM key***. Provides an automated mechanism to discover all your AWS accounts within an organization using a single IAM key. This is the recommended method of discovery when your Data Collector is not an EC2 instance.

- *Automated Discovery when the Data Collector runs as an EC2 instance*. Provides a fully automated mechanism to discover all your AWS accounts when your Data Collectors are running as EC2 instances. Skylar One does not need any AWS credentials in this case. This is the recommended approach when your Data Collectors are EC2 instances.
- *AWS Guided Discovery*. Uses guided workflows in Skylar One. This method is recommended when you want to use a separate IAM key for each AWS account. The guided worfklows provide a more user-friendly version of the manual process. Choose from the following workflows:
  - *AWS EC2*
  - *AWS IAM*
  - *AWS Assume Role*

NOTE: These Guided Discovery Workflows are available in Skylar One version 11.2.0 and later. A basic Guided Discovery Workflow is available in earlier versions of Skylar One.

Before determining your method of discovery, it is recommended to define the minimum permissions policy in AWS. This policy defines the minimum permissions needed to monitor all AWS services and is needed regardless of which of the above methods is used.

You can discover a maximum of 10 accounts with the following requirements on the Data Collector:

- 8 cores
- 32 GB of RAM
- 100 GB of HDD

Preparing your AWS Organization or AWS Account for Discovery

You can discover your AWS organization or individually configured AWS accounts using the "Amazon Web Services" PowerPack. There are two options available for configuration of AWS Organization Discovery :

- IAM User
- EC2 "Credential-less" for Skylar One platforms hosted within Amazon Web Services.

Individual accounts can only be configured for discovery via the IAM User option.

Organization Account

The "Amazon Web Services" PowerPack supports the discovery of individually configured organizations. Each variation leverages the AssumeRole capability to provide monitoring insight to individual accounts. To complete this configuration, you will need to configure on IAM User and corresponding IAM Role within each account you want to monitor. For Skylar One systems hosted in Amazon Web Services/EC2, you do not need an IAM User. For more information about AWS/EC2 configuration, see the *EC2 Credential-less* section.

IAM User

Before you configure an IAM User for Organization Discovery, you must assign the following "Delegated Administrator" Policy to your organization.

1. Replace the <Account ID> with an account ID from withing your AWS organization.

IMPORTANT: ScienceLogic recommends that you do not use your Management Account ID for this step.

```
{
	"Version": "2012-10-17",
	"Statement": [
		{
			"Sid": "Statement",
			"Effect": "Allow",
			"Principal": {
				"AWS": "<Account_ID>"
			},
			"Action": [
				"organizations:ListAccounts",
				"organizations:DescribeAccount",
				"organizations:DescribeOrganization"
			],
			"Resource": "*"
		}
	]
}
```

2. After you have replaced your <Account ID>, you can create the IAM User (e.g. Skylar One-Service-User) within the account you just delegated the above actions.

NOTE: ScienceLogic does not recommend that this IAM User have access to the AWS Management Console.

3. When prompted to add the permissions, select "Attach policies directly".Each account within the organization that you want to monitor should be represented and ScienceLogic does not recommend using Wildcards or non-explicit configurations. Follow the prompts to create a policy with JSON similar to the example below.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "organizations:ListAccounts",
                "organizations:DescribeOrganization",
                "organizations:DescribeAccount"
            ],
            "Resource": "*"
        },
        {
            "Sid": "VisualEditor1",
            "Effect": "Allow",
            "Action": "sts:AssumeRole",
            "Resource": [
                "arn:aws:iam::<Account ID 1>:role/Sciencelogic-Monitor",
                "arn:aws:iam::<Account ID 2>:role/Sciencelogic-Monitor",
                "arn:aws:iam::<Account ID 3>:role/Sciencelogic-Monitor"
            ]
        }
    ]
}
```

4. When you are finished navigating the IAM User creation wizard, you should see your newly created user appear on the IAM User list. Select your user and then select "Security Credentials" and then "Access Key" to customize the role for the new IAM User.
5. Create a role called ScienceLogic-Monitor, which is case sensitive, for each account that you specified in Step 1.
6. Assign the following policy to each role:

```
{
       "Version": "2012-10-17",
       "Statement": [
               {
                       "Sid": "VisualEditor0",
                       "Effect": "Allow",
                       "Action": [
                               "acm-pca:ListCertificateAuthorities",
```

```
                "airflow:ListEnvironments",
                "amplify:ListApps",
                "apigateway:GET",
                "appflow:ListFlows",
                "apprunner:ListServices",
                "appstream:DescribeFleets",
                "appsync:ListGraphqlApis",
                "athena:ListDatabases",
                "athena:ListDataCatalogs",
                "athena:ListWorkGroups",
                "autoscaling:DescribeAutoScalingGroups",
                "autoscaling:DescribeAutoScalingInstances",
                "autoscaling:DescribeLaunchConfigurations",
                "backup:ListBackupPlans",
                "cassandra:Select",
                "chime:ListAccounts",
                "chime:ListVoiceConnectors",
                "cloudfront:GetDistribution",
                "cloudfront:GetStreamingDistribution",
                "cloudfront:ListDistributions",
                "cloudfront:ListInvalidations",
                "cloudfront:ListStreamingDistributions",
                "cloudhsm:DescribeClusters",
                "cloudsearch:DescribeDomains",
                "cloudtrail:DescribeTrails",
                "cloudtrail:GetTrailStatus",
                "cloudwatch:DescribeAlarmHistory",
                "cloudwatch:DescribeAlarms",
                "cloudwatch:GetMetricData",
                "cloudwatch:GetMetricStatistics",
                "cloudwatch:ListMetrics",
                "cloudwatch:ListMetricStreams",
                "codebuild:ListProjects",
                "codeguru-profiler:ListProfilingGroups",
                "codeguru-reviewer:ListRepositoryAssociations",
                "cognito-identity:ListIdentityPools",
                "config:GetDiscoveredResourceCounts",
                "connect:ListInstances",
                "databrew:ListProjects",
                "dataexchange:ListDataSets",
```

```
                "datasync:ListAgents",
                "datapipeline:ListPipelines",
                "dax:DescribeClusters",
                "devops-guru:ListMonitoredResources",
                "directconnect:DescribeConnections",
                "directconnect:DescribeTags",
                "directconnect:DescribeVirtualInterfaces",
                "dlm:GetLifecyclePolicies",
                "dms:DescribeReplicationInstances",
                "dynamodb:DescribeTable",
                "dynamodb:ListTables",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeClientVpnEndpoints",
                "ec2:DescribeFleets",
                "ec2:DescribeElasticGpus",
                "ec2:DescribeImages",
                "ec2:DescribeInstances",
                "ec2:DescribeNatGateways",
                "ec2:DescribeRegions",
                "ec2:DescribeRouteTables",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeSnapshots",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSubnets",
                "ec2:DescribeTransitGatewayAttachments",
                "ec2:DescribeTransitGatewayRouteTables",
                "ec2:DescribeTransitGateways",
                "ec2:DescribeVolumes",
                "ec2:DescribeVpcPeeringConnections",
                "ec2:DescribeVpcs",
                "ec2:DescribeVpcEndpointConnections",
                "ec2:DescribeVpnConnections",
                "ec2:DescribeVpnGateways",
                "ecs:DescribeClusters",
                "ecs:DescribeContainerInstances",
                "ecs:DescribeServices",
                "ecs:DescribeTaskDefinition",
                "ecs:DescribeTasks",
                "ecs:ListClusters",
                "ecs:ListContainerInstances",
```

```
                "ecs:ListServices",
                "ecs:ListTagsForResource",
                "ecs:ListTasks",
                "eks:DescribeCluster",
                "eks:ListClusters",
                "elasticache:DescribeCacheClusters",
                "elasticache:ListTagsForResource",
                "elasticbeanstalk:DescribeConfigurationSettings",
                "elasticbeanstalk:DescribeEnvironmentHealth",
                "elasticbeanstalk:DescribeEnvironmentResources",
                "elasticbeanstalk:DescribeEnvironments",
                "elasticbeanstalk:DescribeInstancesHealth",
                "elasticfilesystem:DescribeFileSystems",
                "elastic-inference:DescribeAccelerators",
                "elasticloadbalancing:DescribeListeners",
                "elasticloadbalancing:DescribeLoadBalancers",
                "elasticloadbalancing:DescribeTags",
                "elasticloadbalancing:DescribeTargetGroups",
                "elasticloadbalancing:DescribeTargetHealth",
                "elasticmapreduce:ListClusters",
                "elasticmapreduce:ListInstances",
                "elastictranscoder:ListPipelines",
                "es:ListDomainNames",
                "events:ListRules",
                "finspace:ListEnvironments",
                "firehose:ListDeliveryStreams",
                "forecast:ListDatasetGroups",
                "frauddetector:GetDetectors",
                "fsx:DescribeFileSystems",
                "gamelift:ListFleets",
                "geo:ListMaps",
                "glacier:GetVaultNotifications",
                "glacier:ListTagsForVault",
                "glacier:ListVaults",
                "glue:ListJobs",
                "greengrass:ListCoreDevices",
                "groundstation:ListSatellites",
                "guardduty:ListDetectors",
                "healthlake:ListFHIRDatastores",
                "iam:GetAccountAuthorizationDetails",
```

```
                "iam:GetUser",
                "inspector2:ListUsageTotals",
                "iot:DescribeThing",
                "iot:ListScheduledAudits",
                "iot:ListSecurityProfiles",
                "iot:ListTagsForResource",
                "iot:ListThings",
                "iotanalytics:ListDatastores",
                "iotevents:ListDetectorModels",
                "iotsitewise:ListAssetModels",
                "iottwinmaker:ListWorkspaces",
                "ivs:ListChannels",
                "ivschat:ListRooms",
                "kafka:ListClusters",
                "kendra:ListIndices",
                "kinesis:ListStreams",
                "kinesisvideo:ListStreams",
                "kms:DescribeKey",
                "kms:ListAliases",
                "kms:ListKeys",
                "kms:ListResourceTags",
                "lakeformation:ListResources",
                "lambda:GetAccountSettings",
                "lambda:GetPolicy",
                "lambda:ListAliases",
                "lambda:ListEventSourceMappings",
                "lambda:ListFunctions",
                "lambda:ListTags",
                "lex:ListBots",
                "lightsail:GetBundles",
                "lightsail:GetInstanceMetricData",
                "lightsail:GetInstances",
                "lightsail:GetRegions",
                "logs:DescribeLogGroups",
                "lookoutequipment:ListDatasets",
                "lookoutmetrics:ListAnomalyDetectors",
                "lookoutvision:ListProjects",
                "mediaconnect:ListFlows",
                "mediaconvert:ListJobs",
                "mediaconvert:ListQueues",
```

```
                "medialive:ListChannels",
                "mediapackage-vod:ListPackagingConfigurations",
                "mediapackage-vod:ListPackagingGroups",
                "mediapackage:ListChannels",
                "mediatailor:ListPlaybackConfigurations",
                "mq:ListBrokers",
                "network-firewall:ListFirewalls",
                "nimble:ListStudios",
                "opsworks:DescribeInstances",
                "opsworks:DescribeStacks",
                "panorama:ListDevices",
                "personalize:ListDatasets",
                "polly:ListLexicons",
                "qldb:ListLedgers",
                "rds:DescribeDBClusters",
                "rds:DescribeDBInstances",
                "rds:DescribeDBSubnetGroups",
                "rds:ListTagsForResource",
                "redshift:DescribeClusters",
                "redshift:DescribeLoggingStatus",
                "rekognition:DescribeProjects",
                "robomaker:ListSimulationJobs",
                "route53:GetHostedZone",
                "route53:ListHealthChecks",
                "route53:ListHostedZones",
                "route53:ListResourceRecordSets",
                "rum:ListAppMonitors",
                "s3:GetBucketLocation",
                "s3:GetBucketLogging",
                "s3:GetBucketTagging",
                "s3:GetBucketWebsite",
                "s3:GetObject",
                "s3:ListAllMyBuckets",
                "s3:ListBucket",
                "sagemaker:ListDomains",
                "states:ListStateMachines",
                "secretsmanager:ListSecrets",
                "ses:ListIdentities",
                "servicecatalog:SearchProductsAsAdmin",
                "shield:DescribeEmergencyContactSettings",
```

```
                "shield:GetSubscriptionState",
                "shield:ListAttacks",
                "shield:ListProtections",
                "sns:ListSubscriptions",
                "sns:ListTopics",
                "sqs:GetQueueAttributes",
                "sqs:ListQueues",
                "ssm:GetInventory",
                "ssm:GetParameters",
                "storagegateway:ListGateways",
                "storagegateway:ListVolumes",
                "sts:GetCallerIdentity",
                "swf:ListDomains",
                "synthetics:DescribeCanaries",
                "tag:GetResources",
                "tag:GetTagKeys",
                "tag:GetTagValues",
                "timestream:ListDatabases",
                "transcribe:ListTranscriptionJobs",
                "transfer:ListServers",
                "translate:ListTextTranslationJobs",
                "waf-regional:GetRateBasedRule",
                "waf-regional:GetRule",
                "waf-regional:GetRuleGroup",
                "waf-regional:GetWebACL",
                "waf-regional:ListResourcesForWebACL",
                "waf-regional:ListWebACLs",
                "waf:GetRateBasedRule",
                "waf:GetRule",
                "waf:GetRuleGroup",
                "waf:GetWebACL",
                "waf:ListWebACLs",
                "workmail:ListOrganizations",
                "workspaces-web:ListPortals",
                "workspaces:DescribeWorkspaceDirectories",
                "workspaces:DescribeWorkspaces"
            ],
            "Resource": "*"
        }
    ]
```

```
}
```

7. Update the trust policy so that the IAM User you created is able to assume this role.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::<Account ID>:user/Skylar One-Service-
User"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {}
        }
    ]
}
```

You should now be able to discover your organization within ScienceLogic.

EC2 "Credential-less"

The following configuration process is designed for ScienceLogic systems (e.g. database servers, data collectors) that are hosted within Amazon Web Services in EC2. As with the IAM User approach, you must assign the following "Delegated Administrator" Policy to your organization.

1. Replace the <Account ID> with an account ID from withing your AWS organization.

IMPORTANT: ScienceLogic recommends that you do not use your Management Account ID for this step.

```
{
       "Version": "2012-10-17",
       "Statement": [
              {
                     "Sid": "Statement",
                     "Effect": "Allow",
                     "Principal": {
                            "AWS": "<Account_ID>"
                     },
```

```
                        "Action": [
                                "organizations:ListAccounts",
                                "organizations:DescribeAccount",
                                "organizations:DescribeOrganization"
                        ],
                        "Resource": "*"
                }
        ]
}
```

2. Within the AWS account that you assigned the above permissions, create the IAM Role with the following policy attached (e.g. Skylar One-Org-Read).

IMPORTANT: ScienceLogic recommends you take note of the Amazon Resource Number (ARN) of this IAM Role as it will be required for the Skylar One credential setup in future steps.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "organizations:ListAccounts",
                "organizations:DescribeOrganization",
                "organizations:DescribeAccount"
            ],
            "Resource": "*"
        }
    ]
}
```

3. Within the AWS Account that hosts the ScienceLogic platform, create the following IAM Role (e.g. "Skylar One-EC2-Collector") and assign the following policy:

IMPORTANT: ScienceLogic recommends you take note of the Amazon Resource Number (ARN) of this IAM Role as it will be required in a future step.

```
{
    "Version": "2012-10-17",
```

```
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": "sts:AssumeRole",
            "Resource": [
                "arn:aws:iam::<Account ID>:role/Skylar One-Org-Read",
                "arn:aws:iam::<Account ID>:role/ScienceLogic-Monitor",
                "arn:aws:iam::<Account ID>:role/ScienceLogic-Monitor",
                "arn:aws:iam::<Account ID>:role/ScienceLogic-Monitor"
            ]
        }
    ]
}
```

4. You must assign the role above to each EC2 instance that hosts a Skylar One component. Navigate to the *Modify IAM role* menu. Select the role you created and assign it to each of the EC2 instances that runs ScienceLogic components.
5. After you create the IAM Role, you must update the trust policy of the Skylar One-Org-Read role. An example follows:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::<Account ID>:role/Skylar One-EC2-Col-
lector"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

6. Within each of the AWS accounts that you want to monitor you must create an IAM Role named ScienceLogic-Monitor, which is case sensitive, and assign the following policy:

NOTE: For GovCloud recommendations, see the *Appendix* at the end of this chapter.

```
{
	"Version": "2012-10-17",
	"Statement": [
		{
			"Sid": "VisualEditor0",
			"Effect": "Allow",
			"Action": [
				"acm-pca:ListCertificateAuthorities",
				"airflow:ListEnvironments",
				"amplify:ListApps",
				"apigateway:GET",
				"appflow:ListFlows",
				"apprunner:ListServices",
				"appstream:DescribeFleets",
				"appsync:ListGraphqlApis",
				"athena:ListDatabases",
				"athena:ListDataCatalogs",
				"athena:ListWorkGroups",
				"autoscaling:DescribeAutoScalingGroups",
				"autoscaling:DescribeAutoScalingInstances",
				"autoscaling:DescribeLaunchConfigurations",
				"backup:ListBackupPlans",
				"cassandra:Select",
				"chime:ListAccounts",
				"chime:ListVoiceConnectors",
				"cloudfront:GetDistribution",
				"cloudfront:GetStreamingDistribution",
				"cloudfront:ListDistributions",
				"cloudfront:ListInvalidations",
				"cloudfront:ListStreamingDistributions",
				"cloudhsm:DescribeClusters",
				"cloudsearch:DescribeDomains",
				"cloudtrail:DescribeTrails",
				"cloudtrail:GetTrailStatus",
				"cloudwatch:DescribeAlarmHistory",
				"cloudwatch:DescribeAlarms",
				"cloudwatch:GetMetricData",
				"cloudwatch:GetMetricStatistics",
				"cloudwatch:ListMetrics",
				"cloudwatch:ListMetricStreams",
```

```
                "codebuild:ListProjects",
                "codeguru-profiler:ListProfilingGroups",
                "codeguru-reviewer:ListRepositoryAssociations",
                "cognito-identity:ListIdentityPools",
                "config:GetDiscoveredResourceCounts",
                "connect:ListInstances",
                "databrew:ListProjects",
                "dataexchange:ListDataSets",
                "datasync:ListAgents",
                "datapipeline:ListPipelines",
                "dax:DescribeClusters",
                "devops-guru:ListMonitoredResources",
                "directconnect:DescribeConnections",
                "directconnect:DescribeTags",
                "directconnect:DescribeVirtualInterfaces",
                "dlm:GetLifecyclePolicies",
                "dms:DescribeReplicationInstances",
                "dynamodb:DescribeTable",
                "dynamodb:ListTables",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeClientVpnEndpoints",
                "ec2:DescribeFleets",
                "ec2:DescribeElasticGpus",
                "ec2:DescribeImages",
                "ec2:DescribeInstances",
                "ec2:DescribeNatGateways",
                "ec2:DescribeRegions",
                "ec2:DescribeRouteTables",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeSnapshots",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSubnets",
                "ec2:DescribeTransitGatewayAttachments",
                "ec2:DescribeTransitGatewayRouteTables",
                "ec2:DescribeTransitGateways",
                "ec2:DescribeVolumes",
                "ec2:DescribeVpcPeeringConnections",
                "ec2:DescribeVpcs",
                "ec2:DescribeVpcEndpointConnections",
                "ec2:DescribeVpnConnections",
```

```
                "ec2:DescribeVpnGateways",
                "ecs:DescribeClusters",
                "ecs:DescribeContainerInstances",
                "ecs:DescribeServices",
                "ecs:DescribeTaskDefinition",
                "ecs:DescribeTasks",
                "ecs:ListClusters",
                "ecs:ListContainerInstances",
                "ecs:ListServices",
                "ecs:ListTagsForResource",
                "ecs:ListTasks",
                "eks:DescribeCluster",
                "eks:ListClusters",
                "elasticache:DescribeCacheClusters",
                "elasticache:ListTagsForResource",
                "elasticbeanstalk:DescribeConfigurationSettings",
                "elasticbeanstalk:DescribeEnvironmentHealth",
                "elasticbeanstalk:DescribeEnvironmentResources",
                "elasticbeanstalk:DescribeEnvironments",
                "elasticbeanstalk:DescribeInstancesHealth",
                "elasticfilesystem:DescribeFileSystems",
                "elastic-inference:DescribeAccelerators",
                "elasticloadbalancing:DescribeListeners",
                "elasticloadbalancing:DescribeLoadBalancers",
                "elasticloadbalancing:DescribeTags",
                "elasticloadbalancing:DescribeTargetGroups",
                "elasticloadbalancing:DescribeTargetHealth",
                "elasticmapreduce:ListClusters",
                "elasticmapreduce:ListInstances",
                "elastictranscoder:ListPipelines",
                "es:ListDomainNames",
                "events:ListRules",
                "finspace:ListEnvironments",
                "firehose:ListDeliveryStreams",
                "forecast:ListDatasetGroups",
                "frauddetector:GetDetectors",
                "fsx:DescribeFileSystems",
                "gamelift:ListFleets",
                "geo:ListMaps",
                "glacier:GetVaultNotifications",
```

```
                "glacier:ListTagsForVault",
                "glacier:ListVaults",
                "glue:ListJobs",
                "greengrass:ListCoreDevices",
                "groundstation:ListSatellites",
                "guardduty:ListDetectors",
                "healthlake:ListFHIRDatastores",
                "iam:GetAccountAuthorizationDetails",
                "iam:GetUser",
                "inspector2:ListUsageTotals",
                "iot:DescribeThing",
                "iot:ListScheduledAudits",
                "iot:ListSecurityProfiles",
                "iot:ListTagsForResource",
                "iot:ListThings",
                "iotanalytics:ListDatastores",
                "iotevents:ListDetectorModels",
                "iotsitewise:ListAssetModels",
                "iottwinmaker:ListWorkspaces",
                "ivs:ListChannels",
                "ivschat:ListRooms",
                "kafka:ListClusters",
                "kendra:ListIndices",
                "kinesis:ListStreams",
                "kinesisvideo:ListStreams",
                "kms:DescribeKey",
                "kms:ListAliases",
                "kms:ListKeys",
                "kms:ListResourceTags",
                "lakeformation:ListResources",
                "lambda:GetAccountSettings",
                "lambda:GetPolicy",
                "lambda:ListAliases",
                "lambda:ListEventSourceMappings",
                "lambda:ListFunctions",
                "lambda:ListTags",
                "lex:ListBots",
                "lightsail:GetBundles",
                "lightsail:GetInstanceMetricData",
                "lightsail:GetInstances",
```

```
                "lightsail:GetRegions",
                "logs:DescribeLogGroups",
                "lookoutequipment:ListDatasets",
                "lookoutmetrics:ListAnomalyDetectors",
                "lookoutvision:ListProjects",
                "mediaconnect:ListFlows",
                "mediaconvert:ListJobs
                "mediaconvert:ListQueues",
                "medialive:ListChannels",
                "mediapackage-vod:ListPackagingConfigurations",
                "mediapackage-vod:ListPackagingGroups",
                "mediapackage:ListChannels",
                "mediatailor:ListPlaybackConfigurations",
                "mq:ListBrokers",
                "network-firewall:ListFirewalls",
                "nimble:ListStudios",
                "opsworks:DescribeInstances",
                "opsworks:DescribeStacks",
                "panorama:ListDevices",
                "personalize:ListDatasets",
                "polly:ListLexicons",
                "qldb:ListLedgers",
                "rds:DescribeDBClusters",
                "rds:DescribeDBInstances",
                "rds:DescribeDBSubnetGroups",
                "rds:ListTagsForResource",
                "redshift:DescribeClusters",
                "redshift:DescribeLoggingStatus",
                "rekognition:DescribeProjects",
                "robomaker:ListSimulationJobs",
                "route53:GetHostedZone",
                "route53:ListHealthChecks",
                "route53:ListHostedZones",
                "route53:ListResourceRecordSets",
                "rum:ListAppMonitors",
                "s3:GetBucketLocation",
                "s3:GetBucketLogging",
                "s3:GetBucketTagging",
                "s3:GetBucketWebsite",
                "s3:GetObject",
```

```
                "s3:ListAllMyBuckets",
                "s3:ListBucket",
                "sagemaker:ListDomains",
                "states:ListStateMachines",
                "secretsmanager:ListSecrets",
                "ses:ListIdentities",
                "servicecatalog:SearchProductsAsAdmin",
                "shield:DescribeEmergencyContactSettings",
                "shield:GetSubscriptionState",
                "shield:ListAttacks",
                "shield:ListProtections",
                "sns:ListSubscriptions",
                "sns:ListTopics",
                "sqs:GetQueueAttributes",
                "sqs:ListQueues",
                "ssm:GetInventory",
                "ssm:GetParameters",
                "storagegateway:ListGateways",
                "storagegateway:ListVolumes",
                "sts:GetCallerIdentity",
                "swf:ListDomains",
                "synthetics:DescribeCanaries",
                "tag:GetResources",
                "tag:GetTagKeys",
                "tag:GetTagValues",
                "timestream:ListDatabases",
                "transcribe:ListTranscriptionJobs",
                "transfer:ListServers",
                "translate:ListTextTranslationJobs",
                "waf-regional:GetRateBasedRule",
                "waf-regional:GetRule",
                "waf-regional:GetRuleGroup",
                "waf-regional:GetWebACL",
                "waf-regional:ListResourcesForWebACL",
                "waf-regional:ListWebACLs",
                "waf:GetRateBasedRule",
                "waf:GetRule",
                "waf:GetRuleGroup",
                "waf:GetWebACL",
                "waf:ListWebACLs",
```

```
                                "workmail:ListOrganizations",
                                "workspaces-web:ListPortals",
                                "workspaces:DescribeWorkspaceDirectories",
                                "workspaces:DescribeWorkspaces"
                            ],
                            "Resource": "*"
                        }
                ]
}
```

7. Once you finished creating the IAM Role, the trust policy should be updated to the following:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::<Account ID>:role/Skylar One-EC2-Col-
lector"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {}
        }
    ]
}
```

Individual Account

To configure an individual AWS account monitored by ScienceLogic, you must create an IAM User within the particular account you want to monitor.

NOTE: ScienceLogic does not recommend that this IAM User have access to the AWS Management Console.

1. After you create the IAM User, you must assign several policies. Assign the first policy below:

```
{
    "Version": "2012-10-17",
```

```
"Statement": [
    {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": [
            "acm-pca:ListCertificateAuthorities",
            "airflow:ListEnvironments",
            "amplify:ListApps",
            "apigateway:GET",
            "appflow:ListFlows",
            "apprunner:ListServices",
            "appstream:DescribeFleets",
            "appsync:ListGraphqlApis",
            "athena:ListDatabases",
            "athena:ListDataCatalogs",
            "athena:ListWorkGroups",
            "autoscaling:DescribeAutoScalingGroups",
            "autoscaling:DescribeAutoScalingInstances",
            "autoscaling:DescribeLaunchConfigurations",
            "backup:ListBackupPlans",
            "cassandra:Select",
            "chime:ListAccounts",
            "chime:ListVoiceConnectors",
            "cloudfront:GetDistribution",
            "cloudfront:GetStreamingDistribution",
            "cloudfront:ListDistributions",
            "cloudfront:ListInvalidations",
            "cloudfront:ListStreamingDistributions",
            "cloudhsm:DescribeClusters",
            "cloudsearch:DescribeDomains",
            "cloudtrail:DescribeTrails",
            "cloudtrail:GetTrailStatus",
            "cloudwatch:DescribeAlarmHistory",
            "cloudwatch:DescribeAlarms",
            "cloudwatch:GetMetricData",
            "cloudwatch:GetMetricStatistics",
            "cloudwatch:ListMetrics",
            "cloudwatch:ListMetricStreams",
            "codebuild:ListProjects",
            "codeguru-profiler:ListProfilingGroups",
```

```
                "codeguru-reviewer:ListRepositoryAssociations",
                "cognito-identity:ListIdentityPools",
                "config:GetDiscoveredResourceCounts",
                "connect:ListInstances",
                "databrew:ListProjects",
                "dataexchange:ListDataSets",
                "datasync:ListAgents",
                "datapipeline:ListPipelines",
                "dax:DescribeClusters",
                "devops-guru:ListMonitoredResources",
                "directconnect:DescribeConnections",
                "directconnect:DescribeTags",
                "directconnect:DescribeVirtualInterfaces",
                "dlm:GetLifecyclePolicies",
                "dms:DescribeReplicationInstances",
                "dynamodb:DescribeTable",
                "dynamodb:ListTables",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeClientVpnEndpoints",
                "ec2:DescribeFleets",
                "ec2:DescribeElasticGpus",
                "ec2:DescribeImages",
                "ec2:DescribeInstances",
                "ec2:DescribeNatGateways",
                "ec2:DescribeRegions",
                "ec2:DescribeRouteTables",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeSnapshots",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSubnets",
                "ec2:DescribeTransitGatewayAttachments",
                "ec2:DescribeTransitGatewayRouteTables",
                "ec2:DescribeTransitGateways",
                "ec2:DescribeVolumes",
                "ec2:DescribeVpcPeeringConnections",
                "ec2:DescribeVpcs",
                "ec2:DescribeVpcEndpointConnections",
                "ec2:DescribeVpnConnections",
                "ec2:DescribeVpnGateways",
                "ecs:DescribeClusters",
```

```
                "ecs:DescribeContainerInstances",
                "ecs:DescribeServices",
                "ecs:DescribeTaskDefinition",
                "ecs:DescribeTasks",
                "ecs:ListClusters",
                "ecs:ListContainerInstances",
                "ecs:ListServices",
                "ecs:ListTasks",
                "ecs:ListTagsForResource",
                "eks:DescribeCluster",
                "eks:ListClusters",
                "elasticache:DescribeCacheClusters",
                "elasticbeanstalk:DescribeConfigurationSettings",
                "elasticbeanstalk:DescribeEnvironmentHealth",
                "elasticbeanstalk:DescribeEnvironmentResources",
                "elasticbeanstalk:DescribeEnvironments",
                "elasticbeanstalk:DescribeInstancesHealth",
                "elasticfilesystem:DescribeFileSystems",
                "elastic-inference:DescribeAccelerators",
                "elasticloadbalancing:DescribeListeners",
                "elasticloadbalancing:DescribeLoadBalancers",
                "elasticloadbalancing:DescribeTags",
                "elasticloadbalancing:DescribeTargetGroups",
                "elasticloadbalancing:DescribeTargetHealth",
                "elasticmapreduce:ListClusters",
                "elasticmapreduce:ListInstances",
                "elastictranscoder:ListPipelines",
                "es:ListDomainNames",
                "events:ListRules",
                "finspace:ListEnvironments",
                "firehose:ListDeliveryStreams",
                "forecast:ListDatasetGroups",
                "frauddetector:GetDetectors",
                "fsx:DescribeFileSystems",
                "gamelift:ListFleets",
                "geo:ListMaps",
                "glacier:GetVaultNotifications",
                "glacier:ListTagsForVault",
                "glacier:ListVaults",
                "glue:ListJobs",
```

```
                "greengrass:ListCoreDevices",
                "groundstation:ListSatellites",
                "guardduty:ListDetectors",
                "healthlake:ListFHIRDatastores",
                "iam:GetAccountAuthorizationDetails",
                "iam:GetUser",
                "inspector2:ListUsageTotals",
                "iot:DescribeThing",
                "iot:ListScheduledAudits",
                "iot:ListSecurityProfiles",
                "iot:ListTagsForResource",
                "iot:ListThings",
                "iotanalytics:ListDatastores",
                "iotevents:ListDetectorModels",
                "iotsitewise:ListAssetModels",
                "iottwinmaker:ListWorkspaces",
                "ivs:ListChannels",
                "ivschat:ListRooms",
                "kafka:ListClusters",
                "kendra:ListIndices",
                "kinesis:ListStreams",
                "kinesisvideo:ListStreams",
                "kms:DescribeKey",
                "kms:ListAliases",
                "kms:ListKeys",
                "kms:ListResourceTags",
                "lakeformation:ListResources",
                "lambda:GetAccountSettings",
                "lambda:GetPolicy",
                "lambda:ListAliases",
                "lambda:ListEventSourceMappings",
                "lambda:ListFunctions",
                "lambda:ListTags",
                "lex:ListBots",
                "lightsail:GetBundles",
                "lightsail:GetInstanceMetricData",
                "lightsail:GetInstances",
                "lightsail:GetRegions",
                "logs:DescribeLogGroups",
                "lookoutequipment:ListDatasets",
```

```
                "lookoutmetrics:ListAnomalyDetectors",
                "lookoutvision:ListProjects"
            ],
            "Resource": "*"
        }
    ]
}
```

2. The second policy, which is a continuation of the first, contains the remaining required permissions. Two separate policies are required due to character limitations on specific types of policies. Assign the second policy below:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "mediaconnect:ListFlows",
                "mediaconvert:ListJobs",
                "mediaconvert:ListQueues",
                "medialive:ListChannels",
                "mediapackage-vod:ListPackagingConfigurations",
                "mediapackage-vod:ListPackagingGroups",
                "mediapackage:ListChannels",
                "mediatailor:ListPlaybackConfigurations",
                "mq:ListBrokers",
                "network-firewall:ListFirewalls",
                "nimble:ListStudios",
                "opsworks:DescribeInstances",
                "opsworks:DescribeStacks",
                "panorama:ListDevices",
                "personalize:ListDatasets",
                "polly:ListLexicons",
                "qldb:ListLedgers",
                "rds:DescribeDBClusters",
                "rds:DescribeDBInstances",
                "rds:DescribeDBSubnetGroups",
                "rds:ListTagsForResource",
```

```
                "redshift:DescribeClusters",
                "redshift:DescribeLoggingStatus",
                "rekognition:DescribeProjects",
                "robomaker:ListSimulationJobs",
                "route53:GetHostedZone",
                "route53:ListHealthChecks",
                "route53:ListHostedZones",
                "route53:ListResourceRecordSets",
                "rum:ListAppMonitors",
                "s3:GetBucketLocation",
                "s3:GetBucketLogging",
                "s3:GetBucketTagging",
                "s3:GetBucketWebsite",
                "s3:GetObject",
                "s3:ListAllMyBuckets",
                "s3:ListBucket",
                "sagemaker:ListDomains",
                "states:ListStateMachines",
                "secretsmanager:ListSecrets",
                "ses:ListIdentities",
                "servicecatalog:SearchProductsAsAdmin",
                "shield:DescribeEmergencyContactSettings",
                "shield:GetSubscriptionState",
                "shield:ListAttacks",
                "shield:ListProtections",
                "sns:ListSubscriptions",
                "sns:ListTopics",
                "sqs:GetQueueAttributes",
                "sqs:ListQueues",
                "ssm:GetInventory",
                "ssm:GetParameters",
                "storagegateway:ListGateways",
                "storagegateway:ListVolumes",
                "sts:GetCallerIdentity",
                "swf:ListDomains",
                "synthetics:DescribeCanaries",
                "tag:GetResources",
                "tag:GetTagKeys",
                "tag:GetTagValues",
                "timestream:ListDatabases",
```

```
                "transcribe:ListTranscriptionJobs",
                "transfer:ListServers",
                "translate:ListTextTranslationJobs",
                "waf-regional:GetRateBasedRule",
                "waf-regional:GetRule",
                "waf-regional:GetRuleGroup",
                "waf-regional:GetWebACL",
                "waf-regional:ListResourcesForWebACL",
                "waf-regional:ListWebACLs",
                "waf:GetRateBasedRule",
                "waf:GetRule",
                "waf:GetRuleGroup",
                "waf:GetWebACL",
                "waf:ListWebACLs",
                "workmail:ListOrganizations",
                "workspaces-web:ListPortals",
                "workspaces:DescribeWorkspaceDirectories",
                "workspaces:DescribeWorkspaces"
            ],
            "Resource": "*"
        }
    ]
}
```

3. After you configured the IAM User policies, you must create the access key. Select the IAM User you created from the IAM User list.
4. Choose the Security credentials and then choose "Create Access Key". You will need the key and secret when you configure your ScienceLogic platform to discover your account.

Discovering your AWS Organization or AWS Account

After you complete the configuration of the AWS Organization or Account, you can monitor them within Skylar One.

IMPORTANT: If you use Guided Discovery and need to upgrade to a new version of the "Amazon Web Services" PowerPack, do not first uninstall the PowerPack on Skylar One versions 11.2.0 and later as you will lose your Guided Discovery workflows. To avoid losing your workflows, see the upgrade instructions in the "Amazon Web Services" PowerPack *Release Notes.*

To discover your organization or account:

1. Navigate to **Devices** and select **[Add Devices]**. The Guided Discovery workflow wizard appears. You can choose from three paths to follow when mapping your organization or account. You should choose the one that corresponds to the configuration options you previously set up:

- AWS Universal EC2 Guided Discovery (Organization Credential-less)
- AWS Universal IAM Guided Discovery (Individual Account)
- AWS Universal Assume Role Guided Discovery (Organization IAM)

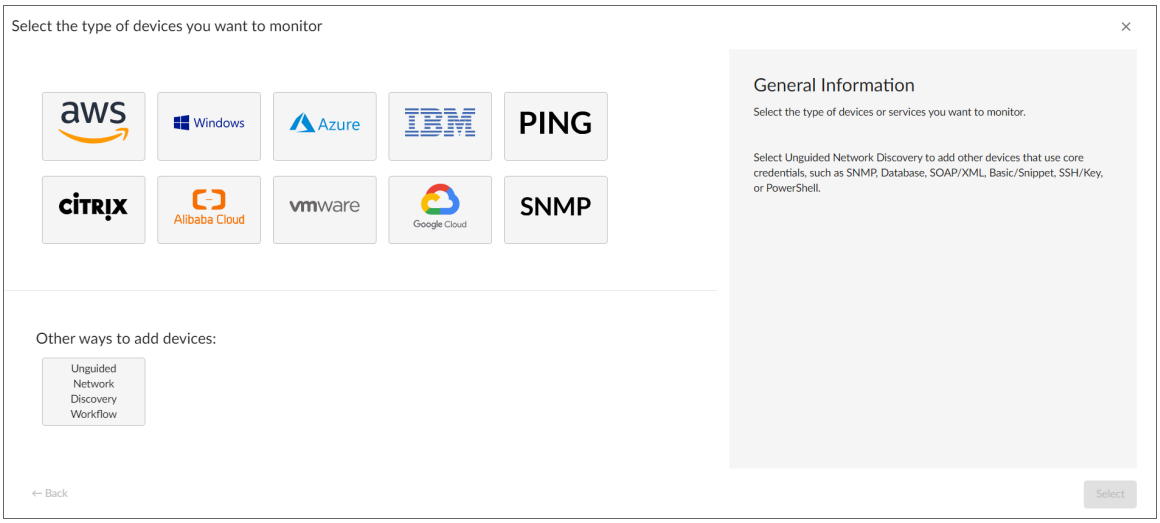

AWS Universal EC2 Guided Discovery (Organization Credential-less)

To define an AWS Universal EC2 Guided Discovery credential:

1. Choose AWS from the Guided Discovery workflow wizard.

2. Select the ***AWS Universal EC2 Guided Discovery*** radio button and click **[Select]**. The credential modal appears.
3. Click **[Create New]** to create the credential.

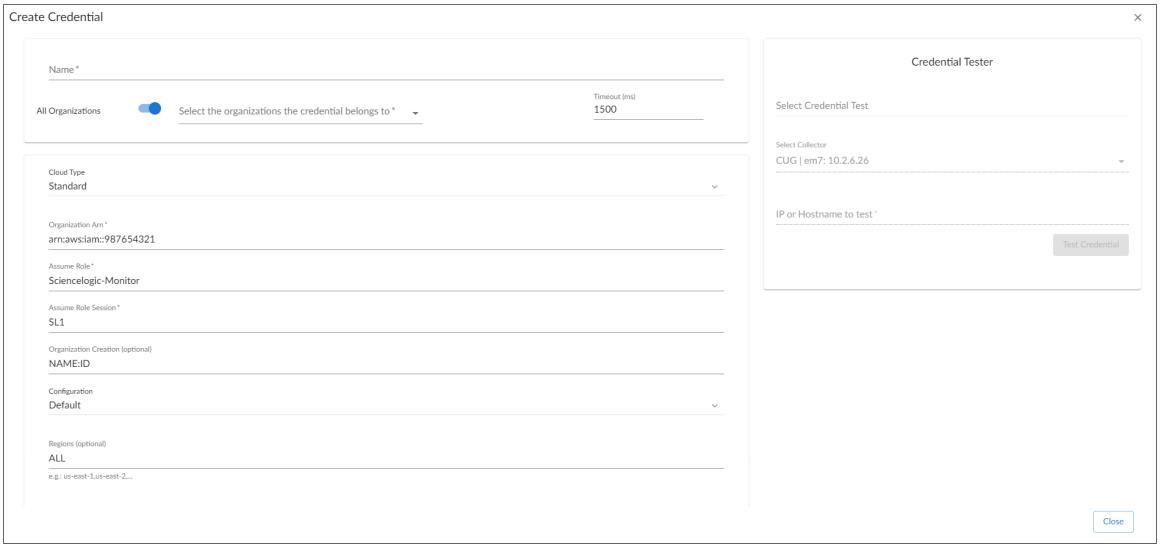

4. Supply the following values:

- ***Name.*** Type a new name for your AWS credential.
- ***Organization.*** This is the ARN of the role you created within the account you delegated adminstrator to. eg. (arn:aws:iam::123456789:role/Skylar One-Org-Read). Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the ***What organization manages this service?*** drop-down field to align the credential with those specific organizations.
- ***Timeout (ms)***. Keep the default value.
- ***AssumeRole.*** The name (case sensitive) of the role you created on each of the accounts with the assigned monitoring policy.
- ***AssumeRole Session.*** Session name used associated with AssumeRole sessions.
- ***Organization-Creation.*** Type in the pattern used to create the ScienceLogic Organization. Your options are:
    - Any combination
    - Name:ID
    - ID:Name
    - Name, ID
- ***Configuration.*** Keep the default value.
- ***Regions.*** You can filter which regions discover. Specify either an allow list (eg. us-east-1,us-east-2) of regions or a block list (eg. !us-east-1,!us-west-1 will prevent us-east-1 and us-west-1 from discovering).
- ***Filter by Tags (optional).*** You can filter some discoveries by tags.

- ***Proxy Hostname IP/Port etc (optional).*** You can configure proxy requests made by the PowerPack.
- ***Proxy Protocol.*** HTTPS by default, optionally HTTP. For HTTPS, SSL Verification is based off of the Skylar One's systems certificate trust bundle so you can configure it.
- ***Enable SSL Verification for Skylar One API.*** Enabled by default. If your Skylar One environment is configured with self signed or self generated certificates you should set this to disabled. Toggle on (blue) to enable, or toggle off (gray) to disable.

5. Click **[Close]**.

> **NOTE:** To discover AWS Gov with credential-less discovery, you must use the classic Skylar One user interface credential option.

AWS Universal IAM Guided Discovery (Individual Account)

To define an AWS Universal IAM Guided Discovery credential:

1. Choose AWS from the Guided Discovery workflow wizard.

2. Select the ***AWS Universal IAM Guided Discovery*** radio button and click **[Select]**. The credential modal appears.
3. Click **[Create New]** to create the credential.

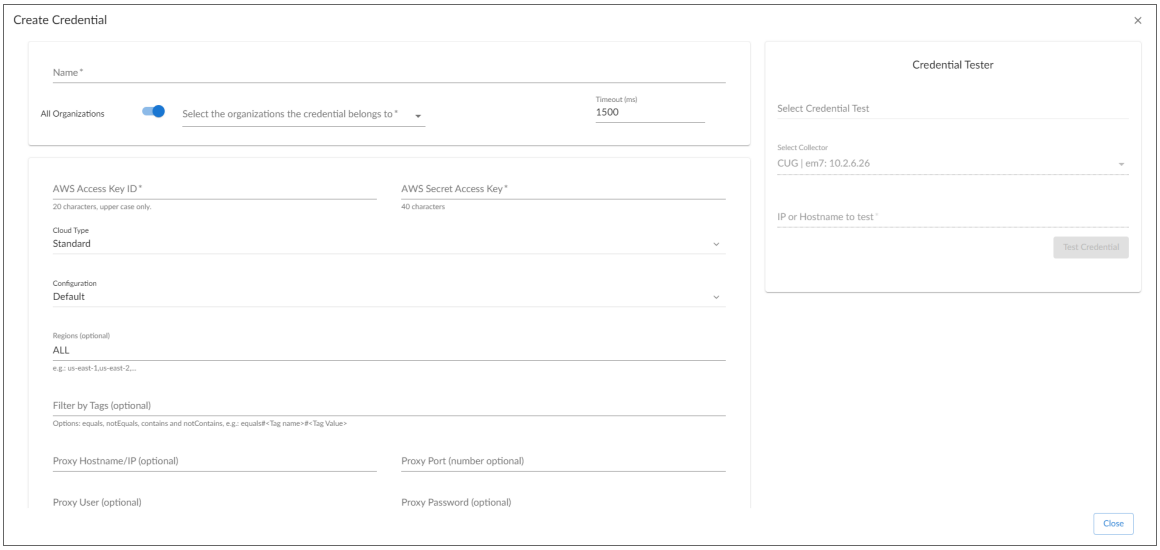

4. Supply the following values:

- ***Name.*** Type a new name for your AWS credential.
- ***Timeout (ms)***. Keep the default value.
- ***AWS Access Key ID.*** Type your ***Access Key ID***.
- ***AWS Secret Access Key.*** Type your ***Secret Access Key***.
- ***Cloud Type.*** Choose Standard, or GovCloud, if you provisioned this user in GovCloud. The Beijing option has been depreciated.
- ***Configuration.*** Keep the default value.
- ***Regions.*** You can filter which regions discover. Specify either an allow list (eg. us-east-1,us-east-2) of regions or a block list (eg. !us-east-1,!us-west-1 will prevent us-east-1 and us-west-1 from discovering).
- ***Filter by Tags (optional).*** You can filter some discoveries by tags.
- ***Proxy Hostname IP/Port etc (optional).*** You can configure proxy requests made by the PowerPack.
- ***Proxy Protocol.*** HTTPS by default, optionally HTTP. For HTTPS, SSL Verification is based off of the Skylar One's systems certificate trust bundle so you can configure it.
- ***Enable SSL Verification for Skylar One API.*** Enabled by default. If your Skylar One environment is configured with self signed or self generated certificates you should set this to disabled. Toggle on (blue) to enable, or toggle off (gray) to disable.

5. Click **[Close]**.

AWS Universal Assume Role Guided Discovery (Organization IAM)

> **IMPORTANT:** If you are upgrading the PowerPack and had previously discovered accounts within an organization separately and now want to use a different discovery method, you must first disable the "AWS: Account Discovery" Dynamic Application in each account that is being upgraded.

To define an AWS Universal Assume Role Guided Discovery credential:

1. Choose AWS from the Guided Discovery workflow wizard.

2. Select the ***AWS Universal Assume Role Guided Discovery*** radio button and click **[Select]**. The credential modal appears.
3. Click **[Create New]** to create the credential.

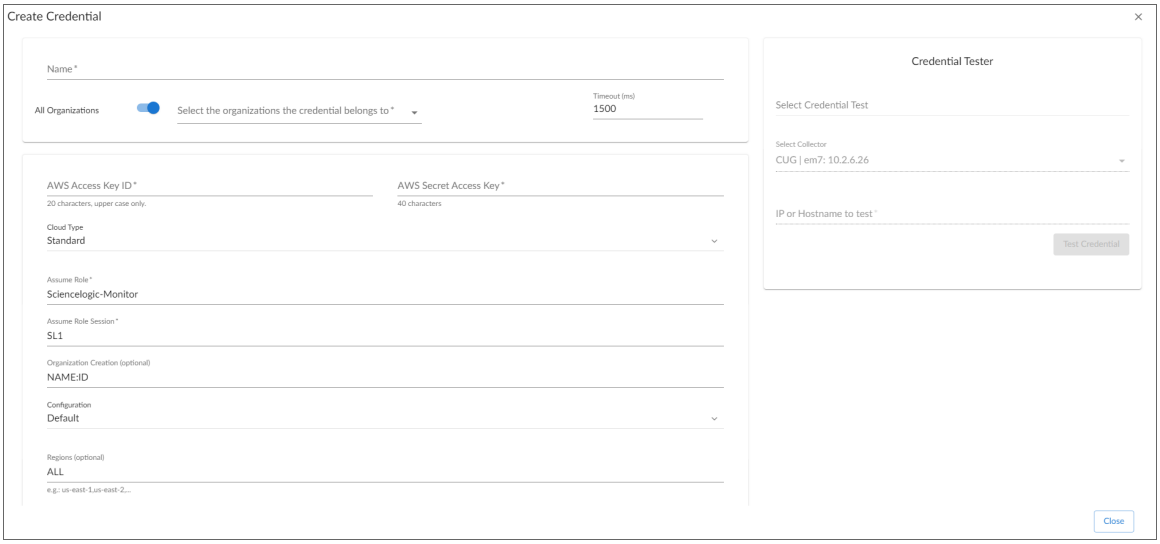

4. Supply the following values:

- ***Name.*** Type a new name for your AWS credential.
- ***Timeout (ms)***. Keep the default value.
- ***AWS Access Key ID.*** Type your ***Access Key ID***.
- ***AWS Secret Access Key.*** Type your ***Secret Access Key***.
- ***Assume Role.*** The name (case sensitive) of the role you created on each of the accounts with the assigned monitoring policy.
- ***Organization-Creation.*** Type in the pattern used to create the ScienceLogic Organization. Your options are:
    - Any combination
    - Name:ID
    - ID:Name
    - Name, ID
- ***Configuration.*** Keep the default value.
- ***Regions.*** You can filter which regions discover. Specify either an allow list (eg. us-east-1, us-east-2) of regions or a block list (eg. !us-east-1, !us-west-1 will prevent us-east-1 and us-west-1 from discovering).
- ***Filter by Tags (optional).*** You can filter some discoveries by tags.
- ***Proxy Hostname IP/Port etc (optional).*** You can configure proxy requests made by the PowerPack.
- ***Proxy Protocol.*** HTTPS by default, optionally HTTP. For HTTPS, SSL Verification is based off of the Skylar One's systems certificate trust bundle so you can configure it.

- ***Enable SSL Verification for Skylar One API.*** Enabled by default. If your Skylar One environment is configured with self signed or self generated certificates you should set this to disabled. Toggle on (blue) to enable, or toggle off (gray) to disable.

5. Click **[Close]**.

After you complete the credential creation, the guided workflow will prompt you to provide a name for the root virtual device and select a collector group to align it to. When the form is complete, the new virtual device will have an "Amazon Web Services" PowerPack Dynamic Application aligned to it.

Discovering your AWS Organization or AWS Account in the Classic Skylar One User Interface

Before you can discover and monitor your AWS organization or account, you must create the credential for the specific account you created.

Creating Credential in the Classic Skylar One User Interface

You can discover your organization or individual account in the classic Skylar One user interface using a number of ample SOAP-type credentials that are shipped in the "Amazon Web Services" PowerPack. In the classic user interface you must override the SOAP-type credentials for each of the following:

- AWS Credential (Individual Account)
- AWS Credential - EC2 Instance (Organization "Credential-less")
- AWS Credential - Master Account (Organization IAM)

AWS Credential (Individual Account)

To discover AWS using the manual discovery method, you must first define an AWS credential in Skylar One.

To define an AWS credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).

2. Locate the **AWS Credential** sample credential and click its wrench icon (🔧). The **Credential Editor** modal page appears.
3. Enter values in the following fields:

   Basic Settings

   - ***Profile Name***. Type a new name for your AWS credential.
   - ***URL***. Not required.
   - ***HTTP Auth User***. Type your ***Access Key ID***.
   - ***HTTP Auth Password***. Type your ***Secret Access Key***. The characters appear as asterisks to protect your password privacy.

   Proxy Settings

   - ***Hostname/IP***. Optional.
   - ***Port***. Optional.
   - ***User***. Optional.
   - ***Password***. Optional.

NOTE: All requests made with the "Amazon Web Services" PowerPack will be performed via HTTPs proxy with provided configuration.

SOAP Options

- ***Embed Value [%1]***. To monitor a GovCloud account, type "us-gov-west-1" or "us-gov-east-1". Otherwise, leave this field blank.

HTTPS Headers

- ***Regions.*** Optional. Type specific regions to discover or which regions you do not want to discover.
    - *Regions:ALL* Automatically discovers regions where your resources have been deployed.
    - *Regions:us-east-1, us-west-2* Discovers only us-east-1 and us-west-2.
    - *Regions:!us-east-1* Does not discover us-east-1, but will discover any other region with resources.

    NOTE: Region Allow List and Block List options cannot be combined.

- ***SL1APISSLVerify.*** Optional. This True or False header can be provided to disable SSL Verification to the Skylar One API if your Skylar One stack is deployed with self signed certificates.
- ***ProxyProtocol.*** Optional. HTTPS or HTTP header can be provided to override the default use of HTTPS for proxy settings. When HTTPS is used, SSL Verification is based on the Skylar One systems trust bundle.

4. Click **[Save]**.

AWS Credential - Master Account (Organization IAM)

To discover AWS using the manual discovery method, you must first define an AWS credential in Skylar One.

To define an AWS credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Locate the **AWS Credential - Master Account** sample credential and click its wrench icon (🔧). The **Credential Editor** modal page appears.

3. Enter values in the following fields:

Basic Settings

- ***Profile Name***. Type a new name for your AWS - Master Account credential.
- ***URL***. For commercial AWS accounts, the URL should be left as default. If you wish to discover a AWS GovCloud Organization the value should be: https://organizations.us-gov-west-1.amazonaws.com
- ***HTTP Auth User***. Type your ***Access Key ID***.
- ***HTTP Auth Password***. Type your ***Secret Access Key***. The characters appear as asterisks to protect your password privacy.

Proxy Settings

- ***Hostname/IP***. Optional.
- ***Port***. Optional.
- ***User***. Optional.
- ***Password***. Optional.

NOTE: All requests made with the "Amazon Web Services" PowerPack will be performed via HTTPs proxy with provided configuration.

HTTPS Headers

- ***Assume Role.*** Type the role which you configured in each of your AWS Organization's accounts that you wish to monitor.
- ***Assume Role Session.*** Type the session name used with the "Amazon Web Services" PowerPack AssumeRole sessions.
- ***Organization-Creation.*** Type in the pattern used to create the ScienceLogic Organization. Your options are:
  - Any combination
  - Name:ID
  - ID:Name
  - Name, ID

- ***Regions.*** Optional. Type specific regions to discover or which regions you do not want to discover.
    - *Regions:ALL* Automatically discovers regions where your resources have been deployed.
    - *Regions:us-east-1, us-west-2* Discovers only us-east-1 and us-west-2.
    - *Regions:!us-east-1* Does not discover us-east-1, but will discover any other region with resources.

    > **NOTE:** Region Allow List and Block List options cannot be combined.

- ***SL1APISSLVerify.*** Optional. This True or False header can be provided to disable SSL Verification to the Skylar One API if your Skylar One stack is deployed with self signed certificates.
- ***ProxyProtocol.*** Optional. HTTPS or HTTP header can be provided to override the default use of HTTPS for proxy settings. When HTTPS is used, SSL Verification is based on the Skylar One systems trust bundle.

4. Click **[Save]**.

AWS Credential - EC2 Instance (Organization "Credential-less")

To discover AWS using the manual discovery method, you must first define an AWS credential in Skylar One.

To define an AWS credential:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Locate the **AWS Credential - EC2 Instance** sample credential and click its wrench icon (🔧). The **Credential Editor** modal page appears.
3. Enter values in the following fields:

Basic Settings

- ***Profile Name***. Type a new name for your AWS - EC2 Instance credential.
- ***URL***. For commercial AWS accounts, the URL should be left as default. If you wish to discover a AWS GovCloud Organization the value should be: https://organizations.us-gov-west-1.amazonaws.com
- ***HTTP Auth User***. Leave as default.
- ***HTTP Auth Password***. Leave blank.

Proxy Settings

- ***Hostname/IP***. Optional.
- ***Port***. Optional.
- ***User***. Optional.
- ***Password***. Optional.

NOTE: All requests made with the "Amazon Web Services" PowerPack will be performed via HTTPs proxy with provided configuration.

HTTPS Headers

- ***OrganizationARN.*** This is the ARN of the role you created within the account you delegated adminstrator to. eg. (arn:aws:iam::123456789:role/Skylar One-Org-Read). Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the ***What organization manages this service?*** drop-down field to align the credential with those specific organizations.
- ***Assume Role.*** Type the role which you configured in each of your AWS Organization's accounts that you wish to monitor.
- ***Assume Role Session.*** Type the session name used with the "Amazon Web Services" PowerPack AssumeRole sessions.
- ***Organization-Creation.*** Type in the pattern used to create the ScienceLogic Organization. Your options are:
  - Any combination
  - Name:ID
  - ID:Name
  - Name, ID

- ***Regions.*** Optional. Type specific regions to discover or which regions you do not want to discover.
    - *Regions:ALL* Automatically discovers regions where your resources have been deployed.
    - *Regions:us-east-1, us-west-2* Discovers only us-east-1 and us-west-2.
    - *Regions:!us-east-1* Does not discover us-east-1, but will discover any other region with resources.

    > **NOTE:** Region Allow List and Block List options cannot be combined.

- ***SL1APISSLVerify.*** Optional. This True or False header can be provided to disable SSL Verification to the Skylar One API if your Skylar One stack is deployed with self signed certificates.
- ***ProxyProtocol.*** Optional. HTTPS or HTTP header can be provided to override the default use of HTTPS for proxy settings. When HTTPS is used, SSL Verification is based on the Skylar One systems trust bundle.

4. Click **[Save]**.

Launching the Discovery Session in the Classic Skylar One User Interface

AWS Organizations

For the both types of the AWS Organization discovery, you should create and launch a Discovery Session to start monitoring your AWS Organization. To discover AWS Accounts in an AWS Organization, perform the following steps:

1. Go to the **Discovery Control Panel** page (System > Manage > Classic Discovery or System > Manage > Discovery in the classic user interface).
2. Click the **[Create]** button. The **Discovery Session Editor** page appears.

3. Supply values in the following fields:
    - ***Name.*** Type the name you want to name your discovery session, such as "AWS Org Discovery".
    - ***IP Address/Hostname***. For commercial AWS type *organizations.us-east-1.amazonaws.com*. For GOVCloud type *organizations.us-gov-west-1.amazonaws.com*.
    - ***Other Credentials***. Select the credential you created.
    - ***Discover Non-SNMP***. Select this checkbox.

> **NOTE:** If your system can not discover a device via PING, please contact ScienceLogic Support for a manual workaround procedure.

4. Click the **[Save]** button to launch the discovery session. Your AWS Organization will now discover over the course of the next several hours.

Individual Accounts

The procedure for the Individual Account configuration for the Classic Skylar One user interface has been simplified to creating a Virtual Device and then aligning the "AWS: Account Discovery" Dynamic Application to your credential.

To create a virtual device that represents your Amazon service:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface in the Skylar One classic user interface).
2. Click the **[Actions]** button, then select *Create Virtual Device*. The **Virtual Device** modal page appears.
3. Enter values in the following fields:
    - ***Device Name.*** Type a name for your virtual device.
    - ***Organization***. Select the organization for this device. The organization the device is associated with limits the users that will be able to view and edit the device.
    - ***Device Class***. Select *Service | AWS Service*.
    - ***Collector***. Select the collector group that will monitor the device.
4. Click the **[Add]** button to create the virtual device.

To discover your AWS account, you must manually align the "AWS: Account Discovery" Dynamic Application with the AWS virtual device. After you do so, the other Dynamic Applications in the "Amazon Web Services" PowerPack will automatically align to discover and monitor all of the components in your AWS account.

To align the "AWS: Account Discovery" Dynamic Application to your virtual device:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. Click the wrench icon (🔧) for your virtual device.
3. In the **Device Administration** panel, click the **[Collections]** tab. The **Dynamic Application Collections** page appears.
4. Click the **[Actions]** button, and then select *Add Dynamic Application* from the menu.
5. In the **Dynamic Application Alignment** modal page, select *AWS: Account Discovery* in the ***Dynamic Applications*** field.
6. In the ***Credentials*** field, select the credential you created for your AWS Individual Account.
7. Click the **[Save]** button to align the Dynamic Application.

IMPORTANT: The only Dynamic Application you should ever need to manually align from the "Amazon Web Services" PowerPack is the "AWS: Account Discovery" Dynamic Application. Aligning any other Dynamic Application from the PowerPack will have unexpected results.

Appendix

To configure the AWS Organization EC2 "Credential-less" account for GovCloud, use the following minimum permissions:

```
{
        "Version": "2012-10-17",
```

```
    "Statement": [
            {
                    "Sid": "VisualEditor0",
                    "Effect": "Allow",
                    "Action": [
                            "acm-pca:ListCertificateAuthorities",
                            "apigateway:GET",
                            "appstream:DescribeFleets",
                            "athena:ListWorkGroups",
                            "athena:ListDatabases",
                            "athena:ListDataCatalogs",
                            "autoscaling:DescribeAutoScalingGroups",
                            "autoscaling:DescribeAutoScalingInstances",
                            "autoscaling:DescribeLaunchConfigurations",
                            "backup:ListBackupPlans",
                            "cloudhsm:DescribeClusters",
                            "cloudtrail:DescribeTrails",
                            "cloudtrail:GetTrailStatus",
                            "cloudwatch:DescribeAlarmHistory",
                            "cloudwatch:DescribeAlarms",
                            "cloudwatch:GetMetricData",
                            "cloudwatch:GetMetricStatistics",
                            "cloudwatch:ListMetrics",
                            "cloudwatch:ListMetricStreams",
                            "codebuild:ListProjects",
                            "cognito-identity:ListIdentityPools",
                            "config:GetDiscoveredResourceCounts",
                            "connect:ListInstances",
                            "databrew:ListProjects",
                            "datasync:ListAgents",
                            "directconnect:DescribeConnections",
                            "directconnect:DescribeTags",
                            "directconnect:DescribeVirtualInterfaces",
                            "dms:DescribeReplicationInstances",
                            "dynamodb:DescribeTable",
                            "dynamodb:ListTables",
                            "ec2:DescribeAvailabilityZones",
                            "ec2:DescribeClientVpnEndpoints",
                            "ec2:DescribeFleets",
                            "ec2:DescribeImages",
```

```
                "ec2:DescribeInstances",
                "ec2:DescribeNatGateways",
                "ec2:DescribeRegions",
                "ec2:DescribeRouteTables",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeSnapshots",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSubnets",
                "ec2:DescribeTransitGatewayAttachments",
                "ec2:DescribeTransitGatewayRouteTables",
                "ec2:DescribeTransitGateways",
                "ec2:DescribeVolumes",
                "ec2:DescribeVpcEndpointConnections",
                "ec2:DescribeVpcPeeringConnections",
                "ec2:DescribeVpcs",
                "ec2:DescribeVpnConnections",
                "ec2:DescribeVpnGateways",
                "ecs:DescribeClusters",
                "ecs:DescribeContainerInstances",
                "ecs:DescribeServices",
                "ecs:DescribeTaskDefinition",
                "ecs:DescribeTasks",
                "ecs:ListClusters",
                "ecs:ListContainerInstances",
                "ecs:ListServices",
                "ecs:ListTasks",
                "ecs:ListTagsForResource",
                "eks:DescribeCluster",
                "eks:ListClusters",
                "elasticache:DescribeCacheClusters",
                "elasticbeanstalk:DescribeConfigurationSettings",
                "elasticbeanstalk:DescribeEnvironmentHealth",
                "elasticbeanstalk:DescribeEnvironmentResources",
                "elasticbeanstalk:DescribeEnvironments",
                "elasticbeanstalk:DescribeInstancesHealth",
                "elasticfilesystem:DescribeFileSystems",
                "elasticloadbalancing:DescribeListeners",
                "elasticloadbalancing:DescribeLoadBalancers",
                "elasticloadbalancing:DescribeTags",
                "elasticloadbalancing:DescribeTargetGroups",
```

```
                "elasticloadbalancing:DescribeTargetHealth",
                "elasticmapreduce:ListClusters",
                "elasticmapreduce:ListInstances",
                "es:ListDomainNames",
                "events:ListRules",
                "firehose:ListDeliveryStreams",
                "fsx:DescribeFileSystems",
                "glacier:GetVaultNotifications",
                "glacier:ListTagsForVault",
                "glacier:ListVaults",
                "glue:ListJobs",
                "greengrass:ListCoreDevices",
                "guardduty:ListDetectors",
                "iam:GetAccountAuthorizationDetails",
                "iam:GetUser",
                "inspector2:ListUsageTotals",
                "iot:DescribeThing",
                "iot:ListScheduledAudits",
                "iot:ListSecurityProfiles",
                "iot:ListTagsForResource",
                "iot:ListThings",
                "iotevents:ListDetectorModels",
                "iotsitewise:ListAssetModels",
                "kafka:ListClusters",
                "kendra:ListIndices",
                "kinesis:ListStreams",
                "kms:DescribeKey",
                "kms:ListAliases",
                "kms:ListKeys",
                "kms:ListResourceTags",
                "lakeformation:ListResources",
                "lambda:GetAccountSettings",
                "lambda:GetPolicy",
                "lambda:ListAliases",
                "lambda:ListEventSourceMappings",
                "lambda:ListFunctions",
                "lambda:ListTags",
                "logs:DescribeLogGroups",
                "mq:ListBrokers",
                "network-firewall:ListFirewalls",
```

```
                "polly:ListLexicons",
                "rds:DescribeDBClusters",
                "rds:DescribeDBInstances",
                "rds:DescribeDBSubnetGroups",
                "rds:ListTagsForResource",
                "redshift:DescribeClusters",
                "redshift:DescribeLoggingStatus",
                "rekognition:DescribeProjects",
                "route53:GetHostedZone",
                "route53:ListHealthChecks",
                "route53:ListHostedZones",
                "route53:ListResourceRecordSets",
                "s3:GetBucketLocation",
                "s3:GetBucketLogging",
                "s3:GetBucketTagging",
                "s3:GetBucketWebsite",
                "s3:GetObject",
                "s3:ListAllMyBuckets",
                "s3:ListBucket",
                "sagemaker:ListDomains",
                "states:ListStateMachines",
                "secretsmanager:ListSecrets",
                "servicecatalog:SearchProductsAsAdmin",
                "ses:ListIdentities",
                "sns:ListSubscriptions",
                "sns:ListTopics",
                "sqs:GetQueueAttributes",
                "sqs:ListQueues",
                "ssm:GetInventory",
                "ssm:GetParameters",
                "storagegateway:ListGateways",
                "storagegateway:ListVolumes",
                "sts:GetCallerIdentity",
                "swf:ListDomains",
                "synthetics:DescribeCanaries",
                "tag:GetResources",
                "tag:GetTagKeys",
                "tag:GetTagValues",
                "transcribe:ListTranscriptionJobs",
                "transfer:ListServers",
```

```
                                "translate:ListTextTranslationJobs",
                                "workspaces:DescribeWorkspaceDirectories",
                                "workspaces:DescribeWorkspaces"
                            ],
                            "Resource": "*"
                        }
                ]
}
```

Chapter 6

The AWS Credential Test and Viewing Component Devices

Overview

The following sections describe how to use the AWS credential test, understanding AWS Dynamic Applications, and how to view AWS component devices with the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

Testing the AWS Credential *104*

Viewing AWS Component Devices *108*

Testing the AWS Credential

NOTE: The Credential Test is for use with the *Manual Discovery* method only.

Skylar One includes a Credential Test for Amazon Web Services. Credential Tests define a series of steps that Skylar One can execute on demand to validate whether a credential works as expected.

The AWS Credential Test can only be used to test a SOAP/XML credential or a SOAP/XML with AWS subtype credential for monitoring AWS using the Dynamic Applications in the "Amazon Web Services" PowerPack. The AWS Credential Test performs the following steps:

- ***Test Reachability***. Performs an ICMP ping request to the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.

- ***Test Port Availability***. Performs an NMAP request to TCP port 443 on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- ***Test Name Resolution***. Performs an nslookup request on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- ***Make connection to AWS account***. Attempts to connect to the AWS service using the account specified in the credential.
- ***Scan AWS services***. Verifies that the account specified in the credential has access to the services.

NOTE: The AWS Credential Test does not support the testing of credentials that connect to AWS through a proxy server.

NOTE: The AWS Credential Test does not support AWS UCF Credentials.

To test the AWS credential:

1. Go to the **Credentials** page (Manage > Credentials).
2. Locate the credential you wish to test, select the **Actions** button (···) next to it and click *Test*.
3. The **Credential Test Form** modal page appears. Fill out the following fields on this page:
    - ***Credential***. This field is read-only and displays the name of the credential you selected.
    - ***Select Credential Test***. Select **AWS Credential Test**.
    - ***Collector***. Select the All-In-One Appliance or Data Collector that will run the test.
    - ***IP or Hostname to Test***. Enter a valid IP address.
4. Click the **[Run Test]** button to run the credential test. The **Testing Credential** window appears:

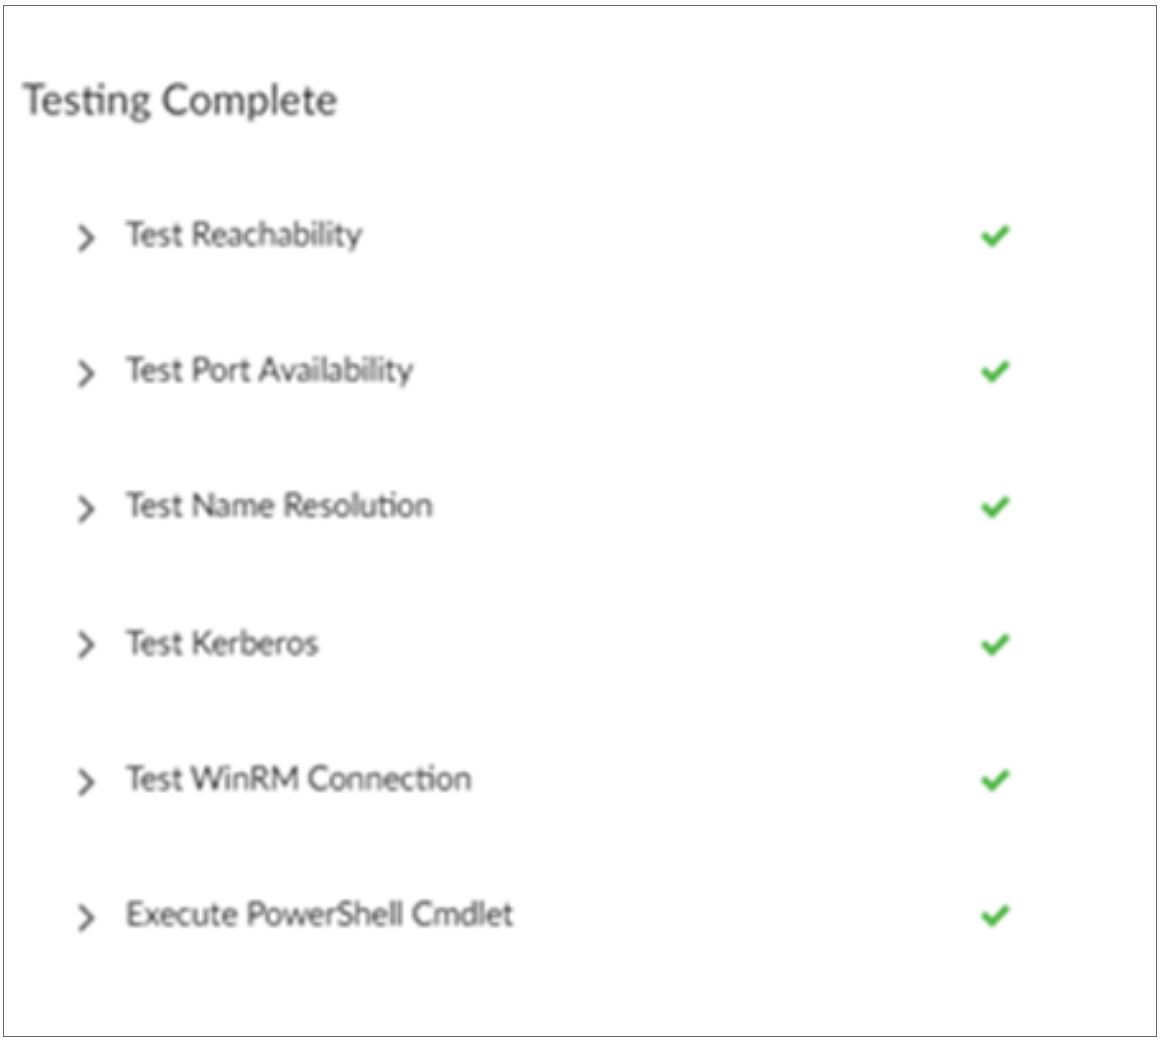

The **Testing Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- ***Step***. The name of the step.
- ***Description***. A description of the action performed during the step.
- ***Log Message***. The result of the step for this execution of the credential test.
- ***Status***. Whether the result of this step indicates the credential and/or the network environment is configured correctly (Passed) or incorrectly (Failed).
- ***Step Tip***. Mouse over the question mark icon (?) to display the tip text. The tip text recommends what to do to change the credential and/or the network environment if the step has a status of "Failed".

Testing the AWS Credential in the Skylar One Classic User Interface

NOTE: The Credential Test is for use with the *Manual Discovery* method only.

Skylar One includes a Credential Test for Amazon Web Services. Credential Tests define a series of steps that Skylar One can execute on demand to validate whether a credential works as expected.

The AWS Credential Test can only be used to test a SOAP/XML credential or a SOAP/XML with AWS subtype credential for monitoring AWS using the Dynamic Applications in the "Amazon Web Services" PowerPack. The AWS Credential Test performs the following steps:

- ***Test Reachability***. Performs an ICMP ping request to the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- ***Test Port Availability***. Performs an NMAP request to TCP port 443 on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- ***Test Name Resolution***. Performs an nslookup request on the URL for the EC2 service in the region specified in the credential. If a region is not specified in the credential, the us-east-1 region is used.
- ***Make connection to AWS account***. Attempts to connect to the AWS service using the account specified in the credential.
- ***Scan AWS services***. Verifies that the account specified in the credential has access to the services.

NOTE: The AWS Credential Test does not support the testing of credentials that connect to AWS through a proxy server.

To test the AWS credential:

1. Go to the **Credential Test Management** page (System > Customize > Credential Tests).
2. Locate the **AWS Credential Test** and click its lightning bolt icon (). The **Credential Tester** modal page appears.
3. Supply values in the following fields:
    - ***Test Type***. This field is pre-populated with the credential test you selected.
    - ***Credential***. Select the credential to test. This drop-down list includes only credentials that you have access to that can be tested using the selected credential test.
    - ***Hostname/IP***. Leave this field blank.
    - ***Collector***. Select the All-In-One Appliance or Data Collector that will run the test.
4. Click the **[Run Test]** button to run the credential test. The **Test Credential** window appears.

The **Test Credential** window displays a log entry for each step in the credential test. The steps performed are different for each credential test. The log entry for each step includes the following information:

- ***Step***. The name of the step.
- ***Description***. A description of the action performed during the step.
- ***Log Message***. The result of the step for this credential test.
- ***Status***. Whether the result of this step indicates the credential or the network environment is configured correctly (Passed) or incorrectly (Failed).
- ***Step Tip***. Mouse over the question mark icon (?) to display the tip text. The tip text recommends what to do to change the credential or the network environment if the step has a status of "Failed".

Viewing AWS Component Devices

When Skylar One performs collection for the AWS virtual device, Skylar One will create component devices that represent each element in your AWS infrastructure and align other Dynamic Applications to those component devices. Some of the Dynamic Applications aligned to the component devices will also be used to create additional component devices. All component devices appear in the **Devices** page.

In addition to the **Devices** page, you can view the AWS service and all associated component devices in the following places in the user interface:

- The **Device Investigator** Map page (click **Map** in the **Device Investigator** page) displays a map of a particular device and all of the devices with which it has parent-child relationships. Double-clicking any of the listed devices reloads the page to make the selected device the primary device

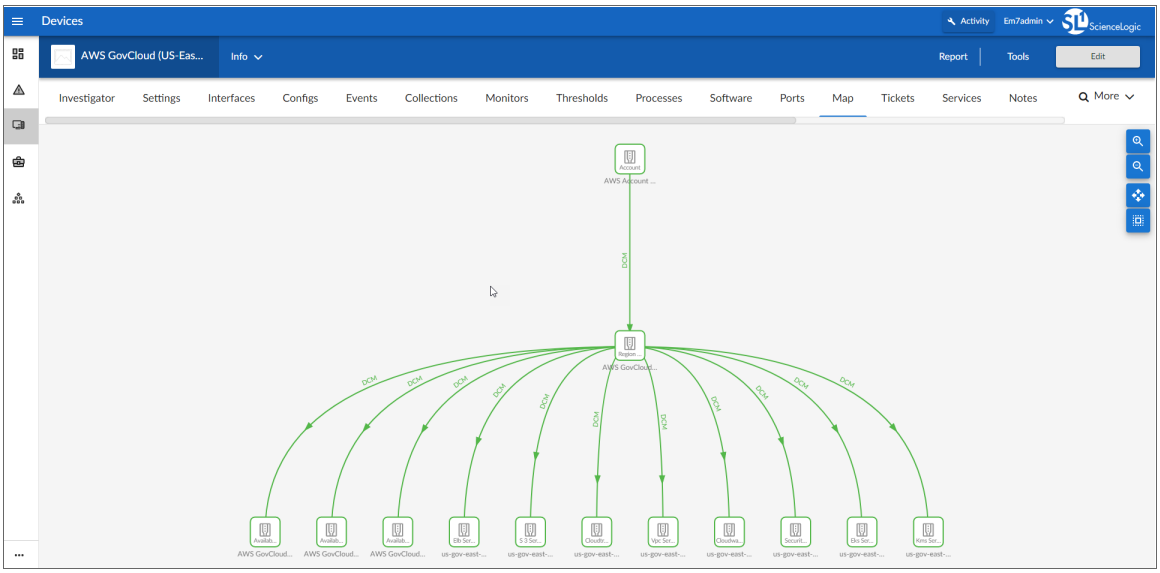

- The **Device Components** page (Devices > Device Components) displays a list of all root devices and component devices discovered by Skylar One in an indented view, so you can easily view the hierarchy and relationships between child devices, parent devices, and root devices. To view the component devices associated with an AWS service, find the AWS virtual device and click its plus icon (**+**).

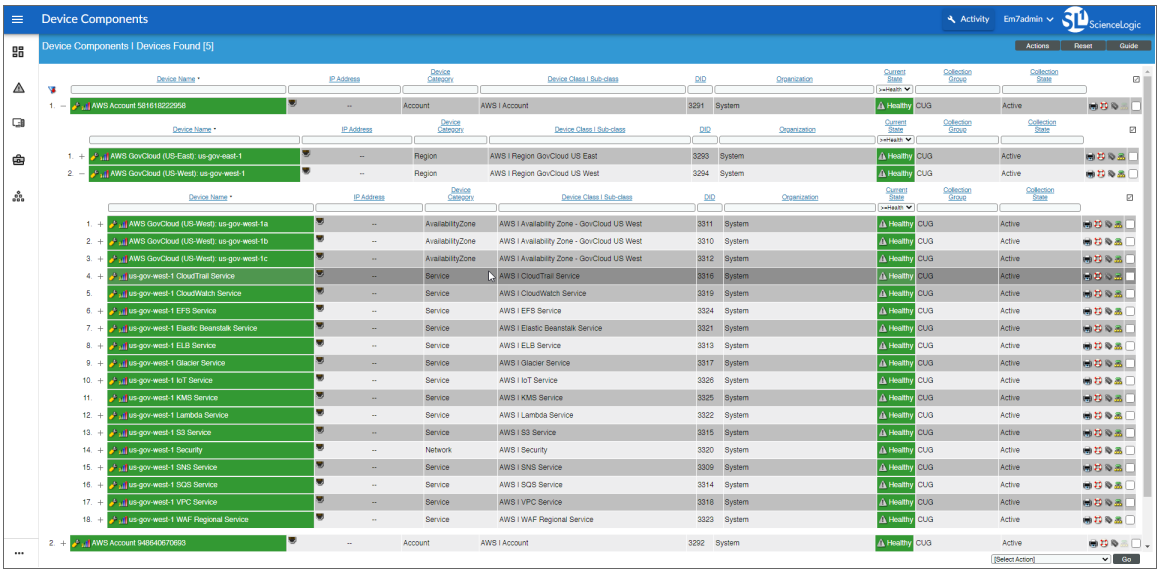

- The **Component Map** page (Classic Maps > Device Maps > Components) allows you to view devices by root node and view the relationships between root nodes, parent components, and child components in a map. This makes it easy to visualize and manage root nodes and their components. Skylar One automatically updates the **Component Map** as new component devices are discovered. Skylar One also updates each map with the latest status and event information. To view the map for an AWS service, go to Classic Maps > Device Maps > Components, and select the map from the list in the left NavBar. To learn more about the **Component Map** page, see the ***Maps*** manual.

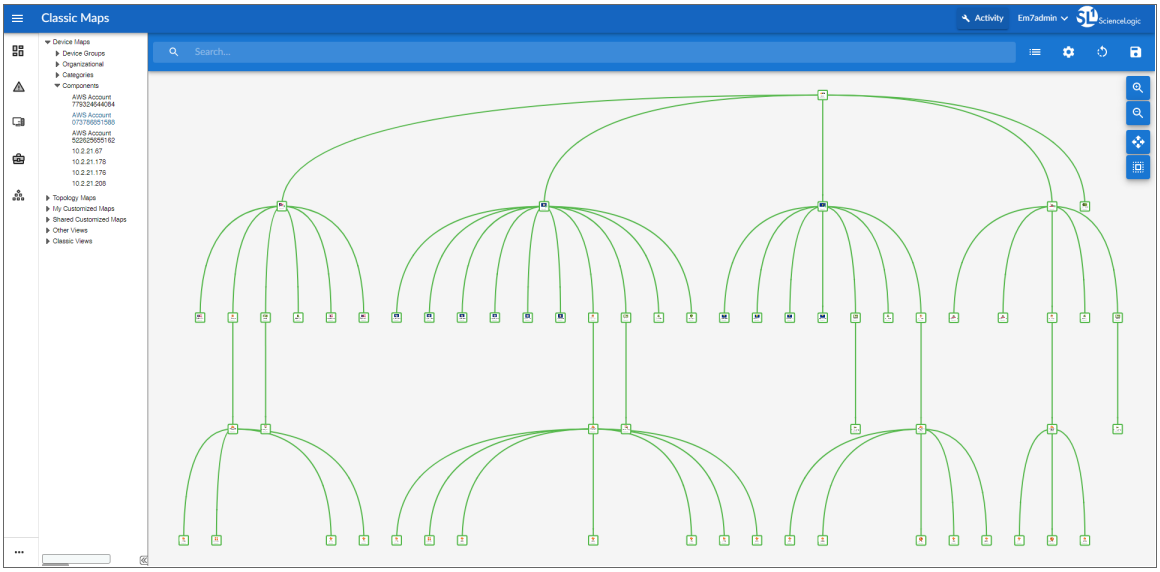

Relationships Between Component Devices

In addition to the parent/child relationships between component devices, relationships are automatically created by the Dynamic Applications in the "Amazon Web Services" PowerPack between the following component devices:

- AWS API Gateway Services and AWS Network Load Balancers
- AWS API Instances and AWS Lambda Functions
- AWS Application ELBs and AWS Availability Zones
- AWS Application ELBs and AWS Route 53-Hosted Zones
- AWS Application ELBs and AWS Security Groups
- AWS Application ELBs and AWS Target Groups
- AWS Application ELBs and AWS VPC Instances
- AWS Auto Scale Groups and AWS Auto Scale Launch Configurations
- AWS Direct Connect Virtual Instances and AWS Virtual Private Gateways
- AWS ECS Instances and AWS EC2 Instances
- AWS ECS Services and AWS Classic Load Balancers
- AWS ECS Services and AWS Security Groups
- AWS ECS Services and AWS Target Groups
- AWS ECS Services and AWS VPC Instances
- AWS ECS Services and AWS VPC Subnets
- AWS EC2 Instances and AWS Auto Scale Groups
- AWS EC2 Instances and AWS EBS Volumes
- AWS EC2 Instances and AWS Elastic Beanstalk Applications
- AWS EC2 Instances and AWS ELB Instances
- AWS EC2 Instances and AWS EMR Instances
- AWS EC2 Instances and AWS OpsWorks Instances
- AWS EC2 Instances and AWS Security Groups
- AWS EC2 Instances and AWS Target Groups
- AWS EC2 Instances and AWS VPC Instances
- AWS EC2 Instances and AWS VPC Subnets
- AWS EC2 Instances and the Cisco Cloud Center application
- AWS Lambda Functions and AWS Security Groups
- AWS Lambda Functions and AWS Simple Notification Services (SNS)
- AWS Lambda Functions and AWS Simple Queue Services (SQS)
- AWS Lambda Functions and AWS VPC Instances
- AWS Lambda Functions and AWS VPC Subnets

- AWS Lambda Function Qualified Services and AWS Security Groups
- AWS Lambda Function Qualified Services and AWS VPC Instances
- AWS Lambda Function Qualified Services and AWS VPC Subnets
- AWS Lambda Function Replicas and their parent AWS Lambda Function Versions
- AWS Network ELBs and AWS Availability Zones
- AWS Network ELBs and AWS Route 53-Hosted Zones
- AWS Network ELBs and AWS Target Groups
- AWS Network ELBs and AWS VPC Instances
- AWS Organizations and AWS Accounts
- AWS RDS Aurora Clusters and AWS RDS DB Instances
- AWS Redshift Instances and AWS Security Groups
- AWS Redshift Instances and AWS VPC Instances
- AWS Route Tables and AWS Virtual Private Gateways
- AWS Route Tables and AWS VPC Subnets
- AWS S3 Instances and AWS CloudTrail Instances
- AWS Security Groups and AWS VPC Instances
- AWS SNS Instances and AWS CloudTrail Instances
- AWS SNS Instances and AWS Glacier Instances
- AWS Transit Gateways and AWS VPC Instances
- AWS VPC Instances and AWS ELB Instances
- AWS VPC Instances and AWS Target Groups
- AWS WorkSpaces and AWS VPC Subnets

Vanishing Component Devices

If Skylar One cannot retrieve information about a component device for the amount of time specified in the ***Component Vanish Timeout*** field (in either the **Global Threshold Settings** page, the **Device Thresholds** page for the component device, or the **Device Thresholds** page for a device higher in the component tree), Skylar One sets the device to "vanished".

When a device is set to "vanished", Skylar One stops trying to collect data about the component device. The vanished device will not appear in reports or views. The vanished device will appear only in the **Vanished Device Manager** page. When a device is set to "vanished", all children of that device are also set to "vanished".

NOTE: This section describes the standard device vanishing behavior that ***does not*** use the "AWS: Vanish Terminated EC2 Instances" Run Book Action and Automation policies. If you use the "AWS: Vanish Terminated EC2 Instances" Run Book Action and Automation policies, see the chapter on **"AWS Run Book Actions and Automations"** for more information about device vanishing.

Most AWS component devices operate using the standard Skylar One vanishing logic: If the device is terminated in AWS, it then becomes unavailable in Skylar One. If the device is unavailable for the amount of time specified in the ***Component Vanish Timeout*** field, then that device is vanished.

However, two AWS component device types operate using slightly different logic:

- ***EC2***. EC2 instances that are deleted in AWS still appear in the AWS portal for one to two hours in a *terminated* state. If Skylar One polls that device and receives a response from AWS that the EC2 is terminated, Skylar One will classify the device as unavailable. If the ***Component Vanish Timeout*** setting has been enabled, then Skylar One will vanish this device automatically. If, however, the EC2 instance has merely been *stopped* rather than terminated, Skylar One will not vanish the device, even if the ***Component Vanish Timeout*** setting has been enabled.
- ***RDS***. RDS instances that have a status of *stopped* or *stopping* in AWS will be classified as unavailable in Skylar One. If the ***Component Vanish Timeout*** setting has been enabled, then Skylar One will vanish this device automatically.

ScienceLogic recommends setting the ***Component Vanish Timeout*** to *120 minutes* when monitoring AWS accounts.

For more information about vanishing devices, see the chapter on "Vanishing & Purging Devices" in the ***Device Management*** manual.

Chapter 7

Configuring Inbound CloudWatch Alarms

Overview

The following sections describe the CloudWatch alarm Event Policies that are included in the "Amazon Web Services" PowerPack and information about configuring CloudWatch and Skylar One to generate events based on CloudWatch alarms:

This chapter covers the following topics:

CloudWatch Alarm Event Policies *113*

Creating Custom CloudWatch Metrics *115*

Configuring CloudWatch to Send Alarms for a Metric *118*

Enabling Custom Metrics Collection in Skylar One *119*

Configuring the "AWS: CloudWatch Alarms Performance" Dynamic Application *119*

Enabling CloudWatch Alarm Events in Skylar One *121*

CloudWatch Alarm Event Policies

Amazon CloudWatch is a service that allows you to monitor your AWS resources and applications in near real-time. You can use CloudWatch to collect and track metrics, and use CloudWatch alarms to send notifications or automatically trigger changes to the resources being monitored based on rules that you define.

In addition to Skylar One collecting metrics for AWS instances, you can configure CloudWatch to send alarm information to Skylar One via API. Skylar One can then generate an event for each alarm.

The "Amazon Web Services" PowerPack includes an "AWS: CloudWatch Alarms Performance" Dynamic Application. This Dynamic Application monitors CloudWatch alarms and associates the alarms with the appropriate AWS component devices, if applicable. If an appropriate component device does not exist in Skylar One or cannot be determined, the alarm is instead associated with the component device for the AWS account.

CAUTION: The performance data collected by the "AWS: CloudWatch Alarms Performance" Dynamic Application is metadata intended to give general insight into the alarm activity the Dynamic Application is processing. This metadata can help identify overall trends, but users should be cautioned that the data presented can be imprecise in certain scenarios, such as when the Dynamic Application is being run in debug mode while data is still being collected.

The "Amazon Web Services" PowerPack also includes several pre-defined event policies for CloudWatch alarms:

Alarm Type	Alarm State	Event Policy Name	Description	Event Source	Severity
Action	Failed	AWS: CloudWatchAlarm_ Action_Failed	An Amazon CloudWatch alarm action has failed.	API	Major
Action	InProgress	AWS: CloudWatchAlarm_ Action_InProgress	An Amazon CloudWatch alarm action is in progress.	API	Notice
Action	Succeeded	AWS: CloudWatchAlarm_ Action_Succeeded	An Amazon CloudWatch alarm action has succeeded.	API	Notice
Configuration Update	Configuration Update	AWS: CloudWatchAlarm_ ConfigurationUpdate	A ConfigurationUpdate alarm type is received.	API	Notice
Status Update	Alarm	AWS: CloudWatchAlarm_ StateUpdate_Alarm	A CloudWatch alarm transitions to an "Alarm" state.	API	Major
Status Update	Insufficient Data	AWS: CloudWatchAlarm_ StateUpdate_ InsufficientData	A CloudWatch alarm transitions to an "Insufficient Data" state.	API	Notice
Status Update	OK	AWS: CloudWatchAlarm_ StateUpdate_OK	A CloudWatch alarm transitions to an "OK" state.	API	Healthy

These events are aligned to AWS Account component devices in the following way:

- If the CloudWatch alarm is configured on a device that is discovered in Skylar One, then the event in Skylar One will be aligned with the component device for that instance.
- If the CloudWatch alarm is configured on a device that is either not discovered or not supported by CloudWatch, or if Skylar One cannot determine a correct component device, then that alarm will be aligned to the Account component device.

The "AWS: CloudWatch Alarms Performance" Dynamic Application and related Event Policies are disabled by default. If you want Skylar One to monitor CloudWatch alarms and generate events about them, you must enable the Dynamic Application and Event Policies. You must also configure the Dynamic Application to specify which types of alarms you want to monitor.

For more information about enabling and configuring the "AWS: CloudWatch Alarms Performance" Dynamic Application, see the *Configuring the "AWS: CloudWatch Alarms Performance" Dynamic Application* section. For more information about enabling the CloudWatch alarms Event Policies, see the *Enabling CloudWatch Alarm Events in the ScienceLogic Platform* section.

NOTE: Because the AWS services make new data points available at varying time intervals, there might be a difference in the data points collected by Skylar One when compared to data presented in CloudWatch at a given time. The difference between Skylar One and CloudWatch is typically less than 1%.

NOTE: If an event expires and the CloudWatch alarm in AWS is still in an "Alarm" state, Skylar One will not generate any additional CloudWatch events unless that CloudWatch alarm changes states in AWS.

Creating Custom CloudWatch Metrics

A CloudWatch alarm watches a single metric and performs one or more actions based on the value of the metric relative to a threshold over a number of time periods. A CloudWatch metric consists of the following elements:

- A ***namespace***, such as *AWS/EC2*
- A ***metric name***, such as *CPUUtilization*
- A ***value***, such as *42*
- A ***dimension*** that identifies a particular resource instance, such as *{'Name': 'InstanceId', 'Value': 'i-0a6a989bb8d57b074'}*

NOTE: For a complete list of supported CloudWatch Metrics and Dimensions, see https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CW_Support_For_AWS.html.

The "Amazon Web Services" PowerPack uses the metric dimensions identified in an alarm to associate the alarm message to a particular ScienceLogic component device. The following table lists the services that are currently supported and the dimensions used to associate an alarm to a component device:

AWS Service	Dimension
API Gateway	'ApiName' | 'ApiName | Stage' **NOTE:** ScienceLogic recommends that you create API Gateways with unique names within the same region.
ApplicationELB	'LoadBalancer' | 'TargetGroup'
CloudFront	'DistributionId'
Direct Connect	'ConnectionID'
DynamoDB	'TableName'
EBS	'VolumeId'
ECS	'ClusterName' | 'ServiceName'
EC2	'InstanceId' | 'AutoScalingGroupName'
EKS Cluster	'ClusterName'
ElasticBeanstalk	'EnvironmentName'
ElastiCache	'CacheClusterId' **NOTE:** Alarms for this service will be associated with the component device for the AWS account.
ElasticMapReduce	'JobFlowId'
ELB	'LoadBalancerName'
Glacier	'VaultId' **NOTE:** This service is not supported by CloudWatch. You must define a custom metric and publish the metric to the CloudWatch service using an agent toolkit or the AWS command-line interface.
Lambda	'FunctionName', 'Resource', 'Version', 'Alias', 'Executed Version' **NOTE:** Alarms "across all functions" for this service will be associated with the component device for the AWS account. Alarms "by function name" will be aligned to a specific Lambda function.
NetworkELB	'LoadBalancer' | 'TargetGroup'
OpsWorks	'StackId' | 'InstanceId'

AWS Service	Dimension
RDS	'DBInstanceIdentifier' **NOTE:** Alarms for this service will be associated with the component device for the AWS account.
Redshift	'ClusterIdentifier' **NOTE:** Alarms for this service will be associated with the component device for the AWS account.
Route53	'HealthCheckId'
Shield	'ShieldService' **NOTE:** CloudWatch alarms are available only for Shield Advanced Services.
SNS	'TopicName'
SQS	'QueueName'
StorageGateway	'GatewayId' | 'VolumeId'
S3	'BucketName'
WAF	'WebACLId'

AWS enables users to create custom metrics for these services and then publish those metrics to CloudWatch using the AWS command-line interface (CLI) or an application programming interface (API). The Dynamic Applications in the "Amazon Web Services" PowerPack can then collect data for these custom AWS metrics (which are not in the "AWS" cloud namespace).

NOTE: For the "Amazon Web Services" PowerPack to collect data for these custom metrics, you must enable certain Dynamic Applications that are disabled by default. For more information, see the *Enabling Custom Metrics Collection in the ScienceLogic Platform* section.

When creating a custom metric, it is important that the metric is correctly formed. For Skylar One to align a custom metric to a particular ScienceLogic component device, the following must be true:

- The metric namespace must include the service being tracked.

  For example, *MyVendorName/EC2* would be a valid namespace that the "Amazon Web Services" PowerPack could use to identify the EC2 service for a tracked metric.

- The dimension must include one or more of the dimensions listed in the preceding table. The dimension enables Skylar One to identify which device to associate with the alarm.

  For example, if the dimension included *{'Name': 'InstanceId', 'Value': 'i-0a6a989bb8d57b074'}*, this would identify the EC2 component. Other dimensions are permitted, but 'InstanceId' is necessary to locate the EC2 instance.

  If the component device was an AutoScaleGroup component that is also under the EC2 service, then the dimension might look like this: *{'Name': 'AutoScalingGroupName', 'Value': 'Y1Z55ZJ390UP'}*.

NOTE: If the CloudWatch event cannot align to a particular ScienceLogic component device, it will instead align to the component device for the AWS account.

Configuring CloudWatch to Send Alarms for a Metric

To configure CloudWatch to send alarms to Skylar One for a metric, perform the following steps:

1. Open a browser session and go to aws.amazon.com.
2. Click **[My Account]** and then select *AWS Management Console*. If you are not currently logged in to the AWS site, you will be prompted to log in.
3. In the **AWS Management Console**, under the **Management Tools** heading, click **[CloudWatch]**.
4. Click the **[Browse Metrics]** button.
5. Select the metric for which you want CloudWatch to send alarms.
6. Select the instances for which you want CloudWatch to send alarms for this metric.
7. Click the **[Create Alarm]** button. The **Create Alarm** page is displayed.
8. Specify a Name and Description for the alarm.
9. If you have previously configured an alarm for Skylar One, select the notification list for Skylar One in the ***Send notification to*** field. Otherwise, select the **[New list]** link to the right of the ***Send notification to*** field and supply values in the following fields:
   - ***Send notification to***. Enter a name for the new notification list. If you add additional alarms, you can select the name you enter in this field instead of re-entering the email address.
   - ***Email list***. Enter the email address to which you want CloudWatch notifications sent.
10. Supply values in the other fields in this page as desired.
11. Click the **[Create Alarm]** button.
12. Log in to the email account you configured to receive email from the email alias.
13. Open the confirmation email from Amazon and click the **[Confirm subscription]** link.

Enabling Custom Metrics Collection in Skylar One

AWS enables users to publish their own custom metrics to CloudWatch using the AWS command-line interface (CLI) or an application programming interface (API). The "Amazon Web Services" PowerPack includes Dynamic Applications that collect data for custom AWS metrics (which are not in the "AWS" cloud namespace). However, these Dynamic Applications are disabled by default and must be enabled for use.

To enable these Dynamic Applications:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Click the wrench icon (🔧) for the "AWS: Custom Metrics" Dynamic Application. The **Dynamic Applications Properties Editor** page appears.
3. In the ***Operational State*** field, select *Enabled*.
4. Click the **[Save]** button.
5. Repeat steps 1 - 4 for the "AWS: Custom Metrics Cache" Dynamic Application.

Configuring the "AWS: CloudWatch Alarms Performance" Dynamic Application

The "Amazon Web Services" PowerPack includes an "AWS: CloudWatch Alarms Performance" Dynamic Application that monitors CloudWatch alarms and associates the alarms with the appropriate AWS component devices, if applicable. This Dynamic Application must be enabled if you want Skylar One to generate CloudWatch alarm events.

> **NOTE:** If an appropriate component device does not exist in Skylar One or cannot be determined, the alarm is instead associated with the "Account" component device.

To enable the "AWS: CloudWatch Alarms Performance" Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Locate the "AWS: CloudWatch Alarms Performance" Dynamic Application and then click its wrench icon (🔧). The **Dynamic Applications Properties Editor** page appears.
3. In the ***Operational State*** field, select *Enabled*.
4. Click **[Save]**.

By default, the "AWS: CloudWatch Alarms Performance" Dynamic Application monitors only the "StateUpdate" type of CloudWatch alarms. If you want the Dynamic Application to also monitor "Action" and "ConfigurationUpdate" alarm types, you must configure the Dynamic Application to do so.

To configure the "AWS: CloudWatch Alarms Performance" Dynamic Application to monitor all CloudWatch alarm types:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Locate the "AWS: CloudWatch Alarms Performance" Dynamic Application and then click its wrench icon (). The **Dynamic Applications Properties Editor** page appears.
3. Click the **[Collections]** tab. The **Collection Objects** page appears.
4. On the **Collection Objects** page, locate the "CloudWatch Alarms Collection Success" collection object and then click its wrench icon ().
5. In the ***Snippet*** field, select one of the following options:
    - *cloudwatch_alarms_performance_StateUpdate_only*. This option is selected by default. This snippet will only trigger events for State Update alarms.
    - *cloudwatch_alarms_performance*.This snippet triggers notifications if any alarm configuration is modified.
    - *cloudwatch_alarms_statistics*. This snippet will trigger events for all CloudWatch alarm types (Action, Configuration Update, and State Update).

NOTE: If you want to monitor all CloudWatch alarm types, then select *cloudwatch_alarms_performance* in the ***Snippet*** field.

6. Click **[Save]**. This Dynamic Application will be executed every 1 minute.

Enabling CloudWatch Alarm Events in Skylar One

The "Amazon Web Services" PowerPack also includes several pre-defined event policies for CloudWatch alarms. These Event Policies must be enabled if you want Skylar One to generate CloudWatch alarm events.

To enable the CloudWatch alarms Event Policies:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic user interface).
2. Perform a search for "CloudWatch".

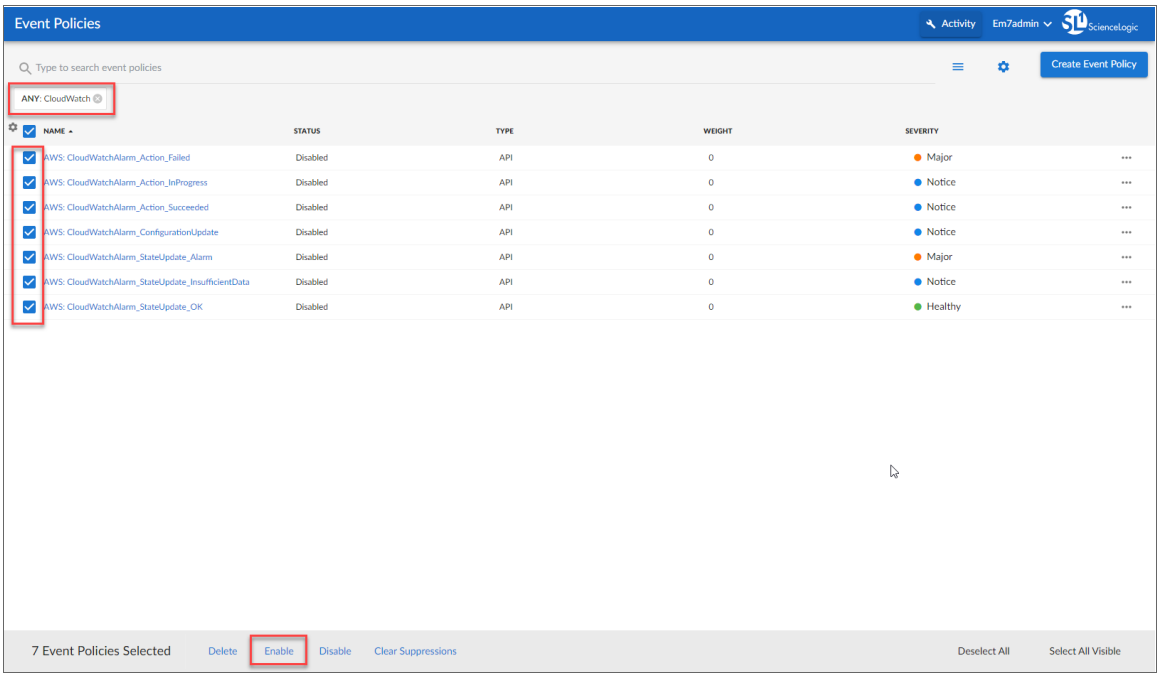

3. Select the check boxes for the events you want to enable.
4. Select **Enable** at the bottom of the screen.

To enable the CloudWatch alarms Event Policies in the Skylar One classic user interface:

1. Go to the **Event Policy Manager** page (Registry > Events > Event Manager).
2. In the ***Event Policy Name*** filter-while-you-type field, type "CloudWatch".
3. Select the check boxes for the events you want to enable.
4. In the ***Select Action*** drop-down field, select *ENABLE these Event Policies*.
5. Click **[Go]**.

Preserving CloudWatch Alarm Event Changes

If you have modified CloudWatch alarm event policies that are included in the "Amazon Web Services" PowerPack, those changes will be overwritten when the PowerPack is updated in your system. If you have modified event policies that are included in the PowerPack, you can:

- Re-implement those changes after each update of the "Amazon Web Services" PowerPack.
- Remove the content from the PowerPack on your system. When the "Amazon Web Services" PowerPack is updated in your system, updated versions of this content will not be installed on your system and your local changes will be preserved.

To remove event policies from the "Amazon Web Services" PowerPack on your system:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the wrench icon ( ) for the "Amazon Web Services" PowerPack. The **Editing PowerPack** page appears.
3. In the left NavBar of the **Editing PowerPack** page, click **[Event Policies]**. The **Embedded Event Policies** and **Available Event Policies** panes appear.
4. In the upper pane, click the delete icon ( ) for each event policy that you want to remove from the "Amazon Web Services" PowerPack on your system.

Chapter 8

AWS Reports

Overview

The following sections describe the reports that are included in the "Amazon Web Services" PowerPack:

This chapter covers the following topics:

AWS Billing Report *124*

AWS Inventory Report *126*

AWS Running Config Report *129*

AWS Billing Report

This report displays service costs for Amazon Web Services. The report includes Total, Monthly, Quarterly, and Annual costs.

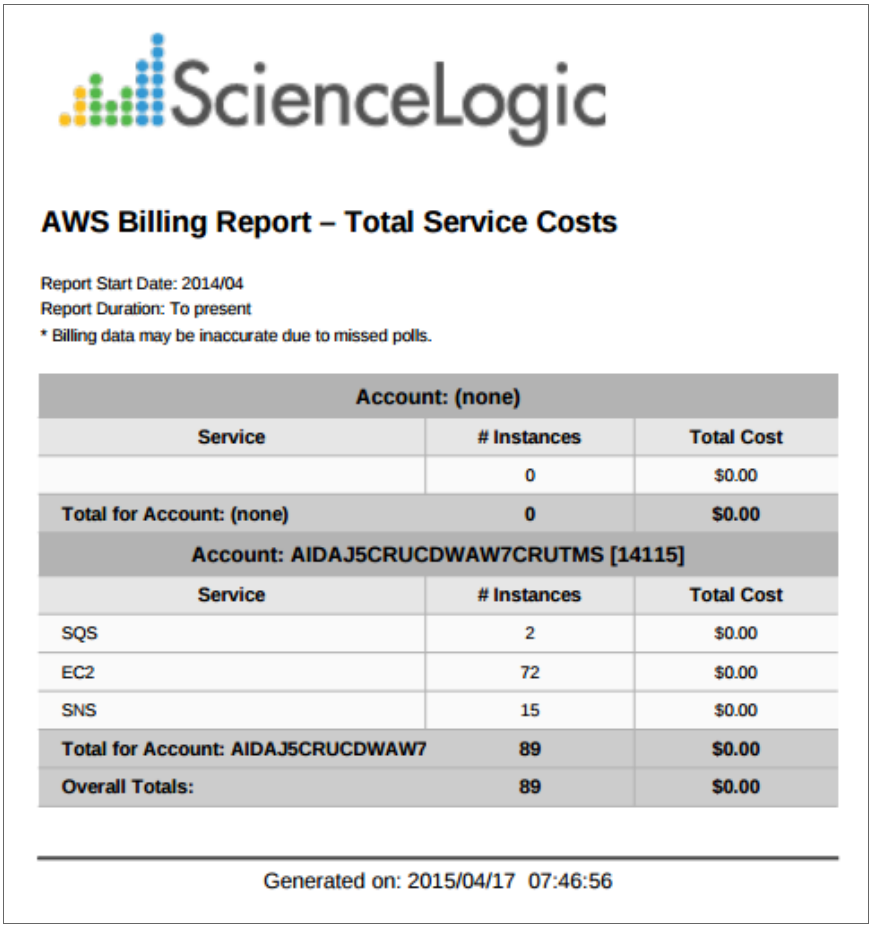

ScienceLogic

AWS Billing Report – Total Service Costs

Report Start Date: 2014/04
Report Duration: To present
* Billing data may be inaccurate due to missed polls.

Account: (none)		
Service	**# Instances**	**Total Cost**
	0	$0.00
Total for Account: (none)	**0**	**$0.00**
Account: AIDAJ5CRUCDWAW7CRUTMS [14115]		
Service	**# Instances**	**Total Cost**
SQS	2	$0.00
EC2	72	$0.00
SNS	15	$0.00
Total for Account: AIDAJ5CRUCDWAW7	**89**	**$0.00**
Overall Totals:	**89**	**$0.00**

Generated on: 2015/04/17 07:46:56

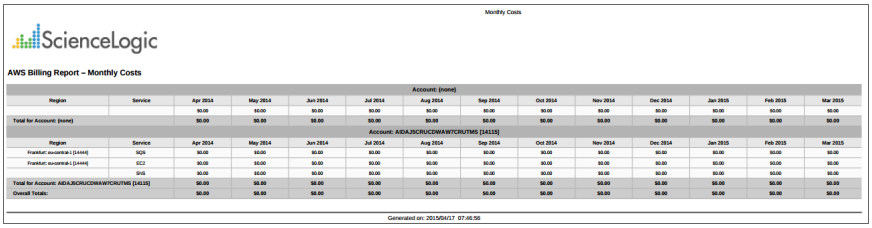

Monthly Costs

ScienceLogic

AWS Billing Report – Monthly Costs

Account: (none)

Region	Service	Apr 2014	May 2014	Jun 2014	Jul 2014	Aug 2014	Sep 2014	Oct 2014	Nov 2014	Dec 2014	Jan 2015	Feb 2015	Mar 2015
		$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00
Total for Account: (none)		$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00

Account: AIDAJ5CRUCDWAW7CRUTMS [14115]

Region	Service	Apr 2014	May 2014	Jun 2014	Jul 2014	Aug 2014	Sep 2014	Oct 2014	Nov 2014	Dec 2014	Jan 2015	Feb 2015	Mar 2015
Frankfurt: eucentral1 [14444]	SQS	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00
Frankfurt: eucentral1 [14444]	EC2	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00
	SNS	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00
Total for Account: AIDAJ5CRUCDWAW7CRUTMS [14115]		$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00
Overall Totals:		$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00	$0.00

Generated on: 2015/04/17 07:46:56

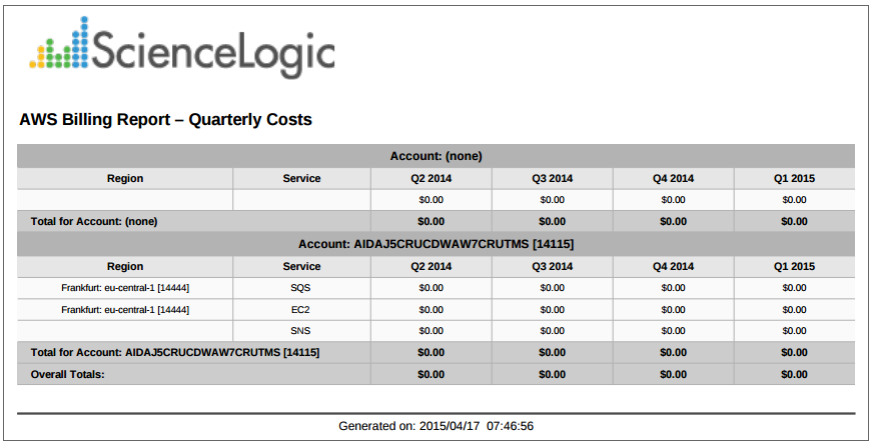

AWS Billing Report – Quarterly Costs

Account: (none)					
Region	**Service**	**Q2 2014**	**Q3 2014**	**Q4 2014**	**Q1 2015**
		$0.00	$0.00	$0.00	$0.00
Total for Account: (none)		**$0.00**	**$0.00**	**$0.00**	**$0.00**
Account: AIDAJ5CRUCDWAW7CRUTMS [14115]					
Region	**Service**	**Q2 2014**	**Q3 2014**	**Q4 2014**	**Q1 2015**
Frankfurt: eu-central-1 [14444]	SQS	$0.00	$0.00	$0.00	$0.00
Frankfurt: eu-central-1 [14444]	EC2	$0.00	$0.00	$0.00	$0.00
	SNS	$0.00	$0.00	$0.00	$0.00
Total for Account: AIDAJ5CRUCDWAW7CRUTMS [14115]		**$0.00**	**$0.00**	**$0.00**	**$0.00**
Overall Totals:		**$0.00**	**$0.00**	**$0.00**	**$0.00**

Generated on: 2015/04/17 07:46:56

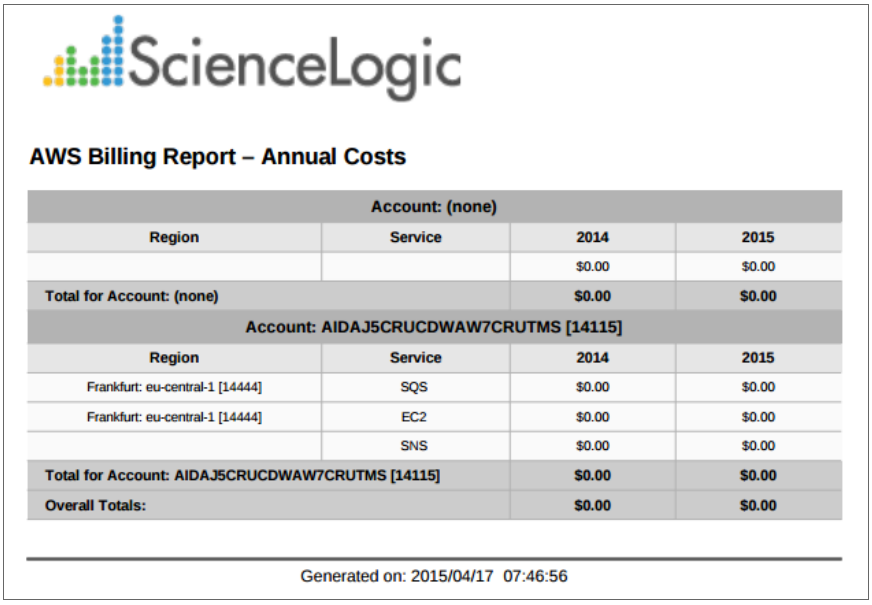

AWS Billing Report – Annual Costs

Account: (none)			
Region	**Service**	**2014**	**2015**
		$0.00	$0.00
Total for Account: (none)		**$0.00**	**$0.00**
Account: AIDAJ5CRUCDWAW7CRUTMS [14115]			
Region	**Service**	**2014**	**2015**
Frankfurt: eu-central-1 [14444]	SQS	$0.00	$0.00
Frankfurt: eu-central-1 [14444]	EC2	$0.00	$0.00
	SNS	$0.00	$0.00
Total for Account: AIDAJ5CRUCDWAW7CRUTMS [14115]		**$0.00**	**$0.00**
Overall Totals:		**$0.00**	**$0.00**

Generated on: 2015/04/17 07:46:56

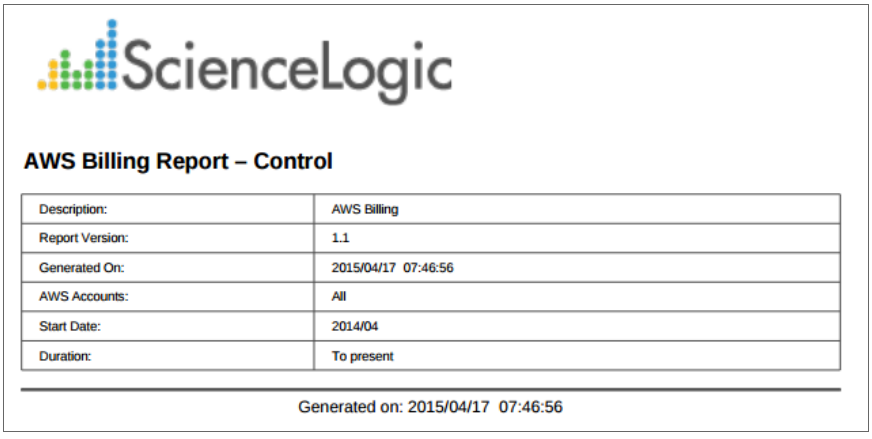

AWS Billing Report – Control

Description:	AWS Billing
Report Version:	1.1
Generated On:	2015/04/17 07:46:56
AWS Accounts:	All
Start Date:	2014/04
Duration:	To present

Generated on: 2015/04/17 07:46:56

The following input options are available when generating the report (Reports > Run Report > Cloud > AWS Billing):

- ***AWS Accounts***. Select the AWS Account(s) for which you want to generate the report. The *All Accounts* checkbox is selected by default. De-selecting this checkbox allows you to select one or more specific accounts for which to generate a report.
- ***Report Span***. Select a span from one to 36 months for the report, or specify a specific starting date for the report.

This description covers the latest version of this report as shipped by ScienceLogic. This report might have been modified on your Skylar One system.

AWS Inventory Report

This report displays an inventory of AWS instance counts. The report includes the number of each kind of instance in every zone associated with the chosen accounts. It also includes a count of each EC2 instance size in each zone.

AWS Inventory Report – Instance Counts

Organization: Pittock [193]										
Account: AIDAJ5CRUCDWAW7CRUTMS [14115]										
Level1: CloudFront Service [14120]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
d12tibk6qbt264.cloudfront.net [14150]	0	0	0	1	1	0	0	0	0	0
Totals for Level1: CloudFront Service [14120]	**0**	**0**	**0**	**1**	**1**	**0**	**0**	**0**	**0**	**0**
Level1: Frankfurt: eu-central-1 [14444]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
eu-central-1 Glacier Service [14467]	1	0	0	0	0	0	0	0	0	0
eu-central-1 VPC Service [14447]	0	0	0	0	0	0	0	2	0	0
eu-central-1a [14446]	0	0	0	0	0	0	0	0	0	1
Totals for Level1: Frankfurt: eu-central-1 [14444]	**1**	**0**	**0**	**0**	**0**	**0**	**0**	**2**	**0**	**1**
Level1: Ireland: eu-west-1 [14117]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
eu-west-1 Glacier Service [14129]	1	0	0	0	0	0	0	0	0	0
eu-west-1 CloudTrail Service [14346]	0	0	0	0	0	1	0	0	0	0
eu-west-1 ELB Service [14124]	0	0	0	0	0	0	1	0	0	7
eu-west-1 SNS Service [14123]	0	0	0	0	0	0	0	0	1	0
eu-west-1 VPC Service [14130]	0	0	0	0	0	0	0	9	0	0
Totals for Level1: Ireland: eu-west-1 [14117]	**1**	**0**	**0**	**0**	**0**	**1**	**1**	**9**	**1**	**7**
Level1: N. Virginia: us-east-1 [14118]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
us-east-1 Auto Scale Service [14138]	0	2	1	0	0	0	2	0	0	38
us-east-1 CloudTrail Service [14139]	0	0	0	0	0	1	0	0	0	0
us-east-1b [14133]	0	0	0	0	0	0	0	0	0	0
us-standard S3 Service [14137]	0	0	0	0	0	0	0	0	0	0
us-east-1 SQS Service [14340]	0	0	0	0	0	0	0	0	8	0
us-east-1 VPC Service [14141]	0	0	0	0	0	0	0	8	0	0
Totals for Level1: N. Virginia: us-east-1 [14118]	**0**	**2**	**1**	**0**	**0**	**1**	**2**	**8**	**8**	**38**
Level1: Oregon: us-west-2 [14119]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
us-west-2 Auto Scale Service [14147]	0	1	1	0	0	0	0	0	0	9
us-west-2 CloudTrail Service [14148]	0	0	0	0	0	1	0	0	0	0
us-west-2 S3 Service [14146]	0	0	0	0	0	0	0	0	0	0
us-west-2 SQS Service [14336]	0	0	0	0	0	0	0	0	4	0
us-west-2 VPC Service [14149]	0	0	0	0	0	0	0	3	0	0
Totals for Level1: Oregon: us-west-2 [14119]	**0**	**1**	**1**	**0**	**0**	**1**	**0**	**3**	**4**	**9**
Level1: Route 53 Service [14116]										
Zone	**Glacier**	**Launch Con**	**AS Group**	**Web Dist**	**udFront Ori**	**CloudTrail**	**ELB**	**Subnet**	**SNS**	**EC2**
mapmycloud.net [14121]	0	0	0	0	0	0	0	0	0	0
Totals for Level1: Route 53 Service [14116]	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
Totals for Account: AIDAJ5CRUCDWAW7CRUTMS [14115]	**2**	**3**	**2**	**1**	**1**	**3**	**3**	**22**	**13**	**55**
Totals for Organization: Pittock [193]	**2**	**3**	**2**	**1**	**1**	**3**	**3**	**22**	**13**	**55**
Overall Totals:	**2**	**3**	**2**	**1**	**1**	**3**	**3**	**22**	**13**	**55**

Generated on: April 17th, 2015 at 7:46am

AWS Inventory Report – EC2 Instance Details

Organization: Pittock [193]								
Account: AIDAJ5CRUCDWAW7CRUTMS [14115]								
Level1: Frankfurt: eu-central-1 [14444]								
Zone	**M1.small**	**M3.large**	**T1.micro**	**T2.small**	**T2.micro**	**C3.large**	**M3.xlarge**	**M3.m**
eu-central-1a [14446]	0	0	0	0	1	0	0	
Totals for Level1: Frankfurt: eu-central-1 [14444]	**0**	**0**	**0**	**0**	**1**	**0**	**0**	
Level1: Ireland: eu-west-1 [14117]								
Zone	**M1.small**	**M3.large**	**T1.micro**	**T2.small**	**T2.micro**	**C3.large**	**M3.xlarge**	**M3.m**
eu-west-1a [14126]	0	1	2	0	0	0	0	
eu-west-1c [14127]	0	0	2	0	0	0	0	
eu-west-1b [14125]	0	0	2	0	0	0	0	
Totals for Level1: Ireland: eu-west-1 [14117]	**0**	**1**	**6**	**0**	**0**	**0**	**0**	
Level1: N. Virginia: us-east-1 [14118]								
Zone	**M1.small**	**M3.large**	**T1.micro**	**T2.small**	**T2.micro**	**C3.large**	**M3.xlarge**	**M3.m**
us-east-1a [14134]	4	4	3	11	1	0	0	
us-east-1e [14135]	0	0	0	0	3	0	1	
us-east-1b [14133]	1	0	4	0	0	0	0	
us-east-1c [14136]	2	0	2	0	0	1	0	
Totals for Level1: N. Virginia: us-east-1 [14118]	**7**	**4**	**9**	**11**	**4**	**1**	**1**	
Level1: Oregon: us-west-2 [14119]								
Zone	**M1.small**	**M3.large**	**T1.micro**	**T2.small**	**T2.micro**	**C3.large**	**M3.xlarge**	**M3.m**
us-west-2c [14145]	0	0	4	0	0	0	0	
us-west-2a [14144]	0	0	3	0	0	0	0	
us-west-2b [14143]	0	0	0	0	0	0	0	
Totals for Level1: Oregon: us-west-2 [14119]	**0**	**0**	**7**	**0**	**0**	**0**	**0**	
Totals for Account: AIDAJ5CRUCDWAW7CRUTMS [14115]	**7**	**5**	**22**	**11**	**5**	**1**	**1**	
Totals for Organization: Pittock [193]	**7**	**5**	**22**	**11**	**5**	**1**	**1**	
Overall Totals:	**7**	**5**	**22**	**11**	**5**	**1**	**1**	

Generated on: April 17th, 2015 at 7:46am

The following input options are available when generating the report (Reports > Run Report > Others > AWS Inventory):

- ***Organizations***. Select the organization for which you want to generate the report. The *All Organizations* checkbox is selected by default. De-selecting this checkbox allows you to select one or more specific organizations for which to generate a report.
- ***AWS Accounts***. Select the AWS Account(s) for which you want to generate the report. The *All Accounts* checkbox is selected by default. De-selecting this checkbox allows you to select one or more specific accounts for which to generate a report.
- ***Filter on EC2 Instance Config Data***. Select the EC2 instances that will be included in the report based on the configuration data reported for each EC2 instance:
  - Choose up to four configuration parameters for EC2 instances.
  - For each selected configuration parameter, enter a value to match against and select how that value should be matched.
  - In the ***Comparison Operator*** field, select whether an EC2 instance must match all configuration parameters (*and*) or only one configuration parameter (*or*) to be included on the report.
- ***Report Options***. Select the *Include Terminated Instances* checkbox to include all terminated instances.

This description covers the latest version of this report as shipped by ScienceLogic. This report might have been modified on your Skylar One system.

AWS Running Config Report

This report displays the running config of all AWS instances for one to all organizations across a number of AWS billing accounts.

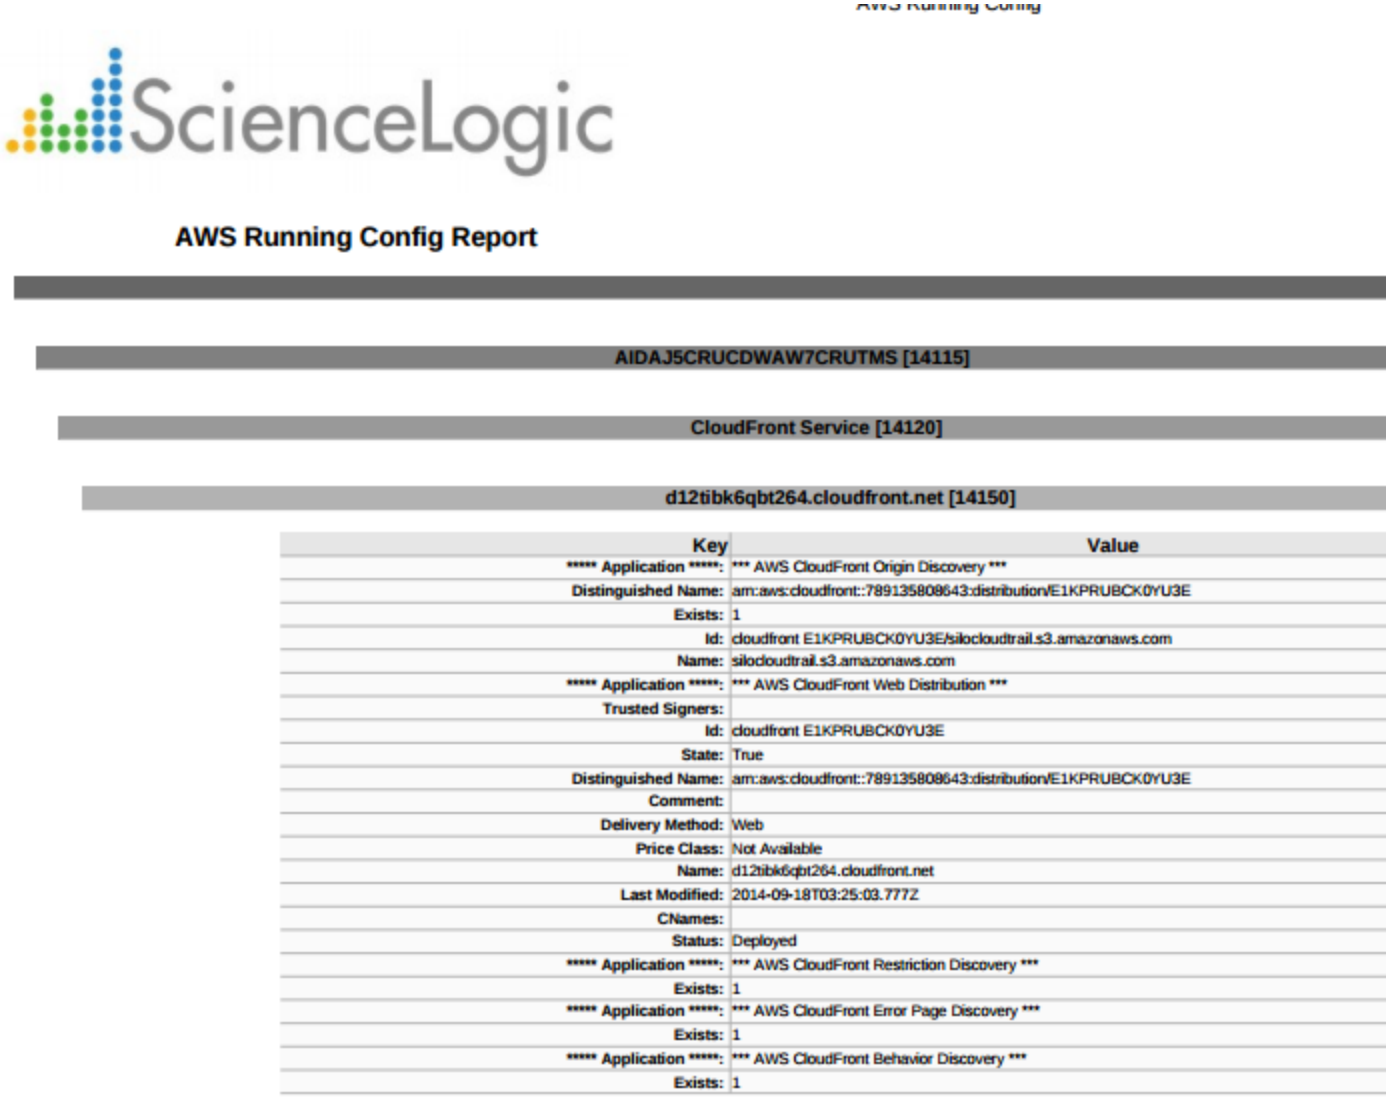

AWS Running Config Report

AIDAJ5CRUCDWAW7CRUTMS [14115]

CloudFront Service [14120]

d12tibk6qbt264.cloudfront.net [14150]

Key	Value
***** Application *****:	*** AWS CloudFront Origin Discovery ***
Distinguished Name:	arn:aws:cloudfront::789135808643:distribution/E1KPRUBCK0YU3E
Exists:	1
Id:	cloudfront E1KPRUBCK0YU3E/silocloudtrail.s3.amazonaws.com
Name:	silocloudtrail.s3.amazonaws.com
***** Application *****:	*** AWS CloudFront Web Distribution ***
Trusted Signers:	
Id:	cloudfront E1KPRUBCK0YU3E
State:	True
Distinguished Name:	arn:aws:cloudfront::789135808643:distribution/E1KPRUBCK0YU3E
Comment:	
Delivery Method:	Web
Price Class:	Not Available
Name:	d12tibk6qbt264.cloudfront.net
Last Modified:	2014-09-18T03:25:03.777Z
CNames:	
Status:	Deployed
***** Application *****:	*** AWS CloudFront Restriction Discovery ***
Exists:	1
***** Application *****:	*** AWS CloudFront Error Page Discovery ***
Exists:	1
***** Application *****:	*** AWS CloudFront Behavior Discovery ***
Exists:	1

The following input options are available when generating the report (Reports > Run Report > Others > AWS Running Config):

- ***Organizations***. Select one, multiple, or all organizations to include in the report.
    - *All Organizations*. This checkbox is selected by default. De-selecting this checkbox allows you to select one or more specific organizations for the report.

    - *Organizations*. If you unchecked the ***All Organizations*** checkbox, select one or more organizations to include in the report.
- ***AWS Accounts***. Select one, multiple, or all AWS Accounts to include in the report.
    - *All Accounts*. This checkbox is selected by default. De-selecting this checkbox allows you to select one or more specific AWS accounts for the report.
    - *Accounts*. If you unchecked the ***All Accounts*** checkbox, select one or more AWS Accounts to include in the report.
- ***Filter on EC2 Instance Config Data***. Select the EC2 instances that will be included on the report based on the configuration data reported for each EC2 instance:
    - Choose up to four configuration parameters for EC2 instances.
    - For each selected configuration parameter, enter a value to match against and select how that value should be matched.
    - In the ***Comparison Operator*** field, select whether an EC2 instance must match all configuration parameters (*and*) or only one configuration parameter (*or*) to be included in the report.
- ***Report Options***. Select the *Include Terminated Instances* checkbox to include all terminated instances.

This description covers the latest version of this report as shipped by ScienceLogic. This report might have been modified on your Skylar One system.

Chapter

9

AWS Dashboards

Overview

The following sections describe how to install the "Amazon Web Services: Dashboards" PowerPack and a description of each dashboard that is included in the PowerPack:

This chapter covers the following topics:

Installing the Amazon Web Services: Dashboards PowerPack 131

AWS Account Billing Dashboard 132

AWS Health Status Dashboard 132

Configuring the AWS Dashboards 133

AWS Service Instance Performance Dashboards 133

Installing the Amazon Web Services: Dashboards PowerPack

To view the Amazon Web Services dashboards in Skylar One, you must first install the "Amazon Web Services: Dashboards" PowerPack.

NOTE: The AWS dashboards have a default ***Access Control*** setting of "Private", which means they can be viewed only by an administrator. For more information about dashboard access settings, see the ***Dashboards*** manual.

To install the PowerPack:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the actions icon, then select *Install PowerPack*. The **Imported PowerPacks** modal page appears.
3. Use the search filter in the ***PowerPack Name*** column heading to locate the PowerPack you want to install. To do so, enter text to match, including special characters, and the **Imported PowerPacks** modal page displays only PowerPacks that have a matching name.
4. Click the lightning-bolt icon ( ) for the PowerPack that you want to install.
5. The **Install PowerPack** modal page appears. To install the PowerPack, click **[Install]**.
6. The PowerPack now appears in the **PowerPack Manager** page. The contents of the PowerPack are automatically installed in your Skylar One System.

AWS Account Billing Dashboard

The AWS Account Billing Dashboard displays:

- A pie chart that shows the estimated billing amount for each service over the selected time period.
- A performance graph that shows the estimated billing amount for the selected service, over time. To select a service, click on the pie-chart segment for that service.
- A table that shows the currently active AWS events.
- A time span selector that controls the amount of data shown in the pie chart and the performance graph.
- An organization selector that limits the data in the pie chart and performance graph to include only instances associated with the selected organizations.

AWS Health Status Dashboard

The AWS Health Status Dashboard displays:

- A traffic light widget that displays a list of AWS services. To populate the other widgets in this dashboard, select a service.
- A tear-sheet widget that displays information and links for the selected service.
- A service health widget, that displays log messages about the health of the service.
- A table that displays currently active events for the service.

- An organization selector and a device group selector that control which services are shown in the traffic light widget.

Configuring the AWS Dashboards

The AWS Account Billing and AWS Health Status dashboards must have their (base) Custom Table widgets manually configured to filter only AWS service-specific events. To do this:

1. Go to Dashboards > Classic Dashboards and select AWS Account Billing, or in the Skylar One classic user interface go to Dashboards and select AWS Account Billing.
2. Click the down-arrow in the upper-right of the AWS Events widget, and then select *Configure* from the **Options** menu. The **Widget Configuration** modal page appears.
3. In the ***Device Class*** filter, enter "AWS" to show only AWS device classes.
4. Control-click on the following items in the ***Device Class*** field:
    - AWS DDB Service
    - AWS EC2 Service
    - AWS ELB Service
    - AWS EMR Service
    - AWS RDS Service
    - AWS SNS Service
    - AWS SQS Service
    - AWS Storage Gateway Service
5. Click the **[Save]** button.
6. Repeat steps 1 - 5 for the AWS Health Status dashboard.

AWS Service Instance Performance Dashboards

The "Amazon Web Services: Dashboards" PowerPack includes a dashboard for each service type. Each dashboard displays performance metrics for instances of an AWS service. The following dashboards are included:

- AWS Application ELB Performance
- AWS Classic ELB Performance
- AWS DDB Performance
- AWS EBS Performance
- AWS EC2 Performance
- AWS EMR Performance
- AWS Network ELB Performance
- AWS RDS Performance

- AWS SQS Performance
- AWS Storage Gateway Performance

Each performance dashboard includes:

- A traffic light widget that shows the status of all instances for the service.
- Four performance graphs that show applicable metrics when you select an instance from the traffic light widget.
- A time span selector that controls the amount of data shown in the performance graphs.
- An organization selector and device group selector that control which instances are shown in the traffic light widget.

Chapter

10

AWS Run Book Actions and Automations

Overview

The following sections describe the Run Book Action and Automation policies that are included in the "Amazon Web Services" PowerPack and how to use them:

This chapter covers the following topics:

About the Run Book Actions and Automations 135

Disabling EC2 and EBS Instances by EC2 Tag 137

Discovering EC2 Instances by Public or Private IP Address 140

Aligning AWS Regions to the AWS Region Device Class 145

Vanishing Terminated or Terminating EC2 Instances 145

About the Run Book Actions and Automations

The "Amazon Web Services" PowerPack includes Run Book Action and Automation policies that can be used to:

- Automatically disable EC2 and EBS devices based on EC2 tags collected from AWS
- Automatically create and start a discovery session for the public or private IP address of an EC2 instance after a component and physical device are merged
- Automatically move an EC2 instance to a vanished state if the EC2 instance is in a terminating or terminated state
- Align AWS region device classes with the correct AWS Region

The following table describes the automation policies and what they do:

Policy Name	Result
AWS: Account Creation	Skylar One creates a virtual device for an AWS account.
AWS: Disable EBS Instances by EC2 Tag	If a component device belongs to the AWS EBS Volumes device group and has an EC2 tag, Skylar One disables the device.
AWS: Disable EC2 and EBS Instances by EC2 Tag	If a component device belongs to either the AWS EBS Volumes or AWS EC2 Instances device group and has an EC2 tag, Skylar One disables the device.
AWS: Disable or Discover EC2 Instances	Skylar One automatically discovers EC2 instances by public or private IP address. Additionally, if a component device belongs to the AWS EC2 Instances device group and has an EC2 tag, Skylar One disables the device.
AWS: Discover EC2 Instances	Skylar One automatically discovers EC2 instances by public or private IP address.
AWS: EKS Cluster Creation	Skylar One automatically discovers EKS Clusters when an AWS EKS Cluster is configured.
AWS: Merge with EC2	If Skylar One determines that the IP address of a physical device matches a custom attribute added to an EC2 Instance component device, Skylar One merges the devices.
AWS: Organization Creation	Skylar One creates a virtual device for an AWS organization.
AWS: RDS DB Instance Device Class Alignment	Skylar One aligns the correct RDS device class the RDS Instance.
AWS: Region Device Class Alignment	If a Region is aligned to an incorrect Region device class, Skylar One will align the Region to the correct device class.
AWS: Vanish Terminated EC2 Instances	If a device belongs to the AWS EC2 Instances device group and is in a terminated or terminating state, Skylar One un-merges the EC2 Instance and physical device (if applicable), clears the device's associated events, and then moves the device to a vanished state.

NOTE: The automation policies in the "Amazon Web Services" PowerPack are disabled by default. To use these automations, you must enable the automation policies and optionally modify the parameters in the automation actions.

NOTE: To use the automation policies in the "Amazon Web Services" PowerPack, the AWS EBS Volumes and AWS EC2 Instances device groups must already be created and populated.

Disabling EC2 and EBS Instances by EC2 Tag

The automation described in this section disables EC2 and EBS devices based on EC2 tags. This can be set up in the "AWS: Disable Instance by Tag" Run Book Automation, so if an EBS or EC2 instance has the tag(s) you specify, Skylar One will disable the device.

The automation for disabling EC2 and EBS instances includes two automation actions that are executed in the following order:

- ***AWS: Get EC2 Instance Configuration***. This action requests information from the AWS API about the EC2 instance that triggered the automation action or the EC2 instance associated with the EBS instance that triggered the automation action. Information about the EC2 instance associated with an EBS instance is returned only if one EC2 instance is associated with the EBS instance.
- ***AWS: Disable Instance By Tag***. This action compares the information collected by the ***AWS: Get EC2 Instance Configuration*** automation action with a pre-defined list of key/value pairs. If an AWS tag matches a key/value pair, the triggering device is disabled.

The "Amazon Web Services" PowerPack includes three automation policies that trigger these actions:

- ***AWS: Disable EC2 and EBS Instances by EC2 Tag***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" or the "AWS: EBS Instance Configuration" Dynamic Applications are aligned (the members of the AWS EC2 Instances and AWS EBS Volumes device groups). The automation policy triggers when the "Component Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to disable EC2 and EBS instances by EC2 tag, but do not want to enable automated discovery of EC2 instances by public or private IP address.
- ***AWS: Disable or Discover EC2 Instances***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" Dynamic Application is aligned (the members of the AWS EC2 Instances). The automation policy triggers when the "Component Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to disable EC2 instances by EC2 tag *and* want to enable automated discovery of EC2 instances by public or private IP address. This automation policy is configured to run both processes in the correct order for EC2 instances. If you enable this automation policy and want to automatically disable associated EBS instances, you must also enable the ***AWS: Disable EBS Instances by EC2 Tag*** automation policy.
- ***AWS: Disable EBS Instances by EC2 Tag***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" Dynamic Application is aligned (the members of the AWS EC2 Instances). The automation policy triggers when the "Component Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to disable EC2 instances by EC2 tag, want to enable automated discovery of EC2 instances by public or private IP address, and want to disable EBS instances by EC2 tag.

To use this automation, you must:

- *Modify the parameters of the automation actions (optional)*
- *Enable the Component Device Record Created event policy*

- *Enable the automation policies*
- *Configure your system to preserve these changes*

Modifying the Parameters of the Automation Actions

The snippet for the ***AWS: Disable Instance by Tag*** automation action includes the pre-defined list of key/value pairs with which the tags collected from the AWS API are compared. You must modify this list to include the key/value pairs that you want to use to disable EC2 instances.

To modify the parameters for the ***AWS: Disable Instance by Tag*** automation action:

1. Go to the **Action Policy Manager** page (Registry > Run Book > Actions).
2. Click the wrench icon (🔧) for the ***AWS: Disable Instance By Tag*** automation action.
3. In the ***Snippet Code*** field, locate and edit the following line:

   ```
   DISABLE_TAGS = [('ExampleKey','ExampleValue')]
   ```

   The line must be in the following format, with each key and each value inside single-quotes and each key/value pair comma-separated inside parentheses, with commas separating each key/value pair.

   ```
   DISABLE_TAGS = [('Key','Value'),('Key','Value'),...,('Key','Value')]
   ```

   For example, suppose you want to disable an EC2 instance where the "Environment" key is either "dev" or "test" or the "Owner" key is "Sales". You would update the line so it looks like this:

   ```
   DISABLE_TAGS = [('Environment','dev'),('Environment','test'),
   ('Owner','Sales')]
   ```

4. Click the **[Save]** button.

Enabling the Component Device Record Created Event Policy

To enable the "Component Device Record Created" event policy:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic user interface).
2. Click the Actions menu (•••) for the "Component Device Record Created" event policy and select *Edit*.
3. In the **Event Policy Editor** page, click on the ***Enable Event Policy*** toggle to enable the event policy.
4. Click **[Save]**.

To enable the "Component Device Record Created" event policy in the Skylar One classic user interface:

1. Go to the **Event Policy Manager** page (Registry > Events > Event Manager).
2. Click the wrench icon (🔧) for the "Component Device Record Created" event policy.
3. In the ***Operational State*** field, select *Enabled*.
4. Click **[Save]**.

To prevent this change from being overwritten when the PowerPacks installed on the system are updated, you can enable the ***Selective PowerPack Field Protection*** option. To enable this option:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. Check the ***Enable Selective PowerPack Field Protection*** checkbox.
3. Click **[Save]**.

Enabling the Automation Policies

To enable one or more automation policies in the "Amazon Web Services" PowerPack:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. Click the wrench icon ( ) for the automation policy you want to enable.
3. In the ***Policy State*** field, select *Enabled*.
4. Click **[Save]**.

Preserving Automation Changes

If you have modified automation actions and policies that are included in the "Amazon Web Services" PowerPack, those changes will be overwritten when the PowerPack is updated in your system. If you have modified automation actions and policies that are included in the PowerPack, you can:

- Re-implement those changes after each update of the "Amazon Web Services" PowerPack.
- Remove the content from the PowerPack on your system before you update it. When the "Amazon Web Services" PowerPack is updated in your system, updated versions of this content will not be installed on your system and your local changes will be preserved.

To remove automation actions or automation policies content from the "Amazon Web Services" PowerPack on your system:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the wrench icon ( ) for the "Amazon Web Services" PowerPack. The **Editing PowerPack** page appears.
3. In the left NavBar of the **Editing PowerPack** page, select the type of content you want to remove:
    - To remove an automation action, click **Run Book Actions**. The **Embedded Run Book Actions** and **Available Run Book Actions** panes appear.
    - To remove an automation policy, click **Run Book Policies**. The **Embedded Run Book Policies** and **Available Run Book Policies** panes appear.
4. In the upper pane, click the delete icon ( ) for each automation action or automation policy that you want to remove from the "Amazon Web Services" PowerPack on your system.

Discovering EC2 Instances by Public or Private IP Address

The automation in this section automatically creates and starts a discovery session for the public or private IP address of an EC2 instance after a component and physical device are merged. If Skylar One determines that the IP address of a physical device matches a custom attribute added to an EC2 instance component device, Skylar One merges the devices.

The automation for discovering EC2 instances by public or private IP addresses includes three automation actions that are executed in the following order:

- ***AWS: Get EC2 Instance Configuration***. This action requests information from the AWS API about the EC2 instance that triggered the automation action.
- ***AWS: Discover from EC2 IP***. This action uses the IP address and port information in the response from the AWS API to create and run a discovery session. This action also adds a custom attribute to the EC2 component device record that can be used to match a newly discovered device to the EC2 instance.
- ***AWS: Merge Physical with Component***. This action matches the IP address of a physical device with the custom attribute added to EC2 component devices by the ***AWS: Discover from EC2 IP*** automation action. If a match is found, the matching EC2 component device is merged with the physical device.

The "Amazon Web Services" PowerPack includes three automation policies that trigger these actions:

- ***AWS: Discover EC2 Instances***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" Dynamic Application is aligned (the members of the AWS EC2 Instances). The automation policy triggers when the "Component Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to enable automated discovery of EC2 instances by public or private IP address but do not want disable EC2 and EBS instances by EC2 tag.
- ***AWS: Disable or Discover EC2 Instances***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" Dynamic Application is aligned (the members of the AWS EC2 Instances). The automation policy triggers when the "Component Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to disable EC2 instances by EC2 tag **and** want to enable automated discovery of EC2 instances by public or private IP address. This automation policy is configured to run both in the correct order for EC2 instances.
- ***AWS: Merge with EC2***. If enabled, this automation policy can trigger for any device. The automation policy triggers when the "Device Record Created" event is active on the matching devices, immediately after the devices are discovered in the system. Enable this automation policy if you want to enable automated discovery of EC2 instances by public or private IP address.

To use this automation, you must:

- *Modify the parameters of the automation actions (optional)*
- *Enable the Component Device Record Created event policy*
- *Enable the Device Record Created event policy*

- *Enable the automation policies*
- *Configure your system to preserve these changes*

Modifying the Parameters of the Automation Actions

The snippet for the ***AWS: Discover from EC2 IP*** automation action includes parameters that define how the automation action creates discovery sessions. You can edit the following lines in the ***Snippet Code*** field of the ***AWS: Discover from EC2 IP*** automation action to change these parameters:

- `EC2_IP_ATTRIBUTE = 'PrivateIpAddress'`

  The attribute returned by the AWS API for EC2 instances that contains the IP address to use in the discovery session. By default, the private IP address is used. To use the public IP address of the EC2 instance, change this line to:

  `EC2_IP_ATTRIBUTE = 'PublicIpAddress'`

- `EXTRA_SCAN_PORTS = ["21","22","23","25","80","443","5985","5986"]`

  The list of TCP ports used in the discovery session includes any TCP ports that are specified explicitly in the security group associated with the EC2 instance, plus any TCP ports included in the EXTRA_SCAN_PORTS parameter. You can add or remove ports from this default list. For example, if you wanted to remove TCP port 21 from this list and add TCP port 53, you would change this line to:

  `EXTRA_SCAN_PORTS = ["22","23","25","53","80","443","5985","5986"]`

  NOTE: The EXTRA_SCAN_PORTS parameter must be populated if there are no rules for specific ports in the security group associated with the EC2 instance.

- `AUTO_INCLUDE_CREDS = True`

  If the AUTO_INCLUDE_CREDS parameter is "True", the automation will automatically add credentials to the discovery session. A credential will be added automatically if it meets one of the following requirements:

  - The credential is an SNMP credential, the Security Group associated with the EC2 instance includes a rule that allows access to UDP port 161, and the credential is explicitly aligned within the organization of the EC2 instance.
  - The credential is an SNMP credential, the Security Group associated with the EC2 instance includes a rule that allows access to UDP port 161, the credential is associated with all organizations in the system, and the INCLUDE_ALL_ORG_CREDS parameter is "True".
  - The credential is not an SNMP credential or an LDAP/AD credential, the TCP port used by the credential is included in the list of TCP ports for the discovery session (the credential is specified explicitly in the security group associated with the EC2 instance or is included in the EXTRA_SCAN_PORTS parameter), and the credential is explicitly aligned with in the organization of the EC2 instance.

- The credential is not an SNMP credential or an LDAP/AD credential, the TCP port used by the credential is included in the list of TCP ports for the discovery session (the credential is specified explicitly in the security group associated with the EC2 instance or is included in the EXTRA_SCAN_PORTS parameter), and the INCLUDE_ALL_ORG_CREDS parameter is "True".

To disable the automatic alignment of credentials to the discovery session, change this line to:

```
AUTO_INCLUDE_CREDS = False
```

- `INCLUDE_ALL_ORG_CREDS = True`

  If INCLUDE_ALL_ORG_CREDS is "True" and the AUTO_INCLUDE_CREDS parameter is "True", credentials that are aligned with all organizations (credentials that do not have an explicit organization alignment) are automatically included in the discovery session when that credential meets the other requirements for being automatically included in the discovery session.

- `EXTRA_CREDS = ""`

  In addition to the credentials that are automatically included in the discovery sessions based on open ports, you can optionally specify a string of comma-separated credential IDs for credentials that will be included in every discovery session created by this automation. For example, if you wanted to include credentials with IDs 10 and 13 in every discovery session created by this automation, you would change this line to:

  ```
  EXTRA_CREDS = "10,13"
  ```

- `DISCOVER_NON_SNMP = "1"`

  If DISCOVER_NON_SNMP is set to "1", discovery sessions created by this automation will be configured to discover non-SNMP devices. If you want the discovery sessions created by this automation to discover only SNMP devices, change this line to:

  ```
  DISCOVER_NON_SNMP = "0"
  ```

- `TEMPLATE_NAME = ""`

  If you specify a device template name in the TEMPLATE_NAME parameter, that device template will be automatically aligned with all discovery sessions created by this automation. For example, if you wanted to align a device template called "Standard Device Template" to every discovery session created by this automation, you would change this line to:

  ```
  TEMPLATE_NAME = "Standard Device Template"
  ```

To modify the parameters for the ***AWS: Discover from EC2 IP*** automation action, perform the following steps:

1. Go to the **Action Policy Manager** page (Registry > Run Book > Actions).
2. Click the wrench icon (🔧) for the ***AWS: Discover from EC2 IP*** automation action.
3. In the ***Snippet Code*** field, locate and edit the line(s) for the parameter(s) you want to change:
4. Click the **[Save]** button.

If you modified the EC2_IP_ATTRIBUTE parameter in the ***AWS: Discover from EC2 IP*** automation action, you must perform the following steps to update the ***AWS: Merge Physical with Component*** automation action:

To modify the parameters for the ***AWS: Discover from EC2 IP*** automation action, perform the following steps:

1. Go to the **Action Policy Manager** page (Registry > Run Book > Actions).
2. Click the wrench icon ( ) for the ***AWS: Discover from EC2 IP*** automation action.
3. In the ***Snippet Code*** field, locate and edit the following line:

   ```
   IP_ATTRIBUTE = 'c-EC2_PrivateIpAddress'
   ```

   If you changed the EC2_IP_ATTRIBUTE parameter in the ***AWS: Discover from EC2 IP*** automation action to 'PublicIpAddress', change this line to:

   ```
   IP_ATTRIBUTE = 'c-EC2_PublicIpAddress'
   ```

4. Click the **[Save]** button.

Enabling the Component Device Record Created Event Policy

To enable the "Component Device Record Created" event policy:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic user interface).
2. Click the Actions menu ( ) for the "Component Device Record Created" event policy and select *Edit*.
3. In the **Event Policy Editor** page, click on the ***Enable Event Policy*** toggle to enable the event policy.
4. Click **[Save]**.

To enable the "Component Device Record Created" event policy in the Skylar One classic user interface:

1. Go to the **Event Policy Manager** page (Registry > Events > Event Manager).
2. Click the wrench icon ( ) for the "Component Device Record Created" event policy.
3. In the ***Operational State*** field, select *Enabled*.
4. Click **[Save]**.

To prevent this change from being overwritten when the PowerPacks installed on the system are updated, you can enable the ***Selective PowerPack Field Protection*** option. To enable this option:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. Check the ***Enable Selective PowerPack Field Protection*** checkbox.
3. Click **[Save]**.

Enabling the Device Record Created Event Policy

To enable the "Device Record Created" event policy:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic user interface).
2. Click the Actions menu ( ) for the "Device Record Created" event policy and select *Edit*.
3. In the **Event Policy Editor** page, click on the ***Enable Event Policy*** toggle to enable the event policy.
4. Click **[Save]**.

To enable the "Device Record Created" event policy in the Skylar One classic user interface:

1. Go to the **Event Policy Manager** page (Registry > Events > Event Manager).
2. Click the wrench icon (🔧) for the "Device Record Created" event policy.
3. In the ***Operational State*** field, select *Enabled*.
4. Click **[Save]**.

To prevent this change from being overwritten when the PowerPacks installed on the system are updated, you can enable the ***Selective PowerPack Field Protection*** option. To enable this option:

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. Check the ***Enable Selective PowerPack Field Protection*** checkbox.
3. Click **[Save]**.

Enabling the Automation Policies

To enable one or more automation policies in the "Amazon Web Services" PowerPack:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. Click the wrench icon (🔧) for the automation policy you want to enable.
3. In the ***Policy State*** field, select *Enabled*.
4. Click **[Save]**.

Preserving Automation Changes

If you have modified automation actions and policies that are included in the "Amazon Web Services" PowerPack, those changes will be overwritten when the PowerPack is updated in your system. If you have modified automation actions and policies that are included in the PowerPack, you can:

- Re-implement those changes after each update of the "Amazon Web Services" PowerPack.
- Remove the content from the PowerPack on your system before you update it. When the "Amazon Web Services" PowerPack is updated in your system, updated versions of this content will not be installed on your system and your local changes will be preserved.

To remove automation actions or automation policies content from the "Amazon Web Services" PowerPack on your system:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the wrench icon (🔧) for the "Amazon Web Services" PowerPack. The **Editing PowerPack** page appears.
3. In the left NavBar of the **Editing PowerPack** page, select the type of content you want to remove:
    - To remove an automation action, click **Run Book Actions**. The **Embedded Run Book Actions** and **Available Run Book Actions** panes appear.
    - To remove an automation policy, click **Run Book Policies**. The **Embedded Run Book Policies** and **Available Run Book Policies** panes appear.
4. In the upper pane, click the delete icon (🗑) for each automation action or automation policy that you want to remove from the "Amazon Web Services" PowerPack on your system.

Aligning AWS Regions to the AWS Region Device Class

The automation for aligning an AWS Region to the correct AWS Region device class includes one automation action:

- ***AWS: Region Device Class Alignment***. This action updates the AWS device class to the correct AWS Region.

NOTE: Device classes for AWS Regions are updated in the second cycle of the "AWS: Region Device Class Discovery" Dynamic Application. Regions will be updated after 24 hours.

The "Amazon Web Services" PowerPack includes an automation policy that triggers this action:

- ***AWS: Region Device Class Alignment***. If enabled, this automation policy can trigger for any device with which the "AWS: Region Device Class Discovery" Dynamic Application is aligned. The automation policy triggers when the "AWS: Device Class Change" event is active on the matching devices, and the automation policy will repeat every 10 minutes until that event is no longer active.

Vanishing Terminated or Terminating EC2 Instances

The automation in this section automatically moves an EC2 instance to a vanished state if the EC2 instance is in a terminating or terminated state. Skylar One unmerges the EC2 instance and physical device, clearing the associated events, and moves the devices to a vanished state.

The automation for vanishing terminated EC2 instances includes one automation action:

- ***AWS: Vanish Terminated EC2 Instances***. If an EC2 instance has been terminated in Amazon, its corresponding device in Skylar One becomes unavailable. This action then requests information from the AWS API about the EC2 instance that triggered the automation action. If the response from the AWS API indicates that the EC2 instance that triggered the automation action is in a terminated or terminating state, the action performs the following steps:
    - If the automation triggers for a physical device that is merged with an EC2 instance, the devices are un-merged.
    - If the automation triggers for a physical device that is merged with an EC2 instance, after being un-merged the physical device is moved to a virtual collector group.
    - If the automation triggers for a physical device that is merged with an EC2 instance, after being unmerged, all events associated with the physical device are cleared.
    - All events associated with the component device are cleared.
    - The component device is vanished.

NOTE: If an EC2 instance is stopped in AWS rather than terminated, then the "AWS Vanish Terminated EC2 Instances" action is not triggered.

The "Amazon Web Services" PowerPack includes an automation policy that triggers this action:

- ***AWS: Vanish Terminated EC2 Instances***. If enabled, this automation policy can trigger for any device with which the "AWS: EC2 Instance Configuration" Dynamic Application is aligned (the members of the AWS EC2 Instances). The automation policy triggers when the "Availability Check Failed" event is active on the matching devices, and the automation policy will repeat every 10 minutes until that event is no longer active.

To use this automation, you must:

- *Enable the AWS: Vanish Terminated EC2 Instances automation policy*
- *Configure your system to preserve this change*

Enabling the Automation Policies

To enable one or more automation policies in the "Amazon Web Services" PowerPack:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. Click the wrench icon ( ) for the automation policy you want to enable.
3. In the ***Policy State*** field, select *Enabled*.
4. Click **[Save]**.

Preserving Automation Changes

If you have modified automation actions and policies that are included in the "Amazon Web Services" PowerPack, those changes will be overwritten when the PowerPack is updated in your system. If you have modified automation actions and policies that are included in the PowerPack, you can:

- Re-implement those changes after each update of the "Amazon Web Services" PowerPack.
- Remove the content from the PowerPack on your system before you update it. When the "Amazon Web Services" PowerPack is updated in your system, updated versions of this content will not be installed on your system and your local changes will be preserved.

To remove automation actions or automation policies content from the "Amazon Web Services" PowerPack on your system:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the wrench icon ( ) for the "Amazon Web Services" PowerPack. The **Editing PowerPack** page appears.

3. In the left NavBar of the **Editing PowerPack** page, select the type of content you want to remove:
    - To remove an automation action, click **Run Book Actions**. The **Embedded Run Book Actions** and **Available Run Book Actions** panes appear.
    - To remove an automation policy, click **Run Book Policies**. The **Embedded Run Book Policies** and **Available Run Book Policies** panes appear.
4. In the upper pane, click the delete icon ( ) for each automation action or automation policy that you want to remove from the "Amazon Web Services" PowerPack on your system.

Appendix

A

Key Metrics Collected by the PowerPack

Overview

This section lists the key metrics for AWS services that the "Amazon Web Services" PowerPack collects by Dynamic Application.

Partial coverage exists for some services for which general information, such as performance counts, is monitored.

This chapter covers the following topics:

Amazon MQ Service 150

API Gateway Service 156

Application ELB Service 159

AppRunner Service 161

AppStream Service 162

AppSync Service 162

Athena Service 165

Auto Scale Service 166

CloudFront Service 168

CloudHSM Service 173

CloudSearch Service 174

CloudTrail Service 174

CloudWatch Service 175
CloudWatch Synthetics Service 176
CodeBuild Service 177
Cognito Service 179
Connect Service 180
DAX Service 183
DDB Service 187
Direct Connect Service 189
DMS Service 193
DocumentDB Service 195
DynamoDB Service 201
EBS Service 202
EC2 Service 204
ECS Service 212
EFS Service 219
EKS Service 221
Elastic Beanstalk Service 223
Elastic Transcoder Service 227
ElastiCache Service 228
ELB Service 234
EMR Service 238
FSx Service 242
Glacier Service 247
Glue Service 248
IoT Service 250
Kafka/MSK Service 251
Kinesis Service 252
KMS Service 259
Lambda Service 260
Lex V2 Service 267
LightSail Service 269

MediaConnect Service 271
MediaConvert Service 280
MediaPackage Service 281
MediaTailor Service 282
Neptune Service 285
Network ELB Service 290
OpenSearch Service 292
OpsWorks Service 303
RDS Service 305
Redshift Service 317
Route 53 Service 320
S3 Service 322
SES Service 325
Shield Standard Service 325
SNS Service 327
SQS Service 329
Storage Gateway Service 331
STS Service 332
SWF Service 333
Transfer Server Service 333
Transit Gateway Service 334
VPC Service 336
WAF Service 343
Workspaces Service 345
WorkMail Service 348
Other Supported Services 348

Amazon MQ Service

These Dynamic Applications are available in the "AWS Extended Services" PowerPack.

AWS: Active MQ Broker-1 Performance	
Object Name	**Object Description**
Amqp Maximum Connections - Avg	The maximum number of clients you can connect to your broker using AMQP. For more information on connection quotas, see Quotas in Amazon MQ.
Burst Balance - Avg	The percentage of burst credits remaining on the Amazon EBS volume used to persist message data for throughput-optimized brokers. If this balance reaches zero, the IOPS provided by the Amazon EBS volume will decrease until the Burst Balance refills.
Cpu Credit Balance - Avg	The number of earned CPU credits that an instance has accrued since it was launched or started (including the number of launch credits). The credit balance is available for the broker instance to spend on bursts beyond the baseline CPU utilization. This metric is available only for the mq.t2.micro broker instance type.
Cpu Utilization - Avg	The percentage of allocated Amazon EC2 compute units that the broker currently uses.
Current Connections Count - Avg	The current number of active connections on the current broker.
Established Connections Count - Avg	The total number of connections, active and inactive, that have been established on the broker.
Heap Usage - Avg	The percentage of the ActiveMQ JVM memory limit that the broker currently uses.
Inactive Durable Topic Subscribers Count - Avg	The number of inactive durable topic subscribers, up to a maximum of 2000.
Job Scheduler Store Percent Usage - Avg	The percentage of disk space used by the job scheduler store.
Journal Files For Fast Recovery - Avg	The number of journal files that will be replayed after a clean shutdown.
Journal Files For Full Recovery - Avg	The number of journal files that will be replayed after an unclean shutdown.
Mqtt Maximum Connections - Avg	The maximum number of clients you can connect to your broker using MQTT. For more information on connection quotas, see Quotas in Amazon MQ.

Network Connector Connection Count - Avg	The number of nodes connected to the broker in a network of brokers using NetworkConnector.
Network In - Avg	The volume of incoming traffic for the broker.
Network Out - Avg	The volume of outgoing traffic for the broker.
Open Transaction Count - Avg	The total number of transactions in progress.
Openwire Maximum Connections - Avg	The maximum number of clients you can connect to your broker using OpenWire. For more information on connection quotas, see Quotas in Amazon MQ.
Stomp Maximum Connections - Avg	The maximum number of clients you can connect to your broker using STOMP. For more information on connection quotas, see Quotas in Amazon MQ.
Store Percent Usage - Avg	The percent used by the storage limit. If this reaches 100, the broker will refuse messages.
Temp Percent Usage - Avg	The percentage of available temporary storage used by non-persistent messages.
Total Consumer Count - Avg	The number of message consumers subscribed to destinations on the current broker.
Total Message Count - Avg	The number of messages stored on the broker.
Total Producer Count - Avg	The number of message producers active on destinations on the current broker.
Volume Read Ops - Avg	The number of read operations performed on the Amazon EBS volume.
Volume Write Ops - Avg	The number of write operations performed on the Amazon EBS volume.
Ws Maximum Connections - Avg	The maximum number of clients you can connect to your broker using WebSocket. For more information on connection quotas, see Quotas in Amazon MQ.

AWS: Active MQ Broker-2 Performance	
Object Name	**Object Description**
Amqp Maximum Connections - Avg	The maximum number of clients you can connect to your broker using AMQP. For more information on connection quotas, see Quotas in Amazon MQ.

Burst Balance - Avg	The percentage of burst credits remaining on the Amazon EBS volume used to persist message data for throughput-optimized brokers. If this balance reaches zero, the IOPS provided by the Amazon EBS volume will decrease until the Burst Balance refills.
Cpu Credit Balance - Avg	The number of earned CPU credits that an instance has accrued since it was launched or started (including the number of launch credits). The credit balance is available for the broker instance to spend on bursts beyond the baseline CPU utilization. This metric is available only for the mq.t2.micro broker instance type.
Cpu Utilization - Avg	The percentage of allocated Amazon EC2 compute units that the broker currently uses.
Current Connections Count - Avg	The current number of active connections on the current broker.
Established Connections Count - Avg	The total number of connections, active and inactive, that have been established on the broker.
Heap Usage - Avg	The percentage of the ActiveMQ JVM memory limit that the broker currently uses.
Inactive Durable Topic Subscribers Count - Avg	The number of inactive durable topic subscribers, up to a maximum of 2000.
Job Scheduler Store Percent Usage - Avg	The percentage of disk space used by the job scheduler store.
Journal Files For Fast Recovery - Avg	The number of journal files that will be replayed after a clean shutdown.
Journal Files For Full Recovery - Avg	The number of journal files that will be replayed after an unclean shutdown.
Mqtt Maximum Connections - Avg	The maximum number of clients you can connect to your broker using MQTT. For more information on connection quotas, see Quotas in Amazon MQ.
Network Connector Connection Count - Avg	The number of nodes connected to the broker in a network of brokers using NetworkConnector.
Network In - Avg	The volume of incoming traffic for the broker.
Network Out - Avg	The volume of outgoing traffic for the broker.
Open Transaction Count - Avg	The total number of transactions in progress.

Openwire Maximum Connections - Avg	The maximum number of clients you can connect to your broker using OpenWire. For more information on connection quotas, see Quotas in Amazon MQ.
Stomp Maximum Connections - Avg	The maximum number of clients you can connect to your broker using STOMP. For more information on connection quotas, see Quotas in Amazon MQ.
Store Percent Usage - Avg	The percent used by the storage limit. If this reaches 100, the broker will refuse messages.
Temp Percent Usage - Avg	The percentage of available temporary storage used by non-persistent messages.
Total Consumer Count - Avg	The number of message consumers subscribed to destinations on the current broker.
Total Message Count - Avg	The number of messages stored on the broker.
Total Producer Count - Avg	The number of message producers active on destinations on the current broker.
Volume Read Ops - Avg	The number of read operations performed on the Amazon EBS volume.
Volume Write Ops - Avg	The number of write operations performed on the Amazon EBS volume.
Ws Maximum Connections - Avg	The maximum number of clients you can connect to your broker using WebSocket. For more information on connection quotas, see Quotas in Amazon MQ.

AWS: Rabbit MQ Broker Performance	
Object Name	**Object Description**
Ack Rate - Sum	The rate at which messages are being acknowledged by consumers. The number produced represents the number of messages per second at the time of sampling.
Channel Count - Avg	The total number of channels established on the broker.
Confirm Rate - Sum	The rate at which the RabbitMQ server is confirming published messages. You can compare this metric with PublishRate to better understand how your broker is performing. The number produced represents the number of messages per second at the time of sampling.

Connection Count - Avg	The total number of connections established on the broker.
Consumer Count - Avg	The total number of consumers connected to the broker.
Exchange Count - Avg	The total number of exchanges configured on the broker.
Message Count - Avg	The total number of messages in the queues. Note The number produced is the total sum of ready and unacknowledged messages on the broker.
Message Ready Count - Avg	The total number of ready messages in the queues.
Message Unacknowledged Count - Avg	The total number of unacknowledged messages in the queues.
Publish Rate - Sum	The rate at which messages are published to the broker. The number produced represents the number of messages per second at the time of sampling.
Queue Count - Avg	The total number of queues configured on the broker.
Rabbit Mqdisk Free - Avg	The total volume of free disk space available in a RabbitMQ broker. When disk usage goes above its limit, the cluster will block all producer connections. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.
Rabbit Mqdisk Free Limit - Avg	The disk limit for a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values. This metric is different per instance size. For more information about Amazon MQ instance types, see Amazon MQ for RabbitMQ instance types.
Rabbit Mqfd Used - Avg	Number of file descriptors used. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.
Rabbit Mqmem Limit - Avg	The RAM limit for a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.

Rabbit Mqmem Used - Avg	The volume of RAM used by a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.
System Cpu Utilization - Avg	The percentage of allocated Amazon EC2 compute units that the broker currently uses. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.

API Gateway Service

AWS: API Gateway HTTP/Websocket Instance Configuration	
Object Name	**Object Description**
Allow Credentials	Specifies if credentials are included in the CORS request. HTTP only.
Allow Headers	Allowed headers. HTTP only.
Allow Methods	Allowed methods. HTTP only.
Allow Origins	Allowed origins. HTTP only.
API Endpoint	The URI of the API, usually in the format: {api-id}.execute-api.{region}.amazonaws.com.
API Gateway Managed	Specifies if the API is managed by API Gateway service.
API Id	The ID of the API.
Api Key Selection Expression	An API key selection expression. Supported only for WebSocket APIs.
Created Date	The date the API was created.
Description	The description of the API.
Disable Execute Api Endpoint	Specifies whether clients can invoke your API by using the default execute-api endpoint. By default, clients can invoke your API with the default https://{api_id}.execute-api.{region}.amazonaws.com endpoint. To require that clients use a custom domain name to invoke your API, disable the default endpoint.

Disable Scheme Validation	Avoid validating models when creating a deployment. Supported only for WebSocket APIs.
Exposed Headers	Collection of exposed headers. HTTP only.
Import Info	The validation information during API import. This may include particular properties of your OpenAPI definition which are ignored during import. Supported only for HTTP APIs.
Key	The key of the key,value tags.
Max Age	Number of seconds browser should cache preflight results. HTTP only.
Name	The name of the API.
Protocol Type	The API protocol.
Route Selection Expression	The route selection expression for the API. For HTTP APIs, the routeSelectionExpression must be ${request.method} ${request.path}. If not provided, this will be the default for HTTP APIs. This property is required for WebSocket APIs.
Value	The value of the the key, value tags.
Version	The version ID for the API.
Warnings	Warnings produced if failonwarnings is turned on during API import.

AWS: API Gateway HTTP Instance Performance	
Object Name	**Object Description**
4XX Errors	Sum of the client side errors (4xx) errors reported over the period.
5XX Errors	The sum of server side errors (5xx) reported over the period.
Count	Sum of total API requests in a period.
Data Processed	The amount of data processed in bytes over the period.
Integration Latency	Time between when API Gateway relays a request to the backend and it responds.
Latency	Time between when API Gateway receives a request from a client and when it provides a response to that client.

AWS: API Gateway Service Configuration	
Object Name	**Object Description**
ID	The identifier of a Usage Plan resource.
Name	The name of a usage plan.
Rate	The API request rate limit.
AWS API Gateway Service/Network Load Balancer	The ELB associated with the API Gateway Service.
Burst	The API request burst limit.
ID	The identifier of the VPC Link.
Name	The name of the VPC link.
Network Load Balancer	The name of ELB associated with the API Gateway.
Quota	The maximum number of requests that can be made in a given time period.
Stages	The associated API stages of a usage plan.
Status	The status of the VPC link.
Target ARNs	The ARNs of network load balancers of the VPC targeted by the VPC link.

AWS: API Gateway Service Health	
Object Name	**Object Description**
Date	The timestamp of this health status update string.
Health	A text description of this AWS API Gateway Service health status.

AWS: API Gateway Websocket Instance Performance	
Object Name	**Object Description**
Client Error	The sum of requests that have a 4xx response returned by API Gateway before the integration is invoked over the period.
Connect Count	The sum of messages sent to the $connect route integration over the period.

Execution Error	The sum of the errors that occurred when calling the integration over the period.
Integration Error	The sum of the request that return an error from the integration over the period.
Integration Latency	The time difference between when API Gateway sending a request to the integration and when the integration provides a response.
Message Count	The sum of messages sent to the Websocket API over the period.

Application ELB Service

AWS: Application ELB Instance Configuration	
Object Name	**Object Description**
AWS Application ELB/Availability Zones	The availability zones for the load balancer.
AWS Application ELB/Security Groups	The unique identifiers of the security groups for the load balancer.
AWS Application ELB/Target Groups	The unique identifiers of the target groups for the application load balancer.
AWS Application ELB/VPC	The unique identifier of the VPC for the load balancer.
DNS Name	The public DNS name of the application load balancer.
Key	The key of the tag, that belongs to the application load balancer.
Listener ARN	The Amazon Resource Name (ARN) of the listener.
Listener Port	The port on which the load balancer is listening.
Listener Protocol	The protocol for connections from clients to the load balancer.
Load Balancer ARN	The Amazon Resource Name (ARN) of the load balancer.
Load Balancer Name	The name of the load balancer.
Load Balancer State	The state of the load balancer.

Scheme	The DNS name of an Internet-facing load balancer is publicly resolvable to the public IP addresses of the nodes.
Type	The type of load balancer.
Value	The value of the tag, that belongs to the application load balancer.

AWS: Application ELB Instance Performance	
Object Name	**Object Description**
Active Connection Count	The total number of concurrent TCP connections active from clients to the load balancer and from the load balancer to targets.
Client TLS Negotiation Error Count	The number of TLS connections initiated by the client that did not establish a session with the load balancer. Possible causes include a mismatch of ciphers or protocols.
HTTPCode_ELB_4XX_Count	The number of HTTP 4XX client error codes that originate from the load balancer. Client errors are generated when requests are malformed or incomplete. These requests have not been received by the target. This count does not include any response codes generated by the targets.
HTTPCode_ELB_5XX_Count	The number of HTTP 5XX server error codes that originate from the load balancer. This count does not include any response codes generated by the targets.
HTTPCode_Target_2XX_Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer.
HTTPCode_Target_3XX_Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer.
HTTPCode_Target_4XX_Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer.
HTTPCode_Target_5XX_Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer.

Rejection Connection Count	The number of connections that were rejected because the load balancer had reached its maximum number of connections.
Request Count	The number of requests received by the load balancer. This includes requests over IPv4 and IPv6.
Target Connection Error Count	The number of connections that were not successfully established between the load balancer and target.
Target Response Time	The time elapsed, in seconds, after the request leaves the load balancer until a response from the target is received. This is equivalent to the target_ processing_time field in the access logs.

AppRunner Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: AppRunner Service Performance	
Object Name	**Object Description**
2XX Status Responses - Sum	The number of HTTP requests that returned each response status, grouped by category 2xx.
4XX Status Responses - Sum	The number of HTTP requests that returned each response status, grouped by category 4xx.
5XX Status Responses - Sum	The number of HTTP requests that returned each response status, grouped by category 5xx.
Active Instances - Avg	The number of instances that are processing HTTP requests for your service. If the ActiveInstances metric displays zero, it means that there are no requests for the service. It does not indicate that the number of instances for your service is zero.
Request Latency - Avg	The time, in milliseconds, that it took your web service to process HTTP requests.
Requests - Sum	The number of HTTP requests that the service received.

AppStream Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: AppStream 2.0 Performance	
Object Name	**Object Description**
Actual Capacity - Avg	The total number of instances that are available for streaming or are currently streaming. ActualCapacity = AvailableCapacity + InUseCapacity
Available Capacity - Avg	The number of idle instances currently available for user sessions. AvailableCapacity = ActualCapacity - InUseCapacity
Capacity Utilization - Avg	Capacity utilization.
Desired Capacity - Avg	The percentage of instances in a fleet that are being used, using the following formula. CapacityUtilization = (InUseCapacity/ActualCapacity) * 100 Monitoring this metric helps with decisions about increasing or decreasing the value of a fleet's desired capacity.
In Use Capacity - Avg	The number of instances currently being used for streaming sessions. One InUseCapacity count represents one streaming session.
Insufficient Capacity Error - Avg	The number of session requests rejected due to lack of capacity.
Pending Capacity - Avg	The number of instances being provisioned by AppStream 2.0. Represents the additional number of streaming sessions the fleet can support after provisioning is complete. When provisioning starts, it usually takes 10-20 minutes for an instance to become available for streaming.
Running Capacity - Avg	The total number of instances currently running. Represents the number of concurrent streaming sessions that can be supported by the fleet in its current state.

AppSync Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: AppSync GraphQLAPI Performance	
Object Name	**Object Description**
4XX Error - Sum	Errors resulting from requests that are not valid due to an incorrect client configuration. Typically, these errors happen anywhere outside of GraphQL processing. For example, these errors can occur when the request includes an incorrect JSON payload or an incorrect query, when the service is throttled, or when the authorization settings are misconfigured.
5XX Error - Sum	Errors encountered during the running of a GraphQL query. For example, this can occur when invoking a query for an empty or incorrect schema. It can also occur when the Amazon Cognito user pool ID or AWS Region is not valid. Alternatively, this could also happen if AWS AppSync encounters an issue during processing of a request.
Active Connections - Sum	The number of concurrent WebSocket connections from clients to AWS AppSync in 1 minute.
Active Subscriptions - Sum	The number of concurrent subscriptions from clients in 1 minute.
Connect Client Error - Sum	The number of WebSocket connections that were rejected by AWS AppSync because of client-side errors. This could imply that the service is throttled or that the authorization settings are misconfigured.
Connect Server Error - Sum	The number of errors that originated from AWS AppSync while processing connections. This usually happens when an unexpected server-side issue occurs.
Connect Success - Sum	The number of successful WebSocket connections to AWS AppSync. It is possible to have connections without subscriptions.
Connection Duration - Avg	The amount of time that the connection stays open.
Disconnect Client Error - Sum	The number of client errors that originated from AWS AppSync while disconnecting WebSocket connections.
Disconnect Server Error - Sum	The number of server errors that originated from AWS AppSync while disconnecting WebSocket connections.

Disconnect Success - Sum	The number of successful WebSocket disconnections from AWS AppSync.
Invalidation Success - Sum	The number of subscriptions successfully invalidated (unsubscribed) by a mutation with $extensions.invalidateSubscriptions().
Latency - Avg	The time between when AWS AppSync receives a request from a client and when it returns a response to the client. This doesn't include the network latency encountered for a response to reach the end devices.
Publish Data Message Client Error - Sum	The number of subscription event messages that failed to publish because of client-side errors.
Publish Data Message Server Error - Sum	The number of errors that originated from AWS AppSync while publishing subscription event messages. This usually happens when an unexpected server-side issue occurs.
Publish Data Message Size - Avg	The size of subscription event messages published.
Publish Data Message Success - Sum	The number of subscription event messages that were successfully published.
Subscribe Client Error - Sum	The number of subscriptions that were rejected by AWS AppSync because of client-side errors. This can occur when a JSON payload is incorrect, the service is throttled, or the authorization settings are misconfigured.
Subscribe Server Error - Sum	The number of errors that originated from AWS AppSync while processing subscriptions. This usually happens when an unexpected server-side issue occurs.
Subscribe Success - Sum	The number of subscriptions that were successfully registered to AWS AppSync through WebSocket. It's possible to have connections without subscriptions, but it's not possible to have subscriptions without connections.
Unsubscribe Client Error - Sum	The number of unsubscribe requests that were rejected by AWS AppSync because of client-side errors.
Unsubscribe Server Error - Sum	The number of errors that originated from AWS AppSync while processing unsubscribe requests. This usually happens when an unexpected server-side issue occurs.
Unsubscribe Success - Sum	The number of unsubscribe requests that were successfully processed.

AWS: AppSync Performance	
Object Name	**Object Description**
Requests - Sum	The number of requests (queries + mutations) that all APIs in your account have processed, by Region.
Tokens Consumed - Sum	Tokens are allocated to requests based on the amount of resources (processing time and memory used) that a request consumes. Usually, each request consumes one token. However, a request that consumes large amounts of resources is allocated additional tokens as needed.

Athena Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Athena WorkGroup Performance	
Object Name	**Object Description**
Engine Execution Time - Sum	The number of milliseconds that the query took to run.
Processed Bytes - Sum	The number of bytes that Athena scanned per DML query. For queries that were canceled (either by the users, or automatically, if they reached the limit), this includes the amount of data scanned before the cancellation time. This metric is not reported for DDL queries.
Query Planning Time - Sum	The number of milliseconds that Athena took to plan the query processing flow. This includes the time spent retrieving table partitions from the data source. Note that because the query engine performs the query planning, query planning time is a subset of EngineExecutionTime.
Query Queue Time - Sum	The number of milliseconds that the query was in the query queue waiting for resources. Note that if transient errors occur, the query can be automatically added back to the queue.

Service Processing Time - Sum	Number of milliseconds that Athena took to process the query results after the query engine finished running the query.
Total Execution Time - Sum	The number of milliseconds that Athena took to run a DDL or DML query. TotalExecutionTime includes QueryQueueTime, QueryPlanningTime, EngineExecutionTime, and ServiceProcessingTime.

Auto Scale Service

AWS: Auto Scale Group Instance Configuration	
Object Name	**Object Description**
.Metric Name	One or more metrics. If you omit this parameter, all metrics are enabled. GroupMinSize GroupMaxSize GroupDesiredCapacity GroupInServiceInstances GroupPendingInstances GroupStandbyInstances GroupTerminatingInstances GroupTotalInstances Note that the GroupStandbyInstances metric is not enabled by default. You must explicitly request this metric.
ARN	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
AWS Auto Scale Group/EC2	The unique identifier for the AutoScale-EC2 relationship.
AWS Auto Scale Group/Launch Configuration	The identifier for this group's LaunchConfig.
Creation Date	The datetime this group was created.
Default Cooldown	The number of seconds allowed for the cooldown period.
Desired Capacity	The desired size of the group.
Health Check Grace Period	The polling interval between health checks.
Health Check Type	The health check status type for the group. This is usually an EC2.
Id	The unique identifier of the Auto Scale Group.
Launch Configuration ARN	The identifier for this group's LaunchConfig.

Load Balancers	A list of load balancers for the group.
Max	The maximum size of the group.
Metric Granularity	The granularity to associate with the metrics to collect. The only valid value is 1Minute.
Min	The minimum size of the group.
Name	The name of the Auto Scale group.
Placement Group	The name of the placement group, if any, for this group.
Subnets	The subnet ID of the VPC where the Group resides.
Suspended Processes	A list of suspended process names.
Termination Policies	Pre-launch environment variables
Zones	All deployable zones in this region.

AWS: Auto Scale Group Performance	
Object Name	**Object Description**
Group Desired Capacity	The number of instances that the Auto Scaling group attempts to maintain.
Group In Service Instances	The number of instances that are running as part of the Auto Scaling group. This metric does not include instances that are pending or terminating.
Group Max Size	The maximum size of the Auto Scaling group.
Group Min Size	The minimum size of the Auto Scaling group.
Group Pending Instances	The number of instances that are pending. A pending instance is not yet in service. This metric does not include instances that are in service or terminating.
Group Standby Instances	The number of instances that are in a Standby state. Instances in this state are still running but are not actively in service.
Group Terminating Instances	The number of instances that are in the process of terminating. This metric does not include instances that are in service or pending.
Group Total Instances	The total number of instances in the Auto Scaling group. This metric identifies the number of instances that are in service, pending, and terminating.

AWS: Auto Scale Launch Config Instance Configuration	
Object Name	**Object Description**
ARN	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Block Devices	Attached EBS filesystems at boot.
Created	The datetime this configuration was created.
EBS Optimized	If this configuration is EBS Optimized.
IAM Instance Profile	The IAM role name associated with this configuration.
Id	The unique identifier of the launch configuration.
Instance Type	The EC2 type of the instances in this configuration.
Kernel ID	The unique identifier for the kernel of this configuration.
Key Name	The key-pair associated with this configuration.
Monitoring	Whether or not detailed monitoring is set for this configuration.
Name	The name of the launch configuration.
RAM Disk Id	The unique ID, if any, for the RAM disk of this configuration.
Security Groups	A list of security groups for this launch configuration.
Spot Price	The spot price in USD/Hour for this configuration.

CloudFront Service

AWS: CloudFront Behavior Configuration	
Object Name	**Object Description**
Forwarded Query Strings	The Forwarded Query Strings of the Behavior.
Id	The unique identifier of the Behavior.
Name	The name of the Behavior.

Origin	The name of the Origin associated with this Behavior.
Path Pattern	A path pattern (for example, images/*.jpg) specifies which requests you want this cache behavior to apply to. When CloudFront receives an end-user request, the requested path is compared with path patterns in the order in which cache behaviors are listed in the distribution.
Trusted Signers	Trusted signers are the AWS accounts that can create signed URLs and signed cookies for a distribution. By default, no account, not even the account that created the distribution, is allowed to create signed URLs or signed cookies. To specify the AWS accounts that you want to use as trusted signers, add the accounts to your distribution (Web distributions and RTMP distributions)
Viewer Protocol Policy	Configure the Viewer Protocol Policy for some or all of your CloudFront cache behaviors either to redirect HTTP requests to HTTPS requests or to require that viewers use only the HTTPS protocol to access your objects in the CloudFront cache.

AWS: CloudFront Error Page Configuration	
Object Name	**Object Description**
Error Caching Minimum TTL	The smallest TTL (in seconds) allowed for this cached element.
HTTP Error Code	The HTTP status code associated with this Error Page.
HTTP Response Code	The HTTP status code used in the response header of the Error Page.
Response Page Path	The URL path associated with the response resource of the Error Page.

AWS: CloudFront Invalidation Configuration	
Object Name	**Object Description**
Date	The date when an invalidation was created.
Id	The unique identifier of the Invalidation.

Object Paths	The object paths of the Invalidation.
Status	The status of the invalidation.

AWS: CloudFront Origin Configuration	
Object Name	**Object Description**
Access Identity	The access identity hash value.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
HTTP Port	The HTTP CloudFront uses only HTTP to access the origin.
HTTPS Port	The HTTPS CloudFront uses only HTTPS to access the origin.
Id	The unique identifier of the Origin.
Name	The name of the Origin.
Origin Protocol Policy	The cache behaviors are routing requests to the origins for which you have configured an Origin Protocol Policy of HTTPS Only or Match Viewer, if applicable.
Type	The type of origin storage.

AWS: CloudFront Restriction Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the Restriction.
Name	The name of the Restriction.
Status	Whether or not this restriction is enabled.
Type	Whether the Restriction is a white-list or a black-list.

AWS: CloudFront RTMP Distribution Configuration

Object Name	Object Description
CNames	Alternative DNS entry names and sub-domains.
Comment	A comment about the specific distribution.
Delivery Method	The transfer protocol of the distribution.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the distribution.
Last Modified	The most recent date and time when the configuration of this distribution was modified.
Log Bucket	The S3 bucket, if any, used to store the access logs of the distribution.
Log Prefix	The log prefix, if any, used for the S3 path of the Log Bucket. Cloudfrontlog is the default value.
Name	The name of the distribution.
Origin	The fully-qualified S3 bucket name for the distribution.
Price Class	The price tiers, if any, for restrictions on the distribution.
State	Whether or not the distribution is enabled for deployment.
Status	Whether or not the distribution is deployed.
Trusted Signers	A list of AWS Account IDs.

AWS: CloudFront Service Health	
Object Name	**Object Description**
Date	A timestamp when the health status of this service was originally written.
Health	A text description of the health status of this AWS service.

AWS: CloudFront Web Distribution Configuration	
Object Name	**Object Description**

AWS CloudFront Web Distribution/WAF Web ACL	The Web ACL associated with the CloudFront.
CNames	Alternative DNS entry names and sub-domains.
Comment	A comment about the specific distribution.
Delivery Method	The transfer protocol of the distribution.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the distribution.
Last Modified	The most recent date and time when the configuration of this distribution was modified.
Log Bucket	The S3 bucket, if any, used to store the access logs of the distribution.
Log Prefix	The log prefix, if any, used for the S3 path of the Log Bucket. Cloudfrontlog is the default value.
Name	The name of the distribution.
Price Class	The price tiers, if any, for restrictions on the distribution.
State	Whether or not the distribution is enabled for deployment.
Status	Whether or not the distribution is deployed.

AWS: CloudFront Web Distribution Performance	
Object Name	**Object Description**
4xx Error Rate	The percentage of all requests for which the HTTP status code is 4xx.
5xx Error Rate	The percentage of all requests for which the HTTP status code is 5xx.
Bytes Downloaded	The number of bytes downloaded by viewers for GET, HEAD, and OPTIONS requests.
Bytes Uploaded	The number of bytes uploaded to your origin with CloudFront using POST and PUT requests.
Requests	The number of requests for all HTTP methods and for both HTTP and HTTPS requests.
Total Error Rate	The percentage of all requests for which the HTTP status code is 4xx or 5xx.

CloudHSM Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: CloudHSM Cluster Performance	
Object Name	**Object Description**
Hsm Keys Session Occupied - Sum	The number of session keys being used by the HSM instance.
Hsm Keys Token Occupied - Sum	The number of token keys being used by the HSM instance and the cluster.
Hsm Session Count - Sum	The number of open connections to the HSM instance. Up to 2,048 are allowed. By default, the client daemon is configured to open two sessions with each HSM instance under one end-to-end encrypted channel. AWS CloudHSM may also have up to 2 connections open with the HSM to monitor the health of the HSMs.
Hsm Ssl Ctxs Occupied - Sum	The number of end-to-end encrypted channels currently established for the HSM instance. Up to 2,048 channels are allowed.
Hsm Temperature - Sum	The junction temperature of the hardware processor. The system shuts down if temperature reaches 110 degrees Centigrade.
Hsm Unhealthy - Sum	The HSM instance is not performing properly. AWS CloudHSM automatically replaces unhealthy instances for you. You may choose to proactively expand cluster size to reduce performance impact while we are replacing the HSM.
Hsm Users Available - Avg	The number of additional users that can be created. This equals the maximum number of users (listed in HsmUsersMax) minus the users created to date.
Hsm Users Max - Avg	The maximum number of users that can be created on the HSM instance. Currently this is 1,024.
Interface Eth2 Octets Input - Sum	The cumulative sum of incoming traffic to the HSM to date.
Interface Eth2 Octets Output - Sum	The cumulative sum of outgoing traffic to the HSM to date.

CloudSearch Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: CloudSearch Domain Performance	
Object Name	**Object Description**
Index Utilization - Avg	The percentage of the search instance's index capacity that has been used. The Maximum value indicates the percentage of the domain's index capacity that has been used.
Partitions - Avg	The number of partitions the index is distributed across.
Searchable Documents - Max	The number of searchable documents in the domain's search index.
Successful Requests - Sum	The number of search requests successfully processed by a search instance.

CloudTrail Service

AWS: CloudTrail Instance Configuration	
Object Name	**Object Description**
AWS CloudTrail/S3	The identifier for the related S3.
AWS CloudTrail/SNS	The unique identifier for the CloudTrail-SNS relationship.
AWS: S3 Relationship	The group label for the CloudTrail-S3 relationship.
AWS: SNS Relationship	The group label for the CloudTrail-SNS relationship.
Bucket Name	The name of the S3 bucket which stores the trail.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).

Latest Delivery Attempt Succeeded	The datetime of the most recent successful SNS delivery.
Latest Delivery Attempt Time	The datetime of the most recent attempt to deliver cloud trail logs from SNS.
Latest Delivery Error	The error type of the most recent, if any, SNS delivery error.
Latest Notification Attempt Succeeded	The most recent datetime when an SNS delivery succeeded.
Latest Notification Attempt Time	The datetime of the most recent SNS delivery attempt whether or not it succeeded.
Latest Notification Error	The most recent notification error.
Log Global Service Events	Whether or not service status events are logged.
Logging	Whether or not the trail is actively logging.
Logging Prefix	The prefix string, if any, attached to the S3 path where trails are stored.
Name	The name of the trail.
SNS Publishing	Whether or not trail logging is published to SNS.
SNS Topic Name	The SNS topic which publishes trail updates.
Time Logging Started	The datetime when logging most recently started.
Time Logging Stopped	The most recent datetime when logging stopped.

CloudWatch Service

AWS: CloudWatch Alarms Performance	
Object Name	**Object Description**
CloudWatch Action (Failed) Alarms	The total number of Action alarms with actionState Failure in the latest polling interval.
CloudWatch Action Alarms	The total number of Action alarms received in the latest polling interval.
CloudWatch Alarms	The total number of alarms received in the latest polling interval.

CloudWatch Alarms Collection Success	Indicates the success (1) or failure (0) of the CloudWatch alarms history collection snippet. Failure (0) occurs if the snippet encounters an unexpected error. NOTE: This collection object is especially important for determining which CloudWatch alarms are to be monitored. 1. To collect all CloudWatch alarm types (ConfigurationUpdate, StateUpdate, Action), select the snippet cloudwatch_alarms_performance. 2. To collect only CloudWatch StateUpdate alarm types, select the snippet cloudwatch_alarms_ performance_StateUpdate_only.
CloudWatch ConfigurationUpdate Alarms	The total number of ConfigurationUpdate alarms received in the latest polling interval.
CloudWatch State Alarms	The total number of StateUpdate alarms received in the latest polling interval.
CloudWatch StateUpdate (ALARM) Alarms	The total number of Alarm StateUpdate alarms in the latest polling interval.
CloudWatch StateUpdate (INSUFFICIENT_DATA) Alarms	The total number of INSUFFICIENT_DATA StateUpdate alarms in the latest polling interval.
CloudWatch StateUpdate (OK) Alarms	The total number of OK StateUpdate alarms in the latest polling interval.
CloudWatch Total Alarms	A running total of all AWS CloudWatch alarms seen since start of collection.

AWS: CloudWatch Service Health	
Object Name	**Object Description**
Date	A timestamp when this service's health status was originally written.
Health	A text description of this AWS Service's health status.

CloudWatch Synthetics Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: CloudWatch Synthetics Canary Performance

Object Name	Object Description
2Xx - Sum	The number of network requests performed by the canary that returned OK responses, with response codes between 200 and 299.
4Xx - Sum	The number of network requests performed by the canary that returned Error responses, with response codes between 400 and 499.
5Xx - Sum	The number of network requests performed by the canary that returned Fault responses, with response codes between 500 and 599.
Duration - Avg	The duration in milliseconds of the canary run.
Failed - Sum	The number of canary runs that failed to execute. These failures are related to the canary itself.
Failed Requests - Sum	The number of HTTP requests executed by the canary on the target website that failed with no response.
Success Percent - Avg	The percentage of the runs of this canary that succeed and find no failures.
Visual Monitoring Success Percent - Avg	The percentage of visual comparisons that successfully matched the baseline screenshots during a canary run.
Visual Monitoring Total Comparisons - Sum	The total number of visual comparisons that happened during a canary run.

CodeBuild Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: CodeBuild Performance	
Object Name	**Object Description**
Build Duration - Avg	Measures the duration of the build's BUILD phase.
Builds - Avg	Measures the number of builds triggered.
Download Source Duration - Avg	Measures the duration of the build's DOWNLOAD_SOURCE phase.
Duration - Avg	Measures the duration of all builds over time.

Failed Builds - Sum	Measures the number of builds that failed because of client error or a timeout.
Finalizing Duration - Avg	Measures the duration of the build's FINALIZING phase.
Install Duration - Avg	Measures the duration of the build's INSTALL phase.
Post Build Duration - Avg	Measures the duration of the build's POST_ BUILD phase.
Pre Build Duration - Avg	Measures the duration of the build's PRE_ BUILD phase.
Provisioning Duration - Avg	Measures the duration of the build's PROVISIONING phase.
Queued Duration - Avg	Measures the duration of the build's QUEUED phase.
Submitted Duration - Avg	Measures the duration of the build's SUBMITTED phase.
Succeeded Builds - Sum	Measures the number of successful builds.
Upload Artifacts Duration - Avg	Measures the duration of the build's UPLOAD_ ARTIFACTS phase.

AWS: CodeBuild Resource Performance	
Object Name	**Object Description**
Cpuutilized - Avg	The number of CPU units of allocated processing used by the build container.
Cpuutilized Percent - Avg	The percentage of allocated processing used by the build container.
Memory Utilized - Avg	The number of megabytes of memory used by the build container.
Memory Utilized Percent - Avg	The percentage of allocated memory used by the build container.
Storage Read Bytes - Avg	The storage read speed used by the build container.
Storage Write Bytes - Avg	The storage write speed used by the build container.

Cognito Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Cognito Resource Instance Performance	
Object Name	**Object Description**
Call Count - Sum	Provides the total number of calls customers made related to a category. This metric includes all the calls, such as throttled calls, failed calls, and successful calls.
Throttle Count - Sum	Provides the total number of throttled calls related to a category.

AWS: Cognito UserPool Instance Performance	
Object Name	**Object Description**
Federation Successes - Sum	Provides the total number of successful identity federation requests to the Amazon Cognito user pool. A successful identity federation request produces a value of 1, whereas an unsuccessful request produces a value of 0. A throttled request is also considered as an unsuccessful request, and hence a throttled request will also produce a count of 0.
Federation Throttles - Sum	Provides the total number of throttled identity federation requests to the Amazon Cognito user pool. A count of 1 is published whenever an identity federation request is throttled.
Sign In Successes - Sum	Provides the total number of successful user authentication requests made to the Amazon Cognito user pool. A user authentication is considered successful when authentication token is issued to the user
Sign In Throttles - Sum	Provides the total number of throttled user authentication requests made to the Amazon Cognito user pool. A count of 1 is published whenever an authentication request is throttled.

Sign Up Successes - Sum	Provides the total number of successful user registration requests made to the Amazon Cognito user pool. A successful user registration request produces a value of 1, whereas an unsuccessful request produces a value of 0. A throttled request is also considered as an unsuccessful request, and hence a throttled request will also produce a count of 0.
Sign Up Throttles - Sum	Provides the total number of throttled user registration requests made to the Amazon Cognito user pool. A count of 1 is published whenever a user registration request is throttled.
Token Refresh Successes - Sum	Provides the total number of successful requests to refresh an Amazon Cognito token that were made to the Amazon Cognito user pool. A successful refresh Amazon Cognito token request produces a value of 1, whereas an unsuccessful request produces a value of 0. A throttled request is also considered as an unsuccessful request, and hence a throttled request will also produce a count of 0.
Token Refresh Throttles - Sum	Provides the total number of throttled requests to refresh an Amazon Cognito token that were made to the Amazon Cognito user pool. A count of 1 is published whenever a refresh Amazon Cognito token request is throttled.

Connect Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Connect Performance	
Object Name	**Object Description**
Call Back Not Dialable Number - Avg	The number of times a queued callback to a customer could not be dialed because the customer's number is in a country for which outbound calls are not allowed for the instance. The countries allowed for an instance are defined by the service quotas.

Call Recording Upload Error - Sum	The number of call recordings that failed to upload to the Amazon S3 bucket configured for your instance. This is the bucket specified in Data Storage > Call Recordings settings for the instance.
Calls Breaching Concurrency Quota - Sum	The total number of voice calls that exceeded the concurrent calls quota for the instance. For the total number of calls that breach the quota, take a look at the Sum statistic. For example, assume your contact center experiences the following volumes, and your service quota is 100 concurrent calls: 0:00 : 125 concurrent calls. This is 25 over the quota. 0:04 : 135 concurrent calls. This is 35 over the quota. 0:10 : 150 concurrent calls. This is 50 over the quota. CallsBreachingConcurrencyQuota = 110: the total number of voice calls that exceeded the quota between 0:00 and 0:10.
Calls Per Interval - Avg	The average number of voice calls, both inbound and outbound, received or placed per second in the instance.
Chats Breaching Active Chat Quota - Sum	The total number of valid requests made to start a chat that exceeded the concurrent active chats quota for the instance. For the total number of chats requests that breach the quota, take a look at the Sum statistic. For example, assume your contact center experiences the following volumes, and your service quota is 2500 concurrent active chats: 0:00 : 2525 concurrent active chats. This is 25 over the quota. 0:04 : 2535 concurrent active chats. This is 35 over the quota. 0:10 : 2550 concurrent active chats. This is 50 over the quota. ChatsBreachingActiveChatsQuota = 110: the total number of chats that exceeded the quota between 0:00 and 0:10.
Concurrent Active Chats - Avg	The number of concurrent active chats in the instance at the time the data is displayed in the dashboard. The value displayed for this metric is the number of concurrent active chats at the time the dashboard is displayed, and not a sum for the entire interval of the refresh interval set. All active chats are included, not only active tasks that are connected to agents.

Concurrent Active Chats Percentage - Sum	The percentage of the concurrent active chats service quota used in the instance. This is calculated by: ConcurrentActiveChats / ConfiguredConcurrentActiveChatsLimit Where ConfiguredConcurrentActiveChatsLimit is the Concurrent active chats per instance configured for your instance.
Concurrent Calls - Avg	The number of concurrent active voice calls in the instance at the time the data is displayed in the dashboard. The value displayed for this metric is the number of concurrent active calls at the time the dashboard is displayed, and not a sum for the entire interval of the refresh interval set. All active voice calls are included, not only active calls that are connected to agents.
Concurrent Calls Percentage - Avg	The percentage of the concurrent active voice calls service quota used in the instance. This is calculated by: ConcurrentCalls / ConfiguredConcurrentCallsLimit.
Concurrent Tasks - Avg	The number of concurrent active tasks in the instance at the time the data is displayed in the dashboard. The value displayed for this metric is the number of concurrent active tasks at the time the dashboard is displayed, and not a sum for the entire interval of the refresh interval set. All active tasks are included, not only active tasks that are connected to agents.
Concurrent Tasks Percentage - Avg	he percentage of the concurrent active tasks service quota used in the instance. This is calculated by: ConcurrentTasks / ConfiguredConcurrentTasksLimit Where ConfiguredConcurrentTasksLimit is the Concurrent tasks per instance configured for your instance. Unit: Percent (output displays as a decimal).
Contact Flow Errors - Sum	The number of times the error branch for a contact flow was run.
Contact Flow Fatal Errors - Sum	The number of times a contact flow failed to execute due to a system error.
Longest Queue Wait Time - Max	The longest amount of time, in seconds, that a contact waited in a queue. This is the length of time a contact waited in a queue during the refresh interval selected in the CloudWatch dashboard, such as 1 minute or 5 minutes.

Misconfigured Phone Numbers - Sum	The number of calls that failed because the phone number is not associated with a contact flow.
Missed Calls - Sum	The number of voice calls that were missed by agents during the refresh interval selected, such as 1 minute or 5 minutes. A missed call is one that is not answered by an agent within 20 seconds.
Public Signing Key Usage - Sum	The number of times a contact flow security key (public signing key) was used to encrypt customer input in a contact flow.
Queue Capacity Exceeded Error - Sum	The number of calls that were rejected because the queue was full.
Queue Size - Sum	The number of contacts in the queue. The value reflects the number of contacts in the queue at the time the dashboard is accessed, not for the duration of the reporting interval.
Successful Chats Per Interval - Sum	The number of chats successfully started in the instance for the defined interval.
Tasks Breaching Concurrency Quota - Sum	The total number of tasks that exceeded the concurrent tasks quota for the instance. For the total number of tasks that breach the quota, take a look at the Sum statistic.
Tasks Expired - Sum	Tasks which have expired after being active for 7 days.
Tasks Expiry Warning Reached - Sum	Tasks that have been active for 6 days 22 hours and reached expiry warning limit.
Throttled Calls - Sum	The number of voice calls that were rejected because the rate of calls per second exceeded the maximum supported quota. To increase the supported rate of calls, request an increase in the service quota for concurrent active calls per instance.
To Instance Packet Loss Rate - Sum	The ratio of packet loss for calls in the instance, reported every 10 seconds. Each data point is between 0 and 1, which represents the ratio of packets lost for the instance.

DAX Service

This Dynamic Application is available in the ***AWS Extended Services***PowerPack.

AWS: DAX Cluster Performance	
Object Name	**Object Description**
Batch Get Item Request Count - Sum	The number of BatchGetItem requests handled by the node or cluster.
Batch Write Item Request Count - Sum	The number of BatchWriteItem requests handled by the node or cluster.
Cache Memory Utilization - Avg	The percentage of available cache memory that is in use by the item cache and query cache on the node or cluster. Cached data starts to be evicted prior to memory utilization reaching 100% (see EvictedSize metric). If CacheMemoryUtilization reaches 100% on any node, write requests will be throttled and you should consider switching to a cluster with a larger node type.
Cache Memory Utilization - Max	The percentage of available cache memory that is in use by the item cache and query cache on the node or cluster. Cached data starts to be evicted prior to memory utilization reaching 100% (see EvictedSize metric). If CacheMemoryUtilization reaches 100% on any node, write requests will be throttled and you should consider switching to a cluster with a larger node type.
Client Connections - Avg	The number of simultaneous connections made by clients to the node or cluster.
CPU Credit Balance - Avg	The number of earned CPU credits that a node has accrued since it was launched or started. Credits are accrued in the credit balance after they are earned, and removed from the credit balance when they are spent. The credit balance has a maximum limit, determined by the DAX node size. After the limit is reached, any new credits that are earned are discarded. The credits in the CPUCreditBalance are available for the node to spend to burst beyond its baseline CPU utilization.

CPU Credit Usage - Avg	The number of CPU credits spent by the node for CPU utilization. One CPU credit equals one vCPU running at 100% utilization for one minute or an equivalent combination of vCPUs, utilization, and time (for example, one vCPU running at 50% utilization for two minutes or two vCPUs running at 25% utilization for two minutes). CPU credit metrics are available at a five-minute frequency only. If you specify a period greater than five minutes, use the Sum statistic instead of the Average.
CPU Surplus Credit Balance - Avg	The number of surplus credits that have been spent by a DAX node when its CPUCreditBalance value is zero. The CPUSurplusCreditBalance value is paid down by earned CPU credits. If the number of surplus credits exceeds the maximum number of credits that the node can earn in a 24-hour period, the spent surplus credits above the maximum incur an additional charge.
CPU Surplus Credits Charged - Avg	The number of spent surplus credits that are not paid down by earned CPU credits, and which thus incur an additional charge. Spent surplus credits are charged when the spent surplus credits exceed the maximum number of credits that the node can earn in a 24-hour period. Spent surplus credits above the maximum are charged at the end of the hour or when the node is terminated.
CPU Utilization - Avg	The percentage of CPU utilization of the node or cluster.
Delete Item Request Count - Sum	The number of DeleteItem requests handled by the node or cluster.
Error Request Count - Sum	Total number of requests that resulted in a user error reported by the node or cluster. Requests that were throttled by the node or cluster are included.
Estimated Db Size - Avg	An approximation of the amount of data cached in the item cache and the query cache by the node or cluster.
Evicted Size - Avg	The amount of data that was evicted by the node or cluster to make room for newly requested data. If the miss rate goes up, and you see this metric also growing, it probably means that your working set has increased. You should consider switching to a cluster with a larger node type.

Failed Request Count - Sum	Total number of requests that resulted in an error reported by the node or cluster.
Fault Request Count - Sum	Total number of requests that resulted in an internal error reported by the node or cluster.
Get Item Request Count - Sum	The number of GetItem requests handled by the node or cluster.
Item Cache Hits - Sum	The number of times an item was returned from the cache by the node or cluster.
Item Cache Misses - Sum	The number of times an item was not in the node or cluster cache, and had to be retrieved from DynamoDB.
Network Bytes In - Avg	The number of bytes received on all network interfaces by the node or cluster.
Network Bytes Out - Avg	The number of bytes sent out on all network interfaces by the node or cluster. This metric identifies the volume of outgoing traffic in terms of the number of bytes on a single node or cluster.
Network Packets In - Avg	The number of packets received on all network interfaces by the node or cluster.
Network Packets Out - Avg	The number of packets sent out on all network interfaces by the node or cluster. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single node or cluster.
Put Item Request Count - Sum	The number of PutItem requests handled by the node or cluster.
Query Cache Hits - Sum	The number of times a query result was returned from the node or cluster cache.
Query Cache Misses - Sum	The number of times a query result was not in the node or cluster cache, and had to be retrieved from DynamoDB.
Query Request Count - Sum	The number of query requests handled by the node or cluster.
Scan Cache Hits - Sum	The number of times a scan result was returned from the node or cluster cache.
Scan Cache Misses - Sum	The number of times a scan result was not in the node or cluster cache, and had to be retrieved from DynamoDB.
Scan Request Count - Sum	The number of scan requests handled by the node or cluster.

Throttled Request Count - Sum	Total number of requests throttled by the node or cluster. Requests that were throttled by DynamoDB are not included, and can be monitored using DynamoDB Metrics.
Total Request Count - Sum	Total number of requests handled by the node or cluster.
Transact Get Items Count - Sum	The number of TransactGetItems requests handled by the node or cluster.
Transact Write Items Count - Sum	The number of TransactWriteItems requests handled by the node or cluster.
Update Item Request Count - Sum	The number of UpdateItem requests handled by the node or cluster.

DDB Service

AWS: DDB Instance Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Read Capacity	The provisioned read capacity of the table
Table Name	The name of the DDB table.
Write Capacity	The provisioned write capacity of the table.

AWS: DDB Instance Performance	
Object Name	**Object Description**
Avg Consumed Read Capacity Units	The average consumed read capacity of the deployment configuration for this table.
Avg Consumed Write Capacity Units	The average consumed write capacity for the table.
Avg Returned Item Count	The average returned item count for the table.
Avg Successful Request Latency	The average latency for successful table operations.

Consumed Read Capacity Units Avg. Prov. Throughput	The average provisioned throughput (sum @ 5 min poll/300) read capacity of the deployment configuration for this table.
Consumed Write Capacity Units Avg. Prov. Throughput	The average provisioned throughput of consumed write capacity for the table.
Max Consumed Read Capacity Units	The maximum consumed read capacity of the deployment configuration for this table.
Max Consumed Write Capacity Units	The maximum consumed write capacity for the table.
Max Returned Item Count	The maximum returned item count for the table.
Max Successful Request Latency	The maximum latency for a successful table operation.
Min Consumed Read Capacity Units	The minimum consumed read capacity of the deployment configuration for this table.
Min Consumed Write Capacity Units	The minimum consumed write capacity for the table.
Min Returned Item Count	The minimum returned item count for the table.
Min Successful Request Latency	The minimum latency for successful table operations.
Returned Item Count	The datum counts of returned items for the table.
Successful Request Latency Counts	The number of successful table operations.
Summed Consumed Read Capacity Units	The sum total consumed read capacity of the deployment configuration for this table.
Summed Consumed Write Capacity Units	The sum total consumed write capacity for the table.
Summed Returned Item Count	The sum total of returned items for the table.
Summed System Errors	The sum of all system error counts.
Summed Throttled Requests	The sum of all table throttling requests.
Summed User Errors	The sum of all user error counts.
System Errors Counts	The datum count of all system errors.
Throttled Requests Count	The datum counts for all table throttling requests.
User Errors Counts	The datum counts for all user errors.

AWS: DDB Service Health

Object Name	Object Description
Date	The timestamp of the health update of this service.
Health	A text description of the health status of this AWS service.

Direct Connect Service

AWS: Direct Connect Instance Configuration	
Object Name	**Object Description**
AWS Device	The Direct Connection endpoint which the physical connection terminates on.
AWS Device v2	The Direct Connect endpoint on which the physical connection terminates.
Bandwidth	Bandwidth of the connection. Example: 1Gbps (for regular connections), or 500Mbps (for hosted connections). A connection represents the physical network connection between the AWS Direct Connect location and the customer.
Connection ID	The ID of the connection.
Connection Name	The name of the connection. A connection represents the physical network connection between the AWS Direct Connect location and the customer.
Has Logical Redundancy	Indicates whether the connection supports a secondary BGP peer in the same address family (IPv4/IPv6). Possible values are yes, no or unknown.
Jumbo Frame Capable	Indicates whether jumbo frames (9001 MTU) are supported. Possible values are True or False.
Key	The key of the tag.
LAG ID	The ID of the LAG.
Location	Where the connection is located. A connection represents the physical network connection between the AWS Direct Connect location and the customer.

Owner Account	The account ID of the owner of the connection. A connection represents the physical network connection between the AWS Direct Connect location and the customer.
Provided By	AWS Direct Connect partner who is a member of the AWS Partner Network (APN). This is optional if your network is collocated with an existing AWS Direct Connect location.
Region	The AWS Region where the connection is located.
State	State of the connection. A connection represents the physical network connection between the AWS Direct Connect location and the customer. Valid Values: ordering | requested | pending | available | down | deleting | deleted | rejected Ordering: The initial state of a hosted connection provisioned on an interconnect. The connection stays in the ordering state until the owner of the hosted connection confirms or declines the connection order. Requested: The initial state of a standard connection. The connection stays in the requested state until the Letter of Authorization (LOA) is sent to the customer. Pending: The connection has been approved, and is being initialized. Available: The network link is up, and the connection is ready for use. Down: The network link is down. Deleted: The connection has been deleted. Rejected: A hosted connection in the 'Ordering' state will enter the 'Rejected' state if it is deleted by the end customer.
Tags	Describes the tags associated with the specified Direct Connect resources.
Time DescribeLoa	The time of the most recent call to DescribeLoa for this connection.
Value	The value of the tag.
VLAN	The ID of the VLAN.

AWS: Direct Connect Instance Performance	
Object Name	**Object Description**
Bit Rate Ingress Data	The bit rate for inbound data to the AWS side of the connection.
Bit Rate Outbound Data	The bit rate for outbound data from the AWS side of the connection.

Connection State	The state of the connection. 0 indicates DOWN and 1 indicates UP.
CRC Error	The number of times cyclic redundancy check (CRC) errors are observed for the data received at the connection.
Health Inbound Traffic	"Indicates the health of the fiber connection for ingress (inbound) traffic to the AWS side of the connection. This metric is available for connections with 10 Gbps port speeds only."
Health Outbound Traffic	"Indicates the health of the fiber connection for egress (outbound) traffic from the AWS side of the connection. This metric is available for connections with 10 Gbps port speeds only."
Packet Rate Ingress Data	The packet rate for inbound data to the AWS side of the connection.
Packet Rate Outbound Data	The packet rate for outbound data from the AWS side of the connection.

AWS: Direct Connect Virtual Interface Configuration	
Object Name	**Object Description**
Account	The AWS account that will own the new virtual interface.
Address Family	Indicates the address family for the BGP peer.
Address Family BGP	Indicates the address family for the BGP peer.
Amazon IP	Indicates the address family for the BGP peer.
Amazon Peer IP	IP address assigned to the Amazon interface. Example: 192.168.1.1/30
Amazon Side ASN	The autonomous system number (ASN) for the Amazon side of the connection.
Authentication Key	The authentication key for BGP configuration.
AWS Device	The Direct Connection endpoint which the virtual interface terminates on.
AWS Direct Connect/Virtual Private Gateway	The unique identifier for the Direct Connect-Virtual Private Gateway relationship.
BGP ASN	Autonomous system (AS) number for Border Gateway Protocol (BGP) configuration.

BGP Authentication Key	Authentication key for BGP configuration.
BGP peers	A list of the BGP peers configured on this virtual interface.
BGP Peers ASN	The autonomous system (AS) number for Border Gateway Protocol (BGP) configuration.
CIDR	CIDR notation for the advertised route.
Connection	The ID of the connection.
Customer IP	IP address assigned to the customer interface.
Device	The Direct Connection endpoint which the BGP peer terminates on.
Direct Connect Gateway	The ID of the direct connect gateway.
ID	The AWS resource ID for the virtual interface.
Jumbo Frame Capable	Indicates whether jumbo frames (9001 MTU) are supported. Possible values are True or False.
Key	The key of the tag.
Location	Where the connection is located.
MTU	The maximum transmission unit (MTU), in bytes. The supported values are 1500 and 9001. The default value is 1500.
Name	The name of the virtual interface assigned by the customer. Example: "My VPC"
Region	The AWS Region where the virtual interface is located.
Router Configuration	Information for generating the customer router configuration.
Routes	A list of routes to be advertised to the AWS network in this region (public virtual interface).

State	State of the virtual interface. One of: Confirming: The creation of the virtual interface is pending confirmation from the virtual interface owner. If the owner of the virtual interface is different from the owner of the connection on which it is provisioned, then the virtual interface will remain in this state until it is confirmed by the virtual interface owner. Verifying: This state only applies to public virtual interfaces. Each public virtual interface needs validation before the virtual interface can be created. Pending: A virtual interface is in this state from the time that it is created until the virtual interface is ready to forward traffic. Available: A virtual interface that is able to forward traffic. Down: A virtual interface that is BGP down. Deleting: A virtual interface is in this state immediately after calling DeleteVirtualInterface until it can no longer forward traffic. Deleted: A virtual interface that cannot forward traffic. Rejected: The virtual interface owner has declined creation of the virtual interface. If a virtual interface in the 'Confirming' state is deleted by the virtual interface owner, the virtual interface will enter the 'Rejected' state.
State BGP	The state of the BGP peer.
Tags	Describes the tags associated with the specified Direct Connect resources.
Type	The type of virtual interface. Example: private (Amazon VPC) or public (Amazon S3, Amazon DynamoDB, and so on.)
Up/Down State	The Up/Down state of the BGP peer
Value	The value of the tag.
Virtual Gateway	The ID of the virtual private gateway to a VPC. This only applies to private virtual interfaces.
VLAN	The VLAN ID. Example: 101
Your Peer IP	IP address assigned to the customer interface. Example: 192.168.1.2/30

DMS Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: DMS Replication Instance Performance	
Object Name	**Object Description**
Available Memory - Avg	An estimate of how much memory is available for starting new applications, without swapping. For more information, see MemAvailable value in /proc/memInfo section of the Linux man-pages.
Cpuallocated - Avg	The percentage of CPU maximally allocated for the task (0 means no limit).
Cpuutilization - Avg	The percentage of allocated vCPU (virtual CPU) currently in use on the instance.
Disk Queue Depth - Avg	The number of outstanding read/write requests (I/Os) waiting to access the disk.
Free Memory - Avg	The amount of physical memory available for use by applications, page cache, and for the kernel's own data structures. For more information, see MemFree value in /proc/memInfo section of the Linux man-pages.
Free Storage Space - Avg	The amount of available storage space.
Freeable Memory - Avg	The amount of available random access memory.
Memory Allocated - Avg	The maximum allocation of memory for the task (0 means no limits).
Network Receive Throughput - Avg	The incoming (Receive) network traffic on the replication instance, including both customer database traffic and AWS DMS traffic used for monitoring and replication.
Network Transmit Throughput - Avg	The outgoing (Transmit) network traffic on the replication instance, including both customer database traffic and AWS DMS traffic used for monitoring and replication.
Read Iops - Avg	The average number of disk read I/O operations per second.
Read Latency - Avg	The average amount of time taken per disk I/O (input) operation.
Read Throughput - Avg	The average number of bytes read from disk per second.
Swap Usage - Avg	The amount of swap space used on the replication instance.

Write Iops - Avg	The average number of disk write I/O operations per second.
Write Latency - Avg	The average amount of time taken per disk I/O (output) operation.
Write Throughput - Avg	The average number of bytes written to disk per second.

DocumentDB Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: DocDB Latency and System Performance	
Object Name	**Object Description**
Buffer Cache Hit Ratio - Avg	The percentage of requests that are served by the buffer cache.
Dbcluster Replica Lag Maximum - Max	The maximum amount of lag, in milliseconds, between the primary instance and each Amazon DocumentDB instance in the cluster.
Dbcluster Replica Lag Minimum - Min	The minimum amount of lag, in milliseconds, between the primary instance and each replica instance in the cluster.
Dbinstance Replica Lag - Avg	The amount of lag, in milliseconds, when replicating updates from the primary instance to a replica instance.
Disk Queue Depth - Avg	The number of concurrent write requests to the distributed storage volume.
Engine Uptime - Avg	The amount of time, in seconds, that the instance has been running.
Index Buffer Cache Hit Ratio - Avg	The percentage of index requests that are served by the buffer cache. You might see a spike greater than 100% for the metric right after you drop an index, collection or database. This will automatically be corrected after 60 seconds. This limitation will be fixed in a future patch update.
Read Latency - Avg	The average amount of time taken per disk I/O operation.
Write Latency - Avg	The average amount of time, in milliseconds, taken per disk I/O operation.

AWS: DocDB Operations Performance	
Object Name	**Object Description**
Documents Deleted - Sum	The number of deleted documents in a one-minute period.
Documents Inserted - Sum	The number of inserted documents in a one-minute period.
Documents Returned - Sum	The number of returned documents in a one-minute period.
Documents Updated - Sum	The number of updated documents in a one-minute period.
Opcounters Command - Sum	The number of commands issued in a one-minute period.
Opcounters Delete - Sum	The number of delete operations issued in a one-minute period.
Opcounters Getmore - Sum	The number of getmores issued in a one-minute period.
Opcounters Insert - Sum	The number of insert operations issued in a one-minute period.
Opcounters Query - Sum	The number of queries issued in a one-minute period.
Opcounters Update - Sum	The number of update operations issued in a one-minute period.
Transactions Aborted - Sum	The number of transactions aborted on an instance in a one-minute period.
Transactions Committed - Sum	The number of transactions committed on an instance in a one-minute period.
Transactions Started - Sum	The number of transactions started on an instance in a one-minute period.
Ttldeleted Documents - Sum	The number of documents deleted by a TTLMonitor in a one-minute period.

AWS: DocDB Resource Performance	
Object Name	**Object Description**

Backup Retention Period Storage Used - Avg	The number of documents deleted by a TTLMonitor in a one-minute period.
Change Stream Log Size - Avg	The amount of storage used by your cluster to store the change stream log in megabytes. This value is a subset of the total storage for the cluster (VolumeBytesUsed) and affects the cost of the cluster. For storage pricing information, see the Amazon DocumentDB product page. The change stream log size is a function of how much change is happening on your cluster and the change stream long retention duration. For more information on change streams, see Using Change Streams with Amazon DocumentDB.
Cpuutilization - Avg	The percentage of CPU used by an instance.
Database Connections - Avg	The number of connections open on an instance taken at a one-minute frequency.
Database Connections Max - Avg	The maximum number of open database connections on an instance in a one-minute period.
Database Cursors - Avg	The number of cursors open on an instance taken at a one-minute frequency.
Database Cursors Max - Avg	The maximum number of open cursors on an instance in a one-minute period.
Database Cursors Timed Out - Avg	The number of cursors that timed out in a one-minute period.
Free Local Storage - Avg	This metric reports the amount of storage available to each instance for temporary tables and logs. This value depends on the instance class. You can increase the amount of free storage space for an instance by choosing a larger instance class for your instance.
Freeable Memory - Avg	The amount of available random access memory, in bytes.
Low Mem Num Operations Throttled - Avg	The number of requests that are throttled due to low available memory in a one-minute period.
Low Mem Throttle Max Queue Depth - Avg	The maximum queue depth for requests that are throttled due to low available memory in a one-minute period.
Low Mem Throttle Queue Depth - Avg	The queue depth for requests that are throttled due to low available memory taken at a one-minute frequency.
Snapshot Storage Used - Avg	The total amount of backup storage in GiB consumed by all snapshots for a given Amazon DocumentDB cluster outside its backup retention window. Included in the total reported by the TotalBackupStorageBilled metric. Computed separately for each Amazon DocumentDB cluster.
Swap Usage - Avg	The amount of swap space used on the instance.

Total Backup Storage Billed - Avg	The total amount of backup storage in GiB for which you are billed for a given Amazon DocumentDB cluster. Includes the backup storage measured by the BackupRetentionPeriodStorageUsed and SnapshotStorageUsed metrics. Computed separately for each Amazon DocumentDB cluster.
Transactions Open - Avg	The number of transactions open on an instance taken at a one-minute frequency.
Transactions Open Max - Avg	The maximum number of transactions open on an instance in a one-minute period.
Volume Bytes Used - Avg	The amount of storage used by your cluster in bytes. This value affects the cost of the cluster. For pricing information, see the Amazon DocumentDB product page.

AWS: DocDB Throughput Performance	
Object Name	**Object Description**
Network Receive Throughput - Avg	The amount of network throughput, in bytes per second, received from clients by each instance in the cluster. This throughput doesn't include network traffic between instances in the cluster and the cluster volume.
Network Throughput - Avg	The amount of network throughput, in bytes per second, both received from and transmitted to clients by each instance in the Amazon DocumentDB cluster. This throughput doesn't include network traffic between instances in the cluster and the cluster volume.
Network Transmit Throughput - Avg	The amount of network throughput, in bytes per second, sent to clients by each instance in the cluster. This throughput doesn't include network traffic between instances in the cluster and the cluster volume.
Read Iops - Avg	The average number of disk read I/O operations per second. Amazon DocumentDB reports read and write IOPS separately, and on one-minute intervals.
Read Throughput - Avg	The average number of bytes read from disk per second.

Volume Read Iops - Avg	The average number of billed read I/O operations from a cluster volume, reported at 5-minute intervals. Billed read operations are calculated at the cluster volume level, aggregated from all instances in the cluster, and then reported at 5-minute intervals. The value is calculated by taking the value of the read operations metric over a 5-minute period. You can determine the amount of billed read operations per second by taking the value of the billed read operations metric and dividing by 300 seconds. For example, if the VolumeReadIOPs returns 13,686, then the billed read operations per second is 45 (13,686 / 300 = 45.62). You accrue billed read operations for queries that request database pages that are not present in the buffer cache and therefore must be loaded from storage. You might see spikes in billed read operations as query results are read from storage and then loaded into the buffer cache.
Volume Write Iops - Avg	The average number of billed write I/O operations from a cluster volume, reported at 5-minute intervals. Billed write operations are calculated at the cluster volume level, aggregated from all instances in the cluster, and then reported at 5-minute intervals. The value is calculated by taking the value of the write operations metric over a 5-minute period. You can determine the amount of billed write operations per second by taking the value of the billed write operations metric and dividing by 300 seconds. For example, if the VolumeWriteIOPs returns 13,686, then the billed write operations per second is 45 (13,686 / 300 = 45.62). Note that VolumeReadIOPs and VolumeWriteIOPs metrics are calculated by the DocumentDB storage layer and it includes IOs performed by the primary and replica instances. The data is aggregated every 20-30 minutes and then reported at 5-minute intervals, thus emitting the same data point for the metric in the time period. If you are looking for a metric to correlate to your insert operations over a 1-minute interval, you can use the instance level WriteIOPS metric. The metric is available in the monitoring tab of your Amazon DocumentDB primary instance.
Write Iops - Avg	The average number of disk write I/O operations per second. When used on a cluster level, WriteIOPs are evaluated across all instances in the cluster. Read and write IOPS are reported separately, on 1-minute intervals.

Write Throughput - Avg	The average number of bytes written to disk per second.

AWS: DocDB Throughput Performance - VOLUME	
Object Name	**Object Description**
Volume Bytes Used - Avg	The average number of bytes used from a cluster volume.
Volume Read Iops - Avg	The average number of billed read I/O operations from a cluster volume, reported at 5-minute intervals. Billed read operations are calculated at the cluster volume level, aggregated from all instances in the cluster, and then reported at 5-minute intervals. The value is calculated by taking the value of the read operations metric over a 5-minute period. You can determine the amount of billed read operations per second by taking the value of the billed read operations metric and dividing by 300 seconds. For example, if the VolumeReadIOPs returns 13,686, then the billed read operations per second is 45 (13,686 / 300 = 45.62). You accrue billed read operations for queries that request database pages that are not present in the buffer cache and therefore must be loaded from storage. You might see spikes in billed read operations as query results are read from storage and then loaded into the buffer cache.

Volume Write Iops - Avg	The average number of billed write I/O operations from a cluster volume, reported at 5-minute intervals. Billed write operations are calculated at the cluster volume level, aggregated from all instances in the cluster, and then reported at 5-minute intervals. The value is calculated by taking the value of the write operations metric over a 5-minute period. You can determine the amount of billed write operations per second by taking the value of the billed write operations metric and dividing by 300 seconds. For example, if the VolumeWriteIOPs returns 13,686, then the billed write operations per second is 45 (13,686 / 300 = 45.62). Note that VolumeReadIOPs and VolumeWriteIOPs metrics are calculated by the DocumentDB storage layer and it includes IOs performed by the primary and replica instances. The data is aggregated every 20-30 minutes and then reported at 5-minute intervals, thus emitting the same data point for the metric in the time period. If you are looking for a metric to correlate to your insert operations over a 1-minute interval, you can use the instance level WriteIOPS metric. The metric is available in the monitoring tab of your Amazon DocumentDB primary instance.

DynamoDB Service

AWS: DynamoDB Performance	
Object Name	**Object Description**
Conditional Check Failed Requests	The number of failed attempts to perform conditional writes.
Online Index Consumed Write Capacity	The number of write capacity units consumed when adding a new global secondary index to a table.
Online Index Percentage Progress	The percentage of completion when a new global secondary index is being added to a table.
Online Index Throttle Events	The number of write throttle events that occur when adding a new global secondary index to a table.
Provisioned Read Capacity Units	The number of provisioned read capacity units for a table or a global secondary index.

Provisioned Write Capacity Units	The number of provisioned write capacity units for a table or a global secondary index
Read Throttle Events	Requests to DynamoDB that exceed the provisioned read capacity units for a table or a global secondary index.
Returned Bytes	The number of bytes returned by GetRecords operations (Amazon DynamoDB Streams) during the specified time period.
Returned Records Count	The number of stream records returned by GetRecords operations (Amazon DynamoDB Streams) during the specified time period.
Successful Request Latency	The number of successful table operations requests.
Throttled Requests	Requests to DynamoDB that exceed the provisioned throughput limits on a table.
Time To Live Deleted Item Count	The number of items deleted by Time To Live (TTL) during the specified time period.
Write Throttle Events	Requests to DynamoDB that exceed the provisioned write capacity units for a table or a global secondary index.

EBS Service

AWS: EBS Instance Configuration	
Object Name	**Object Description**
Volume ID	The unique identifier of this EBS volume.
IO Capacity	If provisioned, the peak IOPS of the volume.
Storage Capacity	The storage capacity of the EBS volume.
Type	Whether or not this EBS volume has been provisioned (type consistent-iops or standard).
Volume Status	The volume state.
EC2 Instance ID	The ID of the EC2 instance that is attached to the EBS Volume if any.
Attached State	Shows whether the EBS volume is attached to an EC2 instance.

Capacity	The size of the snapshot.
Description	The description of the snapshot.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances on the account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Last State Count	Quantity of the last state collected only for creating and deleting states.
Owner	The ID of the snapshot.
Progress	The progress towards completion of the snapshot.
Snapshot ID	The AWS ID of the snapshot.
Snapshots	The volume of the EBS snapshot.
Start Time	The time the snapshot was started.
Status	The status of the snapshot.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances on the account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.

AWS: EBS Instance Performance	
Object Name	**Object Description**
Burst Balance	The percent of General Purpose SSD (gp2) burst-bucket I/O credits available.
EBS Storage Capacity	The storage capacity of the EBS volume.
Volume Consumed Read Write Ops	The consumed total IOPS on the volume.
Volume Idle Time	The average time that the volume has been idle.
Volume Queue Length	The average queue length for the volume.
Volume Read Bytes	The average bytes read from the volume.
Volume Read Ops	The average number of read operations for the volume.
Volume Throughput Percent	The percentage of available throughput currently consumed by the volume.
Volume Total Read Time	The total read times for the volume.
Volume Total Write Time	The total write times for the volume.
Volume Write Bytes	The average bytes written by the volume.
Volume Write Ops	The average write operations for the volume.

EC2 Service

AWS: EC2 Instance Configuration	
Object Name	**Object Description**
AMI Launch Index	The Amazon Machine Image (AMI) index within the EC2's launch reservation list.
Architecture	The architecture of the image.
Architecture	The instance architecture. Type: String. Valid values: i386 | x86_64
Availability Zone	The Availability Zone under which the EC2 instance will be launched.
AWS EC2 Identifier Namespace	The namespace used to link the CCC application.
AWS EC2/EBS Volume	The volume ID of the Amazon EBS volume.

Block Device Attach Time	The attach time for an Amazon EBS volume mapped to the instance (for example, 2010-09-15T17:15:20.000Z)
Block Device Delete On Termination	Indicates whether the Amazon EBS volume is deleted on instance termination.
Block Device Name	The device name exposed to the instance (for example, /dev/sdh, or xvdh).
Block Device Status	The status for the Amazon EBS volume. Valid values: attaching | attached | detaching | detached
CCC Application in AWS EC2 Instance	CCC Applications in EC2 Instances.
Client Token	The idempotency token you provided when you launched the instance. Type: String
Delete On Termination	Indicates whether the EBS volume is deleted on instance termination.
Description	The description of the AMI was provided during image creation.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
EBS Optimized	Indicates whether the instance is optimized for EBS I/O. Type: xsd:boolean
EC2 Unique ID	The unique identifier of the EC2 Instance that matches the identifier given by the CCC application.
Elastic IP	One or more elastic IP addresses assigned to the EC2 instance.
ENA Support	Specifies whether enhanced networking with ENA is enabled.
Encrypted	Indicates whether the encryption state of an EBS volume is changed while being restored from a backing snapshot.
ENI Description	The ENI description.
ENI ID	The ID of the ENI.
ENI MAC Address	The MAC address of the interface.
ENI Owner ID	The ID of the owner of the ENI.
ENI Private IP Address	The IP address of the interface within the subnet.
ENI Source/Dest Check	Flag to indicate whether to validate network traffic to or from this network interface.

ENI Status	The interface's status (available|in-use).
ENI Subnet ID	The ID of the VPC subnet.
ENI VPC ID	The ID of the VPC.
Group	The name of the initial security group to which this EC2 instance is assigned.
Groups	A list of all current groups to which this EC2 instance belongs. Within a VPC, this can change post-launch.
Hypervisor	The hypervisor type of the image.
Hypervisor	The type of root device that the instance uses. Type: String. Valid values: ebs | instance-store
Image Creation Date	The date and time the image was created.
Image ID	The Amazon Machine Image (AMI) id which this instance is booted from.
Image Location	The location of the AMI.
Image Name	The name of the AMI was provided during image creation.
Image Type	The type of image.
Instance ID	The unique identifier of this instance.
Instance Index	The "device name" (if any) of this instance along with its zone and ID. The "device name" is the value of the 'Name' key in the instance's launch-time tag set.
Instance Profile	The instance profile id and arn associated with this instance.
Instance State	The string representation of the instance’s current state.
Instance Type	The cpu, core, and memory capacity type-name of this instance. For example: t1.micro, c1.medium, m1.xlarge.
Kernel ID	The kernel's unique identifier for this Amazon Machine Image (AMI). This is most useful when running a User Provided Kernel (UPK).

Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Key Pair Name	The name of the SSH key associated with the instance.
Launch Time	The UTC time when this instance was launched.
Lifecycle	Whether or not this instance is 'normal'; or 'spot'. A 'spot' lifecycle means this instance's compute-time was purchased in the 'spot market' and has limitations on how, where, and when it can run.
Monitoring Level	Whether or not AWS CloudWatch metrics are reported on a 'basic' (5 min.) or 'detailed' (1 min.) intervals.
Monitoring State	Indicates whether monitoring is enabled for the instance. Type: String. Valid values: disabled | enabled
Persistent Lifecycle	Whether or not Termination Protection is enabled.
Placement Group	The Security Group where the instance is placed upon association with a Virtual Private Cloud (VPC).

Platform	A high-level name of the operating system running on the instance. If such instance is Linux/Unix based the Collection Object won't be visible, only if it is Windows
Platform Details	The platform details associated with the billing code of the AMI.
Private DNS Name	The private AWS Cloud DNS name associated with the instance. Each instance will always have at least one Private DNS name.
Private IP Address	The AWS cloud-local IP address.
Public	Indicates whether the image has public launch permissions.
Public DNS Name	The Public DNS Name (if any) associated with the instance. This DNS name is generally Internet routable and associated with one or more Elastic IPs.
Public IP Address	The public IP address of the instance.
Ramdisk ID	The unique identifier of the ramdisk (if any) associated with the instance.
Region	The AWS Region (datacenter) where the instance is located.
Reservation	The unique identifier associated with the Reservation which launched this instance.
Root Device	The root block device (filesystem) of the instance.
Root Device Name	The device name of the root device volume.
Root Device Type	The type of root device used by the AMI.
Root Device Type	The type of root device that the instance uses. Type: String. Valid values: ebs | instance-store
Security Group Id	The ID of the security group.
Security Group Name	The name of the security group.
Snapshot ID	The ID of the snapshot.
Sriov Net Support	Specifies whether enhanced networking with the Intel 82599 Virtual Function interface is enabled.
State	The current state of the AMI.
State Reason	The reason for the most recent state transition, typically "User initiated", that the state of the instance has changed.

Subnet Component Unique Identifier	The Skylar One component unique identifier for the associated subnet, if any.
Subnet ID	If associated with a VPC, this is the AWS identifier of the subnet within the VPC where the instance resides.
Tag	The key of the tag.
Tenancy	Multi IAM placement status
Value	The value of the tag.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Virtualization Type	The type of virtualization of the AMI.
Virtualization Type	This is always listed as 'paravirtual' or 'hvm' (xen or ovm).
Volume Size	The size of the volume, in GiBs.
Volume Type	The volume type.
VPC ID	If associated with a Virtual Private Cloud (VPC), this is the VPC AWS identifier.

AWS: EC2 Instance Performance	
Object Name	**Object Description**
CPU Utilization	The percentage of allocated EC2 compute units that are currently in use on the instance. This metric identifies the processing power required to run an application upon a selected instance.
Disk Read Bytes	Bytes read from all ephemeral disks available to the instance (if your instance uses Amazon EBS, see Amazon EBS Metrics.) This metric is used to determine the volume of the data the application reads from the hard disk of the instance. This can be used to determine the speed of the application.
Disk Read Ops	Completed read operations from all ephemeral disks available to the instance (if your instance uses Amazon EBS, see Amazon EBS Metrics.) This metric identifies the rate at which an application reads a disk. This can be used to determine the speed in which an application reads data from a hard disk.
Disk Write Bytes	Bytes written to all ephemeral disks available to the instance (if your instance uses Amazon EBS, see Amazon EBS Metrics.) This metric is used to determine the volume of the data the application writes onto the hard disk of the instance. This can be used to determine the speed of the application.
Disk Write Ops	Completed write operations to all ephemeral disks available to the instance (if your instance uses Amazon EBS, see Amazon EBS Metrics.) This metric identifies the rate at which an application writes to a hard disk. This can be used to determine the speed in which an application saves data to a hard disk.
Network In	The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an application on a single instance.
Network Out	The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic to an application on a single instance.

Network Packets In	The number of packets received on all network interfaces by the instance. This metric identifies the volume of incoming traffic in terms of the number of packets on a single instance. This metric is available for basic monitoring only.
Network Packets Out	The number of packets sent out on all network interfaces by the instance. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single instance. This metric is available for basic monitoring only.
Status Check Failed	A combination of StatusCheckFailed_Instance and StatusCheckFailed_System that reports if either of the status checks has failed. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure. Note Status check metrics are available at 1 minute frequency. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.
Status Check Failed for Instance	Reports whether the instance has passed the EC2 instance status check in the last minute. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure. Note Status check metrics are available at 1 minute frequency. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.
Status Check Failed for System	Reports whether the instance has passed the EC2 system status check in the last minute. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure. Note Status check metrics are available at 1 minute frequency. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.

AWS: EC2 Service Health	
Object Name	**Object Description**
Date	The timestamp of this health status update string.
Health	A text description of this AWS Service's health status.

ECS Service

AWS: ECS Cluster Instance Configuration	
Object Name	**Object Description**
Active Services	The number of services that are running on the cluster in an ACTIVE state.
Agent Connected	This parameter returns true if the agent is connected to Amazon ECS. Registered instances with an agent that may be unhealthy or stopped return false.
AWS ECS Instance / EC2	The EC2 instance ID of the container instance.
Cluster ARN	The Amazon Resource Name (ARN) that identifies the cluster.
Cluster Name	The name of the cluster.
Cluster Status	The status of the cluster. The valid values are ACTIVE or INACTIVE .
Container Instance ARN	The Amazon Resource Name (ARN) of the container instance.
Container Instance ARN	The Amazon Resource Name (ARN) of the container instance.
Key	The key of the tag, that belongs to the ECS Cluster instance.
Pending Tasks	The number of tasks in the cluster that are in the PENDING state.
Pending Tasks	The number of tasks on the container instance that are in the PENDING status.

Registered At	The Unix time stamp for when the container instance was registered.
Registered Container Instances	The number of container instances registered into the cluster.
Registered CPU	This value represents the total amount reserved of the resource CPU that can be allocated on the container instance to tasks.
Registered Memory	This value represents the total amount reserved of the resource Memory that can be allocated on the container instance to tasks.
Registered Ports	This value represents the total amount reserved of the resource Ports that can be allocated on the container instance to tasks.
Registered Ports UDP	This value represents the total amount reserved of the resource Ports UDP that can be allocated on the container instance to tasks.
Remaining CPU	This value represents the remaining of resource CPU that has not already been allocated to tasks and is therefore available for new tasks.
Remaining Memory	This value represents the remaining of resource Memory that has not already been allocated to tasks and is therefore available for new tasks.
Remaining Ports	This value represents the remaining of resource Ports that has not already been allocated to tasks and is therefore available for new tasks.
Remaining Ports UDP	This value represents the remaining of resource Ports UDP that has not already been allocated to tasks and is therefore available for new tasks.
Running Tasks	The number of tasks in the cluster that are in the RUNNING state.
Running Tasks	The number of tasks on the container instance that are in the RUNNING status.
Statistic Name	The name of the key value pair. For environment variables, this is the name of the environment variable. (EC2 Launch Type)
Statistic Name	The name of the key value pair. For environment variables, this is the name of the environment variable. (Fargate Launch Type)

Statistic Value	The value of the key value pair. For environment variables, this is the value of the environment variable. (EC2 Launch Type)
Statistic Value	The value of the key value pair. For environment variables, this is the value of the environment variable. (Fargate Launch Type)
Status	The status of the container instance. The valid values are ACTIVE , INACTIVE , or DRAINING.
Value	The value of the tag, that belongs to the ECS Cluster instance.

AWS: ECS Cluster Instance Performance	
Object Name	**Object Description**
CPU Reservation	"The percentage of CPU units that are reserved by running tasks in the cluster. Cluster CPU reservation is measured as the total CPU units that are reserved by Amazon ECS tasks on the cluster, divided by the total CPU units that were registered for all of the container instances in the cluster. This metric is only used for tasks using the EC2 launch type."
CPU Utilization	"The percentage of CPU units that are used in the cluster or service. Cluster CPU utilization is measured as the total CPU units in use by Amazon ECS tasks on the cluster, divided by the total CPU units that were registered for all of the container instances in the cluster. Cluster CPU utilization metrics are only used for tasks using the EC2 launch type."
Memory Reservation	"The percentage of memory that is reserved by running tasks in the cluster. Cluster memory reservation is measured as the total memory that is reserved by Amazon ECS tasks on the cluster, divided by the total amount of memory that was registered for all of the container instances in the cluster. This metric is only used for tasks using the EC2 launch type."

Memory Utilization	"The percentage of memory that is used in the cluster or service. Cluster memory utilization is measured as the total memory in use by Amazon ECS tasks on the cluster, divided by the total amount of memory that was registered for all of the container instances in the cluster. Cluster memory utilization metrics are only used for tasks using the EC2 launch type."

AWS: ECS Cluster Services Configuration	
Object Name	**Object Description**
AWS ECS Service/Classic Load Balancer	The load balancer used for the service if it has created with a classic load balancer type.
AWS ECS Service/Security Groups	The security groups associated with this ECS service.
AWS ECS Service/Subnets	The subnets associated with this ECS service.
AWS ECS Service/Target Group	The target group that is associated with this service if it was created with an application or network load balancer.
AWS ECS Service/VPC	The VPC associated with the service.
Container Name	The container being used with this service on the associated application ELB.
Container Port	The port of the container being used with this service on the associated application ELB.
Created At	The service creation time.
Desired Tasks	The desired number of instantiations of the task definition to keep running on the service.
Key	The key of the tag, that belongs to the ECS Cluster services.
Launch Type	Type of ECS launch (EC2 or Fargate).
Pending Tasks	The number of tasks in the cluster that are in the PENDING state.
Platform Version	The platform version on which your task is running.
Running Tasks	The number of tasks in the cluster that are in the RUNNING state.
Service ARN	The Amazon ID (ARN) of the ECS service.

Service Name	Name of the ECS Service.
Service Status	The status of the service. The valid values are ACTIVE , DRAINING , or INACTIVE.
Service Type	The scheduling strategy to use for the service (Daemon or Replica).
Task Definition	The ARN of the ECS task definition.
Value	The value of the tag, that belongs to the ECS Cluster services.

AWS: ECS Cluster Services Performance	
Object Name	**Object Description**
CPU Utilization	The percentage of CPU units that are used in the cluster or service.
Desired Tasks Count	The desired number of instantiations of the task definition to keep running on the service.
Memory Utilization	The percentage of memory that is used in the cluster or service.
Running Tasks Count	The number of tasks in the services that are in the RUNNING state.
Total Pending Deployments	The total pending deployments of the service.
Total Running Deployments	The total running deployments of the service.

AWS: ECS Cluster Tasks Configuration	
Object Name	**Object Description**
Container Id	The Globally Unique Identifier (GUID) of the container from the Amazon Resource Name (ARN).
Name	The name of the container.
Container Id	The Globally Unique Identifier (GUID) of the container from the Amazon Resource Name (ARN).
Image	The image used to start a container. This string is passed directly to the Docker daemon.
Last Status	The last known status of the container.

Memory	The hard limit (in MiB) of memory to present to the container. If your container attempts to exceed the memory specified here, the container is killed.
Bind IP	The IP address that the container is bound to on the container instance.
Connectivity	The connectivity status of a task.
Container Port	The port number on the container that is used with the network binding.
CPU	The number of CPU units used by the task as expressed in a task definition. It can be expressed as an integer using CPU units. If you are using the EC2 launch type, this field is optional. If you are using the Fargate launch type, this field is required.
CPU	The number of cpu units reserved for the container.
Created At	The Unix timestamp for when the task was created (the task entered the PENDING state).
Desired Status	The desired status of the task.
Essential	If the essential parameter of a container is marked as true , and that container fails or stops for any reason, all other containers that are part of the task are stopped. If the essential parameter of a container is marked as false, then its failure does not affect the rest of the containers in a task.
Group	The name of the task group associated with the task.
Health	The health status of the container. If health checks are not configured for this container in its task definition, then it reports the health status as UNKNOWN.
Health Status	The health status for the task, which is determined by the health of the essential containers in the task. If all essential containers in the task are reporting as HEALTHY , then the task status also reports as HEALTHY . If any essential containers in the task are reporting as UNHEALTHY or UNKNOWN , then the task status also reports as UNHEALTHY or UNKNOWN , accordingly.

Host Port	"The port number on the container instance to reserve for your container. If you are using containers in a task with the awsvpc or host network mode, the hostPort can either be left blank or set to the same value as the containerPort. If you are using containers in a task with the bridge network mode, you can specify a non-reserved host port for your container port mapping, or you can omit the hostPort (or set it to 0) while specifying a containerPort and your container automatically receives a port in the ephemeral port range for your container instance operating system and Docker version."
Key	The key of the tag, that belongs to the ECS Cluster tasks.
Last Status	The last known status of the task.
Launch Type	The launch type on which your task is running. Options are: EC2 or Fargate.
MAC Address	The MAC Address of the network interface.
Memory	The amount of memory (in MiB) used by the task as expressed in a task definition. If you are using the EC2 launch type, this field is optional. If you are using the Fargate launch type, this field is required.
Network Interface Id	The Network Interface Id of the elastic network interface attached if any.
Network Mode	The Network mode of the task. Options are: awsvpc, bridge, host or none.
Platform Version	The platform version on which your task is running. A platform version is only specified for tasks using the Fargate launch type. If one is not specified, the LATEST platform version is used by default.
Private IP	Private IP address of the elastic network interface.
Protocol	The protocol used for the network binding. Options are tpc or udp.
Started At	The Unix timestamp for when the task started (the task transitioned from the PENDING state to the RUNNING state).
Started By	The tag specified when a task is started. If the task is started by an Amazon ECS service, then the startedBy parameter contains the deployment ID of the service that starts it.

Subnet Id	The subnet id of the elastic network interface attachment.
Task ARN	The Amazon Resource Name (ARN) of the task.
Task Definition	The ARN of the task definition that creates the task.
Value	The value of the tag, that belongs to the ECS Cluster tasks.
Version	The version counter for the task. Every time a task experiences a change that triggers a CloudWatch event, the version counter is incremented. If you are replicating your Amazon ECS task state with CloudWatch Events, you can compare the version of a task reported by the Amazon ECS API actions with the version reported in CloudWatch Events for the task (inside the detail object) to verify that the version in your event stream is current.

AWS: ECS Service Health	
Object Name	**Object Description**
Date	The date of the RSS feed entry.
Health	The message of the RSS feed entry.

EFS Service

AWS: EFS File System Configuration	
Object Name	**Object Description**
Creation Time	The time that the file system was created.
Creation Token	The opaque string specified in the request.
Encrypted	A Boolean value that, if true, indicates that the file system is encrypted.
File System Id	The ID of the file system, assigned by Amazon EFS.
Key	The tag key (String).

KMS Key Id	The ID of an AWS Key Management Service (AWS KMS) customer master key (CMK) that was used to protect the encrypted file system.
Last Size Update	The latest known time at which the size of data was determined.
Life Cycle State	The lifecycle phase of the file system.
Name	You can add tags to a file system, including a Name tag. If the file system has a Name tag, Amazon EFS returns the value in this field.
Number of Mount Targets	The current number of mount targets that the file system has.
Owner Id	The AWS account that created the file system. If the file system was created by an IAM user, the parent account to which the user belongs is the owner.
Performance Mode	The performance mode of the file system.
Provisioned Throughput	The throughput, measured in MiB/s, that you want to provision for a file system.
Throughput Mode	The throughput mode for a file system. There are two throughput modes to choose from for your file system: bursting and provisioned.
Value	The value of the tag key.

AWS: EFS File System Performance	
Object Name	**Object Description**
Burst Credit Balance	The number of burst credits that a file system has.
Client Connections	The number of client connections to a file system. When using a standard client, there is one connection per mounted Amazon EC2 instance.
Data Read IO Bytes	The number of bytes for each file system read operation.
Data Write IO Bytes	The number of bytes for each file write operation.
Metadata IO Bytes	The number of bytes for each metadata operation.

Percent IO Limit	Shows how close a file system is to reaching the I/O limit of the General Purpose performance mode. If this metric is at 100% more often than not, consider moving your application to a file system using the Max I/O performance mode.
Permitted Throughput	The maximum amount of throughput a file system is allowed. For file systems in the Provisioned Throughput mode, if the amount of storage allows your file system to drive a higher amount of throughput than you provisioned, this metric will reflect the higher throughput instead of the provisioned amount. For file systems in the Bursting Throughput mode, this value is a function of the file system size and BurstCreditBalance.
Total IO Bytes	The number of bytes for each file system operation, including data read, data write, and metadata operations.

AWS: EFS File System Usage Performance	
Object Name	**Object Description**
Size of Infrequent Access Data	The latest known metered size (in bytes) of data stored in the Infrequent Access storage class.
Size of Standard Access Data	The latest known metered size (in bytes) of data stored in the Standard storage class.
Total File System Size	The latest known metered size (in bytes) of data stored in the file system.

EKS Service

AWS: EKS Cluster Instance Configuration	
Object Name	**Object Description**
ARN	The ARN that AWS provides for this cluster within EKS.
Certificate Authority	The Certificate Authority used by this cluster for signing requests.

Client Request Token	Unique, case-sensitive identifier that you provide to ensure the idempotency of the request.
Created At	The time at which this cluster was created.
Endpoint	The REST endpoint used to make API calls into the Kubernetes cluster.
Keys	Tags keys assigned to this cluster.
Logging Enabled	An object representing the enabled or disabled Kubernetes control plane logs for your cluster.
Name	The name of this cluster within EKS.
Platform Version	The platform version of your Amazon EKS cluster.
Role ARN	The ARN of the AWS role used to manage this cluster.
Security Group IDs	The security groups associated with your cluster.
Status	The active status of this cluster.
Subnet IDs	The subnets associated with your cluster.
Tags	Tags assigned to this cluster.
Types	The available cluster control plane log types.
Values	Tags values assigned to this cluster.
Version	The version of Kubernetes that this cluster is running.
VPC	The VPCs associated with your cluster.
VPC Endpoint Private Access	This parameter indicates whether the Amazon EKS private API server endpoint is enabled. If the Amazon EKS private API server endpoint is enabled, Kubernetes API requests that originate from within your cluster's VPC use the private VPC endpoint instead of traversing the internet.
VPC Endpoint Public Access	This parameter indicates whether the Amazon EKS public API server endpoint is enabled. If the Amazon EKS public API server endpoint is disabled, your clusters Kubernetes API server can receive only requests that originate from within the cluster VPC.

AWS: EKS Token Manager

Object Name	Object Description
Token Created	This token is used to authenticate with the EKS Cluster, True if the token was created otherwise False.
Token Expiration Time	The token expiration time.

Elastic Beanstalk Service

AWS: Elastic Beanstalk Application Instance Health Configuration	
Object Name	**Object Description**
Instance Id	The ID of the Amazon EC2 instance.
Availability Zone	The availability zone in which the instance runs.
Color	Represents the color indicator that gives you information about the health of the EC2 instance.
Health Status	Returns the health status of the instance.
Launched At	The time at which the EC2 instance was launched.

AWS: Elastic Beanstalk Application Performance	
Object Name	**Object Description**
Environment Name Cache	Provides cache of Environment Names for consumption by other collection objects.
P10 Latency	The average latency for the slowest 90 percent of requests over the last 10 seconds.
P50 Latency	The average latency for the slowest 50 percent of requests over the last 10 seconds.
P75 Latency	The average latency for the slowest 25 percent of requests over the last 10 seconds.
P85 Latency	The average latency for the slowest 15 percent of requests over the last 10 seconds.
P90 Latency	The average latency for the slowest 10 percent of requests over the last 10 seconds.

P95 Latency	The average latency for the slowest 5 percent of requests over the last 10 seconds.
P99 Latency	The average latency for the slowest 1 percent of requests over the last 10 seconds.
P999 Latency	The average latency for the slowest 0.1 percent of requests over the last 10 seconds.
Request Count	Average number of requests handled by the web server per second over the last 10 seconds.
Status 2xx	The average number of requests over the last 10 seconds that resulted in a 2xx (200, 201, etc.) status code.
Status 3xx	The average number of requests over the last 10 seconds that resulted in a 3xx (300, 301, etc.) status code.
Status 4xx	The average number of requests over the last 10 seconds that resulted in a 4xx (400, 401, etc.) status code.
Status 5xx	The average number of requests over the last 10 seconds that resulted in a 5xx (500, 501, etc.) status code.

AWS: Elastic Beanstalk Data Tier Configuration	
Object Name	**Object Description**
Engine	The name of the database engine to use for this instance.
Instance Class	The database instance type.
Multi Availability Zone	Specifies whether a database instance Multi-AZ deployment needs to be created.
Storage (GB)	The allocated database storage size, specified in gigabytes.
When Deleted	Decides whether to delete or snapshot the DB instance on environment termination.

AWS: Elastic Beanstalk Network Tier Configuration	
Object Name	**Object Description**

Associate Public IP Address	Specifies whether to launch instances with public IP addresses in your Amazon VPC. Instances with public IP addresses do not require a NAT device to communicate with the Internet. You must set the value to true if you want to include your load balancer and instances in a single public subnet.
Instance Subnets	The IDs of the Auto Scaling group subnet or subnets. If you have multiple subnets, specify the value as a single comma-delimited string of subnet IDs (for example, "subnet-11111111,subnet-22222222").
RDS Subnets	Contains the IDs of the database subnets. This is only used if you want to add an Amazon RDS DB Instance as part of your application. If you have multiple subnets, specify the value as a single comma-delimited string of subnet IDs (for example, "subnet-11111111,subnet-22222222").
VPC ID	The ID for your Amazon VPC.

AWS: Elastic Beanstalk Service Health	
Object Name	**Object Description**
Date	A timestamp when the health status of this service was originally written.
Health	A text description of the health status of this AWS service.

AWS: Elastic Beanstalk Web Tier Configuration	
Object Name	**Object Description**
Add Instance When (>)	If the measurement is higher than this number for the breach duration, a trigger is fired. 0 to 20000000.
Allow URL fopen:	Specifies if PHP file functions are allowed to retrieve data from remote locations, such as websites or FTP servers.
Availability Zones	The availability zones configuration for the beanstalk application.
AWS Elastic Beanstalk/EC2	The ec2 associated with the beanstalk application.

Connection Draining Enabled	If the beanstalk load balancer attempts connections to an Amazon EC2 instance before forcibly closing connections.
Cross Zone Load Balancing Enabled	If the beanstalk load balancer attempts connections to an Amazon EC2 instance in multiple availability zones.
Deployment Batch Size (%)	The size of the set of instances to deploy in each batch.
Display Errors	Information about the errors that are common to all actions.
Environment ID	The beanstalk application environment id.
Environment Name	The beanstalk application environment name.
Environment Type	The beanstalk application environment type.
Instance Port	The beanstalk environment port.
Instance Type	The beanstalk environment type.
Log Publication	Copy the log files for the Amazon EC2 instances of your application into the Amazon S3 bucket associated with your application (valid values are true or false).
Max Execution Time	The maximum number of attempts that Elastic Beanstalk attempts to send the message to the web application that will process it before moving the message to the dead letter queue.
Memory Limit	Amount of memory allocated to the PHP environment.
Notification Protocol	Protocol used to send notifications to your endpoint.
Number Instances	Number of instances.
NumProcesses	The number of daemon processes that should be started for the process group when running WSGI applications.
NumThreads	The number of threads to be created to handle requests in each daemon process within the process group when running WSGI applications.
Remove Instance When (<)	If the measurement falls below this number for the breach duration, a trigger is fired. 0 to 20000000.

Rolling Updates Enabled	Rolling configuration update batches can be processed periodically (time-based), with a delay between each batch, or based on health. The following values are supported for RollingUpdateType: Health | Time | Immutable.
Send To	Endpoint where you want to be notified of important events affecting your application.
StaticFiles	Information about setting of the virtual path and directory mappings in the Static Files section of the Modify software configuration page.
WSGIPath	The file that contains the WSGI application. This file must have an application callable (default: application.py).
Zlib Output Compression	Specifies whether or not PHP should use compression for output.

Elastic Transcoder Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Elastic Transcoder Pipeline Performance	
Object Name	**Object Description**
Billed Audio Output - Sum	The number of billable seconds of audio output for a pipeline.
Billed Hd Output - Sum	The number of billable seconds of HD output for a pipeline.
Billed Sd Output - Sum	The number of billable seconds of SD output for a pipeline.
Jobs Completed - Sum	The number of jobs completed by this pipeline.
Jobs Errored - Sum	The number of jobs that failed because of invalid inputs, such as a request to transcode a file that is not in the given input bucket.
Outputs Per Job - Sum	The number of outputs Elastic Transcoder created for a job.

ElastiCache Service

AWS: ElastiCache Cluster Configuration	
Object Name	**Object Description**
Auto Upgrade	Whether or not the cluster engine can be upgrading during maintenance.
Cluster Node Type	The EC2 class of the cluster nodes.
Creation Time	The time the cluster was created.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Endpoint	The DNS entry for the configuration endpoint.
Engine	The cache engine (either Redis or Memecache) of the cluster.
Id	The unique identifier of the cluster.
Maintenance Window	The datetime when the cluster can undergo maintenance.
Name	The name of the cache cluster.
Nodes	The number of nodes in the cluster.
Notification ARN	The Amazon Resource Name (ARN), if any, of the SNS instance for event notifications.
Parameter Groups	The parameter groups of the cluster.
Preferred Zone	The name of the Availability Zone in which the cluster is located or "Multiple" if the cache nodes are located in different Availability Zones.
Security Groups	The security groups of the cluster.
Status	The status of the cluster.
Version	The version of the engine.

AWS: ElastiCache Cluster Performance	
Object Name	**Object Description**

CPU Utilization	The percentage of CPU utilization.
Freeable Memory	The amount of free memory on the cluster.
Network Bytes In	The bytes the host has read from the network.
Network Bytes Out	The number of bytes the host has written to the network.
Swap Usage	The swap used, if any, on the cluster.

AWS: ElastiCache Memcached Node Configuration	
Object Name	**Object Description**
Availability Zone	The Availability Zone where this node was created and now resides, also referred to as Customer Availability Zone.
Creation Time	The creation time of the node.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Endpoint	The DNS entry of the node.
Id	The node's unique identifier.
Name	The name of the cluster node.
Port	The TCP port of the service running on the node.
Status	The node's EC2 status.

AWS: ElastiCache Memcached Node Performance	
Object Name	**Object Description**
Bytes Read into Memcached	The number of bytes that have been read from the network by the cache node.
Bytes Used for Cache Items	The number of bytes used to store cache items.
Bytes Used For Hash	The number of bytes currently used by hash tables.
Bytes Written Out From Memcached	The number of bytes that have been written to the network by the cache node.

Check and Set Bad Values	The number of CAS (check and set) requests the cache has received where the Cas value did not match the Cas value stored.
Check and Set Misses	The number of Cas requests the cache has received where the key requested was not found.
Check and Set Request Hits	The number of Cas requests the cache has received where the requested key was found and the Cas value matched.
Config Get Requests	The cumulative number of "config get" requests.
Config Set Requests	The cumulative number of "config set" requests.
Current Connections	The number of connections connected to the cache. Note that due to the design of Memcached, this will always return a minimum count of 10.
Current Items	The number of items currently stored in the cache. Note that due to the design of Memcached, this will always return a minimum count of 10.
Decrement Hits	The number of decrement requests the cache has received where the requested key was found.
Decrement Misses	The number of decrement requests the cache has received where the requested key was not found.
Delete Hits	The number of delete requests the cache has received where the requested key was found.
Delete Misses	The number of delete requests the cache has received where the requested key was not found.
Evicted Unfetched Items	The number of valid items evicted from the least recently used cache (LRU) which were never touched after being set.
Evictions	The number of non-expired items the cache evicted to allow space for new writes.
Expired Unfetched Items	The number of expired items reclaimed from the LRU which were never touched after being set.
Flush Commands	The number of flush commands the cache has received.
Get Commands	The number of get commands the cache has received.
Get Hits	The number of get requests the cache has received where the key requested was found.

Get Misses	The number of get requests the cache has received where the key requested was not found.
Increment Hits	The number of increment requests the cache has received where the key requested was found.
Increment Misses	The number of increment requests the cache has received where the key requested was not found.
New Connections	The number of new connections the cache has received. This is derived from the memcached total_connections statistic by recording the change in total_connections across a period of time. This will always be at least 1, due to a connection reserved for an ElastiCache.
New Items	The number of new items the cache has received. This is derived from the memcached total_ connections statistic by recording the change in total_connections across a period of time. This will always be at least 1, due to a connection reserved for a ElastiCache.
Reclaimed	The number of expired items the cache evicted to allow space for new writes.
Set Commands	The number of set commands the cache has received.
Slabs Moved	The total number of slab pages that have been moved.
Stored Configurations	The number of configurations stored.
Touch Commands	The cumulative number of "touch" requests.
Touch Hits	The number of keys that have been touched and were given a new expiration time.
Touch Misses	The number of items that have been touched, but were not found.
Unused Memory	The amount of unused memory the cache can use to store items.

AWS: ElastiCache Node Performance	
Object Name	**Object Description**

CPU Credit Balance	The number of earned CPU credits that an instance has accrued since it was launched or started. CPU credit metrics are available at a five-minute frequency only.
CPU Credit Usage	The number of CPU credits spent by the instance for CPU utilization. CPU credit metrics are available at a five-minute frequency only.
CPU Utilization	The percentage of CPU utilization for the entire host.
Freeable Memory	The amount of free memory available on the host.
Network Bytes In	The number of bytes the host has read from the network.
Network Bytes Out	The number of bytes sent out on all network interfaces by the instance.
Network Packets In	The number of packets received on all network interfaces by the instance. This metric identifies the volume of incoming traffic in terms of the number of packets on a single instance.
Network Packets Out	The number of packets sent out on all network interfaces by the instance. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single instance.
Swap Usage	The amount of swap used on the host.

AWS: ElastiCache Redis Node Configuration	
Object Name	**Object Description**
Availability Zone	The Availability Zone where this node was created and now resides, also referred to as Customer Availability Zone.
Creation Time	The datetime when this node was created.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Endpoint	The DNS entry of this node.
Id	The unique identifier of the Redis node.
Name	The name of the Redis node.

Port	The TCP port of the service cache on the node.
Status	The EC2 status of the node.

AWS: ElastiCache Redis Node Performance	
Object Name	**Object Description**
Bytes Used For Cache	The total number of bytes allocated by Redis.
Cache Hits	The number of successful key lookups.
Cache Misses	The number of unsuccessful key lookups.
Current Connections	The number of client connections, excluding connections from read replicas.
Current Items	The number of items in the cache.
Engine CPU Utilization	Provides more precise visibility into the load of the Redis process itself.
Evictions	The number of keys that have been evicted due to the maxmemory limit.
Get Type Commands	The total number of get types of commands.
Hash Based Commands	The total number of commands that are hash-based.
Hyper Log Log Based Cmds	This is derived from the Redis commandstats statistic by summing all of the pf type of commands (pfadd, pfcount, pfmerge).
Key Based Commands	The total number of commands that are key-based.
List Based Commands	The total number of commands that are list-based.
New Connections	The total number of connections that have been accepted by the server during this period.
Reclaimed	The total number of key expiration events.
Replication Bytes	The number of bytes that the primary is sending to all of its replicas.
Replication Lag	This metric is only applicable for a cache node running as a read replica. It represents how far behind, in seconds, the replica is in applying changes from the primary cache cluster.
Save In Progress	This binary metric returns 1 whenever a background save (forked or forkless) is in progress, and 0 otherwise.

Set Based Commands	The total number of commands that are set-based.
Set Type Commands	The total number of set types of commands.
Sorted Set Based Commands	The total number of commands that are sorted set-based.
String Based Commands	The total number of commands that are string-based.

AWS: ElastiCache Service Health	
Object Name	**Object Description**
Date	A timestamp when this service's health status was originally written.
Health	A text description of this AWS Service's health status.

ELB Service

AWS: ELB Instance Configuration	
Object Name	**Object Description**
Availability Zones	The availability zones where the ELB is deployed.
AWS ELB/EC2	The unique identifier set for the EMR-EC2 relationship.
AWS ELB/VPC	The unique identifier set for the VPC relationship.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Key	The key of the tag, that belongs to the elastic load balancer.
Listener Instance Port	The listener port for the ELB.
Load Balanced Instances	The group label for the EMR-EC2 relationship.
Load Balancer Name	The name of the ELB.
Load Balancer Port	The load-balanced port for the ELB.

Load Balancer Protocol	The protocol supported (HTTP, SMTP, etc.) by the ELB.
Value	The value of the tag, that belongs to the application load balancer.

AWS: ELB Instance Performance	
Object Name	**Object Description**
Backend Connection Errors	The count of the number of connections that were not successfully established between the load balancer and the registered instances. Because the load balancer will retry when there are connection errors, this count can exceed the request rate.
Healthy Host Count Average	The average number of healthy hosts in the ELB.
Healthy Host Count Maximum	The maximum number of healthy hosts in the ELB.
HTTPCode_Backend_2xx	The count of the HTTP 200 status codes on the back-end of the ELB.
HTTPCode_Backend_3xx	The count of the HTTP 300 status codes on the back-end of the ELB.
HTTPCode_Backend_4xx	The count of HTTP 400 status codes (errors) on the back-end of the ELB.
HTTPCode_Backend_5xx	The count of the HTTP 500 status codes (errors) on the back-end of the ELB.
Latency Average	The average latency across the ELB.
Latency Maximum	The maximum latency across the ELB.
Latency Minimum	The minimum latency across the ELB.
Request Count	The total requests datum counts for the ELB.
Spillover Count	A count of the total number of requests that were rejected due to the queue being full.
Surge Queue Length	A count of the total number of requests that are pending submission to a registered instance.
Unhealthy Host Count Average	The average number of unhealthy hosts in the ELB.
Unhealthy Host Count Maximum	The maximum number of unhealthy hosts in the ELB.
Unhealthy Host Count Minimum	The minimum number of unhealthy hosts in the ELB.

AWS: ELB Service Health	
Object Name	**Object Description**
Date	The timestamp when this status was written.
Health	A text description of this AWS Service's health status.

AWS: ELB Target Group Instance Configuration	
Object Name	**Object Description**
Application LoadBalancer ID	The dimension resource identifier of the load balancer that forwards traffic to the target group.
AWS Target Group/EC2	The unique identifiers of EC2s for the target groups.
AWS Target Group/VPC	The unique identifiers of VPC for target groups.
HealthCheck Interval	The approximate amount of time, in seconds, between health checks of an individual target.
HealthCheck Path	The destination for the health check request.
HealthCheck Port	The port to use to connect with the target.
HealthCheck Protocol	The protocol to use to connect with the target.
HealthCheck Timeout	The amount of time, in seconds, during which no response means a failed health check.
Healthy Threshold Count	The number of consecutive health checks successes required before considering an unhealthy target healthy.
Key	The keys of the tag, that belongs to the target group.
Load Balancer ARNs	The Amazon Resource Name (ARN) of the application or network load balancers.
Target Group ARN	The Amazon Resource Name (ARN) of the target group.
Target Group Id	The unique identifier for the target group component.
Target Group Name	The name of the target group component.
Target Group Port	The port on which the targets are listening.

Target Group Protocol	The protocol to use for routing traffic to the targets.
Target Group VPC Id	The ID of the VPC for the targets.
Target Type	The type of target that you must specify when registering targets with this target group. The possible values are instance (targets are specified by instance ID) or ip (targets are specified by IP address).
Unhealthy Threshold Count	The number of consecutive health check failures required before considering the target unhealthy.
Value	The value of the tag, that belongs to the target group.

AWS: ELB Target Group Instance Performance	
Object Name	**Object Description**
Healthy Host Count (Average)	The number of targets that are considered healthy.
Healthy Host Count (Maximum)	The number of targets that are considered healthy.
Healthy Host Count (Minimum)	The number of targets that are considered healthy.
HTTPCode Target 2XX Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
HTTPCode Target 3XX Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
HTTPCode Target 4XX Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.

HTTPCode Target 5XX Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
Request Count Per Target	The average number of requests received by each target in a target group. You must specify the target group using the TargetGroup dimension. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
Target Connection Error Count	The number of connections that were not successfully established between the load balancer and target. This metric does not apply if the target is a Lambda function.
Target Response Time	The time elapsed, in seconds, after the request leaves the load balancer until a response from the target is received. This is equivalent to the target_ processing_time field in the access logs. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
TLS Negotiation Error Count per Target Group	The number of TLS connections initiated by the load balancer that did not establish a session with the target. Possible causes include a mismatch of ciphers or protocols. For now we are not collecting data from amazon api, because it is not supported yet for Network ELB associated with Target Group.
Unhealthy Host Count (Average)	The number of targets that are considered unhealthy.
Unhealthy Host Count (Maximum)	The number of targets that are considered unhealthy.
Unhealthy Host Count (Minimum)	The number of targets that are considered unhealthy.

EMR Service

AWS: EMR Instance Configuration

Object Name	Object Description
AWS EMR/EC2	The unique identifier set for the EMR-EC2 relationship.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
EMR Nodes as EC2 Instances	The group label for the EMR-EC2 relationship.
Instance ID	The unique identifier of EMR master instance.
Name	The name of the cluster.
State	The current state of the cluster.

AWS: EMR Instance Performance	
Object Name	**Object Description**
Apps Completed	The number of applications submitted to YARN that have completed.
Apps Failed	The number of applications submitted to YARN that have failed to complete.
Apps Killed	The number of applications submitted to YARN that have been killed.
Apps Pending	The number of applications submitted to YARN that are in a pending state.
Apps Running	The number of applications submitted to YARN that are in a running state.
Apps Submitted	The number of applications submitted to YARN.
Capacity Remaining (GB)	The amount of remaining HDFS disk capacity.
Container Allocated	The number of resource containers allocated by the ResourceManager.
Container Pending	The number of containers in the queue that have not yet been allocated.

Container Pending Ratio	The ratio of pending containers to containers allocated (ContainerPendingRatio = ContainerPending / ContainerAllocated). If ContainerAllocated = 0, then ContainerPendingRatio = ContainerPending. The value of ContainerPendingRatio represents a number, not a percentage. This value is useful for scaling cluster resources based on container allocation behavior.
Container Reserved	The number of containers reserved.
Core Nodes Pending	The average number of pending nodes in the run.
Core Nodes Running	The average number of core nodes running for this run.
Corrupt Blocks	The number of blocks that HDFS reports as corrupted.
Dfs Pending Replication Blocks	The status of block replication: blocks being replicated, age of replication requests, and unsuccessful replication requests.
HBase Backup Failed	Whether the last backup failed. This is set to 0 by default and updated to 1 if the previous backup attempt failed. This metric is only reported for HBase clusters. Use case: Monitor HBase backups Units: Count
HBase Most Recent Backup Duration	The amount of time it took the previous backup to complete. This metric is set regardless of whether the last completed backup succeeded or failed. While the backup is ongoing, this metric returns the number of minutes after the backup started. This metric is only reported for HBase clusters.
HBase Time Since Last Successful Backup	The number of elapsed minutes after the last successful HBase backup started on your cluster. This metric is only reported for HBase clusters.
HDFS Bytes Read	The average number of bytes read from all Hadoop filesystems.
HDFS Bytes Written	The average number of bytes written to the Hadoop filesystems.
HDFS Utilization	The percentage utilization of the available Hadoop filesystems.
Is Idle	Whether or not this instance is idle at the moment.

Jobs Failed	The number of jobs failed (if any) during the current run.
Jobs Running	The number of current actively running jobs in the current run.
Live Data Nodes	The number of active data nodes with IO for the current run.
Live Task Trackers	The Percentage of active task tracking nodes for the current run.
Map Reduce Lost Nodes	The number of nodes allocated to MapReduce that have been marked in a LOST state.
Map Slots Open	The number of map slots currently open.
MapReduce Active Nodes	The number of nodes presently running MapReduce tasks or jobs. Equivalent to YARN metric
MapReduce Decommissioned Nodes	The number of nodes allocated to MapReduce applications that have been marked in a DECOMMISSIONED state.
MapReduce Rebooted Nodes	The number of nodes available to MapReduce that have been rebooted and marked in a REBOOTED state.
MapReduce Total Nodes	The number of nodes presently available to MapReduce jobs.
MapReduce Unhealthy Nodes	The number of nodes available to MapReduce jobs marked in an UNHEALTHY state.
Memory Allocated (MB)	The amount of memory allocated to the cluster.
Memory Available (MB)	The amount of memory available to be allocated.
Memory Reserved (MB)	The amount of memory reserved.
Memory Total (MB)	The total amount of memory in the cluster.
Missing Blocks	The current number of missing blocks.
Pending Deletion Blocks	The number of blocks marked for deletion.
Reduce Slots Open	The number of open reduction slots available.
Remaining Map Tasks	The number of remaining map tasks for this run.
Remaining Map Tasks Per Slot	The average number of remaining map tasks per available slot.
Remaining Reduce Tasks	The number of remaining reduce tasks for this run.

Running Map Tasks	The number of currently running map tasks for this run.
Running Reduce Tasks	The number of currently running reduce tasks for this run.
S3 Bytes Read	The number of Simple Storage Service (S3) bytes read during this run.
S3 Bytes Written	The number of Simple Storage Service (S3) bytes written during this run.
Task Nodes Pending	The number of pending (inactive) task nodes at the moment.
Task Nodes Running	The number of task nodes running at the moment.
Total Load	The total percentage load across all nodes at the moment.
Under Replicated Blocks	The number of blocks that need to be replicated one or more times.
YARN Memory Available Percentage	The percentage of remaining memory available to YARN (YARNMemoryAvailablePercentage = MemoryAvailableMB / MemoryTotalMB). This value is useful for scaling cluster resources based on YARN memory usage.

AWS: EMR Service Health	
Object Name	**Object Description**
Date	The timestamp when this health status was written.
Health	A text description of the health status of this AWS service.

FSx Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: FSx Lustre Instance Performance	
Object Name	**Object Description**

Data Read Bytes - SC	The number of bytes for file system read operations. The Sum statistic is the total number of bytes associated with read operations during the period. The Minimum statistic is the minimum number of bytes associated with read operations on a single disk. The Maximum statistic is the maximum number of bytes associated with read operations on the disk. The Average statistic is the average number of bytes associated with read operations per disk. The SampleCount statistic is the number of disks.
Data Read Bytes - Sum	The number of bytes for file system read operations. The Sum statistic is the total number of bytes associated with read operations during the period. The Minimum statistic is the minimum number of bytes associated with read operations on a single disk. The Maximum statistic is the maximum number of bytes associated with read operations on the disk. The Average statistic is the average number of bytes associated with read operations per disk. The SampleCount statistic is the number of disks.
Data Read Operations - Sum	The number of read operations. The Sum statistic is the total number of read operations.
Data Write Bytes - Sum	The number of bytes for file system write operations. The Sum statistic is the total number of bytes associated with write operations. The Minimum statistic is the minimum number of bytes associated with write operations on a single disk. The Maximum statistic is the maximum number of bytes associated with write operations on the disk. The Average statistic is the average number of bytes associated with write operations per disk. The SampleCount statistic is the number of disks.
Data Write Operations - Sum	The number of write operations. The Sum statistic is the total number of write operations.

Free Data Storage Capacity - Sum	The amount of available storage capacity. The Sum statistic is the total number of bytes available in the file system. The Minimum statistic is the total number bytes available in the fullest disk. The Maximum statistic is the total number of bytes available in the disk with the most remaining available storage. The Average statistic is the average number of bytes available per disk. The SampleCount statistic is the number of disks.
Logical Disk Usage - Sum	The amount of logical data stored (uncompressed). The Sum statistic is the total number of logical bytes stored in the file system. The Minimum statistic is the least number of logical bytes stored in a disk in the file system. The Maximum statistic is the largest number of logical bytes stored in a disk in the file system. The Average statistic is the average number of logical bytes stored per disk. The SampleCount statistic is the number of disks.
Metadata Operations - Sum	The number of metadata operations. The Sum statistic is the count of metadata operations.
Physical Disk Usage - Sum	The amount of storage physically occupied by file system data (compressed). The Sum statistic is the total number of bytes occupied in disks in the file system. The Minimum statistic is the total number of bytes occupied in the emptiest disk. The Maximum statistic is the total number of bytes occupied in the fullest disk. The Average statistic is the average number of bytes occupied per disk. The SampleCount statistic is the number of disks.

AWS: FSx ONTAP Instance Performance	
Object Name	**Object Description**
Data Read Bytes - Sum	The number of bytes (network I/O) from reads by clients to the file system. The Sum statistic is the total number of bytes associated with read operations during the period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the period.

Data Read Operations - Sum	The count of read operations (network I/O) from reads by clients to the file system. The Sum statistic is the total number of I/O operations that occurred over the specified period. To calculate the average read operations per second for a period, divide the Sum statistic by the number of seconds in the period.
Data Write Bytes - Sum	The number of bytes (network I/O) from writes by clients to the file system. The Sum statistic is the total number of bytes associated with write operations during the period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the period.
Data Write Operations - Sum	The count of write operations (network I/O) from writes by clients to the file system. The Sum statistic is the total number of I/O operations that occurred over the specified period. To calculate the average write operations per second for a period, divide the Sum statistic by the number of seconds in the period.
Logical Data Storage - Avg	The total amount of logical data stored on file system, on both the primary (SSD) tier and the capacity pool tier. This metric does not include storage-efficiency savings. To compute storage-efficiency savings in bytes, take the Average of StorageUsed over a given period and subtract it from the Average of LogicalDataStored over the same period. To compute storage-efficiency savings as a percentage of total logical data size, take the Average of StorageUsed over a given period and subtract it from the Average of LogicalDataStored over the same period. Then divide the difference by the Average of LogicalDataStored over the same period.
Metadata Operations - Sum	The count of metadata operations (network I/O) by clients to the file system. The Sum statistic is the total number of I/O operations that occurred over the specified period. To calculate the average metadata operations per second for a period, divide the Sum statistic by the number of seconds in the period.

Storage Used - Avg	The total amount of physical data stored on the file system, on both the primary (SSD) tier and the capacity pool tier. This metric includes savings from storage-efficiency features, such as data compression and deduplication.

AWS: FSx OpenZFS Instance Performance	
Object Name	**Object Description**
Compression Ratio - Avg	Ratio of compressed storage usage to uncompressed storage usage.
Data Read Bytes - Sum	Number of bytes for read operations.
Data Read Operations - Sum	Number of data read operations.
Data Write Bytes - Sum	Number of bytes for write operations.
Data Write Operations - Sum	Number of data write operations.
Nfs Bad Calls - Sum	Number of calls rejected by NFS server Remote Procedure Call (RPC) mechanism.
Storage Capactiy - Avg	Total storage capacity, equal to the sum of used and available storage capacity.
Used Storage Capacity - Avg	Amount of storage used.

AWS: FSx Window File Server Instance Performance	
Object Name	**Object Description**
Data Read Bytes - Sum	The number of bytes for file system read operations. The Sum statistic is the total number of bytes associated with read operations during the period. To calculate the average throughput (Bytes per second) for a period, divide the Sum statistic by the number of seconds in the period.
Data Read Operations - Sum	The number of read operations. The Sum statistic is the count of read operations during the time period. To calculate the average number of read operations (operations per second) for a period, divide the Sum statistic by the number of seconds in the period.

Data Write Bytes - Sum	The number of bytes for file system write operations. The Sum statistic is the total number of bytes associated with write operations during the period. To calculate the average throughput (Bytes per second) for a period, divide the Sum statistic by the number of seconds in the period.
Data Write Operations - Sum	The number of write operations. The Sum statistic is the count of write operations during the time period. To calculate the average number of write operations (operations per second) for a period, divide the Sum statistic by the number of seconds in the period.
Free Storage Capacity - Avg	The amount of available storage capacity.
Metadata Operations - Sum	The number of metadata operations. The Sum statistic is the count of metadata operations during the time period. To calculate the average number of metadata operations (operations per second) for a period, divide the Sum statistic by the number of seconds in the period.

Glacier Service

AWS: Glacier Instance Configuration	
Object Name	**Object Description**
# of Archives	The number of archives in this vault.
ARN	The unique identifier (ARN) of the vault.
AWS Glacier/SNS	DCMR Relation
AWS: SNS Relationship	Label for the DCMR relationship
Creation Time	The datetime when the vault was created.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Job Type Notifications	A list of job types which trigger notifications.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment.

Last Updated	The most recent, if any, date the inventory of the vault was modified.
Name	The name of the vault.
Notifications	The SNS Topics, if any, which publish events related to this vault.
Region	The region of the vault.
Size	The size of the vault.
Tags	Label for glacier tags.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment.

AWS: Glacier Service Health	
Object Name	**Object Description**
Date	A timestamp when the health status of this service was originally written.
Health	A text description of the health status of this AWS service.

Glue Service

This Dynamic Application is available in the ***AWS Extended Services***PowerPack.

AWS: Glue Job Performance	
Object Name	**Object Description**
Glue All Jvm Heap Usage - Sum	The fraction of memory used by the JVM heap for ALL executors (scale: 0-1).
Glue All Jvm Heap Used - Sum	The number of memory bytes used by the JVM heap for ALL executors.
Glue All S3 Filesystem Read Bytes - Sum	The number of bytes read from Amazon S3 by ALL the executors, since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes read during the previous minute).

Glue All S3 Filesystem Write Bytes - Sum	The number of bytes written to Amazon S3 by ALL the executors ,since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes written during the previous minute).
Glue All System CpuSystemLoad - Sum	The fraction of CPU system load used (scale: 0-1) by ALL the executors.
Glue Driver Aggregate Bytes Read - Sum	The number of bytes read from all data sources by all completed Spark tasks running in all executors.
Glue Driver Aggregate Elapsed Time - Sum	The ETL elapsed time in milliseconds (does not include the job bootstrap times).
Glue Driver Aggregate Num Completed Stages - Sum	The number of completed stages in the job.
Glue Driver Aggregate Num Completed Tasks - Sum	The number of completed tasks in the job.
Glue Driver Aggregate Num Failed Tasks - Sum	The number of failed tasks.
Glue Driver Aggregate Num Killed Tasks - Sum	The number of tasks killed.
Glue Driver Aggregate Records Read - Sum	The number of records read from all data sources by all completed Spark tasks running in all executors.
Glue Driver Aggregate Shuffle Bytes Written - Sum	The number of bytes written by all executors to shuffle data between them since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes written for this purpose during the previous minute).
Glue Driver Aggregate Shuffle Local Bytes Read - Sum	The number of bytes read by all executors to shuffle data between them since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes read for this purpose during the previous minute).
Glue Driver Block Manager Disk DiskSpaceUsed MB - Sum	The number of megabytes of disk space used across all executors.
Glue Driver ExecAllocManager Executors NumberMaxNeededExecutors	The number of maximum (actively running and pending) job executors needed to satisfy the current load.
Glue Driver ExecutorAllocationManagerExecutorsNumberAllExecutors	The number of actively running job executors.

Glue Driver Jvm Heap Usage - Sum	The fraction of memory used by the JVM heap for this driver (scale: 0-1) for driver.
Glue Driver Jvm Heap Used - Sum	The number of memory bytes used by the JVM heap for the driver.
Glue Driver S3 Filesystem Read Bytes - Sum	The number of bytes read from Amazon S3 by the driver, since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes read during the previous minute).
Glue Driver S3 Filesystem Write Bytes - Sum	The number of bytes written to Amazon S3 by the driver, since the previous report (aggregated by the AWS Glue Metrics Dashboard as the number of bytes written during the previous minute).
Glue Driver Streaming Num Records - Sum	The number of records that are received in a micro-batch. This metric is only available for AWS Glue streaming jobs with AWS Glue version 2.0 and above.
Glue Driver StreamingBatchProcessingTimeInMs - Sum	The time it takes to process the batches in milliseconds. This metric is only available for AWS Glue streaming jobs with AWS Glue version 2.0 and above.
Glue Driver System CpuSystemLoad - Sum	The fraction of CPU system load used (scale: 0-1) by the driver.

IoT Service

AWS: IoT Service Performance	
Object Name	**Object Description**
Connect.Success	The number of successful connections to the message broker.
GetThingShadow.Accepted	The number of GetThingShadow requests processed successfully.
Ping.Success	The number of ping messages received by the message broker.
Subscribe.Success	The number of subscribe requests that were successfully processed by the message broker.

UpdateThingShadow.Accepted	The number of UpdateThingShadow requests processed successfully.

AWS: IoT Thing Instance Configuration	
Object Name	**Object Description**
Attribute Keys	The key of the thing attributes assigned to the resource.
Attribute Values	The value of the thing attributes assigned to the resource.
Billing Group	The name of the billing group the thing belongs to.
Key	The key of the tags assigned to the resource.
Thing Name	The name of the thing.
Thing ARN	The ARN of the thing to describe.
Thing Id	The ID of the thing to describe.
Thing Type Name	The thing type name.
Value	The value of the tags assigned to the resource.
Version	The current version of the thing record in the registry.

Kafka/MSK Service

This Dynamic Application is available in the ***AWS Extended Services*** PowerPack.

AWS: MSK (Kafka) Cluster Performance	
Object Name	**Object Description**
Active Controller Count - Avg	Only one controller per cluster should be active at any given time.
Client Connection Count - Avg	The number of active authenticated client connections.

Global Partition Count - Avg	The number of partitions across all topics in the cluster, excluding replicas. Because GlobalPartitionCount doesn't include replicas, the sum of the PartitionCount values can be higher than GlobalPartitionCount if the replication factor for a topic is greater than 1.
Global Topic Count - Avg	Total number of topics across all brokers in the cluster.
Kafka Data Logs Disk Used - Avg	The percentage of disk space used for application logs.
Offline Partitions Count - Avg	Total number of partitions that are offline in the cluster.
Zoo Keeper Request Latency Ms Mean - Avg	The mean latency in milliseconds for Apache ZooKeeper requests from broker.
Zoo Keeper Session State - Avg	Connection status of broker's ZooKeeper session which may be one of the following: NOT_ CONNECTED: '0.0', ASSOCIATING: '0.1', CONNECTING: '0.5', CONNECTEDREADONLY: '0.8', CONNECTED: '1.0', CLOSED: '5.0', AUTH_ FAILED: '10.0'.

Kinesis Service

This Dynamic Application is available in the ***AWS Extended Services***PowerPack.

AWS: Kinesis Data Firehose Delivery Stream Performance	
Object Name	**Object Description**
Data Read From Kinesis Stream Bytes - Avg	When the data source is a Kinesis data stream, this metric indicates the number of bytes read from that data stream. This number includes rereads due to failovers.
Data Read From Kinesis Stream Records - Sum	When the data source is a Kinesis data stream, this metric indicates the number of records read from that data stream. This number includes rereads due to failovers.
Throttled Describe Stream - Sum	The total number of times the DescribeStream operation is throttled when the data source is a Kinesis data stream.

Throttled Get Records - Sum	The total number of times the GetRecords operation is throttled when the data source is a Kinesis data stream.
Throttled Get Shard Iterator - Sum	The total number of times the GetShardIterator operation is throttled when the data source is a Kinesis data stream.

AWS: Kinesis Stream Performance	
Object Name	**Object Description**
Get Records Bytes - Avg	The number of bytes retrieved from the Kinesis stream, measured over the specified time period. Minimum, Maximum, and Average statistics represent the bytes in a single GetRecords operation for the stream in the specified time period.
Get Records Iterator Age Milliseconds - Avg	The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up with the stream.
Get Records Latency - Avg	The time taken per GetRecords operation, measured over the specified time period.
Get Records Records - Avg	The number of records retrieved from the shard, measured over the specified time period. Minimum, Maximum, and Average statistics represent the records in a single GetRecords operation for the stream in the specified time period.
Get Records Success - Avg	The number of successful GetRecords operations per stream, measured over the specified time period.

Incoming Bytes - Avg	The number of bytes successfully put to the Kinesis stream over the specified time period. This metric includes bytes from PutRecord and PutRecords operations. Minimum, Maximum, and Average statistics represent the bytes in a single put operation for the stream in the specified time period.
Incoming Records - Avg	The number of records successfully put to the Kinesis stream over the specified time period. This metric includes record counts from PutRecord and PutRecords operations. Minimum, Maximum, and Average statistics represent the records in a single put operation for the stream in the specified time period.
Put Record Bytes - Sum	The number of bytes put to the Kinesis stream using the PutRecord operation over the specified time period.
Put Record Latency - Avg	The time taken per PutRecord operation, measured over the specified time period.
Put Record Success - Sum	The number of successful PutRecord operations per Kinesis stream, measured over the specified time period. Average reflects the percentage of successful writes to a stream.
Put Records Bytes - Avg	The number of bytes put to the Kinesis stream using the PutRecords operation over the specified time period.
Put Records Failed Records - Sum	The number of records rejected due to internal failures in a PutRecords operation per Kinesis data stream, measured over the specified time period. Occasional internal failures are to be expected and should be retried.
Put Records Latency - Avg	The time taken per PutRecords operation, measured over the specified time period.
Put Records Success - Avg	The number of PutRecords operations where at least one record succeeded, per Kinesis stream, measured over the specified time period.
Put Records Successful Records - Sum	The number of successful records in a PutRecords operation per Kinesis data stream, measured over the specified time period.
Put Records Throttled Records - Sum	The number of records rejected due to throttling in a PutRecords operation per Kinesis data stream, measured over the specified time period.

Put Records Total Records - Sum	The total number of records sent in a PutRecords operation per Kinesis data stream, measured over the specified time period.
Read Provisioned Throughput Exceeded - Avg	The number of GetRecords calls throttled for the stream over the specified time period.
Write Provisioned Throughput Exceeded - Avg	The number of records rejected due to throttling for the stream over the specified time period. This metric includes throttling from PutRecord and PutRecords operations.

AWS: Kinesis Video Stream Performance	
Object Name	**Object Description**
Archived Fragments Consumed Media - Sum	Number of fragment media quota points that were consumed by all of the APIs. For an explanation of the concept of quota points.
Archived Fragments Consumed Metadata - Sum	Number of fragments metadata quota points that were consumed by all of the APIs.
Get Clip Latency - Avg	Latency of the GetClip API calls for the given video stream name.
Get Clip Outgoing Bytes - Avg	Total number of bytes sent out from the service as part of the GetClip API for a given video stream.
Get Clip Requests - Sum	Number of GetClip API requests for a given video stream.
Get Clip Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Dashmanifest Latency - Avg	Latency of the GetDASHManifest API calls for the given stream name.
Get Dashmanifest Requests - Sum	Number of GetDASHManifest API requests for a given stream.
Get Dashmanifest Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Dashstreaming Session Url Latency - Avg	Latency of the GetDASHStreamingSessionURL API calls for the given stream name.
Get Dasstreaming Session Url Requests - Sum	Number of GetDASHStreamingSessionURL API requests for a given stream.
Get Dashstreaming Session Url Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.

Get Hlsmaster Playlist Latency - Avg	Latency of the GetHLSMasterPlaylist API calls for the given stream name.
Get Hlsmaster Playlist Requests - Sum	Number of GetHLSMasterPlaylist API requests for a given stream.
Get Hlsmaster Playlists Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Hlsmedia Playlist Latency - Avg	Latency of the GetHLSMediaPlaylist API calls for the given stream name.
Get Hlsmedia Playlist Requests - Sum	Number of GetHLSMediaPlaylist API requests for a given stream.
Get Hlsmedia Playlist Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Hlsstreaming Session Url Latency - Avg	Latency of the GetHLSStreamingSessionURL API calls for the given stream name.
Get Hlsstreaming Session Url Requests - Sum	Number of GetHLSStreamingSessionURL API requests for a given stream.
Get Hlsstreaming Session Url Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Media For Fragment List Outgoing Bytes - Avg	Total number of bytes sent out from the service as part of the GetMediaForFragmentList API for a given stream.
Get Media For Fragment List Outgoing Fragments - Sum	Total number of fragments sent out from the service as part of the GetMediaForFragmentList API for a given stream.
Get Media For Fragment List Outgoing Frames - Sum	Total number of frames sent out from the service as part of the GetMediaForFragmentList API for a given stream.
Get Media For Fragment List Requests - Sum	Number of GetMediaForFragmentList API requests for a given stream.
Get Media For Fragment List Success - Sum	1 for every fragment successfully sent; 0 for every failure. The average value indicates the rate of success.
Get Media Connection Errors - Sum	The number of connections that were not successfully established.
Get Media Millis Behind Now - Avg	Time difference between the current server timestamp and the server timestamp of the last fragment sent.
Get Media Outgoing Bytes - Avg	Total number of bytes sent out from the service as part of the GetMedia API for a given stream.

Get Media Outgoing Fragments - Sum	Number of fragments sent while doing GetMedia for the stream.
Get Media Outgoing Frames - Sum	Number of frames sent during GetMedia on the given stream.
Get Media Requests - Sum	Number of GetMedia API requests for a given stream.
Get Media Success - Avg	1 for every fragment successfully sent; 0 for every failure. The average value indicates the rate of success.
Get Mp4Init Fragment Latency - Avg	Latency of the GetMP4InitFragment API calls for the given stream name.
Get Mp4Init Fragment Requests - Sum	Number of GetMP4InitFragment API requests for a given stream.
Get Mp4Init Fragment Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Mp4Media Fragment Latency - Avg	Latency of the GetMP4MediaFragment API calls for the given stream name.
Get Mp4Media Fragment Outgoing Bytes - Avg	Total number of bytes sent out from the service as part of the GetMP4MediaFragment API for a given stream.
Get Mp4Media Fragment Requests - Sum	Number of GetMP4MediaFragment API requests for a given stream.
Get Mp4Media Fragment Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
Get Tsfragment Latency - Avg	Latency of the GetTSFragment API calls for the given stream name.
Get Tsfragment Outgoing Bytes - Avg	Total number of bytes sent out from the service as part of the GetTSFragment API for a given stream.
Get Tsfragment Requests - Sum	Number of GetTSFragment API requests for a given stream.
Get Tsfragment Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.
List Fragments Latency - Avg	Latency of the ListFragments API calls for the given stream name.
List Fragments Requests - Sum	Number of ListFragments API requests for a given stream.
List Fragments Success - Avg	1 for every successful request; 0 for every failure. The average value indicates the rate of success.

Put Media Active Connections - Sum	The total number of connections to the service host.
Put Media Buffering Ack Latency - Avg	Time difference between when the first byte of a new fragment is received by Kinesis Video Streams and when the Buffering ACK is sent for the fragment.
Put Media Connection Errors - Sum	Errors while establishing PutMedia connection for the stream.
Put Media Error Ack Count - Sum	Number of Error ACKs sent while doing PutMedia for the stream.
Put Media Fragment Ingestion Latency - Avg	Time difference between when the first and last bytes of a fragment are received by Kinesis Video Streams.
Put Media Fragment Persist Latency - Avg	Time taken from when the complete fragment data is received and archived.
Put Media Incoming Bytes - Avg	Number of bytes received as part of PutMedia for the stream.
Put Media Incoming Fragments - Sum	Number of complete fragments received as part of PutMedia for the stream.
Put Media Incoming Frames - Sum	Number of complete frames received as part of PutMedia for the stream.
Put Media Latency - Avg	Time difference between the request and the HTTP response from InletService while establishing the connection.
Put Media Persisted Ack Latency - Avg	Time difference between when the last byte of a new fragment is received by Kinesis Video Streams and when the Persisted ACK is sent for the fragment.
Put Media Received Ack Latency - Avg	Time difference between when the last byte of a new fragment is received by Kinesis Video Streams and when the Received ACK is sent for the fragment.
Put Media Requests - Sum	Number of PutMedia API requests for a given stream.
Put Media Success - Avg	1 for each fragment successfully written; 0 for every failed fragment. The average value of this metric indicates how many complete, valid fragments are sent.

KMS Service

AWS: KMS Configuration	
Object Name	**Object Description**
Key Id	The unique identifier of every Key from Key Management Service.
Target Key Id	String that contains the key identifier referred to by the alias.
Target Key Id	The key identifier which referred to by the tags.
Alias	String that contains the alias. This value begins with alias/.
ARN	The Amazon Resource Name (ARN) of every Key from Key Management Service.
Creation Date	The date and time when the Key was created.
Key Manager	The manager of the Key. In AWS are either Customer Managed or AWS managed.
Origin	The source of the Key. They could be: AWS_KMS, EXTERNAL, AWS_CLOUDHSM.
Status	The state of the Key.
Tag Key	The key of the tag.
Tag Value	The value of the tag.

AWS: KMS Performance	
Object Name	**Object Description**
Seconds Until Key Material Expiration	This metric tracks the number of seconds remaining until imported key material expires. This metric is valid only for CMKs whose origin is EXTERNAL and whose key material is or was set to expire.

Lambda Service

AWS: Lambda Function Configuration	
Object Name	**Object Description**
Name	The alias name.
State	The state of the event source mapping.
State Transition Reason	The reason the event source mapping is in its current state.
Description	The alias description.
Last Processing Result	The result of the last AWS Lambda invocation of your Lambda function.
Version	List of versions of the Lambda function.
Action	The action that the trigger makes to Lambda Function.
DLQ Resource	The Dead-Letter Queue service. Possible values are "SNS" or "SQS".
Event Source ARN	The Amazon Resource Name (ARN) of the Amazon Kinesis stream that is the source of events.
Function Version	The function version to which the alias points.
Service	The Amazon Service that triggers the Lambda Function.
Version Description	Descriptions of each version of the Lambda function.
ARN	Lambda function ARN that is qualified using the alias name as the suffix.
AWS Lambda Function/Security Group	A list of security group IDs associated with the Lambda function.
AWS Lambda Function/SNS-SQS	The Id of the SQS or SNS service that the lambda is using for the error handling storage.
AWS Lambda Function/Subnet	A list of subnet IDs associated with the Lambda function.
AWS Lambda Function/VPC	The custom VPC that the lambda belongs to.

Batch Size	The largest number of records that AWS Lambda will retrieve from your event source at the time of invoking your function.
Code Size (B)	The size, in bytes, of the function .zip file you uploaded.
Description	The user-provided description for the lambda function.
Event Source ARN	The Amazon Resource Name (ARN) of the Amazon Service that is the source of triggers.
Execution Role	The Amazon Resource Name (ARN) of the IAM role that Lambda assumes when it executes your function to access any other Amazon Web Services (AWS) resources.
Function ARN	Lambda Function Amazon Resource Name identifier.
Function Name	The name of the function.
Handler	The function Lambda calls to begin executing your function.
Key	The key of the lambda function tag.
Last Modified	The time stamp of the last time you updated the function.
Last Modified	The UTC time string indicating the last time the event mapping was updated.
Memory Size (MB)	The memory size, in MB, you configured for the function. Must be a multiple of 64 MB.
Runtime	The runtime environment for the Lambda function.
Timeout (s)	The function execution time in seconds at which Lambda should terminate the function. Because the execution time has cost implications, we recommend you set this value based on your expected execution time. The default is 3 seconds.
Value	The value of the lambda function tag.
Version	The version of the Lambda function.
Version ARN	List of ARNs with the detail of versions of the Lambda function.

AWS: Lambda Function Performance	
Object Name	**Object Description**
Dead Letter Errors	Incremented when Lambda is unable to write the failed event payload to your configured Dead Letter Queues. The major event expires after 90 minutes.
Duration	Elapsed wall clock time from when the function code starts executing as a result of an invocation to when it stops executing.
Errors	Number of invocations that failed due to errors in the function (response code 4XX).
Invocations	Number of times a function is invoked in response to an event or invocation API call.
Iterator Age	Measures the age of the last record for each batch of records processed. Age is the difference between the time Lambda received the batch, and the time the last record in the batch was written to the stream.
Throttles	Number of Lambda function invocation attempts that were throttled due to invocation rates exceeding the customer's concurrent limits (error code 429).

AWS: Lambda Function Qualified Configuration	
Object Name	**Object Description**
Alias ARN	Lambda function ARN that is qualified using the alias name as the suffix.
Alias Description	Alias description.
Alias Function Version	Function version to which the alias points.
Alias Name	Alias name.
Alias Revision ID	Represents the latest updated revision of the function or alias.
AWS Lambda Function Qualified/Security Groups	A list of security group IDs associated with the Lambda function.
AWS Lambda Function Qualified/Subnets	A list of subnet IDs associated with the Lambda function.

AWS Lambda Function Qualified/VPC	The custom VPC that the lambda belongs to.
Version Code Sha256	It is the SHA256 hash of your function deployment package.
Version Code Size (B)	The size, in bytes, of the function .zip file you uploaded.
Version Description	The user-provided description for lambda function version.
Version Function ARN	The Amazon Resource Name (ARN) assigned to a regular lambda function.
Version Function ARN Version	The Amazon Resource Name (ARN) assigned to the lambda function version.
Version Function Name	The name of the function. Note that the length constraint applies only to the ARN. If you specify only the function name, it is limited to 64 characters in length.
Version Handler	The function Lambda calls to begin executing your function.
Version Last Modified	The time stamp of the last time you updated the function.
Version Memory Size (MB)	The memory size, in MB, you configured for the function. Must be a multiple of 64 MB.
Version Revision ID	Represents the latest updated revision of the function or alias.
Version Role	The Amazon Resource Name (ARN) of the IAM role that Lambda assumes when it executes your function to access any other Amazon Web Services (AWS) resources.
Version Runtime	The runtime environment for the Lambda function.
Version Timeout (s)	The function execution time at which Lambda should terminate the function. Because the execution time has cost implications, we recommend you set this value based on your expected execution time. The default is 3 seconds.
Version Version	The version of the Lambda function.
VPC Config Security Group ID	A list of security group IDs associated with the Lambda function.
VPC Config Subnet ID	A list of subnet IDs associated with the Lambda function.
VPC Config VPC ID	The VPC ID associated with you Lambda function.

AWS: Lambda Function Qualified Performance	
Object Name	**Object Description**
Dead Letter Errors	Incremented when Lambda is unable to write the failed event payload to your configured Dead Letter Queues. The major event expires after 90 minutes.
Duration	Elapsed wall clock time from when the function code starts executing as a result of an invocation to when it stops executing.
Errors	Number of invocations that failed due to errors in the function (response code 4XX).
Invocations	Number of times a function is invoked in response to an event or invocation API call.
Iterator Age	Measures the age of the last record for each batch of records processed. Age is the difference between the time Lambda received the batch, and the time the last record in the batch was written to the stream.
Throttles	Number of Lambda function invocation attempts that were throttled due to invocation rates exceeding the customer's concurrent limits (error code 429).

AWS: Lambda Function Replica Configuration	
Object Name	**Object Description**
Code Size (B)	The size, in bytes, of the function .zip file you uploaded.
Description	The user-provided description for the lambda function.
Execution Role	The Amazon Resource Name (ARN) of the IAM role that Lambda assumes when it executes your function to access any other Amazon Web Services (AWS) resources.
Function ARN	Lambda Function Amazon Resource Name identifier.
Function Name	The name of the function.

Handler	The function Lambda calls to begin executing your function.
Lambda Replica Master ARN	The ARN (Amazon Resource Name) of the master function from which the function was replicated.
Last Modified	The time stamp of the last time you updated the function.
Memory Size (MB)	The memory size, in MB, you configured for the function. Must be a multiple of 64 MB.
Parent (Master) Lambda Function/Replica Lambda Function	The master function from which the function was replicated.
Runtime	The runtime environment for the Lambda function.
Timeout (s)	The function execution time in seconds at which Lambda should terminate the function. Because the execution time has cost implications, we recommend you set this value based on your expected execution time. The default is 3 seconds.
Version	The version of the Lambda function.

AWS: Lambda Service Configuration	
Object Name	**Object Description**
Code Storage (B)	Total size, in bytes, of the deployment packages of your account per region.
Memory Allocated (B)	Memory allocated for all lambda functions.
Total Code Size (B)	Maximum size, in bytes, of a code package you can upload per region. The default size is 75 GB.
Code Size Unzipped (B)	Size, in bytes, of code/dependencies that you can zip into a deployment package (uncompressed .zip/.jar size) for uploading. The default limit is 250 MB.
Code Size Zipped (B)	Size, in bytes, of a single zipped code/dependencies package you can upload for your Lambda function(.zip/.jar file). Try using Amazon S3 for uploading larger files. Default limit is 50 MB.
Code Storage Percentage	Total size, in bytes, of the deployment packages of your account per region, it is measured in percentage.

Full Account Concurrency	Number of simultaneous executions of your function per region. The default limit is 1000.
Lambda Functions	The number of existing functions of your account per region.
Unreserved Account Concurrency	The number of concurrent executions available to functions that do not have concurrency limits set.

AWS: Lambda Service Health	
Object Name	**Object Description**
Date	The timestamp of this health status update string.
Health	A text description of this AWS Lambda Service health status.

AWS: Lambda Service Performance	
Object Name	**Object Description**
Concurrent Executions	Sum for all functions within this service instance, during this polling interval, of concurrent executions of functions.
Duration	An average of the elapsed wall clock time from when the function code starts executing as a result of an invocation to when it stops executing. This averaged time value is for all functions within this service instance during this polling interval.
Errors	Sum for all functions within this service instance, during this polling interval, of invocations that failed due to errors (response code 4XX).
Invocations	Sum for all functions within this service instance, during this polling interval, of the times functions are invoked in response to an event or invocation API call.
Throttles	Sum for all functions within this service instance, during this polling interval, of function invocation attempts that were throttled due to invocation rates exceeding the customer's concurrent limits (error code 429).

Unreserved Concurrent Executions	Sum for all functions within this service instance, during this polling interval, for functions invoked or running in concurrency for functions that have no custom concurrency limit specified.

Lex V2 Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Lexv2 PostContent Instance Performance	
Object Name	**Object Description**
Kendra Index Access Error - Sum	The number of times that Amazon Lex V2 could not access your Amazon Kendra index.
Kendra Latency - Avg	The amount of time that it takes Amazon Kendra to respond to a request from the AMAZON.KendraSearchIntent.
Kendra Success - Sum	The number of times that Amazon Lex V2 could not access your Amazon Kendra index.
Kendra System Errors - Sum	The number of times that Amazon Lex V2 could not query the Amazon Kendra index.
Kendra Throttled Events - Sum	The number of times Amazon Kendra throttled requests from the AMAZON.
Missed Utterance Count - Sum	The number of utterances that were not recognized in the specified period
Runtime Concurrency - Sum	The number of concurrent connections in the specified time period. RuntimeConcurrency is reported as a StatisticSet.
Runtime Invalid Lambda Responses - Sum	The number of invalid AWS Lambda responses in the specified period.
Runtime Lambda Errors - Sum	The number of Lambda runtime errors in the specified period.
Runtime Polly Errors - Sum	The number of invalid Amazon Polly responses in the specified period.
Runtime Request Count - Sum	The number of runtime requests in the specified period.

Runtime Successful Request Latency - Sum	The latency for successful requests between the time that the request was made and the response was passed back.
Runtime System Errors - Sum	The number of system errors in the specified period. The response code range for a system error is 500 to 599.
Runtime Throttled Events - Sum	The number of throttled requests. Amazon Lex throttles a request when it receives more requests than the limit of transactions per second set for your account.
Runtime User Errors - Sum	The number of user errors in the specified period. The response code range for a user error is 400 to 499.

AWS: Lexv2 PostText Instance Performance	
Object Name	**Object Description**
Kendra Index Access Error - Sum	The number of times that Amazon Lex V2 could not access your Amazon Kendra index.
Kendra Latency - Avg	The amount of time that it takes Amazon Kendra to respond to a request from the AMAZON.KendraSearchIntent.
Kendra Success - Sum	The number of times that Amazon Lex V2 could not access your Amazon Kendra index.
Kendra System Errors - Sum	The number of times that Amazon Lex V2 could not query the Amazon Kendra index.
Kendra Throttled Events - Sum	The number of times Amazon Kendra throttled requests from the AMAZON.
Missed Utterance Count - Sum	The number of utterances that were not recognized in the specified period
Runtime Concurrency - Sum	The number of concurrent connections in the specified time period. RuntimeConcurrency is reported as a StatisticSet.
Runtime Invalid Lambda Responses - Sum	The number of invalid AWS Lambda responses in the specified period.
Runtime Lambda Errors - Sum	The number of Lambda runtime errors in the specified period.
Runtime Polly Errors - Sum	The number of invalid Amazon Polly responses in the specified period.

Runtime Request Count - Sum	The number of runtime requests in the specified period.
Runtime Successful Request Latency - Sum	The latency for successful requests between the time that the request was made and the response was passed back.
Runtime System Errors - Sum	The number of system errors in the specified period. The response code range for a system error is 500 to 599.
Runtime Throttled Events - Sum	The number of throttled requests. Amazon Lex throttles a request when it receives more requests than the limit of transactions per second set for your account.
Runtime User Errors - Sum	The number of user errors in the specified period. The response code range for a user error is 400 to 499.

LightSail Service

AWS: LightSail Instance Configuration	
Object Name	**Object Description**
Access Direction	The access direction (inbound or outbound).
Access From	The location from which access is allowed (e.g., Anywhere (0.0.0.0/0)).
Access Type	The type of access (Public or Private).
Availability Zone	The Availability Zone.
Blueprint ID	The blueprint ID.
Blueprint Name	The friendly name of the blueprint.
CPU count	The number of vCPUs the instance has.
Created At	The timestamp when the instance was created.
Disk ARN	The Amazon Resource Name (ARN) of the disk.
Disk Attached To	The resources to which the disk is attached.
Disk Attachment State	The attachment state of the disk.
Disk Created At	The date when the disk was created.

Disk Name	The name of the disk.
Disk Path	The disk path.
Disk Resource Type	The resource type of the disk.
Disk Size in GB	The size of the disk in GB.
Distinguished Name	The Amazon Resource Name (ARN) of the instance.
From Port	The first port in the range.
GB per Month Allocation	The amount allocated per month (in GB).
Instance Type	The Amazon EC2 instance type. Note: This collection object may not collect any data if AWS has deprecated the bundle that was used to create the original Lightsail instance.
IPv6 Address	The IPv6 address of the instance.
Is Static Ip?	A Boolean value indicating whether this instance has a static IP assigned to it.
Name	The name the user gave the instance.
Price	The price in US dollars. Note: This collection object may not collect any data if AWS has deprecated the bundle that was used to create the original lightsail instance.
Private IP Address	The private IP address of the instance.
Protocol	The protocol being used, it can contain following values tcp|-1|udp. The -1 value means all protocols (udp and udp).
Public IP Address	The public IP address of the instance.
RAM size in GB	The amount of RAM in GB on the instance (e.g., 1.0).
Region Name	The AWS Region name.
Resource Type	The type of resource (usually Instance).
State	The status code and the state (e.g., running) for the instance.
Support Code	The support code. Include this code in your email to support when you have questions about an instance or another resource in Lightsail. This code enables our support team to look up your Lightsail information more easily.

System Disk	A Boolean value indicating whether this disk is a system disk (has an operating system loaded on it).
To Port	The last port in the range.

AWS: LightSail Instance Performance	
Object Name	**Object Description**
CPU Utilization	This metric refers to the percentage of CPU capacity that your Amazon Lightsail instance is using at the timestamp. This metric identifies the processing power required to run an application on a selected instance.
Network In	The number of bytes received at the timestamp on all network interfaces by the Amazon Lightsail instance. This metric identifies the volume of incoming network traffic to an application on a single instance.
Network Out	The number of bytes sent out at the timestamp on all network interfaces by the Amazon Lightsail instance. This metric identifies the volume of outgoing network traffic to an application on a single instance.
Status Check Failed	Reports whether the instance has passed both the AmazonLightsail instance status check and the system status check in the last minute. This metric can be either 0 (passed) or 1 (failed).
Status Check Failed Instance	Reports whether the instance has passed the Amazon Lightsail instance status check in the last minute. This metric can be either 0 (passed) or 1 (failed).
Status Check Failed System	Reports whether the instance has passed the system status check in the last minute. This metric can be either 0 (passed) or 1 (failed).

MediaConnect Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Media Connect Flow Instance Performance

Object Name	Object Description
Arqrecovered - Sum	The number of dropped packets that were recovered by automatic repeat request (ARQ). This metric applies only to flows that have one source that uses the RIST protocol or the Zixi protocol. It doesn't apply to flows that receive content from entitlement or to flows that have multiple sources. For flows that have multiple sources, use the SourceARQRecovered metric to view data for each source.
Arqrequests - Sum	The number of retransmitted packets that were requested through automatic repeat request (ARQ) and received. This metric applies only to flows that have one source that uses the RIST protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement or to flows that have multiple sources. For flows that have multiple sources, use the SourceARQRequests metric to view data for each source.
Bit Rate - Avg	The bitrate of the incoming (source) video.
Connected - Avg	The status of the source. A value of 1 indicates that the source is connected and a value of 0 (zero) indicates that the source is disconnected. This metric applies only to sources that use the Zixi protocol.
Disconnections - Sum	The number of times that the source status changed from connected to disconnected.
Dropped Packets - Sum	The number of packets that were lost during transit. This value is measured before any error correction takes place.
Failover Switches - Sum	The total number of times the flow switches back and forth between sources when using the Failover mode for source failover.
Fecpackets - Sum	The number of packets that were transmitted using forward error correction (FEC) and received. This metric applies only to flows that have one source that uses the RTP-FEC protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement or to flows that have multiple sources. For flows that have multiple sources, use the SourceFECPackets metric to view data for each source.

Fecrecovered - Sum	The number of packets that were transmitted using forward error correction (FEC), lost during transit, and recovered. This metric applies only to flows that have one source that uses the RTP-FEC protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement or to flows that have multiple sources. For flows that have multiple sources, use the SourceFECRecovered metric to view data for each source.
Merge Active - Avg	The merge status of all sources on the flow. A value of 1 indicates that all sources are merged. A value of 0 (zero) indicates that at least one source is not actively merged with 2022-7.
Merge Latency - Sum	The maximum value for SourceMergeLatency.
Not Recovered Packets - Sum	The number of packets that were lost during transit and were not recovered by error correction.
Overflow Packets - Sum	The number of packets that were lost in transit because the video required more buffer than was available. This metric doesn't apply to flows that receive content from an entitlement or to flows that have multiple sources.
Packet Loss Percent - Avg	The percentage of packets that were lost during transit, even if they were recovered.
Recovered Packets - Sum	The number of packets that were lost during transit, but recovered.
Round Trip Time - Avg	The amount of time it takes for the source to send a signal and receive an acknowledgment from AWS Elemental MediaConnect. This metric applies only to flows that have one source that uses the Zixi protocol. It doesn't apply to flows that receive content from an entitlement or to flows that have multiple sources. For flows that have multiple sources, use the SourceRoundTripTime metric to view data for each source.
Total Packets - Sum	The total number of packets that were received.

AWS: Media Connect Maintenance Flow Instance Performance	
Object Name	**Object Description**
Maintenance Canceled - Avg	Maintenance for this flow is canceled by MediaConnect.

Maintenance Failed - Avg	Maintenance did not complete successfully for this flow.
Maintenance Rescheduled - Avg	MediaConnect is unable to perform maintenance at the previously scheduled date and time. A new date and time has been automatically assigned by MediaConnect for this flow's maintenance.
Maintenance Scheduled - Avg	Maintenance is scheduled for the flow.
Maintenance Started - Avg	Maintenance has started and is currently in progress for this flow.
Maintenance Succeeded - Avg	Maintenance completed successfully for this flow.

AWS: Media Connect Output Instance Performance	
Object Name	**Object Description**
Connected Outputs - Avg	The number of outputs that are currently connected. This metric applies only to outputs that use the Zixi protocol.
Output Connected - Avg	The status of the output. A value of 1 indicates that the output is connected, and a value of 0 (zero) indicates that the output is disconnected. This metric applies only to outputs that use the Zixi protocol.
Output Disconnections - Avg	The number of times that the output status changed from connected to disconnected.
Output Dropped Payloads - Sum	Payloads that were lost during transit from MediaConnect to the output. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.
Output Late Payloads - Sum	Packets of a payload that arrive at the output outside of MediaConnect's internal buffer. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.
Output Total Bytes - Sum	Total amount of bytes transferred from MediaConnect to the output.
Output Total Payloads - Sum	Total amount of payloads delivered from MediaConnect to the output. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.

AWS: Media Connect Source Instance Performance	
Object Name	**Object Description**
Source Arqrecovered - Sum	The number of dropped packets that were recovered by automatic repeat request (ARQ). This metric applies only to sources that use the RIST protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement.
Source Arqrequests - Sum	The number of retransmitted packets that were requested through automatic repeat request (ARQ) and received. This metric applies only to sources that use the RIST protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement.
Source Bit Rate - Avg	The bitrate of the incoming (source) video. Note MediaConnect suppresses null packets in an effort to optimize the data connection between the content originator of flow and the subscribers of flow. This can result in a fluctuating bitrate on the subscribers of flow, or a difference between the bitrate of the content originator of flow and the subscribers flow. We recommend that you monitor source health as a combination of SourceBitRate and other metrics such as SourceContinuityCounter and SourceNotRecoveredPackets.
Source Connected - Avg	The status of the source. A value of 1 indicates that the source is connected and a value of 0 (zero) indicates that the source is disconnected. This metric applies only to sources that use the Zixi protocol.
Source Disconnections - Sum	The number of times that the source status changed from connected to disconnected.
Source Dropped Packets - Sum	The number of packets that were lost during transit. This value is measured before any error correction takes place.
Source Dropped Payloads - Avg	Payloads that were lost during transit to MediaConnect from the source. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.
Source Fecpackets - Sum	The number of packets that were transmitted using forward error correction (FEC) and received. This metric applies only to sources that use the RTP-FEC protocol or the Zixi protocol. It doesn't apply to flows that receive content from an entitlement.
Source Fecrecovered - Sum	The number of packets that were transmitted using forward error correction (FEC), lost during transit, and recovered. This metric applies only to sources that use the RTP-FEC protocol or the Zixi protocol. It does not apply to flows that receive content from an entitlement.

Source Late Payloads - Avg	Packets of a payload that arrive outside of the configured Max sync buffer time frame. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.
Source Merge Active - Avg	An indication of the status of the source with respect to other sources. This metric is useful when the flow has multiple sources for failover and you are using the Merge failover mode. A value of 1 indicates that the flow has multiple sources and that this source is actively in use, with 2022-7 merge. A value of 0 (zero) indicates that the flow is not using the source to form the stream.
Source Merge Latency - Avg	The amount of time that this source trails the primary source. If this source is the primary source, the value is 0 (zero).
Source Merge Status Warn Mismatch - Avg	A status metric warning that the flow is receiving mismatched sources. This means that any dropped packets will not be recovered and will result in poor network reliability. This metric only applies to sources using merge mode failover. Merge mode failover requires both sources to be binary identical. To be binary identical, the sources must originate from the same encoder. This will ensure the sources can share missing packets, as the packets are identical.
Source Merge Status Warn Solo - Avg	A status metric warning that the flow is only receiving one source. This means any dropped packets will not be recovered and will result in poor network reliability. This metric only applies to sources using merge mode failover.
Source Missing Packets - Avg	A packet was missing from both source streams, this means the packet could not be recovered. This metric only applies to sources using merge mode failover.
Source Not Recovered Packets - Avg	The number of packets that were lost during transit and were not recovered by error correction.
Source Overflow Packets - Avg	The number of packets that were lost in transit because the video required more buffer than was available. This metric does not apply to flows that receive content from an entitlement or to flows that have multiple sources.
Source Packet Loss Percent - Avg	The percentage of packets that were lost during transit, even if they were recovered.
Source Recovered Packets - Avg	The number of packets that were lost during transit, but recovered.
Source Round Trip Time - Avg	The amount of time it takes for the source to send a signal and receive an acknowledgment from AWS Elemental MediaConnect. This metric applies only to sources that use the Zixi protocol. It does not apply to flows that receive content from an entitlement.

Source Selected - Avg	An indication if a source is being used as the input for flow ingest. This metric applies if your flow uses source failover, and the failover mode is set to Failover. A value of 1 indicates that the source is being used as the input. A value of 0 (zero) indicates that the flow is being used as the backup stream.
Source Total Bytes - Avg	Total amount of bytes transferred to MediaConnect from the source.
Source Total Packets - Avg	The total number of packets that were received.
Source Total Payouts - Avg	Total amount of payloads delivered to MediaConnect from the source. A payload is a frame of video or an audio sample. Payloads can consist of multiple packets. Payload metrics are only applicable when using CDI.

AWS: Media Connect TR 101 Source Instance Performance	
Object Name	**Object Description**
Source Caterror - Sum	The number of times that a conditional access table (CAT) error occurred. This error indicates that the CAT is not present. The CAT tells the integrated receiver decoder (IRD) where to find management messages for the conditional access (CA) systems that are in use.
Source Continuity Counter - Sum	The number of times that a continuity error occurred. This error indicates an incorrect packet order or lost packets.
Source Crcerror - Sum	The number of times that a cyclic redundancy check (CRC) error occurred. This error happens when a CRC determines that data is corrupted.
Source Paterror - Sum	The number of times that a program association table (PAT) error occurred. This error indicates that the PAT is missing. The PAT lists the programs that are available in a transport stream (TS) and points to the program map tables (PMTs). The decoder needs the PAT to do its job.
Source Pcraccuracy Error - Sum	The number of times that a program clock register (PCR) accuracy error occurred. This error happens when the value of the transmitted PCR differs from what is expected by more than 500 nanoseconds (ns). When a stream is encoded, the encoder assigns periodic PCR values from the encoder of program clock. The decoder relies on these values to ensure that the stream is kept in sync.

Source Pcerror - Sum	The number of times that a PCR error occurred. This error happens when PCR values are not sent frequently enough. The service relies on consistent, frequent PCRs to reset the local 27 MHz system clock. Although the error occurs when the interval exceeds 100 milliseconds (ms), best practices dictate that PCRs should be received at least every 40 ms.
Source Piderror - Sum	The number of times that a packet identifier (PID) error occurred. This error indicates that a PID is missing its associated data stream. The PIDs are identifiers that provide the location of the video, audio, and data streams. This error can occur after the TS has been multiplexed and then remultiplexed.
Source Pmterror - Sum	The number of times that a program map table (PMT) error occurred. This error happens when the PMT is not received at least every 500 milliseconds (ms). Each PMT contains a list of PIDs, which help decoders reassemble data. The decoder needs the PMTs to do its job.
Source Ptserror - Sum	The number of times that a presentation timestamp (PTS) error occurred. This error happens when a presentation timestamp (PTS) is not received at least every 700 ms. This can occur if the PTS is sent less frequently or not at all. The most common cause of this error is when the TS is scrambled.
Source Transport Error - Sum	The number of times that a primary transport error occurred. This error indicates that the TS packet is unusable. When this error occurs, ignore all other TR 101 290 errors for this packet.
Source Tsbyte Error - Sum	The number of times that a TS byte error occurred. This error indicates that the sync byte did not appear after the prescribed number of bytes.
Source Tssync Loss - Sum	The number of times that a TS sync loss error occurred. This error happens after two or more consecutive TS byte errors.

AWS: Media Connect TR101 Flow Instance Performance	
Object Name	**Object Description**

Caterror - Sum	The number of times that a conditional access table (CAT) error occurred. This error indicates that the CAT is not present. The CAT tells the integrated receiver decoder (IRD) where to find management messages for the conditional access (CA) systems that are in use.
Continuity Counter - Sum	The number of times that a continuity error occurred. This error indicates an incorrect packet order or lost packets.
Crcerror - Sum	The number of times that a cyclic redundancy check (CRC) error occurred. This error happens when a CRC determines that data is corrupted.
Paterror - Sum	The number of times that a program association table (PAT) error occurred. This error indicates that the PAT is missing. The PAT lists the programs that are available in a transport stream (TS) and points to the program map tables (PMTs). The decoder needs the PAT to do its job.
Pcraccuracy Error - Sum	The number of times that a program clock register (PCR) accuracy error occurred. This error happens when the value of the transmitted PCR differs from what is expected by more than 500 nanoseconds (ns). When a stream is encoded, the encoder assigns periodic PCR values of the encoder's program clock. The decoder relies on these values to ensure that the stream is kept in sync.
Pcrerror - Sum	The number of times that a PCR error occurred. This error happens when PCR values are not sent frequently enough. The service relies on consistent, frequent PCRs to reset the local 27 MHz system clock. Although the error occurs when the interval exceeds 100 milliseconds (ms), best practices dictate that PCRs should be received at least every 40 ms.
Piderror - Sum	The number of times that a packet identifier (PID) error occurred. This error indicates that a PID is missing its associated data stream. The PIDs are identifiers that provide the location of the video, audio, and data streams. This error can occur after the transport stream has been multiplexed and then remultiplexed.

Pmterror - Sum	The number of times that a program map table (PMT) error occurred. This error happens when the PMT is not received at least every 500 milliseconds (ms). Each PMT contains a list of PIDs, which help decoders reassemble data. The decoder needs the PMTs to do its job.
Ptserror - Sum	The number of times that a presentation timestamp (PTS) error occurred. This error happens when a presentation timestamp (PTS) is not received at least every 700 ms. This can occur if the PTS is sent less frequently or not at all. The most common cause of this error is when the transport stream (TS) is scrambled.
Transport Error - Sum	The number of times that a primary transport error occurred. This error indicates that the TS packet is unusable. When this error occurs, ignore all other TR 101 290 errors for this packet.
Tsbyte Error - Sum	The number of times that a transport stream byte error occurred. This error indicates that the sync byte did not appear after the prescribed number of bytes.
Tssync Loss - Sum	The number of times that a TS sync loss error occurred. This error happens after two or more consecutive TS byte errors.

MediaConvert Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Media Convert Queue Instance Performance	
Object Name	**Object Description**
8KOutput Duration - Avg	The number of milliseconds of 8K output for a queue.
Audio Output Duration - Avg	The number of milliseconds of audio-only output for a queue.
HDOutput Duration - Avg	The number of milliseconds of high-definition (HD) output for a queue.
Jobs Completed Count - Sum	The number of jobs completed in this queue.

Jobs Errored Count - Sum	The number of jobs that failed because of invalid inputs, such as a request to transcode a file that is not in the specified input bucket.
SDOutput Duration - Avg	The number of milliseconds of standard definition (SD) output for a queue.
Standby Time - Avg	The number of milliseconds before AWS Elemental MediaConvert starts transcoding a job.
Transcoding Time - Avg	The number of milliseconds for AWS Elemental MediaConvert to complete transcoding.
UHDOutput Duration - Avg	The number of milliseconds of ultra-high-definition (UHD) output for a queue.

MediaPackage Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Media Package Live Content Instance Performance	
Object Name	**Object Description**
Egress Bytes - Sum	Number of bytes that AWS Elemental MediaPackage successfully sends for each request. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.
Egress Request Count - Sum	Number of content requests that AWS Elemental MediaPackage receives. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.
Egress Response Time - Avg	The time that it takes AWS Elemental MediaPackage to process each output request. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.
Ingress Bytes - Sum	Number of bytes of content that AWS Elemental MediaPackage receives for each input request. If MediaPackage does not receive any requests for input in the specified interval, then no data is given.
Ingress Response Time - Avg	The time that it takes AWS Elemental MediaPackage to process each input request. If MediaPackage does not receive any requests for input in the specified interval, then no data is given.

AWS: Media Package Vod Instance Performance	
Object Name	**Object Description**
Egress Bytes - Sum	Number of bytes that AWS Elemental MediaPackage successfully sends for each request. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.
Egress Request Count - Sum	Number of content requests that AWS Elemental MediaPackage receives. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.
Egress Response Time - Avg	The time that it takes AWS Elemental MediaPackage to process each output request. If MediaPackage does not receive any requests for output in the specified interval, then no data is given.

MediaTailor Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Media Tailor Server Side Ad-Insertion Instance Performance	
Object Name	**Object Description**
Ad Decision Server Ads - Sum	The count of ads included in ad decision server (ADS) responses within the CloudWatch time period that you specified.
Ad Decision Server Duration - Sum	The total duration, in milliseconds, of all ads that MediaTailor received from the ADS within the CloudWatch time period that you specified. This duration can be greater than the Avail.Duration that you specified.
Ad Decision Server Errors - Sum	The number of non-HTTP 200 status code responses, empty responses, and timed-out responses that MediaTailor received from the ADS within the CloudWatch time period that you specified.

Ad Decision Server Fill Rate - Avg	The simple average of the rates at which the responses from the ADS filled the corresponding individual ad avails for the time period that you specified. To get the weighted average, calculate the AdDecisionServer.Duration as a percentage of the Avail.Duration. For more information about simple and weighted averages, see Simple and weighted averages.
Ad Decision Server Latency - Sum	The response time in milliseconds for requests made by MediaTailor to the ADS.
Ad Decision Server Timeouts - Sum	The number of timed-out requests to the ADS in the CloudWatch time period that you specified.
Ad Not Ready - Sum	The number of times that the ADS pointed at an ad that wasn't yet transcoded by the internal transcoder service in the time period that you specified. A high value for this metric might contribute to a low overall Avail.FillRate.
Ads Billed - Sum	The number of ads for which MediaTailor bills customers based on insertion.
Avail Duration - Sum	The planned total number of milliseconds of ad avails within the CloudWatch time period. The planned total is based on the ad avail durations in the origin manifest.
Avail Fill Rate - Avg	The planned simple average of the rates at which MediaTailor will fill individual ad avails within the CloudWatch time period. To get the weighted average, calculate the Avail.FilledDuration as a percentage of the Avail.Duration. For more information about simple and weighted averages, see Simple and weighted averages. The maximum Avail.FillRate that MediaTailor can attain is bounded by the AdDecisionServer.FillRate. If the Avail.FillRate is low, compare it to the AdDecisionServer.FillRate. If the AdDecisionServer.FillRate is low, your ADS might not be returning enough ads for the avail durations.
Avail Filled Duration - Sum	The planned number of milliseconds of ad avail time that MediaTailor will fill with ads within the CloudWatch time period.
Avail Impression - Sum	The number of ads with impression tracking events that MediaTailor sees during server-side beaconing (not the number of impressions).

Avail Observed Duration - Sum	The observed total number of milliseconds of ad avails that occurred within the CloudWatch time period. Avail.ObservedDuration is emitted at the end of the ad avail, and is based on the duration of the segments reported in the manifest during the ad avail.
Avail Observed Fill Rate - Avg	The observed simple average of the rates at which MediaTailor filled individual ad avails within the CloudWatch time period.
Avail Observed Filled Duration - Sum	The observed number of milliseconds of ad avail time that MediaTailor filled with ads within the CloudWatch time period.
Avail Observed Slate Duration - Sum	The observed total number of milliseconds of slate that was inserted within the CloudWatch period.
Cluster Replica Lag - Avg	For a read replica, the amount of lag when replicating updates from the primary instance, in milliseconds.
Cluster Replica Lag Maximum - Max	The maximum amount of lag between the primary instance and each Neptune DB instance in the DB cluster, in milliseconds.
Cluster Replica Lag Minimum - Min	The minimum amount of lag between the primary instance and each Neptune DB instance in the DB cluster, in milliseconds.
Get Manifest Errors - Sum	The number of errors received while MediaTailor was generating manifests in the CloudWatch time period that you specified.
Get Manifest Latency - Sum	The MediaTailor response time in milliseconds for the request to generate manifests.
Origin Errors - Sum	The number of non-HTTP 200 status code responses and timed-out responses that MediaTailor received from the origin server in the CloudWatch time period that you specified.
Origin Latency - Sum	The response time for requests made by MediaTailor to your content origin server.
Origin Manifest File Size Bytes - Sum	The file size of the origin manifest in bytes for both HLS and DASH. Typically this metric is used in conjunction with Origin.ManifestFileSizeTooLarge.
Origin Manifest File Size Too Large - Sum	The number of responses from the origin that have a manifest size larger than the configured amount. Typically this metric is used in conjunction with Origin.ManifestFileSizeBytes.

Origin Timeouts - Sum	The number of timed-out requests to the origin server in the CloudWatch time period that you specified.
Requests - Sum	The number of concurrent transactions per second across all request types. The transaction count depends mainly on the number of players and how often the players request updated manifests. Each player request counts as a transaction.
Skipped Reason Duration Exceeded - Sum	The number of ads that were not inserted into an avail because the ADS returned a duration of ads that was greater than the specified avail duration. A high value for this metric might contribute to a discrepancy between the Avail.Ads and AdDecisionServer.Ads metric.
Skipped Reason Early Cue In - Sum	The number of ads skipped due to an early CUE-IN.
Skipped Reason Internal Error - Sum	The number of ads skipped due to a MediaTailor internal error.
Skipped Reason New Creative - Sum	The number of ads that were not inserted into an avail because it was the first time the asset had been requested by a client. A high value for this metric might temporarily contribute to a low overall Avail.FillRate, until assets can be successfully transcoded.
Skipped Reason No Variant Match - Sum	The number of ads skipped due to there being no variant match between the ad and content.
Skipped Reason Personalization Threshold Exceeded - Sum	The duration of ads exceeding the Personalization Threshold setting in this configuration.
Skipped Reason Profile Not Found - Sum	The number of ads skipped due to the transcoding profile not being found.
Skipped Reason Transcode Error - Sum	The number of ads skipped due to a transcode error.
Skipped Reason Transcode In Progress - Sum	The count of the number of ads that were not inserted into an avail because the ad had not yet been transcoded. A high value for this metric might temporarily contribute to a low overall Avail.FillRate, until the assets can be successfully transcoded.

Neptune Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Neptune Cluster Instance Performance	
Object Name	**Object Description**
Backup Retention Period Storage Used - Sum	The total amount of backup storage, in bytes, used to support from the Neptune DB cluster's backup retention window. Included in the total reported by the TotalBackupStorageBilled metric.
Buffer Cache Hit Ratio - Avg	The percentage of requests that are served by the buffer cache. This metric can be useful in diagnosing query latency, because cache misses induce significant latency. If the cache hit ratio is below 99.9, consider upgrading the instance type to cache more data in memory.
Cpuutilization - Avg	The percentage of CPU utilization.
Engine Uptime - Avg	The amount of time that the instance has been running, in seconds.
Freeable Memory - Min	The amount of available random access memory, in bytes.
Gremlin Requests Per Sec - Avg	Number of requests per second to the Gremlin engine.
Gremlin Web Socket Open Connections - Avg	The number of open WebSocket connections to Neptune.
Loader Requests Per Sec - Avg	Number of loader requests per second.
Main Request Queue Pending Requests - Avg	The number of requests waiting in the input queue pending execution. Neptune starts throttling requests when they exceed the maximum queue capacity.
Neptune Global Dbdata Transfer Bytes - Sum	The number of bytes of redo log data transferred from the primary AWS Region to a secondary AWS Region in a Neptune global database.
Neptune Global Dbprogress Lag - Avg	The number of milliseconds that a secondary cluster is behind the primary cluster for both user transactions and system transactions.
Neptune Global Dbreplicated Write Io - Sum	The number of write I/O operations replicated from the primary AWS Region in the global database to the cluster volume in a secondary AWS Region. The billing calculations for each DB cluster in a Neptune global database use the VolumeWriteIOPS metric to account for writes performed within that cluster. For the primary DB cluster, the billing calculations use NeptuneGlobalDBReplicatedWriteIO to account for the cross-region replication to secondary DB clusters.
Network Throughput - Avg	The amount of network throughput both received from and transmitted to clients by each instance in the Neptune DB cluster, in bytes per second. This throughput does not include network traffic between instances in the DB cluster and the cluster volume.

Network Transmit Throughput - Avg	The amount of outgoing network throughput transmitted to clients by each instance in the Neptune DB cluster, in bytes per second. This throughput does not include network traffic between instances in the DB cluster and the cluster volume.
Num Tx Committed - Avg	The number of transactions successfully committed per second.
Num Tx Opened - Avg	The number of transactions opened on the server per second.
Num Tx Rolled Back - Avg	The number of transactions per second rolled back on the server because of errors.
Open Cypher Bolt Open Connections - Avg	The number of open Bolt connections to Neptune.
Open Cypher Requests Per Sec - Avg	Number of requests per second (both HTTPS and Bolt) to the openCypher engine.
Snapshot Storage Used - Sum	The total amount of backup storage consumed by all snapshots for a Neptune DB cluster outside its backup retention window, in bytes. Included in the total reported by the TotalBackupStorageBilled metric.
Sparql Requests Per Sec - Avg	The number of requests per second to the SPARQL engine.
Stats Num Statements Scanned - Avg	The total number of statements scanned for DFE statistics since the server started. Every time statistics computation is triggered, this number increases, but when no computation is happening, it remains static. As a result, if you graph it over time, you can tell when computation happened and when it didn't: By looking at the slope of the graph in periods where the metric is increasing, you can also tell how quickly the computation was going. If there is no such metric, it means that the statistics feature is disabled on your DB cluster, or that the engine version you're running doesn't have the statistics feature. If the metric value is zero, it means that no statistics computation has occurred.
Total Backup Storage Billed - Avg	The total amount of backup storage for which you are billed for a given Neptune DB cluster, in bytes. Includes the backup storage measured by the BackupRetentionPeriodStorageUsed and SnapshotStorageUsed metrics.
Total Client Errors Per Sec - Avg	The total number per second of requests that errored out because of client-side issues.
Total Requests Per Sec - Avg	The total number of requests per second to the server from all sources.
Total Server Errors Per Sec - Avg	The total number per second of requests that errored out on the server because of internal failures.

Volume Bytes Used - Avg	The total amount of storage allocated to your Neptune DB cluster, in bytes. This is the amount of storage for which you are billed. It is the maximum amount of storage allocated to your DB cluster at any point in its existence, not the amount you are currently using (see Neptune storage billing).
Volume Read Iops - Avg	The average number of billed read I/O operations from a cluster volume, reported at 5-minute intervals. Billed read operations are calculated at the cluster volume level, aggregated from all instances in the Neptune DB cluster, and then reported at 5-minute intervals.
Volume Write Iops - Avg	The average number of write disk I/O operations to the cluster volume, reported at 5-minute intervals.

AWS: Neptune Deprecated Cluster Instance Performance	
Object Name	**Object Description**
Gremlin Errors - Avg	Number of errors in Gremlin traversals.
Gremlin Http1Xx - Avg	Number of HTTP 1xx responses for the Gremlin endpoint per second. We recommend that you use the new Http1xx combined metric instead.
Gremlin Http2Xx - Avg	Number of HTTP 2xx responses for the Gremlin endpoint per second. We recommend that you use the new Http2xx combined metric instead.
Gremlin Http4Xx - Avg	Number of HTTP 4xx errors for the Gremlin endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Gremlin Http5Xx - Avg	Number of HTTP 5xx errors for the Gremlin endpoint per second. We recommend that you use the new Http5xx combined metric instead.
Gremlin Requests - Avg	Number of requests to Gremlin engine.
Gremlin Web Socket Available Connections - Avg	Number of potential WebSocket connections currently available.
Gremlin Web Socket Client Errors - Avg	Number of WebSocket client errors on the Gremlin endpoint per second.
Gremlin Web Socket Server Errors - Avg	Number of WebSocket server errors on the Gremlin endpoint per second.
Gremlin Web Socket Success - Avg	Number of successful WebSocket connections to the Gremlin endpoint per second.
Http100 - Avg	Number of HTTP 100 responses for the endpoint per second. We recommend that you use the new Http1xx combined metric instead.

Http101 - Avg	Number of HTTP 101 responses for the endpoint per second. We recommend that you use the new Http1xx combined metric instead.
Http1Xx - Avg	Number of HTTP 1xx responses for the endpoint per second.
Http200 - Avg	Number of HTTP 200 responses for the endpoint per second. We recommend that you use the new Http2xx combined metric instead.
Http2Xx - Avg	Number of HTTP 2xx responses for the endpoint per second.
Http400 - Avg	Number of HTTP 400 errors for the endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Http403 - Avg	Number of HTTP 403 errors for the endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Http405 - Avg	Number of HTTP 405 errors for the endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Http413 - Avg	Number of HTTP 413 errors for the endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Http429 - Avg	Number of HTTP 429 errors for the endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Http4Xx - Avg	Number of HTTP 4xx errors for the endpoint per second.
Http500 - Avg	Number of HTTP 500 errors for the endpoint per second. We recommend that you use the new Http5xx combined metric instead.
Http501 - Avg	Number of HTTP 501 errors for the endpoint per second. We recommend that you use the new Http5xx combined metric instead.
Http5Xx - Avg	Number of HTTP 5xx errors for the endpoint per second.
Loader Errors - Avg	Number of errors from Loader requests.
Loader Requests - Avg	Number of Loader Requests.
Sparql Errors - Avg	Number of errors in the SPARQL queries.

Sparql Http1Xx - Avg	Number of HTTP 1xx responses for the SPARQL endpoint per second. We recommend that you use the new Http1xx combined metric instead.
Sparql Http2Xx - Avg	Number of HTTP 2xx responses for the SPARQL endpoint per second. We recommend that you use the new Http2xx combined metric instead.
Sparql Http4Xx - Avg	Number of HTTP 4xx errors for the SPARQL endpoint per second. We recommend that you use the new Http4xx combined metric instead.
Sparql Http5Xx - Avg	Number of HTTP 5xx errors for the SPARQL endpoint per second. We recommend that you use the new Http5xx combined metric instead.
Sparql Requests - Avg	Number of requests to the SPARQL engine.
Status Errors - Avg	Number of errors from the status endpoint.
Status Requests - Avg	Number of requests to the status endpoint.

Network ELB Service

AWS: Network ELB Instance Configuration	
Object Name	**Object Description**
AWS Network ELB/Availability Zones	The availability zones for the load balancer.
AWS Network ELB/Target Groups	The unique identifiers of the target groups for the network load balancer.
AWS Network ELB/VPC	The unique identifier of the VPC for the load balancer.
DNS Name	The public DNS name of the network load balancer.
Key	The key of the tag, that belongs to the network load balancer.
Listener ARN	The Amazon Resource Name (ARN) of the listener.
Listener Port	The port on which the load balancer is listening.
Listener Protocol	The protocol for connections from clients to the load balancer.
Load Balancer ARN	The Amazon Resource Name (ARN) of the load balancer.

Load Balancer Name	The name of the load balancer.
Load Balancer State	The state of the load balancer.
Scheme	The DNS name of an Internet-facing load balancer is publicly resolvable to the public IP addresses of the nodes.
Type	The type of load balancer.
Value	The value of the tag, that belongs to the network load balancer.

AWS: Network ELB Instance Performance	
Object Name	**Object Description**
Active Flow Count	The total number of concurrent TCP flows (or connections) from clients to targets.
Active Flow Count TLS	The total number of concurrent TLS flows (or connections) from clients to targets. This metric includes only connections in the ESTABLISHED states.
Client TLS Negotation Error Count	The total number of TLS handshakes that failed during negotiation between a client and a TLS listener.
Consumed LCUst	The number of load balancer capacity units (LCU) used.
New Flow Count	The total number of new TCP flows (or connections) established from clients to targets in the time period.
New Flow Count TLS	The total number of new TLS flows (or connections) established from clients to targets in the time period.
Processed Bytes	The total number of bytes processed by the load balancer, including TCP/IP headers.
Processed Bytes TLS	The total number of bytes processed by TLS listeners.
Target TLS Negotiation Error Count	The total number of TLS handshakes that failed during negotiation between a TLS listener and a target.
TCP Client Reset Count	The total number of reset (RST) packets sent from a client to a target.

TCP ELB Reset Count	The total number of reset (RST) packets generated by the load balancer.
TCP Target Reset Count	The total number of reset (RST) packets sent from a target to a client.

OpenSearch Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: OpenSearch Cluster Performance	
Object Name	**Object Description**
2Xx - Sum	The number of requests to the domain that resulted in the given HTTP response code (2xx).
3Xx - Sum	The number of requests to the domain that resulted in the given HTTP response code (3xx).
4Xx - Sum	The number of requests to the domain that resulted in the given HTTP response code (4xx).
5Xx - Sum	The number of requests to the domain that resulted in the given HTTP response code (5xx).
Automated Snapshot Failure - Sum	The number of failed automated snapshots for the cluster. A value of 1 indicates that no automated snapshot was taken for the domain in the previous 36 hours.
Cluster Index Writes Blocked - Max	Indicates whether your cluster is accepting or blocking incoming write requests. A value of 0 means that the cluster is accepting requests. A value of 1 means that it is blocking requests.
Cluster Status Green - Max	A value of 1 indicates that all index shards are allocated to nodes in the cluster.
Cluster Status Red - Max	A value of 1 indicates that the primary and replica shards for at least one index are not allocated to nodes in the cluster.
Cluster Status Yellow - Max	A value of 1 indicates that the primary shards for all indexes are allocated to nodes in the cluster, but replica shards for at least one index are not.

Cluster Used Space - Max	The total used space for the cluster. You must leave the period at one minute to get an accurate value.
Cpucredit Balance - Min	The remaining CPU credits available for data nodes in the cluster. A CPU credit provides the performance of a full CPU core for one minute.
Cpuutilization - Avg	The percentage of CPU usage for data nodes in the cluster. Maximum shows the node with the highest CPU usage. Average represents all nodes in the cluster.
Cpuutilization - Max	The percentage of CPU usage for data nodes in the cluster. Maximum shows the node with the highest CPU usage. Average represents all nodes in the cluster.
Deleted Documents - Avg	The total number of documents marked for deletion across all data nodes in the cluster. These documents no longer appear in search results, but OpenSearch only removes deleted documents from disk during segment merges. This metric increases after delete requests and decreases after segment merges.
Free Storage Space - Sum	The free space for data nodes in the cluster. Sum shows total free space for the cluster, but you must leave the period at one minute to get an accurate value. Minimum and Maximum show the nodes with the least and most free space, respectively.
Invalid Host Header Requests - Sum	The number of HTTP requests made to the OpenSearch cluster that included an invalid (or missing) host header. Valid requests include the domain hostname as the host header value. OpenSearch Service rejects invalid requests for public access domains that don't have a restrictive access policy. We recommend applying a restrictive access policy to all domains. If you see large values for this metric, confirm that your OpenSearch clients include the domain hostname (and not, for example, its IP address) in their requests.

Jvmmemory Pressure - Max	The maximum percentage of the Java heap used for all data nodes in the cluster. OpenSearch Service uses half of an instance's RAM for the Java heap, up to a heap size of 32 GiB. You can scale instances vertically up to 64 GiB of RAM, at which point you can scale horizontally by adding instances.
Kibana Reporting Failed Request Sys Err Count - Sum	The number of requests to generate OpenSearch Dashboards reports that failed due to server problems or feature limitations.
Kibana Reporting Gailed Request User Err Count - Sum	The number of requests to generate OpenSearch Dashboards reports that failed due to client issues.
Kibana Reporting Request Count - Sum	The total number of requests to generate OpenSearch Dashboards reports.
Kibana Reporting Success Count - Sum	The number of successful requests to generate OpenSearch Dashboards reports.
Kmskey Error - Max	A value of 1 indicates that the AWS KMS key used to encrypt data at rest has been disabled. To restore the domain to normal operations, re-enable the key. The console displays this metric only for domains that encrypt data at rest.
Kmskey Inaccessible - Max	A value of 1 indicates that the AWS KMS key used to encrypt data at rest has been deleted or revoked its grants to OpenSearch Service. You can't recover domains that are in this state. If you have a manual snapshot, though, you can use it to migrate the domain's data to a new domain. The console displays this metric only for domains that encrypt data at rest.
Nodes - Avg	The number of nodes in the OpenSearch Service cluster, including dedicated master nodes and UltraWarm nodes.
Open Search Dashboards Healthy Nodes - Avg	A health check for OpenSearch Dashboards. If the minimum, maximum, and average are all equal to 1, Dashboards is behaving normally. If you have 10 nodes with a maximum of 1, minimum of 0, and average of 0.7, this means 7 nodes (70%) are healthy and 3 nodes (30%) are unhealthy. Note that this object was previously KibanaHealthyNodes in AWS).
Open Search Requests - Sum	The number of requests made to the OpenSearch cluster.

Searchable Documents - Avg	The total number of searchable documents across all data nodes in the cluster.
Shards Active - Avg	The total number of active primary and replica shards.
Shards Active Primary - Avg	The number of active primary shards.
Shards Delayed Unassigned - Avg	The number of shards whose node allocation has been delayed by the timeout settings.
Shards Initializing - Avg	The number of shards that are under initialization.
Shards Relocating - Avg	The number of shards that are under relocation.
Shards Unassigned - Avg	The number of shards that are not allocated to nodes in the cluster.

AWS: OpenSearch EBS Performance	
Object Name	**Object Description**
Disk Queue Depth - Avg	The number of pending input and output (I/O) requests for an EBS volume.
Disk Queue Depth - Max	The number of pending input and output (I/O) requests for an EBS volume.
Disk Queue Depth - Min	The number of pending input and output (I/O) requests for an EBS volume.
Read Iops - Avg	The number of input and output (I/O) operations per second for read operations on EBS volumes.
Read Iops - Max	The number of input and output (I/O) operations per second for read operations on EBS volumes.
Read Iops - Min	The number of input and output (I/O) operations per second for read operations on EBS volumes.
Read Latency - Avg	The latency, in seconds, for read operations on EBS volumes.
Read Latency - Max	The latency, in seconds, for read operations on EBS volumes.
Read Latency - Min	The latency, in seconds, for read operations on EBS volumes.
Read Throughput - Avg	The throughput, in bytes per second, for read operations on EBS volumes.

Read Throughput - Max	The throughput, in bytes per second, for read operations on EBS volumes.
Read Throughput - Min	The throughput, in bytes per second, for read operations on EBS volumes.
Write Iops - Avg	The number of input and output (I/O) operations per second for write operations on EBS volumes.
Write Iops - Max	The number of input and output (I/O) operations per second for write operations on EBS volumes.
Write Iops - Min	The number of input and output (I/O) operations per second for write operations on EBS volumes.
Write Latency - Avg	The latency, in seconds, for write operations on EBS volumes.
Write Latency - Max	The latency, in seconds, for write operations on EBS volumes.
Write Latency - Min	The latency, in seconds, for write operations on EBS volumes.
Write Throughput - Avg	The throughput, in bytes per second, for write operations on EBS volumes.
Write Throughput - Max	The throughput, in bytes per second, for write operations on EBS volumes.
Write Throughput - Min	The throughput, in bytes per second, for write operations on EBS volumes.

AWS: OpenSearch Instances Performance	
Object Name	**Object Description**
Coordinating Write Rejected - Sum	The total number of rejections happened on the coordinating node due to indexing pressure since the last OpenSearch Service process startup.
Indexing Latency - Avg	The average time, in milliseconds, that it takes a shard to complete an indexing operation.
Indexing Latency - Max	The average time, in milliseconds, that it takes a shard to complete an indexing operation.

Indexing Rate - Avg	The number of indexing operations per minute. A single call to the _bulk API that adds two documents and updates two counts as four operations, which might be spread across one or more nodes. If that index has one or more replicas, other nodes in the cluster also record a total of four indexing operations. Document deletions do not count towards this metric.
Jvmgcold Collection Count - Avg	The number of times that "old generation" garbage collection has run. In a cluster with sufficient resources, this number should remain small and grow infrequently.
Jvmgcold Collection Count - Max	The number of times that "old generation" garbage collection has run. In a cluster with sufficient resources, this number should remain small and grow infrequently.
Jvmgcold Collection Time - Avg	The amount of time, in milliseconds, that the cluster has spent performing "old generation" garbage collection.
Jvmgcyoung Collection Count - Avg	The number of times that "young generation" garbage collection has run. A large, ever-growing number of runs is a normal part of cluster operations.
Jvmgcyoung Collection Time - Avg	The amount of time, in milliseconds, that the cluster has spent performing "young generation" garbage collection.
Jvmgcyoung Collection Time - Max	The amount of time, in milliseconds, that the cluster has spent performing "young generation" garbage collection.
Open Search Dashboards Concurrent Connections - Avg	The number of active concurrent connections to OpenSearch Dashboards. If this number is consistently high, consider scaling your cluster. (previously KibanaConcurrentConnections).
Open Search Dashboards Concurrent Connections - Max	The number of active concurrent connections to OpenSearch Dashboards. If this number is consistently high, consider scaling your cluster. (previously KibanaConcurrentConnections).
Open Search Dashboards Healthy Node - Avg	A health check for the individual OpenSearch Dashboards node. A value of 1 indicates normal behavior. A value of 0 indicates that Dashboards is inaccessible. (previously KibanaHealthyNode).
Open Search Dashboards Healthy Node - Min	A health check for the individual OpenSearch Dashboards node. A value of 1 indicates normal behavior. A value of 0 indicates that Dashboards is inaccessible. (previously KibanaHealthyNode).
Open Search Dashboards Heap Total - Avg	The amount of heap memory allocated to OpenSearch Dashboards in MiB. Different EC2 instance types can impact the exact memory allocation. (previously KibanaHealthyNode).

Open Search Dashboards Heap Used - Avg	The absolute amount of heap memory used by OpenSearch Dashboards in MiB. (previously KibanaHeapUsed).
Open Search Dashboards Heap Utilization - Avg	The maximum percentage of available heap memory used by OpenSearch Dashboards. If this value increases above 80%, consider scaling your cluster. (previously KibanaHeapUtilization).
Open Search Dashboards Heap Utilization - Max	The maximum percentage of available heap memory used by OpenSearch Dashboards. If this value increases above 80%, consider scaling your cluster.
Open Search Dashboards Os1Minute Load - Avg	The one-minute CPU load average for OpenSearch Dashboards. The CPU load should ideally stay below 1.00. While temporary spikes are fine, we recommend increasing the size of the instance type if this metric is consistently above 1.00. (previously KibanaOS1MinuteLoad).
Open Search Dashboards Os1Minute Load - Max	The one-minute CPU load average for OpenSearch Dashboards. The CPU load should ideally stay below 1.00. While temporary spikes are fine, we recommend increasing the size of the instance type if this metric is consistently above 1.00. (previously KibanaOS1MinuteLoad).
Open Search Dashboards Request Total - Sum	The total count of HTTP requests made to OpenSearch Dashboards. If your system is slow or you see high numbers of Dashboards requests, consider increasing the size of the instance type. (previously KibanaRequestTotal).
Open Search Dashboards Response Times Max In Millis - Avg	The maximum amount of time, in milliseconds, that it takes for OpenSearch Dashboards to respond to a request. If requests consistently take a long time to return results, consider increasing the size of the instance type. (previously KibanaResponseTimesMaxInMillis).
Open Search Dashboards Response Times max In Millis - Max	The maximum amount of time, in milliseconds, that it takes for OpenSearch Dashboards to respond to a request. If requests consistently take a long time to return results, consider increasing the size of the instance type. (previously KibanaResponseTimesMaxInMillis).
Primary Write Rejected - Sum	The total number of rejections happened on the primary shards due to indexing pressure since the last OpenSearch Service process startup.

Replica Write Rejected - Sum	The total number of rejections happened on the replica shards due to indexing pressure since the last OpenSearch Service process startup. This metric is available in version 7.9 and above.
Search Latency - Avg	The average time, in milliseconds, that it takes a shard on a data node to complete a search operation.
Search Latency - Max	The average time, in milliseconds, that it takes a shard on a data node to complete a search operation.
Search Rate - Avg	The total number of search requests per minute for all shards on a data node. A single call to the _search API might return results from many different shards. If five of these shards are on one node, the node would report 5 for this metric, even though the client only made one request.
Segment Count - Avg	The number of segments on a data node. The more segments you have, the longer each search takes. OpenSearch occasionally merges smaller segments into a larger one.
Segment Count - Max	The number of segments on a data node. The more segments you have, the longer each search takes. OpenSearch occasionally merges smaller segments into a larger one.
Sys Memory Utilization - Avg	The percentage of the instance's memory that is in use. High values for this metric are normal and usually do not represent a problem with your cluster. For a better indicator of potential performance and stability issues, see the JVMMemoryPressure metric.
Threadpool Bulk Queue - Sum	The number of queued tasks in the bulk thread pool. If the queue size is consistently high, consider scaling your cluster.
Threadpool Bulk Rejected - Sum	The number of rejected tasks in the bulk thread pool. If this number continually grows, consider scaling your cluster.
Threadpool Bulk Threads - Sum	The size of the bulk thread pool.
Threadpool Force_Merge Queue - Sum	The number of queued tasks in the force merge thread pool. If the queue size is consistently high, consider scaling your cluster.
Threadpool Force_Merge Rejected - Sum	The number of rejected tasks in the force merge thread pool. If this number continually grows, consider scaling your cluster.
Threadpool Force_Merge Threads - Sum	The size of the force merge thread pool.

Threadpool Index Queue - Sum	The number of queued tasks in the index thread pool. If the queue size is consistently high, consider scaling your cluster. The maximum index queue size is 200.
Threadpool Index Rejected - Sum	The number of rejected tasks in the index thread pool. If this number continually grows, consider scaling your cluster.
Threadpool Index Threads - Sum	The size of the index thread pool.
Threadpool Search Queue - Sum	The number of queued tasks in the search thread pool. If the queue size is consistently high, consider scaling your cluster. The maximum search queue size is 1,000.
Threadpool Search Rejected - Sum	The number of rejected tasks in the search thread pool. If this number continually grows, consider scaling your cluster.
Threadpool Search Threads - Sum	The size of the search thread pool.
Threadpool Write Queue - Sum	The number of queued tasks in the write thread pool.
Threadpool Write Rejected - Sum	The number of rejected tasks in the write thread pool. Note Because the default write queue size was increased from 200 to 10000 in version 7.9, this metric is no longer the only indicator of rejections from OpenSearch Service. Use the CoordinatingWriteRejected, PrimaryWriteRejected, and ReplicaWriteRejected metrics to monitor rejections in versions 7.9 and later.
Threadpool Write Threads - Sum	The size of the write thread pool.
Threadpoolsql-Worker Queue - Sum	The number of queued tasks in the SQL search thread pool. If the queue size is consistently high, consider scaling your cluster.
Threadpoolsql-Worker Rejected - Sum	The number of rejected tasks in the SQL search thread pool. If this number continually grows, consider scaling your cluster.
Threadpoolsql-Worker Threads - Sum	The size of the SQL search thread pool.

AWS: OpenSearch k-NN Performance	
Object Name	**Object Description**
Knncache Capacity Reached - Max	Per-node metric for whether cache capacity has been reached. This metric is only relevant to approximate k-NN search.

Knncircuit Breaker Triggered - Max	Per-cluster metric for whether the circuit breaker is triggered. If any nodes return a value of 1 for KNNCacheCapacityReached, this value will also return 1. This metric is only relevant to approximate k-NN search.
Knneviction Count - Sum	Per-node metric for the number of graphs that have been evicted from the cache due to memory constraints or idle time. Explicit evictions that occur because of index deletion are not counted. This metric is only relevant to approximate k-NN search.
Knngraph Index Errors - Sum	Per-node metric for the number of requests to add the knn_vector field of a document to a graph that produced an error.
Knngraph Index Requests - Sum	Per-node metric for the number of requests to add the knn_vector field of a document to a graph.
Knngraph Memory Usage - Avg	Per-node metric for the current cache size (total size of all graphs in memory) in kilobytes. This metric is only relevant to approximate k-NN search.
Knngraph Query Errors - Sum	Per-node metric for the number of graph queries that produced an error.
Knngraph Query Requests - Sum	Per-node metric for the number of graph queries.
Knnhit Count - Sum	Per-node metric for the number of cache hits. A cache hit occurs when a user queries a graph that is already loaded into memory. This metric is only relevant to approximate k-NN search.
Knnload Exception Count - Sum	Per-node metric for the number of times an exception occurred while trying to load a graph into the cache. This metric is only relevant to approximate k-NN search.
Knnload Success Count - Sum	Per-node metric for the number of times the plugin successfully loaded a graph into the cache. This metric is only relevant to approximate k-NN search.
Knnmiss Count - Sum	Per-node metric for the number of cache misses. A cache miss occurs when a user queries a graph that is not yet loaded into memory. This metric is only relevant to approximate k-NN search.
Knnquery Requests - Sum	Per-node metric for the number of query requests the k-NN plugin received.
Knnscript Compilation Errors - Sum	Per-node metric for the number of errors during script compilation. This statistic is only relevant to k-NN score script search.

Knnscript Compilations - Sum	Per-node metric for the number of times the k-NN script has been compiled. This value should usually be 1 or 0, but if the cache containing the compiled scripts is filled, the k-NN script might be recompiled. This statistic is only relevant to k-NN score script search.
Knnscript Query Errors - Sum	Per-node metric for the number of errors during script queries. This statistic is only relevant to k-NN score script search.
Knnscript Query Requests - Sum	Per-node metric for the total number of script queries. This statistic is only relevant to k-NN score script search.
Knntotal Load Time - Sum	The time in nanoseconds that k-NN has taken to load graphs into the cache. This metric is only relevant to approximate k-NN search.

AWS: OpenSearch Master Nodes Performance	
Object Name	**Object Description**
Master Cpucredit Balance - Min	The remaining CPU credits available for dedicated master nodes in the cluster. A CPU credit provides the performance of a full CPU core for one minute. For more information, see CPU credits in the Amazon EC2 Developer Guide. This metric is available only for the T2 instance types.
Master Cpuutilization - Max	The maximum percentage of CPU resources used by the dedicated master nodes. We recommend increasing the size of the instance type when this metric reaches 60 percent.
Master Jvmmemory Pressure - Max	The maximum percentage of the Java heap used for all dedicated master nodes in the cluster. We recommend moving to a larger instance type when this metric reaches 85 percent.
Master Reachable From Node - Min	A health check for MasterNotDiscovered exceptions. A value of 1 indicates normal behavior. A value of 0 indicates that /_cluster/health/ is failing. Failures mean that the master node stopped or is not reachable. They are usually the result of a network connectivity issue or AWS dependency problem.
Master Sys Memory Utilization - Max	The percentage of the master node's memory that is in use.

AWS: OpenSearch SQL Performance	
Object Name	**Object Description**
Sqldefault Cursor Request Count - Sum	Similar to SQLRequestCount but only counts pagination requests.
Sqlfailed Request Count By Cus Err - Sum	The number of requests to the _sql API that failed due to a client issue. For example, a request might return HTTP status code 400 due to an IndexNotFoundException.
Sqlfailed Request Count By Sys Err - Sum	The number of requests to the _sql API that failed due to a server problem or feature limitation. For example, a request might return HTTP status code 503 due to a VerificationException.
Sqlrequest Count - Sum	The number of requests to the _sql API.
Sqlunhealthy - Max	A value of 1 indicates that, in response to certain requests, the SQL plugin is returning 5xx response codes or passing invalid query DSL to OpenSearch. Other requests should continue to succeed. A value of 0 indicates no recent failures. If you see a sustained value of 1, troubleshoot the requests your clients are making to the plugin.

OpsWorks Service

AWS: OpsWorks Instance Configuration	
Object Name	**Object Description**
AWS OpsWorks/EC2	The unique identifier for the OpsWorks-EC2 relationship.
Chef Version	The versionof Chef used to deploy the stack.
Default Operating System	The default Operating System of the Stack.
Default Root Device Type	The type of the root file system.
Default SSH Key	The SSH keyname associated with the key-pair of the stack.
Default Subnet	The default subnet of the Stack.

Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Hostname Theme	The default naming scheme, if any, for the instances of the stack.
Id	The unique identifier of the Stack.
Name	The common name of the Stack.
Region	The region where the stack is deployed.
Use Custom Chef Cookbooks	Whether or not the stack was deployed with a custom Chef cookbook.
Use OpsWorks Security Groups	Whether or not to use OpsWorks Security Groups for instances within the stack.
VPC	The virtual private cloud, if any, associated with the Stack.

AWS: OpsWorks Service Health	
Object Name	**Object Description**
Date	A timestamp when this service's health status was originally written.
Health	A text description of the health status of this AWS service.

AWS: OpsWorks Stack Performance	
Object Name	**Object Description**
Active Processes	The number of active processes.
CPU Idle	The percentage of time the CPU is idle.
CPU Nice	The percentage of time that the CPU is handling processes with a positive nice value, which have lower scheduling priority. For further information, see nice (Unix).
CPU Steal	The percentage of time that an instance is waiting for the hypervisor to allocate physical CPU resources.

CPU System	The percentage of time that the CPU is handling system operations.
CPU User	The percentage of time the CPU is handling user operations.
CPU Wait I/O	The percentage of time the CPU is waiting on I/O operations.
Load 1 Minute	The Unix load average over a 1 minute window.
Load 15 Minutes	The Unix load averaged over a 15 minute window.
Load 5 Minutes	The Unix load averaged over a 5 minute window.
Memory Buffers	The amount of buffered memory.
Memory Cached	The amount of cached memory.
Memory Free	The amount of free memory.
Memory Swap	The amount of swap space used.
Memory Total	The total amount of memory.
Memory Used	The total amount of memory used.

RDS Service

AWS: RDS Aurora Cluster Instance Configuration	
Object Name	**Object Description**
Allocated Storage	The amount of allocated storage available for a database cluster. For AWS Aurora clusters this will always return a value of 1 as Aurora DB Cluster sizes are not fixed.
Availablity Zones	Availability zones where database instances in the cluster can be created.
Backup Retention Period	Number of days which automatic DB snapshots are stored.
Capacity	Current capacity active AWS Aurora Serverless DB Clusters. Will return 0 when cluster is paused.
Cluster Creation Time	When the DB cluster was created in UTC
Custom Endpoints	Custom endpoints associated with the cluster.

Database Name	The name of the Database Cluster that was provided at creation time.
DB Cluster ARN	The ARN associated with the DB cluster.
DB Cluster Identifier	DBClusterIdentifier provided for the DB Cluster
DB Cluster Parameter Group	Name of the DB Cluster Parameter Group.
DB Cluster Resource ID	Region unique immutable identifier for the cluster.
DB Subnet Group	Specifies information about the DB Clusters subnet group.
Deletion Protection	Specifies if deletion protection is enabled.
Endpoint	Connection endpoint for the primary instance of the DB Cluster
Engine	Name of the database engine used for the DB cluster.
Engine Mode	Mode of the DB engine for the cluster.
Engine Version	Version of the database engine used for the DB cluster.
Hosted Zone ID	Specifies the ID that Amazon Route 53 assigns when you create a hosted zone.
IAM Database Authentication Enabled	Specifies if IAM database authentication is enabled.
Key	Tag key value.
KMS Key ID	If database is encrypted the KMS Key identifier for the encrypted DB cluster.
Master Username	Master username for the DB cluster.
Multiple Availability Zones	Identifies if the database cluster has instances in multiple availability zones.
Port	The port that the database engine is listening on.
Read Replica Identifier	Contains one or more of the identifiers for the read replicas associated with the DB cluster.
Reader Endpoint	Reader endpoint for the DB Cluster.
Replication Source Identifier	Contains the identifier of the source DB cluster if the DB cluster is a read replica.
Status	The status of the database cluster.
Storage Encrypted	Specifies if the DB Cluster is encrypted.
Value	Tag value.

AWS: RDS Aurora Cluster Instance Performance	
Object Name	**Object Description**
Aborted Clients	The number of client connections that have not been closed properly.
Active Transactions	The average number of current transactions executing on an Aurora database instance per second.
Aurora Bin Log Replica Lag	The amount of time a replica DB cluster running on Aurora with MySQL compatibility lags behind the source DB cluster.
Aurora Global DB Data Transfer Bytes	The amount of redo log data transferred from the master AWS Region to a secondary AWS Region.
Aurora Global DB Replicated Write IO	The number of write I/O operations replicated from the primary AWS Region to the cluster volume in a secondary AWS Region.
Aurora Global DB Replication Lag	The amount of lag when replicating updates from the primary AWS Region.
Aurora Parallel Query Bytes Returned	The number of bytes for the tuple data structures transmitted to the head node during parallel queries.
Aurora Parallel Query Max Concurrent Requests	The maximum number of parallel query sessions that can run concurrently on this Aurora DB instance.
Aurora Parallel Query Pages Pushed Down	The number of data pages (each with a fixed size of 16 KiB) where parallel query avoided a network transmission to the head node.
Aurora Parallel Query Request Attempted	The number of parallel query sessions requested.
Aurora Parallel Query Request Executed	The number of parallel query sessions run successfully.
Aurora Parallel Query Request Failed	The number of parallel query sessions that returned an error to the client.
Aurora Parallel Query Request In Progress	The number of parallel query sessions currently in progress.
Aurora Parallel Query Request Not Chosen	The number of times parallel query was not chosen to satisfy a query.
Aurora Parallel Query Request Not Chosen Below Min Rows	The number of times parallel query was not chosen due to the number of rows in the table.

Aurora Parallel Query Request Not Chosen Column Bit	The number of parallel query requests that use the nonparallel query processing path because of an unsupported data type in the list of projected columns.
Aurora Parallel Query Request Not Chosen Column Geometry	The number of parallel query requests that use the nonparallel query processing path because the table has columns with the GEOMETRY data type.
Aurora Parallel Query Request Not Chosen Column Lob	The number of parallel query requests that use the nonparallel query processing path because the table has columns with a LOB data type, or VARCHAR columns that are stored externally due the declared length.
Aurora Parallel Query Request Not Chosen Column Virtual	The number of parallel query requests that use the nonparallel query processing path because the table contains a virtual column.
Aurora Parallel Query Request Not Chosen Custom Charset	The number of parallel query requests that use the nonparallel query processing path because the table has columns with a custom character set.
Aurora Parallel Query Request Not Chosen Fast DDL	The number of parallel query requests that use the nonparallel query processing path because the table is currently being altered by a fast DDL ALTER statement.
Aurora Parallel Query Request Not Chosen Full Text Index	The number of parallel query requests that use the nonparallel query processing path because the table has full-text indexes.
Aurora Parallel Query Request Not Chosen High Buffer Pool Pct	The number of times parallel query was not chosen because a high percentage of the table data (currently, greater than 95 percent) was already in the buffer pool.
Aurora Parallel Query Request Not Chosen Index Hint	The number of parallel query requests that use the nonparallel query processing path because the query includes an index hint.
Aurora Parallel Query Request Not Chosen InnoDB Table Format	The number of parallel query requests that use the nonparallel query processing path because the table uses an unsupported InnoDB row format.
Aurora Parallel Query Request Not Chosen Long Trx	The number of parallel query requests that used the nonparallel query processing path, due to the query being started inside a long-running transaction.
Aurora Parallel Query Request Not Chosen No Where Clause	The number of parallel query requests that use the nonparallel query processing path because the query does not include any WHERE clause.

Aurora Parallel Query Request Not Chosen Range Scan	The number of parallel query requests that use the nonparallel query processing path because the query uses a range scan on an index.
Aurora Parallel Query Request Not Chosen Row Length Too Long	The number of parallel query requests that use the nonparallel query processing path because the total combined length of all the columns is too long.
Aurora Parallel Query Request Not Chosen Small Table	The number of times parallel query was not chosen due to the overall size of the table, as determined by number of rows and average row length.
Aurora Parallel Query Request Not Chosen Temporary Table	The number of parallel query requests that use the nonparallel query processing path because the query refers to temporary tables that use the unsupported MyISAM or memory table types.
Aurora Parallel Query Request Not Chosen Tx Isolation	The number of parallel query requests that use the nonparallel query processing path because query uses an unsupported transaction isolation level.
Aurora Parallel Query Request Not Chosen Unsupported Access	The number of parallel query requests that use the nonparallel query processing path because the WHERE clause does not meet the criteria for parallel query.
Aurora Parallel Query Request Not Chosen Update Delete Stmts	The number of parallel query requests that use the nonparallel query processing path because the query is part of an UPDATE or DELETE statement.
Aurora Parallel Query Request Throttled	The number of times parallel query was not chosen due to the maximum number of concurrent parallel queries already running on a particular Aurora DB instance.
Aurora PQ Request Not Chosen Few Pages Outside Buffer Pool	The number of times parallel query was not chosen, even though less than 95 percent of the table data was in the buffer pool, because there was not enough unbuffered table data to make parallel query worthwhile.
Aurora Replica Lag	For an Aurora replica, the amount of lag when replicating updates from the primary instance.
Aurora Replica Lag Maximum	The maximum amount of lag between the primary instance and each Aurora DB instance in the DB cluster.
Aurora Replica Lag Minimum	The minimum amount of lag between the primary instance and each Aurora DB instance in the DB cluster.

Aurora Volume Bytes Left Total	The remaining available space for the cluster volume. As the cluster volume grows, this value decreases. If it reaches zero, the cluster reports an out-of-space error.
Backtrack Change Records Creation Rate	The number of backtrack change records created over 5 minutes for your DB cluster.
Backtrack Change Records Stored	The number of backtrack change records used by your DB cluster.
Backtrack Window Actual	The difference between the target backtrack window and the actual backtrack window.
Backtrack Window Alert	The number of times that the actual backtrack window is smaller than the target backtrack window for a given period of time.
Backup Retention Period Storage Used	The total amount of backup storage used to support the point-in-time restore feature within the Aurora DB cluster's backup retention window.
Bin Log Disk Usage	The amount of disk space occupied by binary logs on the primary instance.
Blocked Transactions	The average number of transactions in the database that are blocked per second.
Buffer Cache Hit Ratio	The percentage of requests that are served by the buffer cache.
Commit Latency	The average duration of commit operations.
Commit Throughput	The average number of commit operations per second.
CPU Credit Balance	The number of CPU credits that an instance has accumulated, reported at 5-minute intervals. This metric applies only to db.t2.small and db.t2.medium instances for Aurora MySQL, and to db.t3 instances for Aurora PostgreSQL.
CPU Credit Usage	The number of CPU credits consumed during the specified period, reported at 5-minute intervals. This metric applies only to db.t2.small and db.t2.medium instances for Aurora MySQL, and to db.t3 instances for Aurora PostgreSQL.
CPU Utilization	The percentage of CPU used by an Aurora DB instance.
Database Connections	The current number of connections to an Aurora DB instance.

DDL Latency	The average duration of requests such as example, create, alter, and drop requests.
DDL Throughput	The average number of DDL requests per second.
Deadlocks	The average number of deadlocks in the database per second.
Delete Latency	The average duration of delete operations.
Delete Throughput	The average number of delete queries per second.
Disk Queue Depth	The number of outstanding read/write requests waiting to access the disk.
DML Latency	The average duration of inserts, updates, and deletes.
DML Throughput	The average number of inserts, updates, and deletes per second.
Engine Uptime	The amount of time that the instance has been running.
Free Local Storage	The amount of local storage available.
Freeable Memory	The amount of available random access memory.
Insert Latency	The average duration of insert operations.
Insert Throughput	The average number of insert operations per second.
Login Failures	The average number of failed login attempts per second.
Maximum Used Transaction IDs	The age of the oldest unvacuumed transaction ID, in transactions.
Network Receive Throughput	The amount of network throughput received from clients by each instance in the Aurora MySQL DB cluster.
Network Throughput	The amount of network throughput both received from and transmitted to clients by each instance in the Aurora MySQL DB cluster.
Network Transmit Throughput	The amount of network throughput sent to clients by each instance in the Aurora DB cluster.
Num Binary Log Files	The number of binlog files generated.
Queries	The average number of queries executed per second.

Read IOPS	The average number of disk I/O operations per second.
Read Latency	The average amount of time taken per disk I/O operation.
Read Throughput	The average number of bytes read from disk per second.
Result Set Cache Hit Ratio	The percentage of requests that are served by the Resultset cache.
Rollback Segment History List Length	The undo logs that record committed transactions with delete-marked records.
Row Lock Time	The total time spent acquiring row locks for InnoDB tables.
Select Latency	The average amount of time for select operations.
Select Throughput	The average number of select queries per second.
Serverless Database Capacity	The current capacity of an Aurora Serverless v1 DB cluster.
Snapshot Storage Used	The total amount of backup storage consumed by all Aurora snapshots for an Aurora DB cluster outside its backup retention window.
Sum Binary Log Size	The total size of the binlog files.
Swap Usage	The amount of swap space used. This metric is available for the Aurora PostgreSQL instance classes db.t3.medium, db.r5.large, db.r5.xlarge, db.r4.large, and db.r4.xlarge. For Aurora MySQL, this metric applies only to db.t* instances.
Total Backup Storage Billed	The total amount of backup storage in bytes for which you are billed for a given Aurora DB cluster.
Transaction Logs Disk Usage	The amount of disk space consumed by transaction logs on the Aurora PostgreSQL DB instance.
Update Latency	The average amount of taken taken for update operations.
Update Throughput	The average number of updates per second.
Volume Bytes Used	The amount of storage used by your Aurora DB instance.
Volume Read IOPs	The number of billed read I/O operations from a cluster volume within a 5-minute interval.
Volume Write IOPs	The number of write disk I/O operations to the cluster volume, reported at 5-minute intervals.

Write IOPS	The average number of disk I/O operations per second.
Write Latency	The average amount of time taken per disk I/O operation.
Write Throughput	The average number of bytes written to persistent storage every second.

AWS: RDS Instance Configuration	
Object Name	**Object Description**
DB Subnet Group Name	The name of the DB subnet group.
Is Clustered	Boolean value to know if a cluster has members or not.
RDS Instance/RDS Aurora Cluster	The Amazon Resource Name (ARN) for the DB cluster.
VPC ID	For instances associated with a Virtual Private Cloud (VPC), this is the unique identifier of the VPC.
Created Time	The timestamp the instance was created.
Database Class	The size of the database. For example: db.t1.micro, db.m2.2xlarge, db.m5.xlarge, etc.
DB Cluster Parameter Group Status	The status of the DB cluster parameter group for this member of the DB cluster.
DB Default Max Connections	The default maximum number of simultaneous database connections varies by the DB engine type and the memory allocation for the DB instance class.
DB Instance Identifier	The unique identifier for this instance.
DB Parameter Groups	Provides the list of DB parameter groups applied to this DB instance.
DB Security Group Name	Security group associated with RDS instance.
DB Security Group Status	Security group description associated with RDS instance.
DB Subnet Group Description	Provides the description of the DB subnet group.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).

Engine Name	A high-level name for the database engine such as MySQL, or SQLServer.
Is Cluster Writer	Value that is true if the cluster member is the primary instance for the DB cluster and false otherwise.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for Amazon EC2 instances of your account that helps you track the stack level and owner of each instance. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Multi-AZ	A boolean specifying whether or not this instance exists in more than one zone.
Pending Values	A list of pending events such as replication, class change, or group changes.
Primary Zone	The primary availability zone where this instance is addressable.
Promotion Tier	Value that specifies the order in which an Aurora Replica is promoted to the primary instance after a failure of the existing primary instance.
Status	Whether this instance is terminated, starting, stopped, running, or replicating.
Storage Size	The file system size of the local instance used for table spaces.

Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for Amazon EC2 instances of your account that helps you track the stack level and owner of each instance. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
VPC Security Group ID	The unique identifier of the VPC security group.
VPC Security Group Name	The name of the VPC security group.
VPC Security Group Status	The status of the VPC security group.

AWS: RDS Instance Performance	
Object Name	**Object Description**
Allocated Storage	Specifies the allocated storage size specified in gigabytes.
Bin Log Disk Usage	The bytes consumed by the binary log of RDS instance.
CPU Credit Balance	The number of CPU credits available for the instance to burst beyond its base CPU utilization.
CPU Credit Usage	The number of CPU credits consumed by the instance.
CPU Utilization	The utilization percentage across the virtualized CPU cores of database.
Database Connections	The number of current database connections.
Disk Queue Depth	The number of outstanding IOs (read/write requests) waiting to access the disk.
Free Storage Space	The available free storage space on the mounted filesystem of database.
Freeable Memory	The available memory which has yet to be freed by the the processes of database.

Max Database Connections	The maximum number of database connections, this value represents the default max_connections defined in AWS documentation. If Skylar One administrator wants to customize this max_ connections value, just edit the snippet and set the new value for an DB Engine (MAX_MYSQL_ CONNECTIONS, MAX_MARIADB_ CONNECTIONS, MAX_POSTGRESQL_ CONNECTIONS, MAX_ORACLE_ CONNECTIONS or MAX_SQLSERVER_ CONNECTIONS).
Network Receive Throughput	The incoming (Receive) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.
Network Transmit Throughput	The outgoing (Transmit) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.
Read IOPS	The read Input/Output Operations per Second for the database instance.
Read Latency	The average read latency on the database instance.
Read Throughput	The average read throughput on the database instance.
Replica Lag	The current lag time (latency) for replication.
Swap Usage	The usage in bytes of the swap space.
Write IOPS	The write Input/Output Operations per Second for the database instance.
Write Latency	The average latency of write operations on the database instance.
Write Throughput	The average write throughput in Bytes per Second for the database instance.

AWS: RDS Service Health	
Object Name	**Object Description**
Date	The timestamp when the health status was written.
Health	A text description of the health status of this AWS service.

Redshift Service

AWS: Redshift Instance Configuration	
Object Name	**Object Description**
Allow Version Upgrade	Whether or not to allow automatic version upgrades to the database.
Audit Logging Enabled	Whether or not audit logs are kept.
Automated Snapshot Retention Period	The number of days a snapshot is maintained.
AWS Redshift Instance / Security Groups	The security groups associated with this cluster.
AWS Redshift Instance / VPC	The VPC associated with this cluster.
Cluster Id	The unique ID of the cluster.
Cluster Name	The name of the cluster.
Cluster Parameter Group	The parameter group of the cluster.
Cluster Public key	The public key of the cluster in RSA base-64 encoded text format.
Cluster Security Groups	The names of the security groups of the cluster.
Cluster Version	The version ID of the Amazon Redshift engine that is running on the cluster.
Created Time	The creation date of the cluster.
Cross-Region Snapshots Enabled	Whether or not database snapshots are automatically copied to another region.
Current Node Type	The EC2 type used for the database nodes.
Database Name	The name of the database.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Encrypted	Whether or not (yes/no) the database is encrypted.
Endpoint	The DNS name of the cluster.
In Maintenance Mode	A Boolean (yes or no); if maintenance mode is active.
Key	Redshift cluster tag key.

Maintenance Track	The name of the maintenance track for the cluster.
Maintenance Window	The timespan chosen for maintenance.
Master Username	The username of the master (i.e. root) database user.
Nodes	The number of nodes in the cluster.
Number of Nodes	The number of nodes in the cluster.
Port	The TCP port of the cluster.
Publicly Accessible	Indicates if the cluster can be accessed from a public network.
Value	Redshift cluster tag value.

AWS: Redshift Instance Performance	
Object Name	**Object Description**
CPU Utilization	The utilization of the cluster's CPU.
Database Connections	The number of current database connections.
Health Status	A Boolean duty-cycle for cluster health.
Maintenance Mode	A Boolean duty-cycle for the maintenance of a cluster.
Network Receive Throughput	The total network throughput received from the cluster.
Network Transmit Throughput	The total transmitted network throughput for the cluster.
Percentage Disk Space Used	The percentage of disk space used by the cluster.
Read IOPS	The total read IOPs for the cluster.
Read Latency	The average read latency for the cluster.
Read Throughput	The total read throughput for the cluster.
Write IOPS	The total write IOPS for the cluster.
Write Latency	The average write latency for the cluster.
Write Throughput	The average write throughput for the cluster.

AWS: Redshift Node Instance Configuration

Object Name	Object Description
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the node.
Name	The node name.
Private IP Address	The private IP address of the node.
Public IP Address	The public IP address of the node.

AWS: Redshift Node Performance	
Object Name	**Object Description**
Commit Queue Length	Number of transactions in Leader Node where the queueing started.
CPU Utilization	The CPU utilization of the cluster.
Database Connections	The number of current database connections.
Health Status	A Boolean duty-cycle for Leader Node health.
Maintenance Mode	A Boolean duty-cycle for the maintenance of Leader Node.
Network Receive Throughput	The total network throughput received from the cluster.
Network Transmit Throughput	The total transmitted network throughput for the cluster.
Percentage Disk Space Used	The percentage of disk space used by the cluster.
Read IOPS	The total read IOPs for the cluster.
Read Latency	The average read latency for the cluster.
Read Throughput	The total read throughput for the cluster.
Total Table Count	The number of user tables open at a particular point in time.
Write IOPS	The total write IOPS for the cluster.
Write Latency	The average write latency for the cluster.
Write Throughput	The average write throughput for the cluster.

AWS: Redshift Service Health	
Object Name	**Object Description**
Date	A timestamp when the health status of this service was originally written.
Health	A text description of the health status of this AWS service.

Route 53 Service

AWS: Route 53 Health Check Instance Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Failure Threshold	The minimum count of Health Check request failures before marking a route as unhealthy.
Host	The name of the health check instance.
Id	The unique identifier of the health check.
IP Address	The IP address of the health check.
Port	The port of the health check.
Protocol	The protocol of the health check.
Request Interval	The period of time between health check network requests.
Search String	The URL path component query string, if any, for the Health Check request.
URL	The resource referenced by the URL for the check.

AWS: Route 53 Health Check Performance	
Object Name	**Object Description**

Health Check Status	This metric represents the duty cycle of the Health Check endpoint divided by 100. For example, 0 means the endpoint was repsonsive 0% of the time during the monitoring interval, 0.5 = 50% of the time, and 1 = responsive 100% of the time.
Health checkers that report the endpoint healthy	The percentage of Amazon Route 53 health checkers that consider the selected endpoint to be healthy.
Number of healthy child health checks	For a calculated health check, the number of health checks that are healthy among the health checks that Amazon Route 53 is monitoring.
TCP connection time	The average time, in milliseconds, that it took Amazon Route 53 health checkers to establish a TCP connection with the endpoint.
Time to complete SSL handshake	The average time, in milliseconds, that it took Amazon Route 53 health checkers to complete the SSL handshake.
Time to first byte	The average time, in milliseconds, that it took Amazon Route 53 health checkers to receive the first byte of the response to an HTTP or HTTPS request.

AWS: Route 53 Hosted Zone Instance Configuration	
Object Name	**Object Description**
AWS Route 53-Hosted Zone/Application ELB	The unique identifiers of Application ELBs in the Route 53.
AWS Route 53-Hosted Zone/Network ELB	The unique identifiers of Network ELBs in the Route 53.
Comment	Comments, if any, regarding the zone.
Delegation Set	The delegation set of the zone.
Id	The unique identifier of the hosted zone.
Name	The name of the hosted zone.
Record Set Count	The number of zone records.

AWS: Route 53 Service Health

Object Name	Object Description
Date	A timestamp when this service's health status was originally written.
Health	A text description of this AWS Service's health status.

S3 Service

AWS: S3 Cache Configuration	
Object Name	**Object Description**
s3_bucket_cache	Whether or not the S3 cache is collected.

AWS: S3 Instance Configuration	
Object Name	**Object Description**
Bucket Name	The name of the bucket.
Creation Time	The datetime the bucket was created.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the bucket.
Key	Name of the tag.
Logging Bucket	If logging is enabled, the name of the S3 bucket where the logs are stored.
Logging Prefix	If logging is enabled, the value, if any, of the S3 bucket's object prefix.
Logging Status	Whether or not this bucket's activities are logged.
Owner	The owner of the bucket.
Region	The AWS region where this bucket is deployed.
Value	Value of the tag.
Web Hosting	The website index and error document configurations, if any, for this bucket.

AWS: S3 Request Performance	
Object Name	**Object Description**
4xx Errors	The number of HTTP 4xx client error status code requests made to an Amazon S3 bucket with a value of either 0 or 1.
5xx Errors	The number of HTTP 5xx server error status code requests made to an Amazon S3 bucket with a value of either 0 or 1.
All Requests	The total number of HTTP requests made to an Amazon S3 bucket, regardless of type.
Bytes Downloaded	The number bytes downloaded for requests made to an Amazon S3 bucket, where the response includes a body.
Bytes Uploaded	The number bytes uploaded that contain a request body, made to an Amazon S3 bucket.
Delete Requests	The number of HTTP DELETE requests made for objects in an Amazon S3 bucket.
First Byte Latency	The per-request time from the complete request being received by an Amazon S3 bucket to when the response starts to be returned.
Get Requests	The number of HTTP GET requests made for objects in an Amazon S3 bucket.
Head Requests	The number of HTTP HEAD requests made to an Amazon S3 bucket.
List Requests	The number of HTTP requests that list the contents of a bucket.
Post Requests	The number of HTTP POST requests made to an Amazon S3 bucket.
Put Requests	The number of HTTP PUT requests made for objects in an Amazon S3 bucket.
Select Bytes Returned	The number of bytes of data returned with Amazon S3 SELECT Object Content requests in an Amazon S3 bucket.
Select Bytes Scanned	The number of bytes of data scanned with Amazon S3 SELECT Object Content requests in an Amazon S3 bucket.

Select Requests	The number of Amazon S3 SELECT Object Content requests made for objects in an Amazon S3 bucket.
Total Request Latency	The elapsed per-request time from the first byte received to the last byte sent to an Amazon S3 bucket.

AWS: S3 Service Health	
Object Name	**Object Description**
Date	A timestamp when this service's health status was originally written.
Health	A text description of this AWS Service's health status.

AWS: S3 Service Performance	
Object Name	**Object Description**
Total Bucket Size Bytes	The total collected metrics for BucketSizeBytes metric and its Storage Types: Standard, StandardIA, ReducedRedundancy, OneZoneIAStorage, IntelligentTieringFAStorage, GlacierStorage.
Total Number Of Objects	The total collected metrics for NumberOfObjects metric and its Storage Types: AllStorageTypes.

AWS: S3 Storage Performance	
Object Name	**Object Description**
Bucket Size Bytes	The amount of data in bytes stored in a bucket in the Standard storage class, Standard - Infrequent Access (Standard_IA) storage class, or the Reduced Redundancy Storage (RRS) storage class.
Number Of Objects	The total number of objects stored in a bucket for all storage classes.
OI Average Object Size	The average object size in a bucket. If bucket object versioning is enabled then the calculations do include all the object versions.

SES Service

AWS: SES Service Performance	
Object Name	**Object Description**
Bounces	The number of emails that were rejected by email servers of the recipient. This count also includes soft bounces, meaning AWS attempted to deliver the email but was unable to over a period of time.
Clicks	The number of recipients who clicked one or more links in emails.
Complaints	The number of sent emails marked as spam by recipients.
Deliveries	The number of emails Amazon SES successfully delivered to recipients mail servers.
Opens	The number of emails opened by recipients.
Rejects	Calls to Amazon SES that were not successful and Amazon SES will not attempt to deliver the email.
Rendering Failure	The Email was not sent because there was a template rendering error. This can only occur with Template and Bulk Template emails.
Reputation.BounceRate	The rate at which sent emails are bounced.
Reputation.ComplaintRate	The Rate at which sent emails are receiving complaints.
Sends	The number of successful calls to Amazon SES meaning Amazon SES will attempt to deliver the email.

Shield Standard Service

AWS: Shield Advanced Service Performance	
Object Name	**Object Description**

DDoS Attack Bits Per Second	The number of bytes observed during a DDoS event for a particular Amazon Resource Name (ARN). Reporting criteria: Non-zero value during an attack. Zero when there is no attack.
DDoS Attack Packets Per Second	The number of packets observed during a DDoS event for a particular Amazon Resource Name (ARN). Reporting criteria: Non-zero value during an attack. Zero when there is no attack.
DDoS Attack Requests Per Second	The number of requests observed during a DDoS event for a particular Amazon Resource Name (ARN). Reporting criteria: Non-zero value during an attack. Zero when there is no attack.
DDoS Detected	Indicates a DDoS event for a particular Amazon Resource Name (ARN). Reporting criteria: Non-zero value indicates a DDoS event. Zero when there is no DDoS event detected.

AWS: Shield Service Configuration	
Object Name	**Object Description**
Attack Id	The unique identifier (ID) of the attack.
Emergency Contact Emails	Email address that the DRT can use to contact you during a suspected attack.
End Time	The end time of the attack.
Name	The name of the protection.
Protection ID	The unique identifier (ID) of the protection.
Resource ARN	The ARN (Amazon Resource Name) of the resource that was attacked.
Resource ARN	The ARN (Amazon Resource Name) of the AWS resource that is protected.
Start Time	The start time of the attack.
Subscription Status	The status of the subscription.

SNS Service

AWS: SNS Instance Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Subscription ARN	The ARN of the subscription.
Subscription Endpoint	The location of the endpoint for the subscriber. * For email, this would be a valid email address * For email-json, this would be a valid email address * For http, this would be a URL beginning with http * For https, this would be a URL beginning with https * For sqs, this would be the ARN of an SQS Queue * For sms, this would be a phone number of an SMS-enabled device. For application, the endpoint is the EndpointArn of a mobile app and device.
Subscription Owner	The endpoint owner.
Subscription Protocol	The protocol used to communicate with the subscriber. Current choices are: email|email-json|http|https|sqs|sms|application.
Topic Name	The topic name of this instance.

AWS: SNS Instance Performance	
Object Name	**Object Description**
Avg Published Size	The average published size of messages for this topic.
Max Published Size	The maximum published size of a message for this topic.
Min Published Size	The minimum published size of a message for this topic.
Number of Messages Published	A count of the number of published messages for this topic.

Number of Notifications Delivered	A count of the number of delivered notifications for this instance.
Number of Notifications Failed	The number of notifications which failed delivery.
Number Of Notifications Filtered Out	The number of messages that were rejected by subscription filter policies. A filter policy rejects a message when the message attributes do not match the policy attributes.
Number Of Notifications Filtered Out Invalid Attributes	The number of messages that were rejected by subscription filter policies because the messages attributes are invalid. For example, because the attribute JSON is incorrectly formatted.
Number Of Notifications Filtered Out No Message Attributes	The number of messages that were rejected by subscription filter policies because the messages have no attributes.
Published Size Count	The datum count of published size messages for this topic.

AWS: SNS Service Health	
Object Name	**Object Description**
Date	A timestamp when this service's health status was originally written.
Health	A text description of this AWS Service's health status.

AWS: SNS Service Performance	
Object Name	**Object Description**
SMS Month To Date Spent USD	Month to Date Spent on SNS in USD.
SMS Success Rate	The success rate of messages delivered to subscribers per country.
SMS Success Rate Label	Country Label for SMS Success Rate Metrics.

SQS Service

AWS: SQS Instance Configuration	
Object Name	**Object Description**
ARN	The Amazon Resource Name (ARN) (a unique identifier) of this queue.
Created	The timestamp when this queue was created.
Delivery Delay	The built-in delivery delay for messages from the initial time of queueing.
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Last Updated	The timestmap when this queue was last modified.
Max Message Size	The maximum message size, in bytes, that this queue has contained.
Message Delayed	An estimate for the number of delayed messages (incorporating the built-in delay) in the queue.
Messages Available	An estimate for the number of messages still in the queue.
Messages In Flight	An estimate for the number of messages that have been dequeued, but not yet delivered.
Name	The name of the queue.
Retention Period	The maximum time (a kind of TTL) that a message may remain in the queue.
URL	The Uniform Resource Locator (URL) of the queue.
Visibility Timeout	The time period when AWS prevents queue clients from processing messages which has been delivered, but not not explicitly deleted by prior clients.

AWS: SQS Instance Performance	
Object Name	**Object Description**

Approximate Age of Oldest Message	The approximate age of the oldest non-deleted message in the queue.
Approximate Number of Messages Delayed	The number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter.
Approximate Number of Messages Not Visible	The approximate number of messages not visible in this queue.
Approximate Number of Messages Visible	The approximate number of visible messages in this queue.
Avg Sent Message Size	The average size of messages sent through this queue.
Max Sent Message Size	The maximum size of all messages sent through this queue.
Min Sent Message Size	The minimum size of all messages sent through this queue.
Number of Empty Receives	The number of empty messages received by this queue.
Number of Messages Deleted	The number of messages deleted from this queue.
Number of Messages Received	The number of messages received by this queue.
Number of Messages Sent	The number of messages send through this queue.
Sent Message Size Count	The datum counts of all message sizes sent through this queue.

AWS: SQS Service Health	
Object Name	**Object Description**
Date	The timestamp when this health status was written.
Health	A text description of the health status of the AWS Service.

Storage Gateway Service

AWS: Storage Gateway Instance Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Gateway ID	The unique identifier of this gateway containing Region information.
Gateway Name	The name of this gateway.

AWS: Storage Gateway Instance Performance	
Object Name	**Object Description**
Cloud Bytes Downloaded	The average number of bytes downloaded from the cloud to this gateway.
Cloud Bytes Uploaded	The average number of bytes uploaded to the cloud from this gateway.
Cloud Download Latency	The average latency when downloading from the cloud to this gateway.
Queued Writes	The average number of queued write operations on this gateway.
Read Bytes	The average number of bytes read by this gateway.
Read Time	The average read times for this gateway.
Working Storage Free	The average free/available working storage on this gateway.
Working Storage Percent Used	The percentage of the local Working Storage buffer currently being used.
Working Storage Used	The average working storage used for this gateway.
Write Bytes	The average number of bytes written by this gateway.
Write Time	The average write times for this storage gateway.

AWS: Storage Gateway Service Health	
Object Name	**Object Description**
Date	The timestamp when this health status update was written.
Health	A text description of the health status of the AWS Service.

AWS: Storage Gateway Volume Configuration	
Object Name	**Object Description**
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Volume ID	The unique identifier for this volume.
Volume Name	The name of this filesystem volume.

AWS: Storage Gateway Volume Performance	
Object Name	**Object Description**
Queued Writes	The average number of queued (pending) write operations on this gateway volume.
Read Bytes	The average bytes read from this gateway volume.
Read Time	The average time for read operations for this gateway volume.
Write Bytes	The average bytes written to this gateway volume.
Write Time	The average time for write operations for this gateway volume.

STS Service

AWS: STS Session Manager	
Object Name	**Object Description**

Account ID	The unique identifier (ID) of the account.
Organization Namespace	Namespace defined to create relationship between organization and its accounts.
Response Status	Response Status as value starting with Succes or Error. Once you run the collect method of the sts_ session_manager class it will change its value.

SWF Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: SWF General Limit Instance Performance	
Object Name	**Object Description**
Consumed Limit - Sum	The amount of general limit that has been consumed.
Provisioned Limit - Sum	The amount of general limit that is provisioned to the account.

Transfer Server Service

This Dynamic Application is available in the "AWS Extended Services" PowerPack.

AWS: Transfer Server Performance	
Object Name	**Object Description**
Bytes In - Sum	The total number of bytes transferred into the server.
Bytes Out - Sum	The total number of bytes transferred out of the server.
Files In - Sum	The total number of files transferred into the server.
Files Out - Sum	The total number of files transferred out of the server.
On Upload Executions Failed - Sum	The total number of unsuccessful workflow executions on the server.

On Upload Executions Started - Sum	The total number of workflow executions started on the server.
On Upload Executions Success - Sum	The total number of successful workflow executions on the server.

Transit Gateway Service

AWS: Transit Gateway Instance Configuration	
Object Name	**Object Description**
Attachment Id	The ID of the attachment.
Attachment Id	The ID of the attachment.
Attachment Name	The name of the attachment.
Attachment Name	The name of the attachment.
ID	The ID of transit gateway.
Name	The name of transit gateway.
Transit Gateway Route Table Id	The ID of the transit gateway route table.
Transit Gateway/VPC	The ID of the resource.
ARN	The ARN of transit gateway.
Default Association Route Table	Indicates whether this is the default association route table for the transit gateway (true | false).
Default Propagation Route Table	Indicates whether this is the default propagation route table for the transit gateway (true | false).
Owner Id	The ID of the AWS account that owns the transit gateway.
Resource Id	The ID of the resource.
Resource Type	The resource type. Valid values are vpc | vpn | direct-connect-gateway | peering | connect.
Resource Type	The resource type. Valid values are vpc | vpn | direct-connect-gateway | peering | connect.
State	The state of transit gateway (available | deleted | deleting | modifying | pending).

State	The state of the attachment. Valid values are available | deleted | deleting | failed | failing | initiatingRequest | modifying | pendingAcceptance | pending | rollingBack | rejected | rejecting.
State	The state of the route table (available | deleting | deleted | pending).
State	The state of the attachment. Valid values are available | deleted | deleting | failed | failing | initiatingRequest | modifying | pendingAcceptance | pending | rollingBack | rejected | rejecting.
Creation Time	The creation time of transit gateway.
Creation Time	The creation time of the attachment.
Creation Time	The creation time of the route table.
Creation Time	The creation time of the attachment.
Key	The tag key (String).
Value	The value of the tag key.

AWS: Transit Gateway Instance Performance	
Object Name	**Object Description**
Bytes Drop Count Blackhole	The number of bytes dropped because they matched a blackhole route.
Bytes Drop Count No Route	The number of bytes dropped because they did not match a route.
Bytes In	The number of bytes received by the transit gateway.
Bytes Out	The number of bytes sent from the transit gateway.
Packet Drop Count Blackhole	The number of packets dropped because they matched a blackhole route.
Packet Drop Count No Route	The number of packets dropped because they did not match a route.
Packets In	The number of packets received by the transit gateway.
Packets Out	The number of packets sent by the transit gateway.

VPC Service

AWS: VPC Instance Configuration	
Object Name	**Object Description**
VPC Peer ID	The ID of the VPC peering connection.
VPC Peer Name	The name of the VPC peering connection.
Accepter VPC ID	The ID of the accepter VPC.
Status	The status of the VPC peering connection.
AWS VPC/EC2	The unique identifier for the VPC-EC2 relationship.
AWS VPC/Security Group	The unique identifier for the VPC-Security Group relationship.
AWS VPC/VPC Inter-Account	The relationship between Inter-Account VPCs. Contains the relationship namespace value. A relationship between accounts is drawn when the relationship namespace has the same value as the identity namespace and their respective group indices also have the same value.
AWS VPC/VPC Intra-Account	The relationship between Intra-Account VPCs.
DHCP Options Set	The resource ID of the DHCP option set for the VPC. DHCP options are associated with your AWS account so that you can use them across all of your virtual private clouds (VPC).
Distinguished Name	The internal Skylar One distinguished name of the AWS component. This follows closely the format of the AWS Amazon Resource Name (ARN).
Id	The unique identifier of the VPC.
Inter-Account Requester VPC ID	The Inter-Account Requester VPC ID. Value populated only when the current VPC is the Requester VPC.
Inter-Account Requester VPC ID	The Inter-Account Requester VPC ID.

Is Default VPC	Is this a default VPC. True | False. A default VPC combines the benefits of the advanced features provided by EC2-VPC with the ease of use of EC2-Classic. If you have a default VPC and do not specify a subnet when you launch an instance, the instance is launched into your default VPC. You can launch instances into your default VPC without needing to know anything about Amazon VPC.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Name	The Name property, if any, of the VPC as set in its tags.
Requester VPC ID	The Requester VPC ID.
State	Whether the VPC is available or not.
Tenancy	One of default | dedicated. If default, your instance runs on shared hardware. If dedicated, your instance runs on single-tenant hardware.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
VPC CIDR	The Classless Inter-Domain Routing(CIDR) block for the VPC.
vpc_identity_namespace	The identity namespace used to identify Inter-Account relationships. Value populated only when the current VPC is the Requester VPC.

AWS: VPC NAT Gateway Instance Configuration	
Object Name	**Object Description**
Allocated ID	The allocation ID of the Elastic IP address that is associated with the NAT gateway.
AWS NAT Gateway/Subnet	The unique identifier for the VPC NAT Geteway-Subnet relationship.
Create Time	The date and time the NAT gateway was created.
Key	The key of a tag assigned to the resource.
NAT Gateway ID	The ID of the NAT gateway.
Network Interface ID	The ID of the network interface associated with the NAT gateway.
Private Ip	The private IP address associated with the Elastic IP address.
Public Ip	The Elastic IP address associated with the NAT gateway.
State	The state of the NAT gateway. Possible values are: pending|failed|failureMessage|available|deleting|deleted.
Value	The value of a tag assigned to the resource.
VPC ID	The ID of the VPC in which the NAT gateway is located.

AWS: VPC NAT Gateway Instance Performance	
Object Name	**Object Description**
Active Connection Count	The total number of concurrent active TCP connections through the NAT gateway. A value of zero indicates that there are no active connections through the NAT gateway.
Bytes In From Destination	The number of bytes received by the NAT gateway from the destination. If the value for BytesOutToSource is less than the value for BytesInFromDestination, there may be data loss during NAT gateway processing, or traffic being actively blocked by the NAT gateway.

Bytes In From Source	The number of bytes received by the NAT gateway from clients in your VPC. If the value for BytesOutToDestination is less than the value for BytesInFromSource, there may be data loss during NAT gateway processing.
Bytes Out To Destination	The number of bytes sent out through the NAT gateway to the destination. A value greater than zero indicates that there is traffic going to the internet from clients that are behind the NAT gateway. If the value for BytesOutToDestination is less than the value for BytesInFromSource, there may be data loss during NAT gateway processing.
Bytes Out To Source	The number of bytes sent through the NAT gateway to the clients in your VPC. A value greater than zero indicates that there is traffic coming from the internet to clients that are behind the NAT gateway. If the value for BytesOutToSource is less than the value for BytesInFromDestination, there may be data loss during NAT gateway processing, or traffic being actively blocked by the NAT gateway.
Connection Attempt Count	The number of connection attempts made through the NAT gateway. If the value for ConnectionEstablishedCount is less than the value for ConnectionAttemptCount, this indicates that clients behind the NAT gateway attempted to establish new connections for which there was no response.
Connection Established Count	The number of connections established through the NAT gateway. If the value for ConnectionEstablishedCount is less than the value for ConnectionAttemptCount, this indicates that clients behind the NAT gateway attempted to establish new connections for which there was no response.
Error Port Allocation	The number of times the NAT gateway could not allocate a source port. A value greater than zero indicates that too many concurrent connections are open through the NAT gateway.

Idle Timeout Count	The number of connections that transitioned from the active state to the idle state. An active connection transitions to idle if it was not closed gracefully and there was no activity for the last 350 seconds. A value greater than zero indicates that there are connections that have been moved to an idle state. If the value for IdleTimeoutCount increases, it may indicate that clients behind the NAT gateway are re-using stale connections.
Packets Drop Count	The number of packets dropped by the NAT gateway. A value greater than zero may indicate an ongoing transient issue with the NAT gateway.
Packets In From Destination	The number of packets received by the NAT gateway from the destination. If the value for PacketsOutToSource is less than the value for PacketsInFromDestination, there may be data loss during NAT gateway processing, or traffic being actively blocked by the NAT gateway.
Packets In From Source	The number of packets received by the NAT gateway from clients in your VPC. If the value for PacketsOutToDestination is less than the value for PacketsInFromSource, there may be data loss during NAT gateway processing.
Packets Out To Destination	The number of packets sent out through the NAT gateway to the destination. A value greater than zero indicates that there is traffic going to the internet from clients that are behind the NAT gateway. If the value for PacketsOutToDestination is less than the value for PacketsInFromSource, there may be data loss during NAT gateway processing.
Packets Out To Source	The number of packets sent through the NAT gateway to the clients in your VPC. A value greater than zero indicates that there is traffic coming from the internet to clients that are behind the NAT gateway. If the value for PacketsOutToSource is less than the value for PacketsInFromDestination, there may be data loss during NAT gateway processing, or traffic being actively blocked by the NAT gateway.

AWS: VPC Route Table Configuration

Object Name	Object Description
AWS Route Table/Subnet	The unique identifier for the Route Table-Subnet relationship.
AWS Route Table/Virtual Private Gateway	The unique identifier for the Route Table-Virtual Private Gateway relationship.
Destination	The Destination in a route.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Main	Whether the Route Table is the default (main) or otherwise.
Route Table ID	The unique identifier of the VPC Route Table.
Status	Whether a Route in the VPC Route Table is active or not.
Target	The Gateway of a Route defined in a VPC Route Table.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
VPC ID	The unique identifier of the VPC.

AWS: VPC Service Health

Object Name	Object Description
Date	The timestamp of this service's health update.
Health	A text description of this AWS Service's health status.

AWS: VPC Subnet Configuration	
Object Name	**Object Description**
Available IP Address Count	The number of IPv4 addresses in the subnet that are available.
AWS VPC Subnet/EC2	The unique identifier for the VPC Subnet-EC2 relationship.
Id	The unique identifier of the subnet.
Key	The tag key. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.
Name	The Name property, if any, of the subnet as set in its tags.
Subnet CIDR	The IPv4 CIDR block of the subnet.
Value	The tag value. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. For example, you could define a set of tags for your Amazon EC2 instances account that helps you track each owner of the instance and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.

AWS: VPC Virtual Private Gateway Configuration	
Object Name	**Object Description**
Key	The key of a tag assigned to the resource.
State	The current state of the attachment between the gateway and the VPC (attaching | attached | detaching | detached).
Value	The value of a tag assigned to the resource.
VPC ID	The unique identifier of an attached VPC.
VPN Gateway ID	The unique identifier of the virtual private gateway.
VPN Gateway Name	The name of the virtual private gateway.
VPN Gateway State	The state of the virtual private gateway (pending | available | deleting | deleted).
VPN Gateway Type	The type of virtual private gateway. Currently the only supported type is ipsec.1.

WAF Service

AWS: WAF Global WebACL Instance Configuration	
Object Name	**Object Description**
Action	The action that CloudFront or AWS WAF takes when a web request matches the conditions in the Rule.
Default Action	The WebACL instance component default action.
Id	The Id for a Rule.
Name	The WebACL instance component name.
Name	The name for a Rule.
Priority	The order in which the Rules in a WebACL are evaluated.
Type	Type of AWS WAF response to requests that match the settings in a Rule.
Web ACL ARN	The ARN for the WebACL component.
Web ACL Id	The WebACL instance component Unique Identifier.

AWS: WAF Global WebACL Instance Performance	
Object Name	**Object Description**
Allowed (By Rule)	The number of allowed web requests by rule.
Blocked (By Rule Group)	The number of blocked web requests by rule group.
Blocked (By Rule)	The number of blocked web requests by rule.
Counted (By Rule Group)	The number of counted web requests by rule group.
Counted (By Rule)	The number of counted web requests by rule.

AWS: WAF Regional WebACL Instance Configuration	
Object Name	**Object Description**
Action	Specifies the action that API Gateway or Application ELB or AWS WAF takes when a web request matches the conditions in the Rule.
AWS WAF/API Gateway	The API IDs of resources associated with the specified web ACL.
AWS WAF/API Stage	The API Stage IDs of resources associated with the specified web ACL.
AWS WAF/Application ELB	The Application ELB IDs of resources associated with the specified web ACL.
Default Action	The action to perform if none of the Rules contained in the WebACL match.
Metric Name	A friendly name or description for the metrics for this WebACL.
Name	The name for a Rule.
Priority	Specifies the order in which the Rules in a WebACL are evaluated.
Rule Id	The RuleId for a Rule.
Type	Specifies how you want AWS WAF to respond to requests that match the settings in a Rule.
WebACL ARN	The ARN for the WebACL component.
WebACL Id	The WebACL instance component Unique Identifier.
WebACL Name	A friendly name or description of the WebACL.

AWS: WAF Regional WebACL Instance Performance	
Object Name	**Object Description**
Allowed (By Rule)	The number of allowed web requests by rule.
Blocked (By Rule Group)	The number of blocked web requests by rule group.
Blocked (By Rule)	The number of blocked web requests by rule.
Counted (By Rule Group)	The number of counted web requests by rule group.
Counted (By Rule)	The number of counted web requests by rule.

AWS: WAF Service Health	
Object Name	**Object Description**
Date	The date of the RSS feed entry.
Health	The message of the RSS feed entry.

Workspaces Service

AWS: Workspaces Directory Configuration	
Object Name	**Object Description**
Alias	The directory alias.
Change Compute Type	Specifies whether users can change the compute type (bundle) for their WorkSpace.
Customer User Name	The user name for the service account.
Device Type Android	Indicates whether users can use Android and Android-compatible Chrome OS devices to access their WorkSpaces.
Device Type Chrome OS	Indicates whether users can use Chromebooks to access their WorkSpaces.
Device Type IOS	Indicates whether users can use iOS devices to access their WorkSpaces.
Device Type Linux	Indicates whether users can use Linux clients to access their WorkSpaces.

Device Type OSX	Indicates whether users can use OSX clients to access their WorkSpaces.
Device Type Web	Indicates whether users can access their WorkSpaces through a web browser.
Device Type Windows	Indicates whether users can use Windows clients to access their WorkSpaces.
Device Type Zero Client	Indicates whether users can use zero client devices to access their WorkSpaces.
Directory Id	The unique identifier of the directory
Directory Name	The directory name or directory id if name is not present, prefixed by the region where it is located.
Directory Type	The directory type.
DNS IP Addresses	The IP addresses of the DNS servers for the directory.
Enable Internet Access	Specifies whether to automatically assign an Elastic public IP address to WorkSpaces in this directory by default.
Enable Maintenance Mode	Specifies whether maintenance mode is enabled for WorkSpaces.
Enable Work Docs	Specifies whether the directory is enabled for Amazon WorkDocs.
IAM Role Id	The identifier of the IAM role. This is the role that allows Amazon WorkSpaces to make calls to other services, such as Amazon EC2, on your behalf.
Increase Volume Size	Specifies whether users can increase the volume size of the drives on their WorkSpace.
Rebuild Workspace	Specifies whether users can rebuild the operating system of a WorkSpace to its original state.
Registration Code	The registration code for the directory. This is the code that users enter in their Amazon WorkSpaces client application to connect to the directory.
Restart Workspace	Specifies whether users can restart their WorkSpace.
State	The state of the directory registration with Amazon WorkSpaces.
Subnet Ids	The identifiers of the subnets used with the directory.

Switch Running Mode	Specifies whether users can switch the running mode of their WorkSpace.
Tenancy	Specifies whether the directory is dedicated or shared.
User Enabled As Local Administrator	Specifies whether WorkSpace users are local administrators on their WorkSpaces.

AWS: Workspaces Directory Performance	
Object Name	**Object Description**
Available	The number of WorkSpaces that returned a healthy status.
Connection Attempt	The number of connection attempts.
Connection Failure	The number of failed connections.
Connection Success	The number of successful connections.
In Session Latency	The round trip time between the WorkSpaces client and the WorkSpace.
Maintenance	The number of WorkSpaces that are under maintenance.
Session Disconnect	The number of connections that were closed, including user-initiated and failed connections.
Session Launch Time	The amount of time it takes to initiate a WorkSpaces session.
Stopped	The number of WorkSpaces that are stopped.
Trusted Device Validation Attempt	The number of device authentication signature validation attempts.
Trusted Device Validation Failure	The number of failed device authentication signature validations.
Trusted Device Validation Success	The number of successful device authentication signature validations.
Unhealthy	The number of WorkSpaces that returned an unhealthy status.
User Connected	The number of WorkSpaces that have a user connected.

WorkMail Service

This Dynamic Application is available in the ***AWS Extended Services***PowerPack.

AWS: WorkMail Organization Performance	
Object Name	**Object Description**
Incoming Email Bounced - Avg	The number of incoming emails that bounced due to full mailboxes. This metric is counted for each intended recipient. For example, if 1 email is sent to 10 recipients in your organization, and 2 of the recipients have full mailboxes resulting in a bounce response, the IncomingEmailBounced count is 2.
Mailbox Email Delivered - Avg	The number of emails delivered to individual mailboxes in your Amazon WorkMail organization. If 1 email is successfully delivered to 10 recipients in your organization, the MailboxEmailDelivered count is 10.
Organization Email Received - Avg	The number of emails received by your Amazon WorkMail organization. If 1 email is addressed to 10 recipients in your organization, the OrganizationEmailReceived count is 1.
Outgoing Email Bounced - Avg	The number of outgoing emails that could not be delivered, counted for each intended recipient. For example, if 1 email is sent to 10 recipients, and 2 emails could not be delivered, the OutgoingEmailBounced count is 2.
Outgoing Email Sent - Avg	The number of emails successfully sent from your Amazon WorkMail organization. This metric is counted for each recipient of a successfully sent email. For example, if 1 email is sent to 10 recipients, and the email was successfully delivered to 8 of the recipients, the OutgoingEmailSent count is 8 .

Other Supported Services

The following services have partial coverage with configuration information, health status, and AWS alerts:

AWS Service
Amplify
AppFlow
Backup
Chime
Chime SDK
Client VPN
CloudWatch Logs
CloudWatch Metric Streams
CloudWatch RUM
CodeGuru Profiler
CodeGuru Reviewer
Data Exchange
Data Pipeline
DataSync
DevOps Guru
DLM (Data Life Cycle Manager)
EC2 Fleet
EC2 Spot
EC2 Spot Fleet
Elastic Inference
EventBridge
FinSpace
Forecast
Fraud Detector
GameLift
Gateway Load Balancer (ELB)
Glue DataBrew
Ground Station

GuardDuty
HealthLake
Inspector
Interactive Video Service Channel
Interactive Video Service Chat
IoT Analytics
IoT Device Defender
IoT Events
IoT FleetWise
IoT Greengrass
IoT RoboRunner
IoT Sitewise
IoT Twinmaker
IoT Things Graph
Kendra
Keyspaces
Lake Formation
Location Service
Lookout for Equipment
Lookout for Metrics
Lookout for Vision
MWAA (Managed Workflows for Apache Airflow)
Network Firewall
Network Manager
Panorama
Personalize
Polly
Private Certificate Authority
PrivateLink
QLDB (Quantum Ledger DB)

QuickSight
Rekognition
RoboMaker
SageMaker
Secrets Manager
Service Catalog
Step Functions
Systems Manager
Textract
Timestream
Transcribe
Translate
WorkSpaces Web

© 2003 - 2026, ScienceLogic, Inc.

All rights reserved.

ScienceLogic™, the ScienceLogic logo, and ScienceLogic's product and service names are trademarks or service marks of ScienceLogic, Inc. and its affiliates. Use of ScienceLogic's trademarks or service marks without permission is prohibited.

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information herein, the information provided in this document may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described herein at any time without notice.

800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010