



AWS Incident Detection and Response Concepts and Procedures

AWS Incident Detection and Response User Guide



Version June 12, 2024

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Incident Detection and Response User Guide: AWS Incident Detection and Response Concepts and Procedures

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is AWS Incident Detection and Response?	1
Product terms	2
Availability	2
RACI	3
Architecture	6
Get started with Incident Detection and Response	8
Onboard a workload	8
Workload onboarding	9
Alarm ingestion	9
Account subscription	10
Workload discovery	13
Alarm configuration	13
Create CloudWatch alarms that fit your business	15
Use AWS CloudFormation templates to build CloudWatch alarms	18
Example uses cases for CloudWatch alarms	21
Ingest alerts into AWS Incident Detection and Response	26
Provision access	27
Integrate with CloudWatch	27
Ingest alarms from APMs with EventBridge integration	28
Example: Integrating notifications from Datadog and Splunk	29
Ingest alarms from APMs without direct integration with Amazon EventBridge	40
Develop runbooks	41
Test onboarded workloads	47
CloudWatch alarms	48
Third party APM alarms	49
Key outputs	49
Workload onboarding and alarm ingestion questionnaires	49
Workload onboarding questionnaire - General questions	49
Workload onboarding questionnaire - Architecture questions	50
Workload onboarding questionnaire - AWS Service Event questions	52
Alarm Ingestion Questionnaire	53
Alarm matrix	54
Request changes to a workload	58
Offboard a workload	61

Monitoring and observability	63
Implementing observability	64
Incident management	65
Provision access for application teams	68
Incident management for service events	68
Incident Response Request	70
Reporting	75
Security and resiliency	76
Access to your accounts	77
Your alarm data	77
Document history	78
AWS Glossary	82

What is AWS Incident Detection and Response?

AWS Incident Detection and Response offers eligible AWS Enterprise Support customers proactive incident engagement to reduce the potential for failure and accelerates recovery of critical workloads from disruption. Incident Detection and Response facilitates your collaboration with AWS to develop runbooks and response plans customized to each onboarded workload. A team of Incident Management Engineers (IMEs) monitor your onboarded workloads 24x7 and engages you on a call bridge within 5 minutes of a critical alarm.

Incident Detection and Response offers the following key features:

- **Improved observability:** AWS experts provide guidance to help you define and correlate metrics and alarms between the application and infrastructure layers of your workload to detect disruptions early.
- **5-minute response time:** IMEs monitor your onboarded workloads 24x7 to detect critical incidents. The IMEs respond within 5-minutes of an alarm trigger or in response to a business-critical Support case that you raise to Incident Detection and Response.
- **Faster resolution:** IMEs use pre-defined and custom runbooks developed for your workloads to respond within 5-minutes, create a Support case on your behalf, and manage incidents on your workload. IMEs provide single-threaded ownership for incidents and keep you engaged with the right AWS experts until the incident is resolved.
- **Incident management for AWS events:** Because we understand the context of your critical workload (for example, accounts, services and instances), we can detect and proactively notify you of a potential impact to your workload during an AWS service event. If requested, IMEs engage you during the AWS service events and provide updates on the events. While Incident Detection and Response cannot prioritize you for recovery during a service event, Incident Detection and Response does provide Support guidance to help you implement your mitigation plan.
- **Reduced potential for failure:** After resolution, the IMEs provide you with a post-incident review (upon request). And, AWS experts work with you to apply lessons learned to improve the incident response plan and runbooks. You can also leverage AWS Resilience Hub for continuous resiliency tracking on your workloads.

Incident Detection and Response product terms

- AWS Incident Detection and Response is available to direct and Partner-resold Enterprise Support accounts.
- AWS Incident Detection and Response is not available to accounts on Partner Led Support.
- You must maintain AWS Enterprise Support at all times during the term of your Incident Detection and Response service. For information, see [Enterprise Support](#). Termination of Enterprise Support results in concurrent removal from the AWS Incident Detection and Response service.
- All workloads on AWS Incident Detection and Response must go through the workload onboarding process.
- The minimum duration to subscribe an account to AWS Incident Detection and Response is ninety (90) days. All cancellation request must be submitted thirty (30) days prior to the intended effective date of cancellation.
- AWS handles your information as described in the [AWS Privacy Notice](#).

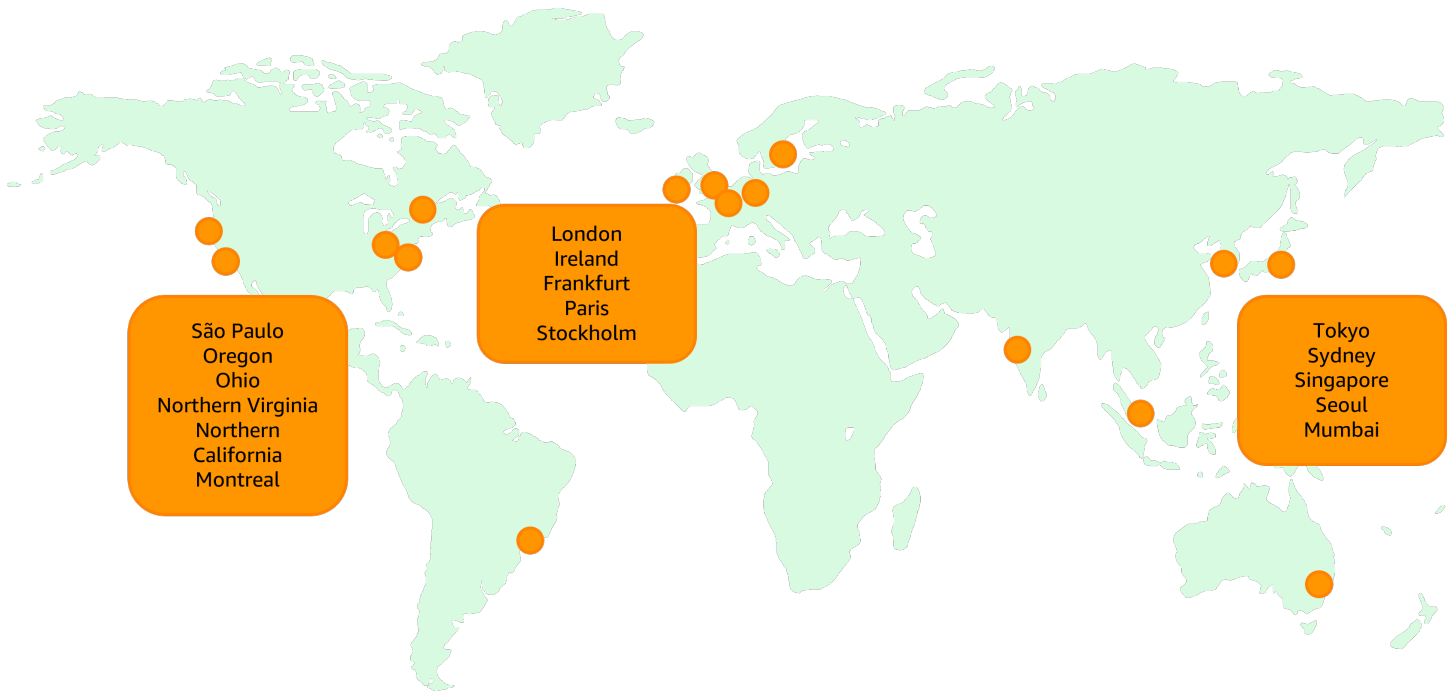
Note

For Incident Detection and Response billing related questions, see [Getting help with AWS Billing](#).

Incident Detection and Response availability

AWS Incident Detection and Response is currently available in the English language for Enterprise Support accounts hosted in any of the following regions:

US East (N. Virginia), US East (Ohio), US West (Oregon), US West (N. California), Canada (excluding Quebec), EU (Frankfurt), EU (Ireland), EU (London), EU (Paris), EU (Stockholm), Asia Pacific (Mumbai), Asia Pacific (Tokyo), Asia Pacific (Singapore), Asia Pacific (Seoul), Asia Pacific (Sydney), South America (Sao Paulo). Shown here:



AWS Incident Detection and Response RACI

The following table shows the AWS Incident Detection and Response responsible, accountable, consulted, and informed or RACI.

Activity	Customer	Incident Detection and Response
Data collection		
Customer and workload introduction	C	R
Architecture	R	A
Operations	R	A
Determine CloudWatch alarms to be configured	R	A
Define incident response plan	R	A
Completing on-boarding questionnaire	R	A


Activity	Customer	Incident Detection and Response
Operations readiness review		
Conduct well architected review (WAR) on workload	C	R
Validate incident response	C	R
Validate alarm matrix	C	R
Identify key AWS services being used by the workload	A	R
Account configuration		
Create IAM role in customer account	R	I
Install managed EventBridge rule using created role	I	R
Test CloudWatch alarms	R	A
Verify that customer alarms engage the incident detection and response	I	R
Update alarms	R	C
Update runbooks	C	R
Incident management		
Proactively notify Incidents detected by Incident Detection and Response	I	R
Provide incident response	I	R
Provide Incident resolution / infrastructure restore	R	C
Post incident review		

Activity	Customer	Incident Detection and Response
Request post incident review	R	I
Provide post incident review	I	R

AWS Incident Detection and Response architecture

AWS Incident Detection and Response integrates with your existing environment as shown in the following graphic. The architecture includes the following services:


- **Amazon EventBridge:** Amazon EventBridge serves as the sole integration point between your workloads and AWS Incident Detection and Response. Alarms are ingested from your monitoring tools, such as Amazon CloudWatch, through Amazon EventBridge using predefined rules managed by AWS. To allow Incident Detection and Response to build and manage the EventBridge rule, you install a service-linked role. To learn more about these services, see [What is Amazon EventBridge](#) and [Amazon EventBridge rules](#), [What is Amazon CloudWatch](#), and [Using service-linked roles for AWS Health](#).
- **AWS Health:** AWS Health provides ongoing visibility into your resource performance and the availability of your AWS services and accounts. Incident Detection and Response uses AWS Health to track events on the AWS services used by your workloads and to notify you when an alert has been received from your workload. To learn more about AWS Health, see [What is AWS Health](#).
- **AWS Systems Manager:** Systems Manager provides a unified user interface for automation and task management across your AWS resources. AWS Incident Detection and Response hosts information about your workloads including workload architecture diagrams, alarm details and their corresponding incident management runbooks in AWS Systems Manager documents (for details, see [AWS Systems Manager Documents](#)). To learn more about AWS Systems Manager, see [What is AWS Systems Manager](#).
- **Your specific runbooks:** An incident management runbook defines the actions that AWS Incident Detection and Response performs during incident management. Your specific runbooks tell AWS Incident Detection and Response who to contact, how to contact them, and what information to share.



Get started with AWS Incident Detection and Response

You can select specific workloads for monitoring and critical incident management using AWS Incident Detection and Response. A workload is a collection of resources and code that work together to deliver a business value. A workload might be all the resources and code that make up your banking payment portal or a customer relationship management (CRM) system. You can host a workload in a single AWS account or multiple AWS accounts.

For example, you might have a monolithic application hosted in a single account (for example, Employee Performance App in Fig.1). Or, you might have an application (for example, Storefront Webapp in Fig. 1) broken into microservices that stretch across different accounts. A workload might share resources, such as a database, with other applications or workloads as shown in Fig. 1.



Note

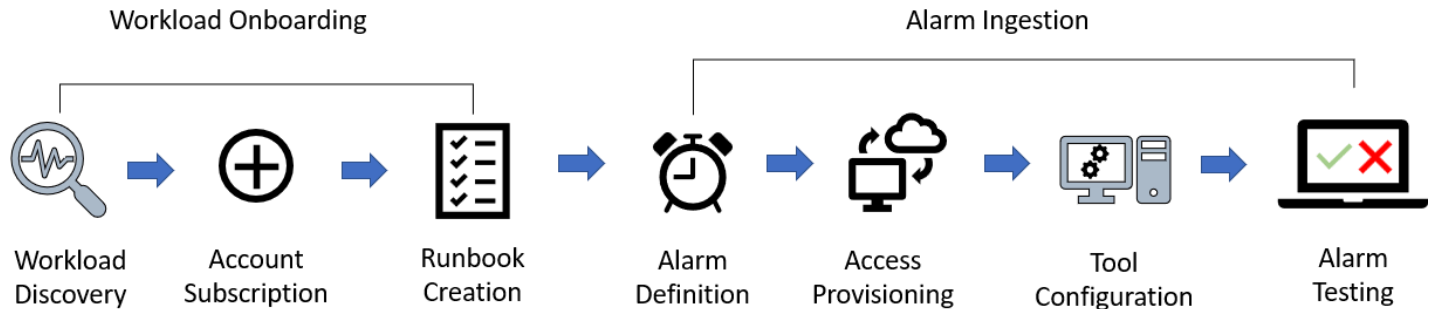
To make changes to your runbooks, workload information, or the alarms monitored on AWS Incident Detection and Response, create an [Request changes to an onboarded workload](#).

Onboarding

AWS works with you to onboard your workload and alarms to AWS Incident Detection and Response. You provide key information to AWS in the [Workload onboarding and alarm ingestion](#)

[questionnaires](#). It's a best practice that you also register your workloads on AppRegistry. For more information, see the [AppRegistry User Guide](#).

The following diagram shows the flow for workload onboarding and alarm ingestion in Incident Detection and Response:



Workload onboarding

During workload onboarding, AWS works with you to understand your workload and how to support you during incidents and AWS Service Events. You provide key information about your workload that assists with impact mitigation.

Key outputs:

- General workload information
- Architecture details including diagrams
- Runbook Information
- Customer-initiated incidents
- AWS Service Events

Alarm ingestion

AWS works with you to onboard your alarms. AWS Incident Detection and Response can ingest alarms from Amazon CloudWatch and Third-Party Application Performance Monitoring (APM) tools through Amazon EventBridge. Onboarding alarms allows for proactive incident detection and automated engagement. For more information, see [Ingest alarms from APMs that have direct integration with Amazon EventBridge](#).

Key outputs:

- Alarm matrix

The following table lists the steps required to onboard a workload to AWS Incident Detection and Response. This table shows example durations of each task. The actual dates for each task are defined based on the availability of your team and schedule.

Phase	Task	Customer	AWS	Duration	Meeting Suggested?	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Workload Onboarding	Complete questionnaire	R	C	2		■	■												
	Workload Discovery	R	C	1				■	■										
	Subscription	R	C	1	Yes					■									
	Runbook Creation	C	R	1						■									
Alarm Ingestion	Complete questionnaire	R	C	2							■	■							
	Alarm Definition	R	C	3									■	■	■				
	Access Provisioning	R	C	1												■			
	Tool Configuration	C	R	2													■	■	
	Alarm Testing	R	C	1															■

R - Responsible, C-Consulted

Account subscription

To subscribe a workload to AWS Incident Detection and Response, create a new support case for each workload. When you create the support case, keep the following in mind:

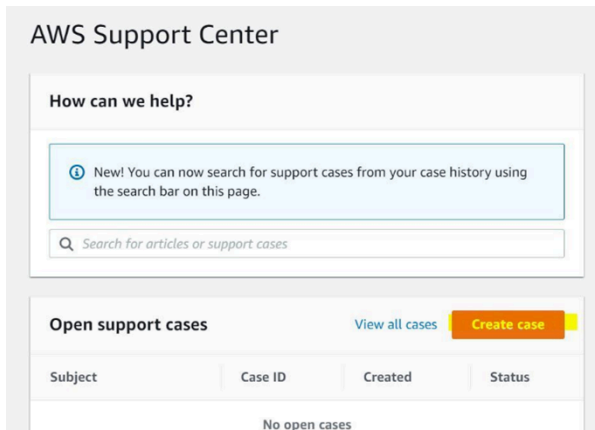
- To onboard a workload that's in a single AWS account, create the support case either from the workload's account or from your payer account.
- To onboard a workload that spans multiple AWS accounts, create the support case from your **payer account**. In the body of the support case, list all account IDs to onboard.

Important

If you create a support case to subscribe a workload to Incident Detection and Response from the incorrect account, then you might experience delays and requests for additional information before your workloads can be subscribed.

To subscribe a workload

- Go to the [AWS Support Center](#), and then select **Create case** as shown in the following example. You can only subscribe workloads from accounts that are enrolled in Enterprise Support.



2. Complete the support case form:

- Select **Technical support**.
- For **Service**, choose **Incident Detection and Response**.
- For **Category**, choose **Onboard New Workload**.
- For **Severity**, choose **General guidance**.

3. Enter a **Subject** for this change. For example:

[Onboard] AWS Incident Detection and Response - *workload_name*

4. Enter a **Description** for this change. For example, enter "This request is to onboard a workload to AWS Incident Detection and Response". Make sure that you include the following information in your request:

- **Workload name:** Your workload name.
- **Account ID(s):** ID1, ID2, ID3, and so on. These are the accounts that you want to onboard to AWS Incident Detection and Response.
- **Subscription start date:** The date that you want to start the AWS Incident Detection and Response subscription.

5. In the **Additional contacts - optional** section, enter any email IDs that you want to receive correspondence about this request.

The following is an example of the **Additional contacts - optional** section:

⚠ Important

Failure to add email IDs in the **Additional contacts - optional** section might delay the AWS Incident Detection and Response onboarding process.

6. Choose Submit.

After you submit the request, you can add additional emails from your organization. To add emails, reply to the case, and then add the email IDs in the **Additional contacts - optional** section.

The following is an example of the **Additional contacts - optional** section:

After you create a support case for the subscription request, keep the following two documents ready to proceed with the workload onboarding process:

- AWS workload architecture diagram.
- [Workload onboarding and alarm ingestion questionnaires](#): Complete all of the information in the questionnaire that's related to the workload that you're onboarding. If you have multiple workloads to be onboarded, then create a new onboarding questionnaire for each workload. If you have questions about completing the onboarding questionnaire, then contact your Technical Account Manager (TAM).

📘 Note

DO NOT attach these two documents to the case using the **Attach files** option. AWS Incident Detection and Response team will reply to the case with an Amazon Simple Storage Service Uploader link for you to upload the documents.

For information on how to create a case with AWS Incident Detection and Response to request changes to an existing onboarded workload, see [Request changes to an onboarded workload](#). For information on how to offboard a workload, see [Offboard a workload](#).

Workload discovery

AWS works with you to understand as much context about your workload as possible. AWS Incident Detection and Response uses this information to create runbooks to support you during incidents and AWS Service Events. The required information is captured in the [Workload onboarding and alarm ingestion questionnaires](#). It's a best practice to register your workloads on AppRegistry. For more information, see the [AppRegistry User Guide](#).

Key outputs:

- Workload information, such as the workload's description, architecture diagrams, contact, and escalation details.
- Details of how the workload employs AWS services in each AWS Region.
- Specific information on how AWS supports you during a Service Event.
- Alarms used by your team that detect critical workload impact.

Alarm configuration

AWS works with you to define metrics and alarms to provide visibility into the performance of your applications and their underlying AWS infrastructure. We ask that alarms adhere to the following criteria when defining and configuring thresholds:

- Alarms only enter the "Alarm" state when there is critical impact to the monitored workload (loss of revenue or degraded customer experience that significantly reduces performance) that requires immediate operator attention.
- Alarms must also engage your specified resolvers for the workload at the same time, or prior to, engaging the incident management team. Incident management engineers should be collaborating with your specified resolvers in the mitigation process, not serve as a first line responder and then escalate to you.
- Alarm thresholds must be set to an appropriate threshold and duration so that any time an alarm fires an investigation must take place. If an alarm is flapping between "Alarm" and "OK" state, sufficient impact is occurring to warrant operator response and attention.

Types of alarms:

- Alarms that portray the level of business impact and pass relevant information for simple fault detection.
- Amazon CloudWatch canaries. For more information, see [Canaries and X-Ray tracing](#), and [X-Ray](#).
- Aggregate alarming (monitoring of dependencies)

Example alarm, all using the CloudWatch monitoring system

Metric name / Alarm threshold	Alarm ARN or resource ID	If this alarm fires	If engaged, cut a Premium Support Case for these services
API errors / # of errors >= 10 for 10 datapoints	arn:aws:cloudwatch:us-west-2:000000000000:alarm:E2MPmimLambda-Errors	Ticket cut to database administrator (DBA) team	Lambda, API Gateway
ServiceUnavailable (Http status code 503) # of errors >=3 for 10 datapoints (different clients) in a 5 minute window	arn:aws:cloudwatch:us-west-2:xxxxx:alarm:httperrorcode503	Ticket cut to Service team	Lambda, API Gateway
ThrottlingException (Http status code 400)	arn:aws:cloudwatch:us-west-2:xxxxx:alarm:httperrorcode400	Ticket cut to Service team	EC2, Amazon Aurora

Metric name / Alarm threshold	Alarm ARN or resource ID	If this alarm fires	If engaged, cut a Premium Support Case for these services
# of errors ≥ 3 for 10 datapoints (different clients) in a 5 minute window			

For more details, see [AWS Incident Detection and Response monitoring and observability](#).

Key outputs:

- Definition and configuration of alarms on your workloads.
- Completion of the alarm details on the onboarding questionnaire.

Create CloudWatch alarms that fit your business needs in Incident Detection and Response

When you create Amazon CloudWatch alarms, there are several steps that you can take to make sure your alarms best fit your business needs.

Review your proposed CloudWatch alarms

Review your proposed alarms to make sure that they only enter the "Alarm" state when there is critical impact to the monitored workload (loss of revenue or degraded customer experience that significantly reduces performance). For example, do you consider this alarm critical enough that you must react immediately if it goes into the "Alarm" state?

The following are suggested metrics that might represent critical business impact, such as affecting your end users' experience with an application:

- **CloudFront:** For more information, see [Viewing CloudFront and edge function metrics](#).

- **Application Load Balancers:** It's a best practice that you create the following alarms for Application Load Balancers, if possible:
 - HTTPCode_ELB_5XX_Count
 - HTTPCode_Target_5XX_Count

The preceding alarms allow you to monitor responses from targets that are behind the Application Load Balancer, or behind other resources. This makes it easier to identify the source of 5XX errors. For more information, see [CloudWatch metrics for your Application Load Balancer](#).

- **Amazon API Gateway:** If you use WebSocket API in Elastic Beanstalk, then consider using the following metrics:
 - Integration error rates (filtered to 5XX errors)
 - Integration latency
 - Execution errors

For more information, see [Monitoring WebSocket API execution with CloudWatch metrics](#).

- **Amazon Route 53:** Monitor the **EndPointUnhealthyENICount** metric. This metric is the number of elastic network interfaces in the **Auto-recovering** status. This status indicates attempts by the resolver to recover one or more of the Amazon Virtual Private Cloud network interfaces that are associated with the endpoint (specified by **EndpointId**). In the recovery process, the endpoint functions with limited capacity. The endpoint can't process DNS queries until it's fully recovered. For more information, see [Monitoring Route 53 Resolver endpoints with Amazon CloudWatch](#).

Validate your alarm configurations

After you confirm that your proposed alarms fit your business needs, validate the configuration and history of the alarms:

- Validate the **Threshold** for the metric to enter the "Alarm" state against the metric's graph trend.
- Validate the **Period** used for polling data points. Polling data points at 60 seconds assist in early incident detection.
- Validate the **DatapointToAlarm** configuration. In most cases, it's a best practice to set this to 3 out of 3 or 5 out of 5. In an incident, the alarm triggers after 3 minutes when set as [60 second metrics with 3 out of 3 DatapointToAlarm] or 5 minutes when set as [60 second metrics with 5 out of 5 DatapointToAlarm]. Use this combination to eliminate noisy alarms.

Note

The preceding recommendations might vary depending on how you use a service. Each AWS service operates differently within a workload. And, the same service might operate differently when used in multiple places. You must be sure that you understand how your workload utilizes the resources that feed the alarm, as well as the upstream and downstream effects.

Validate how your alarms handle missing data

Some metric sources don't send data to CloudWatch at regular intervals. For these metrics, it's a best practice to treat missing data as **notBreaching**. For more information, see [Configuring how CloudWatch alarms treat missing data](#) and [Avoiding premature transitions to alarm state](#).

For example, if a metric monitors an error rate, and there are no errors, then the metric reports no data (nil) data points. If you configure the alarm to treat missing data as **Missing**, then a single breaching data point followed by two no data (nil) data points causes the metric to go into the "Alarm" state (for 3 out of 3 data points). This is because the missing data configuration evaluates the last known data point in the evaluation period.

In cases where metrics monitor an error rate, in the absence of service degradation you can assume that no data is a good thing. It's a best practice to treat missing data as **notBreaching** so that missing data is treated as "OK" and the metric doesn't enter the "Alarm" state on a single data point.

Review the history of each alarm

If an alarm's history shows that it frequently enters the "Alarm" state and then recovers quickly, then the alarm might become an issue for you. Make sure that you tune the alarm to prevent noise or false alarms.

Validate metrics for underlying resources

Make sure that your metrics look at valid underlying resources and use the correct statistics. If an alarm is configured to review invalid resource names, then the alarm might not be able to track the underlying data. This might cause the alarm to enter the "Alarm" state.

Create composite alarms

If you provide Incident Detection and Response operations with a large number of alarms for onboarding, you might be asked to create composite alarms. Composite alarms reduce the total number of alarms that need to be onboarded.

Use AWS CloudFormation templates to build CloudWatch alarms in Incident Detection and Response

To accelerate onboarding to AWS Incident Detection and Response, and to reduce the effort needed to build alarms, AWS provides you with AWS CloudFormation templates. These templates include optimized alarm settings for commonly onboarded services, such as Application Load Balancer, Network Load Balancer, and Amazon CloudFront.

Build CloudWatch alarms with CloudFormation templates

1. Download a template using the provided links:

NameSpace	Metrics	ComparisonOperator (Threshold)	Period	DatapointsToAlarm	TreatMissingData	Statistic	Template Link
Application Elastic Load Balancer	$(m1+m2)/(m1+m2+n4)*100$ m1=HTTPCode_Target_2XX_Count m2=HTTPCode_Target_3XX_Count m3=HTTPCode_Target_4XX_Count	LessThanThreshold(95)	60	3 out of 3	missing	Sum	Template

NameSpace	Metrics	ComparisonOperator (Threshold)	Period	DatapointsToAlarm	TreatingData	Statistic	Template Link
	HTTPCode_Target_5XX_Count						
Amazon CloudFront	TotalErrorRate	GreaterThanThreshold(5)	60	3 out of 3	notBreaching	Average	Template
Application Elastic Load Balancer	UnHealthyHostCount	GreaterThanOrEqualToThreshold(2)	60	3 out of 3	notBreaching	Maximum	Template
Network Elastic Load Balancer	UnHealthyHostCount	GreaterThanOrEqualToThreshold(2)	60	3 out of 3	notBreaching	Maximum	Template

- Review the downloaded JSON file to make sure that it meets your organization's operation and security processes.
- Create a CloudFormation stack:

Note

The following steps use the standard CloudFormation stack creation process. For detailed steps, see [Creating a stack on the AWS CloudFormation console](#).

- Open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation>.

- b. Choose **Create stack**.
- c. Choose **Template is ready**, and then upload the template file from your local folder.

The following is an example of the **Create stack** screen.

Create stack

Prerequisite - Prepare template

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

Template is ready

Use a sample template

Specify template
A template is a JSON or YAML file that describes your stack's resources and properties.

Template source
Selecting a template generates an Amazon S3 URL where it will be stored.

Amazon S3 URL

Upload a template file

Upload a template file

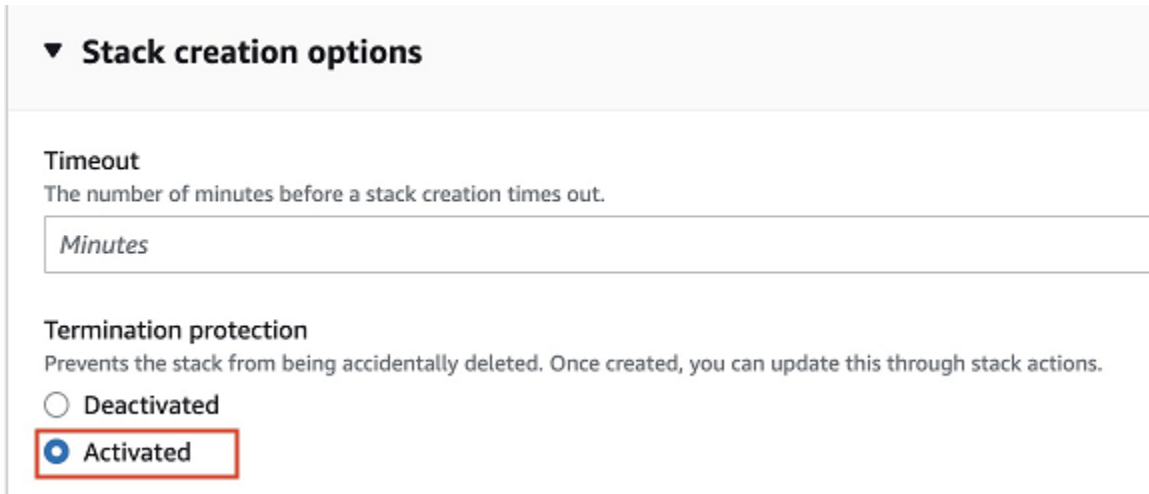
TotalErrorRateFromCloudFront.json

JSON or YAML formatted file

S3 URL: <https://s3.ap-southeast-2.amazonaws.com/cf-templates-1vgfb61k8amki-ap-southeast-2/2023-05-27T090143.078Z5pg-TotalErrorRateFromCloudFront.json>

- d. Choose **Next**.
- e. Enter the following required information:
 - **AlarmNameConfig** and **AlarmDescriptionConfig**: Enter a name and description for your alarm.
 - **ThresholdConfig**: Revise the threshold value to meet your application's requirements.
 - **DistributionIDConfig**: Make sure that the distribution ID point to the correct resources in the account tht you're creating theAWS CloudFormation stack in.
- f. Choose **Next**.
- g. Review the default values in the **PeriodConfig**, **EvalutionPeriodConfig**, and **DatapointsToAlarmConfig** fields. It's a best practice to use the default values for these fields. You can make adjustments, if needed, to meet your application's requirements.
- h. Optionally enter tags and SNS notification information as needed. It's a best practice to turn on **Termination protection**to prevent accidental deletion of the alarm. To turn

on termination protection, select the **Activated** radio button, as shown in the following example:



▼ Stack creation options

Timeout
The number of minutes before a stack creation times out.

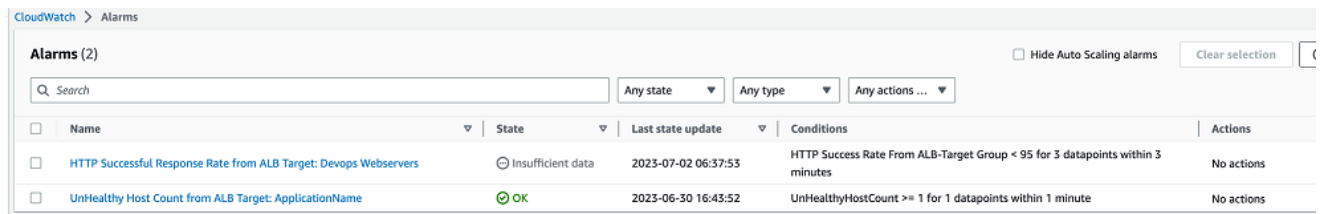
Minutes

Termination protection
Prevents the stack from being accidentally deleted. Once created, you can update this through stack actions.

Deactivated

Activated

- i. Choose **Next**.
- j. Review your stack settings, and then choose **Create stack**.
- k. After you create the stack, you see the alarm listed in the Amazon CloudWatch **Alarm list**, as shown in the following example:



<input type="checkbox"/>	Name	State	Last state update	Conditions	Actions
<input type="checkbox"/>	HTTP Successful Response Rate from ALB Target: Devops Webservers	⊖ Insufficient data	2023-07-02 06:37:53	HTTP Success Rate From ALB-Target Group < 95 for 3 datapoints within 3 minutes	No actions
<input type="checkbox"/>	UnHealthy Host Count from ALB Target: ApplicationName	⊕ OK	2023-06-30 16:43:52	UnHealthyHostCount >= 1 for 1 datapoints within 1 minute	No actions

4. After you create all of your alarms in the correct account and AWS Region, notify your Technical Account Manager (TAM). The AWS Incident Detection and Response team reviews the status of your new alarms, and then continues your onboarding.

Example uses cases for CloudWatch alarms in Incident Detection and Response

Review the following uses cases for examples of how you can use Amazon CloudWatch alarms in Incident Detection and Response.

Example Use Case A: Application Load Balancer

Create the following CloudWatch alarm that signals potential workload impact. You can create a metric math that alarms when successful connections drop below a certain threshold. For the available CloudWatch metrics, see [CloudWatch metrics for your Application Load Balancer](#)

Metric: HTTPCode_Target_3XX_Count;HTTPCode_Target_4XX_Count;HTTPCode_Target_5XX_Count
 $(m1+m2)/(m1+m2+m3+m4)*100$ m1 = HTTP Code 2xx || m2 = HTTP Code 3xx || m3 = HTTP Code 4xx || m4 = HTTP Code 5xx

NameSpace: AWS/ApplicationELB

ComparisonOperator(Threshold): Less than x (x = customer's threshold).

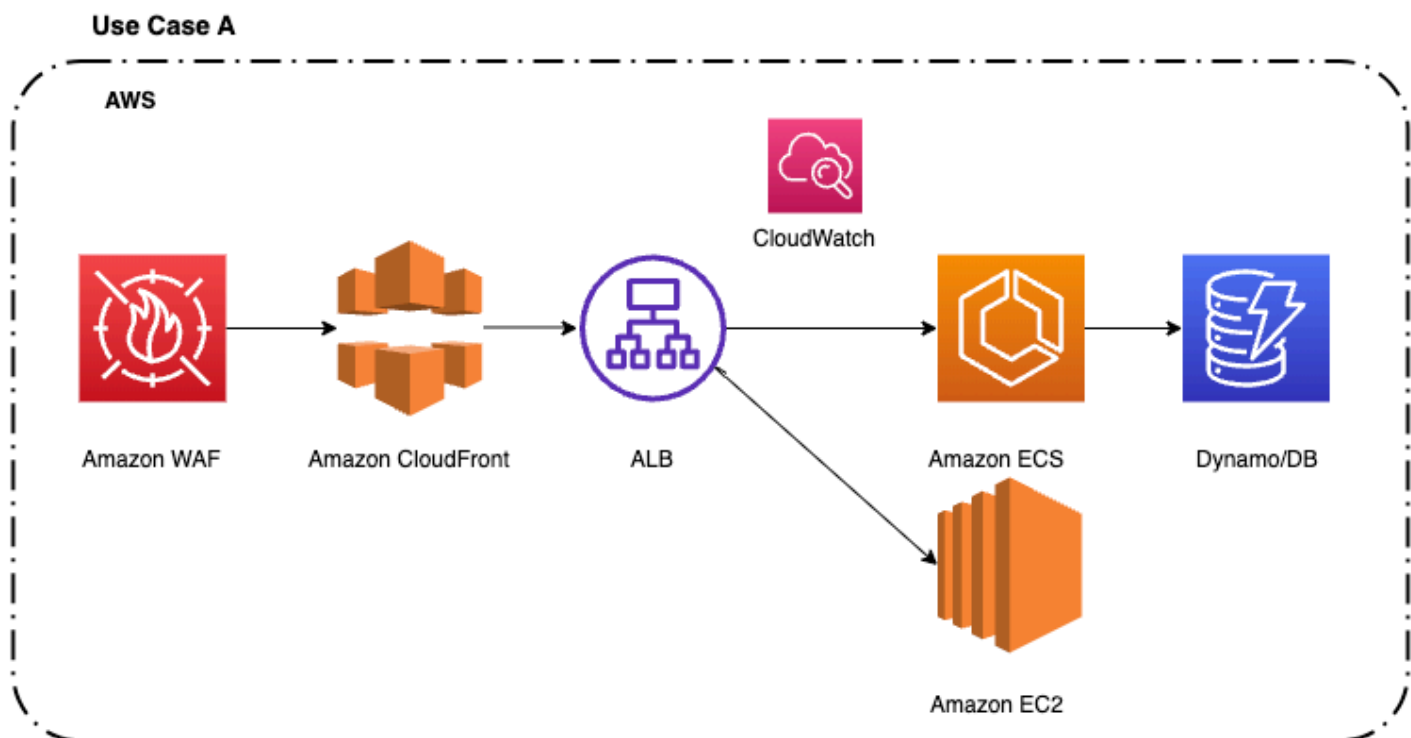
Period: 60 seconds

DatapointsToAlarm: 3 out of 3

Missing data treatment: Treat missing data as [breaching](#).

Statistic: Sum

The following diagram shows the flow for Use Case A:



Example Use Case B: Amazon API Gateway

Create the following CloudWatch alarm that signals potential workload impact. You can create a composite metric that alarms when there is high latency or a high average number of 4XX errors in the API Gateway. For the available metrics, see [Amazon API Gateway dimensions and metrics](#)

Metric: compositeAlarmAPI Gateway (ALARM(error4XXMetricApiGatewayAlarm)) OR (AALARM(latencyMetricApiGatewayAlarm))

NameSpace: AWS/API Gateway

ComparisonOperator(Threshold): Greater than (x or y customer's thresholds)

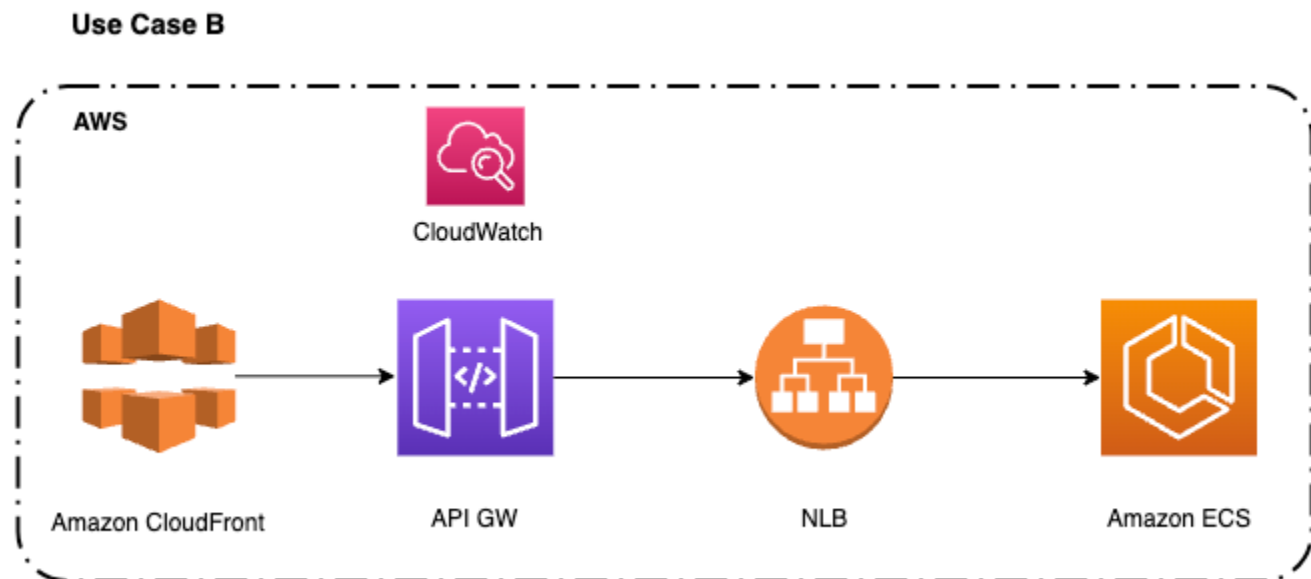
Period: 60 seconds

DatapointsToAlarm: 1 out of 1

Missing data treatment: Treat missing data as [not breaching](#).

Statistic:

The following diagram shows the flow for Use Case B:



Example Use Case C: Amazon Route 53

You can monitor your resources by creating Route 53 health checks that use CloudWatch to collect and process raw data into readable, near real-time metrics. You can create the following

CloudWatch alarm that signals potential workload impact. You can use the CloudWatch metrics to create an alarm that triggers when it breaches the established threshold. For the available CloudWatch metrics, see [CloudWatch metrics for Route 53 health checks](#)

Metric:R53-HC-Success

NameSpace: AWS/Route 53

Threshold HealthCheckStatus: HealthCheckStatus < x for 3 datapoints within 3 minutes (being x customer's threshold)

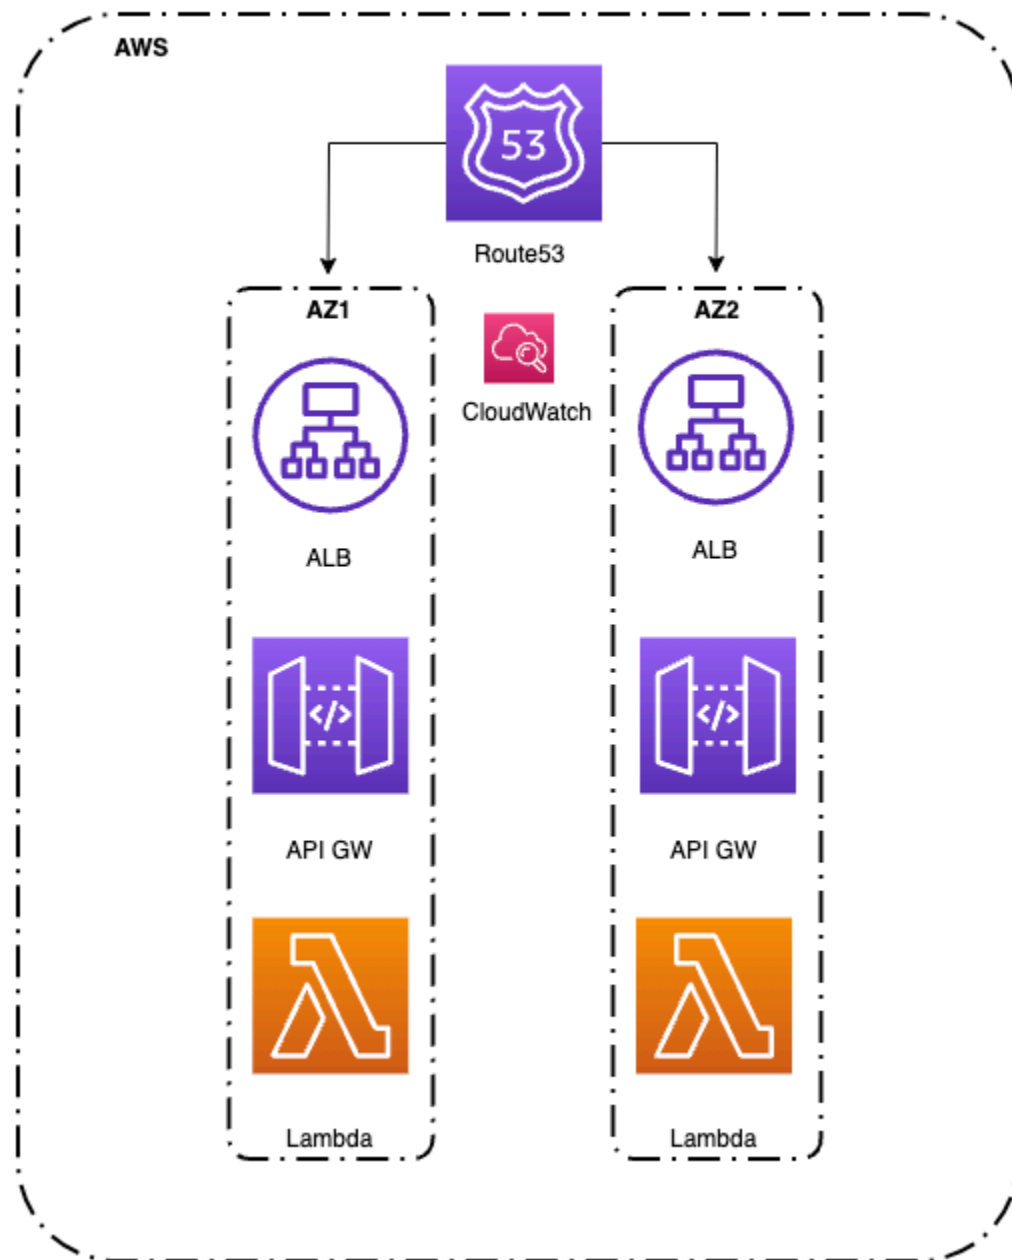
Period: 1 minute

DatapointsToAlarm: 3 out of 3

Missing data treatment: Treat missing data as [breaching](#).

Statistic: Minimum

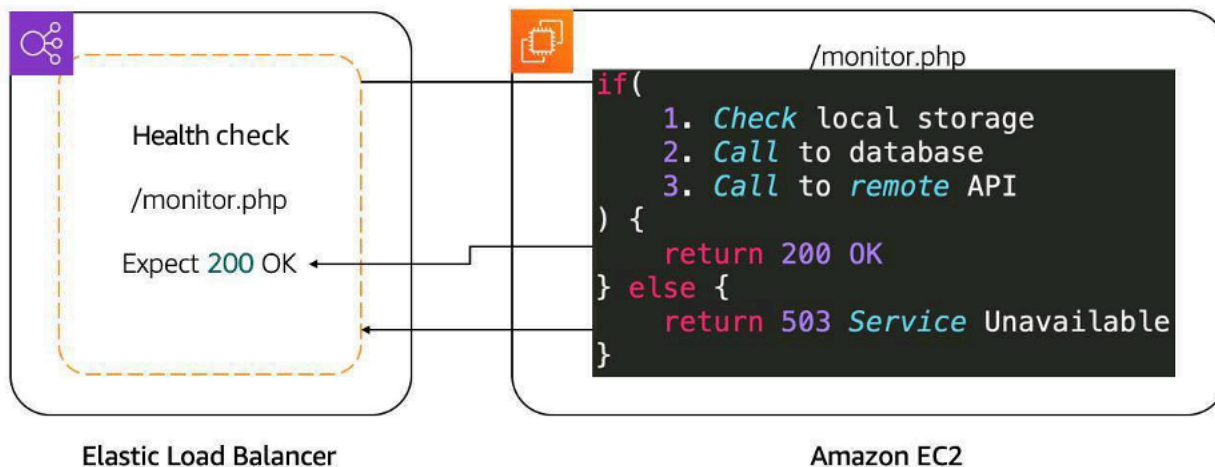
The following diagram shows the flow for Use Case C:

Use Case C**Example Use Case D: Monitor a workload with a custom app**

It's critical that you take the time to define an appropriate health check in this scenario. If you only verify that an application's port is open, then you haven't verified that the application is working. Additionally, making a call to the home page of an application is not necessarily the correct way to determine if the app is working. For instance, if an application depends on a database AND Amazon Simple Storage Service, then the health check must validate all of the elements. One way to do

that is to create a monitoring webpage, such as `/monitor`. The monitoring webpage makes a call to the database to make sure that it can connect and get data. And, the monitoring webpage makes a call to Amazon S3. Then, you point the health check on the load balancer to the `/monitor` page.

The following diagram shows the flow for Use Case D:



Ingest alerts into AWS Incident Detection and Response

AWS Incident Detection and Response supports alarm ingestion through [Amazon EventBridge](#). This section describes how to integrate AWS Incident Detection and Response with different Application Performance Monitoring (APM) tools, including Amazon CloudWatch, APMs with direct integration with Amazon EventBridge (for example, DataDog and New Relic), and APMs without direct integration with Amazon EventBridge. For a complete list of APMs with direct integration to Amazon EventBridge, see [Amazon EventBridge integrations](#).

Topics

- [Provision access for alert ingestion to Incident Detection and Response](#)
- [Integrate Incident Detection and Response with Amazon CloudWatch](#)
- [Ingest alarms from APMs that have direct integration with Amazon EventBridge](#)
- [Example: Integrate notifications from Datadog and Splunk](#)
- [Use webhooks to ingest alarms from APMs without direct integration with Amazon EventBridge](#)

Provision access for alert ingestion to Incident Detection and Response

To allow AWS Incident Detection and Response to ingest alarms from your account, install the `AWSServiceRoleForHealth_EventProcessor` service-linked role (SLR). AWS assumes the SLR to create an Amazon EventBridge-managed rule. The managed rule sends notifications from your accounts to AWS Incident Detection and Response. For information about this SLR, including the associated AWS managed policy, see [Using service-linked roles](#) in the *AWS Health User Guide*.

You can install this service-linked role in your account by following the instructions in [Create service-linked role](#) in the *AWS Identity and Access Management User Guide*. Or, you can use the following AWS Command Line Interface (AWS CLI) command:

```
aws iam create-service-linked-role --aws-service-name event-processor.health.amazonaws.com
```

Key outputs

- Successful installation of the Service Linked Role in your account.

Related information

For more information, see the following topics:

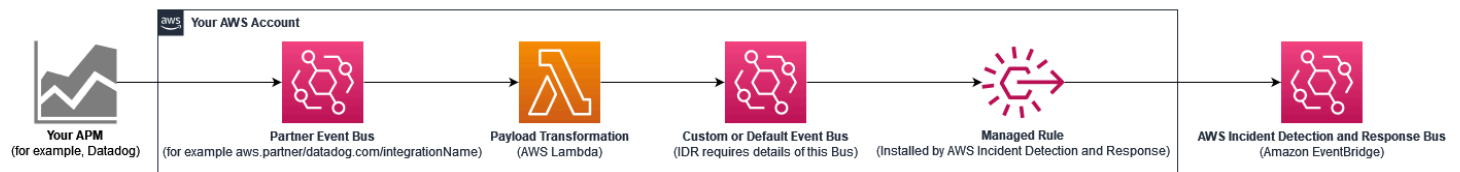
- [Using service-linked roles for AWS Health](#)
- [Creating a service-linked role](#)
- [AWS managed policy: AWSHealth_EventProcessorServiceRolePolicy](#)

Integrate Incident Detection and Response with Amazon CloudWatch

AWS Incident Detection and Response uses the service-linked role (SLR) that you turned on during access provisioning to create an Amazon EventBridge-managed rule in your AWS account named `AWSHealthEventProcessor-DO-NOT-DELETE`. Incident Detection and Response uses this rule to ingest Amazon CloudWatch alarms from your accounts. Additional steps aren't required to ingest alarms from CloudWatch.

Ingest alarms from APMs that have direct integration with Amazon EventBridge

The following illustration shows the process for sending notifications to AWS Incident Detection and Response from Application Performance Monitoring (APM) tools that have direct integration with Amazon EventBridge, such as Datadog and Splunk. For a complete list of APMs that have direct integration with EventBridge, see [Amazon EventBridge integrations](#)



Use the following steps to set up integration with AWS Incident Detection and Response. Before performing these steps, verify that the AWS service-linked role (SLR) `AWSServiceRoleForHealth_EventProcessor`, is [installed](#) in your accounts.

Set up integration with AWS Incident Detection and Response

You must complete the following steps for each AWS account and AWS Region. Alerts must come from the AWS account and AWS Region where the application resources reside.

1. Set up each of your APMs as Amazon EventBridge partner event sources (for example, `aws.partner/my_apm/integrationName`). For guidelines on setting up your APM as an event source, see [Receiving events from a SaaS partner with Amazon EventBridge](#). This creates a partner event bus in your account.
2. Do one of the following:
 - (Recommended method) Create a custom EventBridge event bus. AWS Incident Detection and Response installs a managed rule (`AWSHealthEventProcessorEventSource-DO-NOT-DELETE`) bus through the `AWSServiceRoleForHealth_EventProcessor` SLR. The rule source is the custom event bus. The rule destination is AWS Incident Detection and Response. The rule matches the pattern for ingesting 3rd party APM events.
 - (Alternative method) Use the default event bus instead of a custom event bus. The default event bus requires the managed rule to send APM alerts to AWS Incident Detection and Response.
3. Create an [AWS Lambda](#) function (for example, `My_APM-AWSIncidentDetectionResponse-LambdaFunction`) to transform your partner event bus events. The transformed events matches the managed rule `AWSHealthEventProcessorEventSource-DO-NOT-DELETE`.

- a. Transformed events include a unique AWS Incident Detection and Response identifier, and sets the source and detail type of the event to the required values. The pattern matches the managed rule.
 - b. Set the target of the Lambda function to either the custom event bus created in Step 2 (Recommended method) or to your default event bus.
4. Create an EventBridge rule and define the event patterns that match the list of events that you want to push to AWS Incident Detection and Response. The source of the rule is the partner event bus that you define in step 1 (for example, `aws.partner/my_apm/integrationName`). The target of the rule is the Lambda function that you define in step 3 (for example, `My_APM-AWSIncidentDetectionResponse-LambdaFunction`). For guidelines on defining your EventBridge rule, see [Amazon EventBridge rules](#).

For examples on how to set up a partner event bus integration for use with AWS Incident Detection and Response, see [Example: Integrate notifications from Datadog and Splunk](#).

Example: Integrate notifications from Datadog and Splunk

This example provides detailed steps for integrating notifications from Datadog and Splunk to AWS Incident Detection and Response.

1. Set up your APM as an event source in Amazon EventBridge in your AWS account.
2. Create a custom event bus.
3. Create an AWS Lambda function for transformation.
4. Create your custom EventBridge rule.

Step 1: Set up your APM as an event source in Amazon EventBridge

Set up each of your APMs as an event source in Amazon EventBridge in your AWS account. For instructions on setting up your APM as an event source, see the [event source set up instructions for your tool in Amazon EventBridge partners](#).

By setting up your APM as an event source, you can ingest notifications from your APM to an event bus in your AWS account. After setup, AWS Incident Detection and Response can start the incident management process when the event bus receives an event. This process adds Amazon EventBridge as a destination in your APM.

Step 2: Create a custom event bus

It's a best practice to use a custom event bus. AWS Incident Detection and Response uses the custom event bus to ingest transformed events. An AWS Lambda function transforms the partner event bus event and sends it to the custom event bus. AWS Incident Detection and Response installs a managed rule to ingest events from the custom event bus.

You can use the default event bus instead of a custom event bus. AWS Incident Detection and Response modifies the managed rule to ingest from the default event bus instead of a custom one.

Create a custom event bus in your AWS account:

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>
2. Choose **Buses, Event bus**.
3. Under **Custom event bus**, choose **Create**.
4. Provide a name for your event bus under **Name**. The recommended format is **APMName-AWSIncidentDetectionResponse-EventBus**.

As an example, use one of the following if you use Datadog or Splunk:

- **Datadog:** Datadog-AWSIncidentDetectionResponse-EventBus
- **Splunk:** Splunk-AWSIncidentDetectionResponse-EventBus

Step 3: Create an AWS Lambda function for transformation

The Lambda function transforms events between the partner event bus in Step 1 and the custom (or default) event bus from Step 2. The Lambda function transformation matches the AWS Incident Detection and Response managed rule.

Create an AWS Lambda function in your AWS account

1. Open the [Functions page](#) on the AWS Lambda console.
2. Choose **Create function**.
3. Choose the **Author from scratch** tab.
4. For **Function name**, enter a name using the format **APMName-AWSIncidentDetectionResponse-LambdaFunction**.

The following are examples for Datadog and Splunk:

- **Datadog:** Datadog-AWSIncidentDetectionResponse-LambdaFunction
 - **Splunk:** Splunk-AWSIncidentDetectionResponse-LambdaFunction
5. For **Runtime**, enter **Python 3.10**.
 6. Leave the remaining fields at the default values. Choose **Create function**.
 7. On the **Code edit** page, replace the default Lambda function content with the function in the following code examples.

Note the comments starting with # in the following code examples. These comments indicate which values to change.

Datadog transformation code template:

```
import logging
import json
import boto3

logger = logging.getLogger()
logger.setLevel(logging.INFO)

# Change the EventBusName to the custom event bus name you created previously or
# use your default event bus which is called 'default'.
# Example 'Datadog-AWSIncidentDetectionResponse-EventBus'
EventBusName = "Datadog-AWSIncidentDetectionResponse-EventBus"

def lambda_handler(event, context):
    # Set the event["detail"]["incident-detection-response-identifier"] value to
    # the name of your alert that is coming from your APM. Each APM is different and
    # each unique alert will have a different name.
    # Replace the dictionary path, event["detail"]["meta"]["monitor"]["name"], with
    # the path to your alert name based on your APM payload.
    # This example is for finding the alert name for Datadog.
    event["detail"]["incident-detection-response-identifier"] = event["detail"]
    ["meta"]["monitor"]["name"]
    logger.info(f"We got: {json.dumps(event, indent=2)}")

    client = boto3.client('events')
    response = client.put_events(
        Entries=[
            {
                'Detail': json.dumps(event["detail"], indent=2),
```

```

        'DetailType': 'ams.monitoring/generic-apm', # Do not modify. This
DetailType value is required.
        'Source': 'GenericAPMEvent', # Do not modify. This Source value is
required.
        'EventBusName': EventBusName # Do not modify. This variable is set
at the top of this code as a global variable. Change the variable value for your
eventbus name at the top of this code.
    }
]
)
print(response['Entries'])

```

Splunk transformation code template:

```

import logging
import json
import boto3

logger = logging.getLogger()
logger.setLevel(logging.INFO)

# Change the EventBusName to the custom event bus name you created previously or
use your default event bus which is called 'default'.
# Example Splunk-AWSIncidentDetectionResponse-EventBus
EventBusName = "Splunk-AWSIncidentDetectionResponse-EventBus"

def lambda_handler(event, context):
    # Set the event["detail"]["incident-detection-response-identifier"] value to
the name of your alert that is coming from your APM. Each APM is different and
each unique alert will have a different name.
    # replace the dictionary path event["detail"]["ruleName"] with the path to your
alert name based on your APM payload.
    # This example is for finding the alert name in Splunk.
    event["detail"]["incident-detection-response-identifier"] = event["detail"]
["ruleName"]
    logger.info(f"We got: {json.dumps(event, indent=2)}")

    client = boto3.client('events')
    response = client.put_events(
        Entries=[
            {
                'Detail': json.dumps(event["detail"], indent=2),

```

```

        'DetailType': 'ams.monitoring/generic-apm', # Do not modify. This
DetailType value is required.
        'Source': 'GenericAPMEvent', # Do not modify. This Source value is
required.
        'EventBusName': EventBusName # Do not modify. This variable is set
at the top of this code as a global variable. Change the variable value for your
eventbus name at the top of this code.
    }
]
)
print(response['Entries'])

```

8. Choose **Deploy**.
9. Add **PutEvents** permission to the Lambda execution role for the event bus that you're sending the transformed data to:
 - a. Open the [Functions page](#) on the AWS Lambda console.
 - b. Select the function, and then choose **Permissions** on the **Configuration** tab.
 - c. Under **Execution role**, select the **Role name** to open the execution role in the AWS Identity and Access Management console.
 - d. Under **Permissions policies**, select the existing policy name to open the policy.
 - e. Under **Permissions defined in this policy**, choose **Edit**.
 - f. On the **Policy editor** page, select **Add new statement**:
 - g. The **Policy editor** adds a new blank statement similar to the following

```

{
  "Sid": "Statement1",
  "Effect": "Allow",
  "Action": [],
  "Resource": []
}

```

- h. Replace the new auto-generated statement with the following:

```

{
  "Sid": "AWSIncidentDetectionResponseEventBus0",
  "Effect": "Allow",
  "Action": "events:PutEvents",
  "Resource": "arn:aws:events:{region}:{accountId}:event-bus/{custom-eventbus-
name}"
}

```

```
}

```

- i. The **Resource** is the ARN of the custom event bus that you created in [Step 2: Create a custom event bus](#) or the ARN of your default event bus if you are using the default event bus in your Lambda code.

10. Review and confirm that the required permission are added to the role.

11. Choose **Set this new version as the default**, and then choose **Save changes**.

What's required from a payload transformation?

The following JSON key:value pairs are required in event bus events ingested by AWS Incident Detection and Response.

```
{
  "detail-type": "ams.monitoring/generic-apm",
  "source": "GenericAPMEvent"
  "detail" : {
    "incident-detection-response-identifier": "Your alarm name from your APM",
  }
}
```

The following examples show an event from a partner event bus before and after it is transformed.

```
{
  "version": "0",
  "id": "a6150a80-601d-be41-1a1f-2c5527a99199",
  "detail-type": "Datadog Alert Notification",
  "source": "aws.partner/datadog.com/Datadog-aaa111bbbc",
  "account": "123456789012",
  "time": "2023-10-25T14:42:25Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "alert_type": "error",
    "event_type": "query_alert_monitor",
    "meta": {
      "monitor": {
        "id": 222222,
        "org_id": 3333333333,
        "type": "query alert",
        "name": "UnHealthyHostCount",

```

```
    "message": "@awseventbridge-Datadog-aaa111bbbc",
    "query":
"max(last_5m):avg:aws.applicationelb.un_healthy_host_count{aws_account:123456789012}
\u003c\u003d 1",
    "created_at": 1686884769000,
    "modified": 1698244915000,
    "options": {
      "thresholds": {
        "critical": 1.0
      }
    },
  },
  "result": {
    "result_id": 7281010972796602670,
    "result_ts": 1698244878,
    "evaluation_ts": 1698244868,
    "scheduled_ts": 1698244938,
    "metadata": {
      "monitor_id": 222222,
      "metric": "aws.applicationelb.un_healthy_host_count"
    }
  },
  "transition": {
    "trans_name": "Triggered",
    "trans_type": "alert"
  },
  "states": {
    "source_state": "OK",
    "dest_state": "Alert"
  },
  "duration": 0
},
"priority": "normal",
"source_type_name": "Monitor Alert",
"tags": [
  "aws_account:123456789012",
  "monitor"
]
}
}
```

Note that before the event is transformed, `detail-type` indicates the APM that the alert came from, the source is from a partner APM, and the `incident-detection-response-identifier` key is not present.

The Lambda function transforms the above event and puts it in to the target custom or default event bus. The transformed payload now includes the required key:value pairs.

```
{
  "version": "0",
  "id": "7f5e0fc1-e917-2b5d-a299-50f4735f1283",
  "detail-type": "ams.monitoring/generic-apm",
  "source": "GenericAPMEvent",
  "account": "123456789012",
  "time": "2023-10-25T14:42:25Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "incident-detection-response-identifier": "UnHealthyHostCount",
    "alert_type": "error",
    "event_type": "query_alert_monitor",
    "meta": {
      "monitor": {
        "id": 222222,
        "org_id": 3333333333,
        "type": "query alert",
        "name": "UnHealthyHostCount",
        "message": "@awseventbridge-Datadog-aaa111bbbc",
        "query":
"max(last_5m):avg:aws.applicationelb.un_healthy_host_count{aws_account:123456789012}
\u003c\u003d 1",
        "created_at": 1686884769000,
        "modified": 1698244915000,
        "options": {
          "thresholds": {
            "critical": 1.0
          }
        },
      },
    },
    "result": {
      "result_id": 7281010972796602670,
      "result_ts": 1698244878,
      "evaluation_ts": 1698244868,
      "scheduled_ts": 1698244938,
    }
  }
}
```



```
    "metadata": {
      "monitor_id": 222222,
      "metric": "aws.applicationelb.un_healthy_host_count"
    }
  },
  "transition": {
    "trans_name": "Triggered",
    "trans_type": "alert"
  },
  "states": {
    "source_state": "OK",
    "dest_state": "Alert"
  },
  "duration": 0
},
"priority": "normal",
"source_type_name": "Monitor Alert",
"tags": [
  "aws_account:123456789012",
  "monitor"
]
}
```

Note that `detail-type` is now `aws.monitoring/generic-apm`, `source` is now `GenericAPMEvent`, and under `detail` there is new key:value pair: `incident-detection-response-identifier`.

In the preceding example, the `incident-detection-response-identifier` value is taken from the alert name under the path `$.detail.meta.monitor.name`. APM alert name paths are different from one APM to another. The Lambda function must be modified to take the alarm name from the correct partner event JSON path and use it for the `incident-detection-response-identifier` value.

Each unique name that is set on the `incident-detection-response-identifier` is provided to the AWS Incident Detection and Response team during on-boarding. Events that have an unknown name for the `incident-detection-response-identifier` aren't processed.

Step 4: Create a custom Amazon EventBridge rule

The partner event bus created in Step 1 requires an EventBridge rule that you create. The rule sends the desired events from the partner event bus to the Lambda function created in Step 3.

For guidelines on defining your EventBridge rule, see [Amazon EventBridge rules](#).

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>
2. Choose **Rules**, and then select the partner event bus associated with your APM. The following are examples of partner event buses:
 - **Datadog:** aws.partner/datadog.com/eventbus-name
 - **Splunk:** aws.partner/signalfx.com/RandomString
3. Choose **Create rule** to create a new EventBridge rule.
4. For rule name, enter a name in the following format `APMName-AWS Incident Detection and Response-EventBridgeRule`, and then choose **Next**. The following are example names:
 - **Datadog:** Datadog-AWSIncidentDetectionResponse-EventBridgeRule
 - **Splunk:** Splunk-AWSIncidentDetectionResponse-EventBridgeRule
5. For **Event source**, select **AWS events or EventBridge partner events**.
6. Leave **Sample event** and **Creation method** as the default values.
7. For **Event pattern**, choose the following:
 - a. **Event source:** EventBridge partners.
 - b. **Partner:** Select your APM Partner.
 - c. **Event Type:** All events.

The following are example event patterns:

Example Datadog event pattern

Event pattern [Info](#)

Event source
AWS service or EventBridge partner as source

EventBridge partners ▼

Partner
EventBridge Partner as the event source

Datadog ▼

Event type
The type of events as the source of the matching pattern

All Events ▼

Event pattern
Event pattern, or filter to match the events

```
1 {  
2   "source": [{  
3     "prefix": "aws.partner/datadog.com"  
4   }]  
5 }
```

[Copy](#) [Test pattern](#) [Edit pattern](#)

Example Splunk event pattern

Event pattern [Info](#)

Event source
AWS service or EventBridge partner as source

EventBridge partners ▼

Partner
EventBridge Partner as the event source

SignalFx ▼

Event type
The type of events as the source of the matching pattern

All Events ▼

Event pattern
Event pattern, or filter to match the events

```
1 {  
2   "source": [{  
3     "prefix": "aws.partner/signalfx.com"  
4   }]  
5 }
```

[Copy](#) [Test pattern](#) [Edit pattern](#)

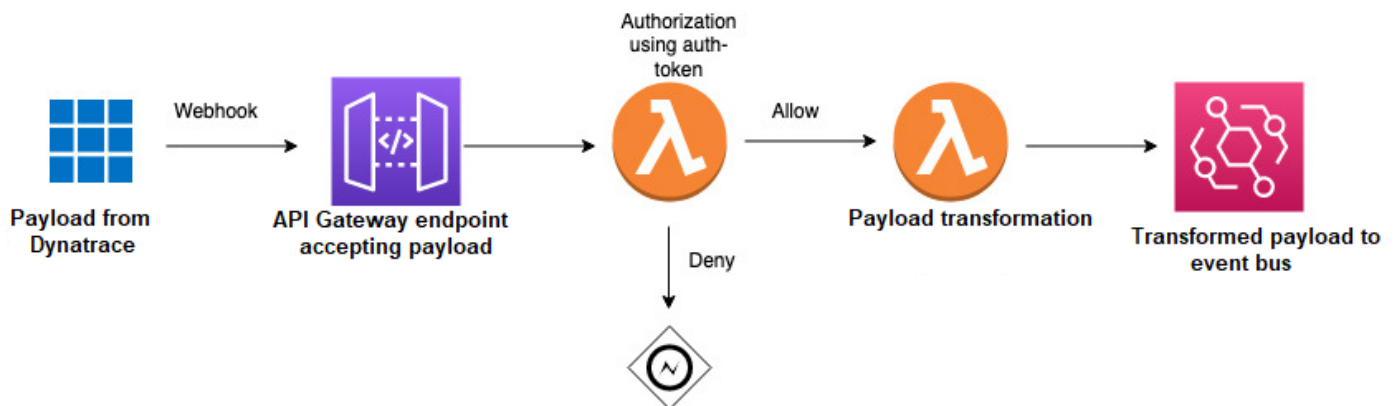
8. For **Targets**, choose the following:

- a. **Target types:** AWS service
 - b. **Select a target:** Choose Lambda function.
 - c. **Function:** The name of the Lambda function that you created in Step 2.
9. Choose **Next, Save rule**.

Use webhooks to ingest alarms from APMs without direct integration with Amazon EventBridge

AWS Incident Detection and Response supports using webhooks for alarm ingestion from third party APMs that don't have direct integration with Amazon EventBridge.

For a list of APMs with direct integrations with Amazon EventBridge, see [Amazon EventBridge integrations](#).



Use the following steps to set up integration with AWS Incident Detection and Response. Before performing these steps, verify that the AWS Managed Rule, *AWSHealthEventProcessorEventSource-DO-NOT-DELETE*, is installed in your accounts

Ingest events using webhooks

1. Define an Amazon API Gateway to accept the payload from your APM.
2. Define an AWS Lambda function for authorization using an authentication token, as displayed in the preceding illustration.
3. Define a second Lambda function to transform and append the AWS Incident Detection and Response identifier to your payload. You can also use this function to filter for the events that you want to send to AWS Incident Detection and Response.
4. Set up your APM to send notifications to the URL generated from the API Gateway.

Develop runbooks for AWS Incident Detection and Response

You can download an example Incident Detection and Response runbook: [aws-idr-runbook-example.zip](#).

Incident Detection and Response uses information captured from your onboarding questionnaire to develop runbooks and response plans for the management of incidents affecting your workloads. Runbooks document steps Incident Managers take when responding to an incident. A response plan is mapped to at least one of your workloads. The incident management team creates these templates from the information provided by you during workload discovery, described previously. Response plans are AWS Systems Manager (SSM) document templates used to trigger incidents. To learn more about SSM documents, see [AWS Systems Manager Documents](#), to learn more about Incident Manager, see [What Is AWS Systems Manager Incident Manager?](#)

Key outputs:

- Completion of your workload definition on AWS Incident Detection and Response.
- Completion of alarms, runbooks and response plan definition on AWS Incident Detection and Response.

You can also download an AWS Incident Detection and Response Runbook example: [aws-idr-runbook-example.zip](#).

Example runbook:

Runbook template for AWS Incident Detection and Response

Description

This document is intended for [CustomerName] [WorkloadName].

[Insert short description of what the workload is intended for].

Step: Priority

Priority actions

1. When a case is created with Incident Detection and Response, lock the case to yourself, verify the Customer Stakeholders in the Case from *Engagement Plans - Initial Engagement*.
2. Send the first correspondence on the support case to the customer as below. If there is no support case or if it is not possible to use the support case then backup communication details are listed in the steps that follow.

```
...
Hello,

This is <<Engineer's name>> from AWS Incident Detection and Response. An alarm has
triggered for your workload <<application name>>. I am currently investigating and
will update you in a few minutes after I have finished initial investigation.

Alarm Identifier - <insert CloudWatch Alarm ARN or APM Response Identifier>
...

**Compliance and regulatory requirements for the workload**
<<e.g. The workload deals with patient health records which must be kept secured and
confidential. Information not to be shared with any third parties.>>

**Actions required from Incident Detection and Response in complying**
<<e.g Incident Management Engineers must not shared data with third parties.>>

## Step: Information
**Review of common information**

* This section provides a space for defining common information which may be needed
through the life of the incident.
* The target user of this information is the Incident Management Engineer and
Operations Engineer.
* The following steps may reference this information to complete an action (for
example, execute the "Initial Engagement" plan).

---
**Engagement plans**

Describe the engagement plans applicable to this runbook. This section contains
only contact details. Engagement plans will be referenced in the step by step
**Communication Plans**.

* **Initial engagement**

AWS Incident Detection and Response Team will add customer stakeholder addresses below
to the Support Case. AWS Stakeholders are for additional stakeholders that may need to
be made aware of any issues.
When updating customer stakeholders details in this plan also update the Backup Mailto
links.
* ***Customer Stakeholders***: customeremail1; customeremail2; etc
* ***AWS Stakeholders***: aws-idr-oncall@amazon.com; tam-team-email; etc.
```

* **One Time Only Contacts**: [These are email contacts that are included on only the first communication. Remove these contacts after the first communication has gone out. These could be customer paging email addresses such as pager-duty that must not be paged for every correspondence]

* **Backup Mailto Impact Template**: <Insert Impact Template Mailto Link here>

* Use the backup Mailto when communication over cases is not possible.

* **Backup Mailto No Impact Template**: <Insert No Impact Mailto Link here>

* Use the backup Mailto when communication over cases is not possible.

* **Engagement Escalation**

AWS Incident Detection and Response will reach out to the following contacts when the contacts from the **Initial engagement** plan do not respond to incidents.

For each Escalation Contact indicate if they must be added to the support case, phoned or both.

* **First Escalation Contact**: [escalationEmailAddress#1] / [PhoneNumber] - Wait XX Minutes before escalating to this contact.

* [add Contact to Case / phone] this contact.

* **Second Escalation Contact**: [escalationEmailAddress#2] / [PhoneNumber] - Wait XX Minutes before escalating to this contact.

* [add Contact to Case / phone] this contact.

* Etc;

* **Communication plans**

Describe how Incident Management Engineer communicates with designated stakeholders outside the incident call and communication channels.

* **Impact Communication plan**

This plan is initiated when Incident Detection and Response have determined from step **Triage** that an alert indicates potential impact to a customer.

Incident Detection and Response will request the customer to join the predetermined bridge (Chime Bridge/Customer Provided Bridge / Customer Static Bridge) as indicated in **Engagement plans - Incident call setup**.

All backup email templates for use when cases can't be used are in **Engagement plans - Initial engagement**.

* 1 - Before sending the impact notification, verify then remove and/or add customer contacts from the Support Case CC based on the contacts listed in the **Initial engagement** Engagement plan.

* 2 - Send the engagement notification to the customer based the following Template:

(choose one and remove the rest)

Impact Template - Chime Bridge

...

The following alarm has engaged AWS Incident Detection and Response to an Incident bridge:

Alarm Identifier - <insert CloudWatch Alarm ARN or APM Response Identifier>

Alarm State Change Reason - <insert state change reason>

Alarm Start Time - <Example: 1 January 2023, 3:30 PM UTC>

Please join the Chime Bridge below so we can start the steps outlined in your Runbook:

<insert Chime Meeting ID>

<insert Link to Chime Bridge>

International dial-in numbers: <https://chime.aws/dialinnumbers/>

...

Impact Template - Customer Provided Bridge

...

The following alarm has engaged AWS Incident Detection and Response:

Alarm Identifier - <insert CloudWatch Alarm ARN or APM Response Identifier>

Alarm State Change Reason - <insert state change reason>

Alarm Start Time - <Example: 1 January 2023 3:30 PM UTC>

Please respond with your internal bridge details so we can join and start the steps outlined in your Runbook.

...

Impact Template - Customer Static Bridge

...

The following alarm has engaged AWS Incident Detection and Response to an Incident bridge:

Alarm Identifier - <insert CloudWatch Alarm ARN or APM Response Identifier>

Alarm State Change Reason - <insert state change reason>

Alarm Start Time - <Example: 1 January 2023, 3:30 PM UTC>

Please join the Bridge below so we can start the steps outlined in your Runbook:

Conference Number: <insert conference number>

Conference URL : <insert bridgeURL>

...

- * 3 - Set the Case to Pending Customer Action
- * 4 - Follow **Engagement Escalation** plan as mentioned above.
- * 5 - If the customer does not respond within 30 minutes, disengage and continue to monitor until the alarm recovers.

* **No Impact Communication plan**

This plan is initiated when an alarm recovers before Incident Detection and Response have completed initial **Triage**.

- * 1 - Before sending the no impact notification, verify then remove and/or add customer contacts from the Support Case CC based on the contacts listed in the **Engagement plans - Initial engagement** Engagement plan.

* 2 - Send a no engagement notification to the customer based on the below template:

No Impact Template

...

AWS Incident Detection and Response received an alarm that has recovered for your workload.

Alarm Identifier - <insert CloudWatch Alarm ARN or APM Response Identifier>

Alarm State Change Reason - <insert state change reason>

Alarm Start Time - <Example: 1 January 2023, 3:30 PM UTC>

Alarm End Time - <Example: 1 January 2023, 3:35 PM UTC>

This may indicate a brief customer impact that is currently not ongoing.

If there is an ongoing impact to your workload, please let us know and we will engage to assist.

...

* 3 - Put the case in to Pending Customer Action.

* 4 - If the customer does not respond within 30 minutes Resolve the case.

* **Updates**

If AWS Incident Detection and Response is expected to provide regular updates to customer stakeholders, list those stakeholders here. Updates must be sent via the same support case.

Remove this section if not needed.

* Update Cadence: Every XX minutes

* External Update Stakeholders: customeremailaddress1; customeremailaddress2; etc

* Internal Update Stakeholders: awsemailaddress1; awsemailaddress2; etc

Application architecture overview

This section provides an overview of the application/workload architecture for Incident Management Engineer and Operations Engineer awareness.

* **AWS Accounts and Regions with key services** - list of AWS accounts with regions supporting this application. Assists Engineers in assessing underlying infrastructure supporting the application.

* 123456789012

* US-EAST-1 - brief desc as appropriate

* EC2 - brief desc as appropriate

* DynamoDB - brief desc as appropriate

* etc.

* US-WEST-1 - brief desc as appropriate

* etc.

* another-account-etc.

```
* Resource identification - describe how engineers determine resource association with application
  * Resource groups: etc.
  * Tag key/value: AppId=123456
```

```
* CloudWatch Dashboards - list dashboards relevant to key metrics and services
  * 123456789012
  * us-east-1
    * some-dashboard-name
  * etc.
  * some-other-dashboard-name-in-current-acct
```

Step: Triage

Evaluate incident and impact

This section provides instructions for triaging of the incident to determine correct impact, description, and overall correct runbook being executed.

* **Evaluation of initial incident information**

- * 1 - Review Incident Alarm, noting time of first detected impact as well as the alarm start time.
- * 2 - Identify which service(s) in the customer application is seeing impact.
- * 3 - Review AWS Service Health for services listed under **AWS Accounts and Regions with key services**.
- * 4 - Review any customer provided dashboards listed under **CloudWatch Dashboards**

* **Impact**

Impact is determined when either the customer's metrics do not recover, appear to be trending worse or if there is indication of AWS Service Impact.

- * 1 - Start **Communication plans - Impact Communication plan**
- * 2 - Start **Engagement plans - Engagement Escalation** if no response is received from the **Initial Engagement** contacts.
- * 3 - Start **Communication plans - Updates** if specified in **Communication plans**

* **No Impact**

No Impact is determined when the customer's alarm recovers before Triage is complete and there are no indications of AWS service impact or sustained impact on the customer's CloudWatch Dashboards.

- * 1 - Start **Communication plans - No Impact Communication plan**

Step: Investigate

Investigation

This section describes performing investigation of known and unknown symptoms.

****Known issue****

- * *List all known issues with the application and their standard actions here*

****Unknown issues****

- * Investigate with the customer and AWS Premium Support.
- * Escalate internally as required.

Step: Mitigation****Collaborate****

- * Communicate any changes or important information from the ****Investigate**** step to the members of the incident call.

****Implement mitigation****

- * ***List customer failover plans / Disaster Recovery plans / etc here for implementing mitigation.


Step: Recovery****Monitor customer impact****

- * Review metrics to confirm recovery.
- * Ensure recovery is across all Availability Zones / Regions / Services
- * Get confirmation from the customer that impact is over and the application has recovered.

****Identify action items****

- * Record key decisions and actions taken, including temporary mitigation that might have been implemented.
- * Ensure outstanding action items have assigned owners.
- * Close out any Communication plans that were opened during the incident with a final confirmation of recovery notification.

Test onboarded workloads

 Note

The AWS Identity and Access Management user or role that you use for alarm testing must have `cloudwatch:SetAlarmState` permission.

The last step in the onboarding process is to perform a gameday for your new workload. After alarm ingestion completes, AWS Incident Detection and Response confirms a date and time of your choosing to start your gameday.

Your gameday serves two main purposes:

- **Functional Validation:** Confirms that AWS Incident Detection and Response can correctly receive your alarm events. And, functional validation confirms that your alarm events trigger the appropriate runbooks and any other desired actions, such as auto case creation if you selected it during alarm ingestion.
- **Simulation:** The gameday is an end to end simulation of what might happen during a real incident. AWS Incident Detection and Response follows your prescribed runbook steps to give you insight into how a real incident might unfold. The gameday is an opportunity for you to ask questions or refine instructions to improve the engagement.

During the alarm test, AWS Incident Detection and Response works with you to remediate any issues identified.

CloudWatch alarms

AWS Incident Detection and Response tests your Amazon CloudWatch alarms by monitoring the state change of your alarm. To do this, manually change the alarm to the **Alarm** state using the AWS Command Line Interface. You can also access the AWS CLI from AWS CloudShell. AWS Incident Detection and Response provides you with a list of AWS CLI commands for you to use during testing.

Example AWS CLI command to set an alarm state:

```
aws cloudwatch set-alarm-state --alarm-name "ExampleAlarm" --state-value ALARM --state-reason "Testing AWS Incident Detection and Response Configuration" --region us-east-1
```

To learn more about manually changing the state of CloudWatch alarms, see [SetAlarmState](#).

To learn more about the permissions required for CloudWatch API operations, see [Amazon CloudWatch permissions reference](#).

Third party APM alarms

Workloads that utilize a third party Application Performance Monitoring (APM) tool such as DataDog, Splunk, NewRelic, or Dynatrace require different instructions to simulate an alarm. At the start of the GameDay, AWS Incident Detection and Response requests that you temporarily change your alarm thresholds or comparison operators to force the alarm into the **ALARM** status. This status triggers a payload to AWS Incident Detection and Response.

Key outputs

Key outputs:

- Alarm ingestion is successful and your alarm configuration is correct.
- Alarms are successfully created and received by AWS Incident Detection and Response.
- A support case is created for your engagement and your prescribed contacts are notified.
- AWS Incident Detection and Response can engage with you by your prescribed conference means.
- All alarms and support cases generated as part of the Gameday are resolved.
- A Go-Live email is sent confirming your workload is now being monitored by AWS Incident Detection and Response.

Workload onboarding and alarm ingestion questionnaires

Download the [Workload onboarding questionnaire](#).

Download the [Alarm ingestion questionnaire](#).

Workload onboarding questionnaire - General questions

General questions

Question	Example Response
Enterprise Name	Amazon Inc.
Name of this workload (include any abbreviations)	Amazon Retail Operations (ARO)


Question	Example Response
Primary end user and the function of this workload.	This workload is an e-commerce application that allows end users to purchase various items. This workload is the primary revenue generator for our business.
Applicable compliance and/or regulatory requirements for this workload and any actions required from AWS after an incident.	The workload deals with patient health records which must be kept secured and confidential.

Workload onboarding questionnaire - Architecture questions

Architecture questions

Question	Example Response
<p>A list of AWS resource tags used to define resources that are part of this workload. AWS uses these tags to identify this workload's resources to expedite support during incidents.</p> <p>.</p> <div data-bbox="113 1207 792 1522" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Tags are case sensitive. If you provide multiple tags, all resources used by this workload must have the same tags.</p> </div>	<p>appName: Optimax</p> <p>environment: Production</p>
A list of AWS Services utilized by this workload and the AWS Account and Regions that they're in.	<p>Route 53: Routes internet traffic to the ALB.</p> <p>Account:123456789101</p> <p>Region: US-EAST-1, US-WEST-2</p>

Question	Example Response
<p>Note Create a new row for each service.</p>	
<p>A list of AWS Services utilized by this workload and the AWS Account and Regions that they're in.</p> <p>Note Create a new row for each service.</p>	<p>ALB: Routes incoming traffic to a target group of ECS containers.</p> <p>Account: 123456789101</p> <p>Region: N/A</p>
<p>A list of AWS Services utilized by this workload and the AWS Account and Regions that they're in.</p> <p>Note Create a new row for each service.</p>	<p>ECS: Compute infrastructure for main business logic fleet. Responsible for handling incoming user requests and making queries to persistence layer.</p> <p>Account: 123456789101</p> <p>Region: US-EAST-1</p>
<p>A list of AWS Services utilized by this workload and the AWS Account and Regions that they're in.</p> <p>Note Create a new row for each service.</p>	<p>RDS: Amazon Aurora cluster stores user data accessed by ECS business logic layer.</p> <p>Account: 123456789101</p> <p>Region: US-EAST-1</p>

Question	Example Response
<p>A list of AWS Services utilized by this workload and the AWS Account and Regions that they're in.</p> <div data-bbox="115 401 792 569" style="border: 1px solid #0070C0; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note Create a new row for each service.</p> </div>	<p>S3: Stores website static assets.</p> <p>Account: 123456789101</p> <p>Region: N/A</p>
<p>Detail any upstream/downstream component s not being onboarded that could affect this workload if experiencing an outage.</p>	<p>Authentication Microservice: Will prevent users from loading their health records as they will be unauthenticated.</p>
<p>Are there any on-premise or non-AWS components for this workload? If so, what are they and what functions are performed?</p>	<p>All internet based traffic in/out of AWS is routed via our on-prem proxy service.</p>
<p>Provide details of any manual or automated failover/disaster recovery plans at the Availability Zone and regional level.</p>	<p>Warm standby. Automated failover to US-WEST-2 during sustained drop in success rate.</p>

Workload onboarding questionnaire - AWS Service Event questions

AWS Service Event questions

Question	Example Response
<p>Provide the contact details (name/email/ phone) of your company's internal major incident/IT crisis management team.</p>	<p>Major Incident Management Team</p> <p>mim@example.com</p> <p>+61 2 3456 7890</p>
<p>Provide details of any static incident/crisis management bridge established by your company. If you utilize non-static bridges, then specify your preferred application and</p>	<p>Amazon Chime</p> <p>https://chime.aws/1234567890</p>

Question	Example Response
<p>AWS will request these details during an incident.</p> <div data-bbox="115 331 792 602" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>If one isn't provided, then AWS will reach out during an incident and provide a Chime bridge for you to join.</p> </div>	

Alarm Ingestion Questionnaire

Runbook Questions

Question	Example Response
<p>AWS will engage workload contacts through the AWS Support Case. Who is the primary contact when an alarm triggers for this workload?</p> <p>Specify your preferred conferencing application and AWS will request these details during an incident.</p> <div data-bbox="115 1346 792 1661" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>If a preferred conferencing application isn't provided, then AWS will reach out during an incident and provide a Chime bridge for you to join.</p> </div>	<p>Application Team</p> <p>app@example.com</p> <p>+61 2 3456 7890</p>
<p>If the primary contact is unavailable during an incident, please provide escalation contacts and timeline in the preferred communication order.</p>	<p>1. After 10 minutes, if no response from Primary Contact, engage:</p> <p>John Smith - Application Supervisor</p>

Question	Example Response
	<p>john.smith@example.com</p> <p>+61 2 3456 7890</p> <p>2. After 10 minutes, if no response from John Smith, contact:</p> <p>Jane Smith - Operations Manager</p> <p>jane.smith@example.com</p> <p>+61 2 3456 7890</p>
<p>AWS communicates updates through the support case at regular intervals throughout the incident. Are there additional contacts that should receive these updates?</p>	<p>john.smith@example.com, jane.smith@example.com</p>

Alarm matrix

Alarm Matrix

Provide the following information to identify the set of alarms that will engage AWS Incident Detection and Response to create incidents on behalf of your workload. Once engineers from AWS Incident Detection and Response have reviewed your alarms additional onboarding steps will be delivered.

AWS Incident Detection and Response Critical Alarm Criteria:

- AWS Incident Detection and Response alarms should only enter "Alarm" state upon significant business impact to the monitored workload (loss of revenue/degraded customer experience) that requires immediate operator attention.
- AWS Incident Detection and Response alarms must also engage your resolvers for the workload at the same time or prior to engagement. AWS Incident Managers collaborate with your resolvers in the mitigation process, and do not serve as a first-line responders who then escalate to you.
- AWS Incident Detection and Response alarm thresholds must be set to an appropriate threshold and duration so that any time an alarm fires an investigation must take place. If an alarm is

moving between the "Alarm" and "OK" state, sufficient impact is occurring to warrant operator response and attention.

AWS Incident Detection and Response Policy for Criteria Violations:

These criteria can only be evaluated on a case-by-case basis as events occur. The Incident Management team works with your technical account managers (TAMs) to adjust alarms and in rare cases disable monitoring if it is suspected that customer alarms do not adhere to this criteria and is engaging the Incident Management team unnecessarily at a regular rate.

Important

Provide a group distribution email addresses when supplying contact addresses, so that you can control recipient additions and deletions without runbook updates.

Provide the contact phone number for your site reliability engineering (SRE) team if you would like the AWS Incident Detection and Response team to call them after sending an initial engagement email.

Alarm Matrix table

Metric name / ARN / Threshold	Description	Notes	Actions requested
Workload volume / <i>CW Alarm ARN</i> / CallCount < 100000 for 5 datapoints within 5 minute , treat missing data as missing	This metric represents the number of incoming requests coming to the workload, measured at the Application Load Balancer level. This alarm is important because significant drops in incoming requests may indicate issues with upstream	The alarm has entered the "Alarm" state 10 times in the last week. This alarm is at risk of false positives. Threshold review is planned. Issues? No or Yes (if No, leave blank): This alarm flips frequently during a particular batch job execution.	Engage the Site Reliability Engineering team by sending an email to SRE@xyz.com Create an AWS Premium Support case for our ELB, and Route 53 services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the

Metric name / ARN / Threshold	Description	Notes	Actions requested
	network connectivity, or issues with our DNS implementation that result in users not being able to access the workload.	Resolvers: Site Reliability Engineers	<i>XYZ</i> Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)
Workload Request Latency / <i>CW Alarm ARN /</i> p90 Latency > 100ms for 5 datapoints within 5 minutes , treat missing data as missing	This metric represents the p90 latency for HTTP requests to be fulfilled by the workload. This alarm represents latency (important measure of customer experience for the website).	The alarm has entered the "Alarm" state 0 times in the last week. Issues? No or Yes (if No, leave blank): This alarm flips frequently during a particular batch job execution. Resolvers: Site Reliability Engineers	Engage the Site Reliability Engineering team by sending an email to <i>SRE@xyz.com</i> Create an AWS Premium Support case for our ECW, and RDS services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the <i>XYZ</i> Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)

Metric name / ARN / Threshold	Description	Notes	Actions requested
<p>Workload Request Availability /</p> <p><i>CW Alarm ARN /</i></p> <p>Availability < 95% for 5 datapoints within 5 minutes , treat missing data as missing.</p>	<p>This metric represents the availability for HTTP requests to be fulfilled by the workload. (# of HTTP 200 / # of Requests) per period.</p> <p>This alarm represents the availability of the workload.</p>	<p>The alarm has entered the "Alarm" state 0 times in the last week.</p> <p>Issues? No or Yes (if No, leave blank): This alarm flips frequently during a particular batch job execution.</p> <p>Resolvers: Site Reliability Engineers</p>	<p>Engage the Site Reliability Engineering team by sending an email to <i>SRE@xyz.com</i></p> <p>Create an AWS Premium Support case for our ELB, and Route 53 services.</p> <p>If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the <i>XYZ</i> Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)</p>

New Relic Alarm Example

Metric name / ARN / Threshold	Description	Notes	Actions requested
<p>End to End Integration test /</p> <p><i>CW Alarm ARN /</i></p> <p>3% failure rate for 1 minute metrics over 3 minutes duration , treat missing data as missing</p> <p>Workload Identifier: End to End Test Workflow, AWS Region: US-EAST-1 , AWS Account ID: 012345678910</p>	<p>This metric tests if a request can traverse each layer of the workload. If this test fails, it represents a critical failure to process business transactions.</p> <p>This alarm represents the ability to process business transactions for the workload.</p>	<p>The alarm has entered the "Alarm" state 0 times in the last week.</p> <p>Issues? No or Yes (if No, leave blank): This alarm flips frequently during a particular batch job execution.</p> <p>Resolvers: Site Reliability Engineers</p>	<p>Engage the Site Reliability Engineering team by sending an email to <i>SRE@xyz.com</i></p> <p>Create an AWS Premium Support case for our ECS, and DynamoDB services.</p> <p>If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the <i>XYZ</i> Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)</p>

Request changes to an onboarded workload

To request changes to an onboarded workload, complete the following steps to create a support case with AWS Incident Detection and Response.

1. Go to the [AWS Support Center](#), and then select **Create case**, as shown in the following example:
2. Choose **Technical**.
3. For **Service**, choose **Incident Detection and Response**.
4. For **Category**, choose **Workload change request**.

5. For **Severity**, choose **General Guidance**.
6. Enter a **Subject** for this change. For example:

AWS Incident Detection and Response - *workload_name*

7. Enter a **Description** for this change. For example, enter "This request is for changes to an existing workload onboarded into AWS Incident Detection and Response". Make sure that you include the following information in your request:
 - **Workload name:** Your workload name.
 - **Account ID(s):** ID1, ID2, ID3, and so on.
 - **Change details:** Enter the details for your requested change.
8. In the **Additional contacts - optional** section, enter any email IDs that you want to receive correspondence about this change.

The following is an example of the **Additional contacts - optional** section.

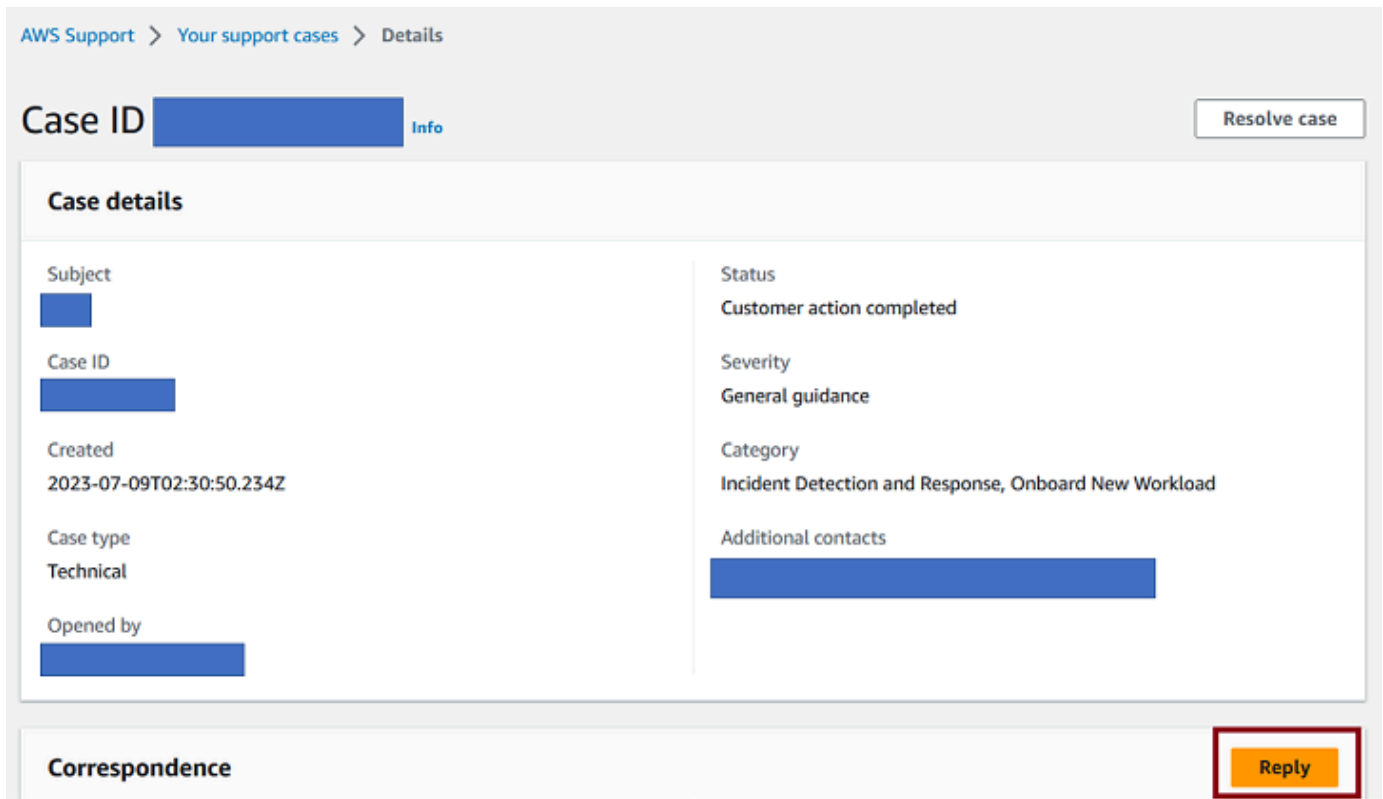
The screenshot shows a web form titled "Solve now or contact us". At the top right, there is a green status message: "Case draft saved". Below the title, there are two tabs: "Solve now" and "Contact us", with the "Contact us" tab selected and highlighted with a red box. Under the "Contact us" tab, there is a "Preferred contact language" dropdown menu set to "English". Below that is a "Web" contact method button with a radio button and the text "We'll get back to you within 24 hours." The "Additional contacts - optional" section is highlighted with a red box and contains the text "Additional contacts - optional Info" and "We will copy email notifications for this case to the following email addresses." Below this is a text input field labeled "Email addresses". At the bottom of the form, there are three buttons: "Cancel", "Previous", and "Submit", with the "Submit" button highlighted with a red box.

⚠ Important

Failure to add email IDs in the **Additional contacts - optional** section might delay the change process.

9. Choose **Submit**.

After you submit the change request, you can add additional emails from your organization. To add emails, choose **Reply** in **Case details**, as shown in the following example:



The screenshot displays the AWS Support console interface for a case. At the top, the breadcrumb navigation shows 'AWS Support > Your support cases > Details'. Below this, the 'Case ID' is shown as a redacted blue box with an 'Info' link to its right. A 'Resolve case' button is located in the top right corner. The main content area is titled 'Case details' and is divided into two columns. The left column contains the following fields: 'Subject' (redacted), 'Case ID' (redacted), 'Created' (2023-07-09T02:30:50.234Z), 'Case type' (Technical), and 'Opened by' (redacted). The right column contains: 'Status' (Customer action completed), 'Severity' (General guidance), 'Category' (Incident Detection and Response, Onboard New Workload), and 'Additional contacts' (redacted). At the bottom of the console, there is a 'Correspondence' section with a red-bordered 'Reply' button highlighted in orange.

Then, add the email IDs in the **Additional contacts - optional** section.

The following is an example of the **Reply** page showing where you can enter additional emails.


Reply

Do not share any sensitive information in case correspondences, such as credentials, credit cards, signed URLs, or personally identifiable information. Find more information [here](#) .

Adding additional email IDs

Maximum 8000 characters (8000 remaining)

Attachments

 Choose files

Up to 3 attachments, each less than 5MB.

Contact methods [Info](#)

Web 

We'll respond by email and Support Center.

Additional contacts - optional [Info](#)

When we contact you via email, we will copy the correspondence to the following email addresses

Use commas or semicolons to separate email addresses - Maximum 10 email addresses (8 remaining) or 254 characters (213 remaining)

Offboard a workload

To offboard a workload from AWS Incident Detection and Response, create a new support case for each workload. When you create the support case, keep the following in mind:

- To offboard a workload that's in a single AWS account, create the support case either from the workload's account or from your payer account.
- To offboard a workload that spans multiple AWS accounts, then create the support case from your **payer account**. In the body of the support case, list all account IDs to offboard.

Important

If you create a support case to offboard a workload from the incorrect account, then you might experience delays and requests for additional information before your workloads can be offloaded.

Request to offboard a workload

1. Go to the [AWS Support Center](#), and then select **Create case**.
2. Choose **Technical**.
3. For **Service**, choose **Incident Detection and Response**.
4. For **Category**, choose **Workload Offboarding**.
5. For **Severity**, choose **General Guidance**.
6. Enter a **Subject** for this change. For example:

[Offboard] AWS Incident Detection and Response - *workload_name*

7. Enter a **Description** for this change. For example, enter "This request is for offboarding an existing workload onboarded into AWS Incident Detection and Response". Make sure that you include the following information in your request:
 - **Workload name:** Your workload name.
 - **Account ID(s):** ID1, ID2, ID3, and so on.
 - **Reason for offboarding:** Provide a reason for offboarding the workload.
8. In the **Additional contacts - optional** section, enter any email IDs that you want to receive correspondence about this offboarding request.
9. Choose **Submit**.

AWS Incident Detection and Response monitoring and observability

AWS Incident Detection and Response offers you expert guidance on defining observability across your workloads from the application layer to the underlying infrastructure. Monitoring tells you that something is wrong. Observability uses data collection to tell you what is wrong and why it happened.

The Incident Detection and Response system monitors your AWS workloads for failures and performance degradation by leveraging native AWS services such as Amazon CloudWatch and Amazon EventBridge to detect events that may impact your workload. Monitoring provides you notification of imminent, on-going, receding, or potential failures or of performance degradation. When you onboard your account to Incident Detection and Response, you select which alarms in your account should be monitored by the Incident Detection and Response monitoring system and you associate those alarms with an application and a runbook used during incident management.

Incident Detection and Response uses Amazon CloudWatch and other AWS services to build your observability solution. AWS Incident Detection and Response helps you with observability in two ways:

- **Business Outcome metrics:** Observability on AWS Incident Detection and Response starts with defining the key metrics that monitor the outcomes of your workloads or end-user experience. AWS experts work with you to understand the objectives of your workload, the key outputs or factors that may impact user-experience, and to define the metrics and alerts that capture any degradation in those key metrics. For example a key business metric for a mobile calling application is the *Call Setup Success Rate* (monitors the success rate of user call attempts), and a key metric for a website is *page speed*. Incident engagement is triggered based on business outcome metrics.
- **Infrastructure level metrics:** At this stage, we identify the underlying AWS services and infrastructure supporting your application and define metrics and alarms to track the performance of these infrastructure services. These may include metrics such as `ApplicationLoadBalancerErrorCount` for Application Load Balancer instances. This starts after the workload has been onboarded and monitoring set up.

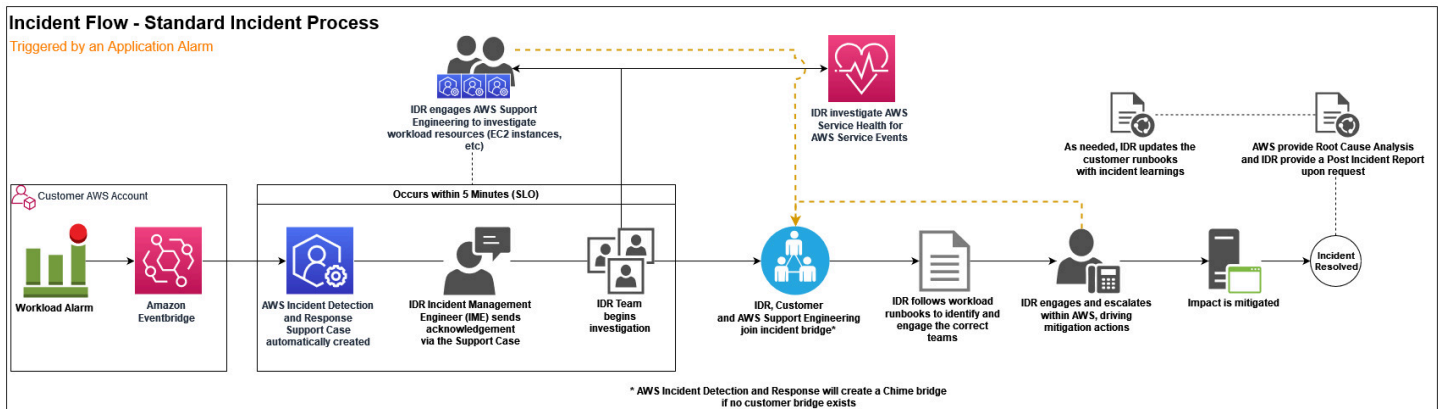
Implementing observability on AWS Incident Detection and Response

Because observability is a continuous process that may not be completed in one exercise or time frame, AWS Incident Detection and Response implements observability in two phases:

- **Onboarding phase:** Observability during onboarding is focused on detecting when the business outcomes of your application are impaired. To this end, observability during the onboarding phase is focused on defining the key business outcome metrics at the application layer to notify AWS of disruptions to your workloads. This way AWS can promptly respond to these disruption and provide you help toward recovery.
- **Post Onboarding phase:** AWS Incident Detection and Response offers a number of proactive services for observability including the definition of infrastructure level metrics, metric tuning, and setting up traces and logs depending, on the maturity level of the customer. The implementation of these services may span several months and involve multiple teams. AWS Incident Detection and Response provides guidance on observability setup and customers are required to implement the required changes in their workload environment. For help with hands-on implementation of observability features, raise a request to your technical account managers (TAMs).

Incident management with AWS Incident Detection and Response

AWS Incident Detection and Response offers you 24x7 proactive monitoring and incident management delivered by a designated team of incident managers.



- 1. Alarm Generation:** Alarms triggered on your workloads are pushed through Amazon EventBridge to AWS Incident Detection and Response. AWS Incident Detection and Response automatically pulls up the runbook associated with your alarm and notifies an incident manager. If a critical incident occurs on your workload that isn't detected by alarms monitored by AWS Incident Detection and Response, then you can create a support case to request an Incident Response. For more information on requesting an Incident Response, see [Incident Response Request](#).
- 2. AWS Incident Manager Engagement:** The incident manager responds to the alarm and engages you on a conference call or as otherwise specified in the runbook. The incident manager verifies the health of the AWS services to determine if the alarm is related to issues with AWS services used by the workload and advises on the status of the underlying services. If required, the incident manager then creates a case on your behalf and engages the right AWS experts for support.

Because AWS Incident Detection and Response monitors AWS services specifically for your applications, AWS Incident Detection and Response might determine that the incident is related to an AWS service issue even before an AWS service event is declared. In this scenario, the incident manager advises you on the status of the AWS service, triggers the AWS Service Event Incident Management flow, and follows up with the service team on resolution. The information provided gives you the opportunity to implement your recovery plans or workarounds early to

mitigate the impact of the AWS Service Event. For more information, see [Incident management for service events](#).

- 3. Incident Resolution:** The incident manager coordinates the incident across the required AWS teams and makes sure that you remain engaged with the right AWS experts until the incident is mitigated or resolved.
- 4. Post Incident Review (if requested):** After an incident, AWS Incident Detection and Response can perform a post incident review at your request and generate a Post Incident Report. The Post Incident Report includes a description of the issue, the impact, which teams were engaged, and workarounds or actions taken to mitigate or resolve the incident. The Post Incident Report might contain information that can be used to reduce the likelihood of incident recurrence, or to improve the management of a future occurrence of a similar incident. The Post Incident Report isn't a Root Cause Analysis (RCA). You can request a RCA in addition to the Post Incident Report. An example of a Post Incident Report is provided in the following section.

Important

The following report template is an example only.

Post ** Incident ** Report ** Template

Post Incident Report - 0000000123

Customer: Example Customer

AWS Support case ID(s): 0000000000

Customer internal case ID (if provided): 1234567890

Incident start: 2023-02-04T03:25:00 UTC

Incident resolved: 2023-02-04T04:27:00 UTC

Total Incident time: 1:02:00 s

Source Alarm ARN: arn:aws:cloudwatch:us-east-1:000000000000:alarm:alarm-prod-workload-impaired-useast1-P95

Problem Statement:

Outlines impact to end users and operational infrastructure impact.

Starting at 2023-02-04T03:25:00 UTC, the customer experienced a large scale outage of their workload that lasted one hour and two minutes and spanning across all Availability Zones where the application is deployed. During impact, end users were unable to connect to the workload's Application Load Balancers (ALBs) which service inbound communications to the application.

Incident Summary:

Summary of the incident in chronological order and steps taken by AWS Incident Managers to direct the incident to a path to mitigation.

At 2023-02-04T03:25:00 UTC, the workload impairments alarm triggered a critical incident for the workload. AWS Incident Detection and Response Incident Managers responded to the alarm, checking AWS service health and steps outlined in the workload's runbook.

At 2023-02-04T03:28:00 UTC, ** per the runbook, the alarm had not recovered and the Incident Management team sent the engagement email to the customer's Site Reliability Team (SRE) team, created a troubleshooting bridge, and an AWS Support support case on behalf of the customer.

At 2023-02-04T03:32:00 UTC, ** the customer's SRE team, and AWS Support Engineering joined the bridge. The Incident Manager confirmed there was no on-going AWS impact to services the workload depends on. The investigation shifted to the specific resources in the customer account.

At 2023-02-04T03:45:00 UTC, the Cloud Support Engineer discovered a sudden increase in traffic volume was causing a drop in connections. The customer confirmed this ALB was a newly provisioned to handle an increase in workload traffic for an on-going promotional event.

At 2023-02-04T03:56:00 UTC, the customer instituted back off and retry logic. The Incident Manager worked with the Cloud Support Engineer to raise an escalation a higher support level to quickly scale the ALB per the runbook.

At 2023-02-04T04:05:00 UTC, ALB support team initiates scaling activities. The back-off/retry logic yields mild recovery but timeouts are still being seen for some clients.

By 2023-02-04T04:15:00 UTC, scaling activities complete and metrics/alarms return to pre-incident levels. Connection timeouts subside.

At 2023-02-04T04:27:00 UTC, per the runbook the call was spun down, after 10 minutes of recovery monitoring. Full mitigation is agreed upon between AWS and the customer.

Mitigation:

Describes what was done to mitigate the issue. NOTE: this is not an Root Cause Analysis (RCA).

Back-off and retries yielded mild recovery. Full mitigation happened after escalation to ALB support team (per runbook) to scale the newly provisioned ALB.

Follow up action items (if any):

Action items to be reviewed with your Technical Account Manager (TAM), if required. Review alarm thresholds to engage AWS Incident Detection and Response closer to the time of impact.

Work with AWS Support and TAM team to ensure newly created ALBs are pre-scaled to accommodate expected spikes in workload traffic.

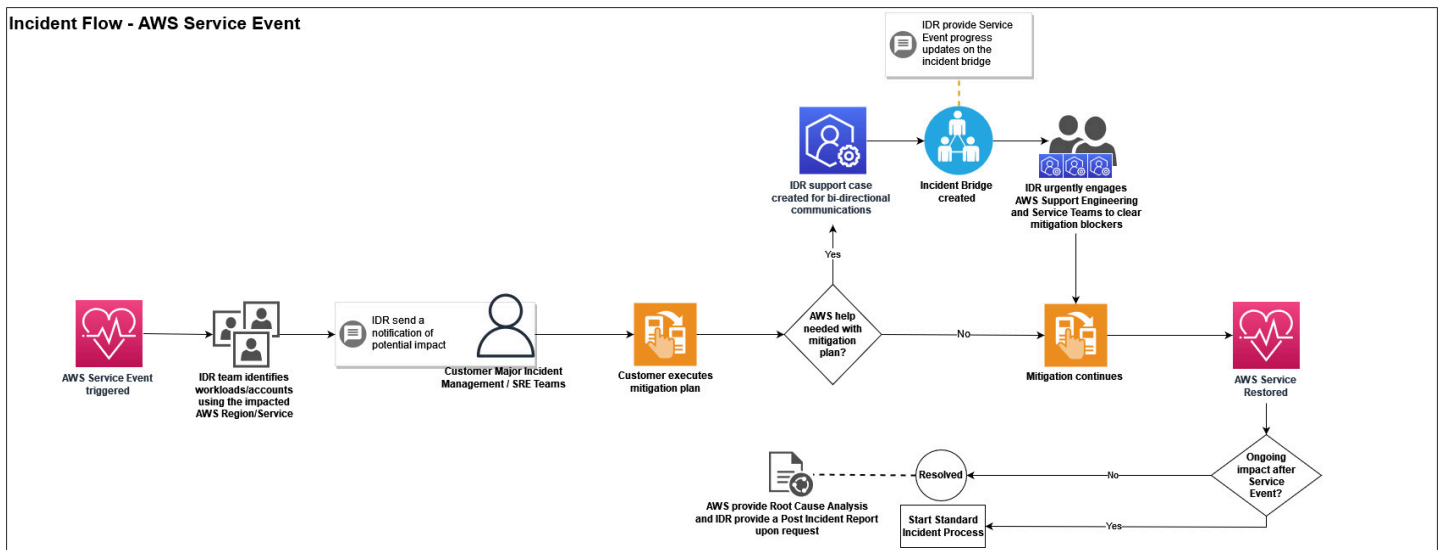
Provision access for application teams

AWS Incident Detection and Response communicates with you through AWS Support cases during the lifecycle of an incident. To correspond with Incident Managers, your teams must have access to the AWS Support Center.

For more information on provisioning access, see [Manage access to AWS Support Center](#) in the *AWS Support User Guide*.

Incident management for service events

AWS Incident Detection and Response notifies you of an ongoing service event in your AWS Regions, whether or not your workload is impacted. During an AWS service event, AWS Incident Detection and Response creates an AWS Support case, joins your conference call bridge to receive feedback on impact and sentiment, and provides guidance to invoke your recovery plans during the event. You also receive a notification through AWS Health containing details of the event. Customers who aren't affected by the AWS owned service event (for example, operating in a different AWS Region, don't use the AWS service that's impaired, and so on) continue to be supported by the standard engagement. For more information about AWS Health, see [What is AWS Health?](#)



Post Incident Report for Service Events (if requested): If a service event causes an incident, then you can request AWS Incident Detection and Response to perform a post incident review and generate a Post Incident Report. The Post Incident Report for service events includes the following:

- A description of the issue

- The incident's impact
- Information shared on the AWS Health dashboard
- The teams that were engaged during the incident
- Workarounds and actions taken to mitigate or resolve the incident

The Post Incident Report for service events might contain information that can be used to reduce the likelihood of incident recurrence, or to improve the management of a future occurrence of a similar incident. The Post Incident Report for service events isn't a Root Cause Analysis (RCA). You can request a RCA in addition to the Post Incident Report for service events.

The following is an example of a Post Incident Report for service event:

 **Note**

The following report template is an example only.

Post Incident Report - LSE000123

Customer: Example Customer

AWS Support Case ID(s): 0000000000

Incident Start: Example: 1 January 2024, 3:30 PM UTC

Incident Resolved: Example: 1 January 2024, 3:30 PM UTC

Incident Duration: 1:02:00

Service(s) Impacted: Lists the impacted services such as EC2, ALB

Region(s): Lists the impacted AWS Regions, such as US-EAST-1

Alarm Identifiers: Lists any customer alarms that triggered during the Service Level Event

Problem Statement:

Outlines impact to end users and operational infrastructure impact during the Service Level Event.

Starting at 2023-02-04T03:25:00 UTC, the customer experienced a service outage...

Impact Summary for Service Level Event:

(This section is limited to approved messaging available on the AWS Health Dashboard)
Outline approved customer messaging as provided on the AWS Health Dashboard.

Between 1:14 PM and 4:33 PM UTC, we experienced increased error rates for the Amazon SNS Publish, Subscribe, Unsubscribe, Create Topic, and Delete Topic APIs in the EU-WEST-1 Region. The issue has been resolved and the service is operating normally.

Incident Summary:

Summary of the incident in chronological order and steps taken by AWS Incident Managers during the Service Level Event to direct the incident to a path to mitigation.

At 2024-01-04T01:25:00 UTC, the workload alarm triggered a critical incident...

At 2024-01-04T01:27:00 UTC, customer was notified via case 000000000 about the triggered alarm

At 2024-01-04T01:30:00 UTC, IDR team identified an ongoing service event which was related to the customer triggered alarm

At 2024-01-04T01:32:00 UTC, IDR team sent an impact case correspondence requesting for the incident bridge details

At 2024-01-04T01:32:00 UTC, customer provided the incident bridge details

At 2024-01-04T01:32:00 UTC, IDR team joined the incident bridge and provided information about the ongoing service outage

By 2024-01-04T02:35:00 UTC, customer failed over to the secondary region (EU-WEST-1) to mitigate impact...

At 2024-01-04T03:27:00 UTC, customer confirmed recovery, the call was spun down...

Mitigation:

Describes what was done to mitigate the issue. NOTE: this is not an Root Cause Analysis (RCA).

Back-off and retries yielded mild recovery. Full mitigation happened ...

Follow up action items (if any):

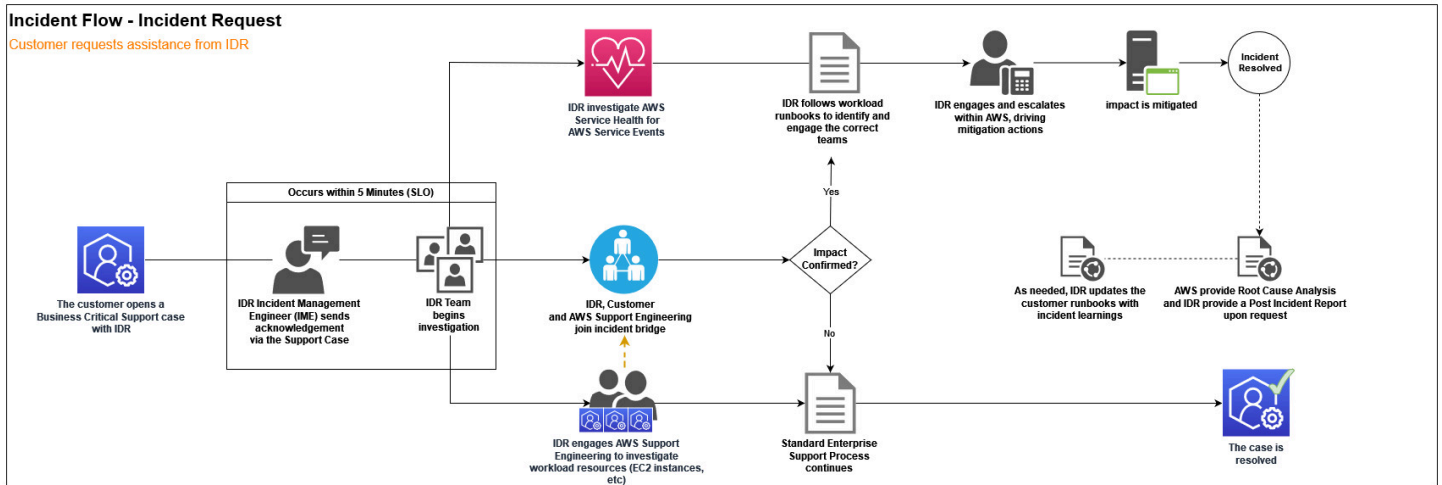
Action items to be reviewed with your Technical Account Manager (TAM), if required.

Review alarm thresholds to engage AWS Incident Detection and Response closer ...

Work with AWS Support and TAM team to ensure ...

Incident Response Request

If a critical incident occurs on your workload that isn't detected by alarms monitored by AWS Incident Detection and Response, then you can create a support case to request an Incident Response. You can request an Incident Response for any workload that's subscribed to AWS Incident Detection and Response, including workloads in the process of onboarding.



To request an Incident Response for an incident that's actively impacting your workload, create an AWS Support Case. After the support case is raised, AWS Incident Detection and Response will engage you on a conference bridge with the AWS experts required to accelerate the recovery of your workload.

Request an Incident Response using the AWS Support Center Console

1. Open the [AWS Support Center Console](#), and then choose **Create case**.
2. Choose **Technical**.
3. For **Service**, choose **Incident Detection and Response**.
4. For **Category**, choose **Active Incident**.
5. For **Severity**, choose **Business-critical system down**.
6. Enter a **Subject** for this incident. For example:

AWS Incident Detection and Response - Active Incident - workload_name

7. Enter the **Problem Description** for this incident. Add the following details:

- **Technical Information:**

Affected Service(s):

Affected Resource(s):

Affected Region(s):

Workload Name:

- **Business Information:**

Description of impact to the business:

[Optional] Customer Bridge Details:

8. In the **Additional contacts** section, enter any email addresses that you want to receive correspondences about this incident.

The following illustration shows the console screen with the **Additional contacts** field highlighted.

Solve now or contact us ✔ Case draft saved

Solve now | **Contact us**

Preferred contact language
English

Web
We'll get back to you within 15 minutes.

Additional contacts - optional [Info](#)
We will copy email notifications for this case to the following email addresses.

Use commas (,) or semicolons (;) to separate email addresses. You can enter up to 10 email addresses (10 remaining) or 254 characters (254 remaining).

Cancel Previous **Submit**

9. Choose **Submit**.

After submitting an Incident Response request, you can add additional email addresses from your organization. To add additional addresses, reply to the case, and then add the email addresses in the **Additional contacts** section.

The following illustration shows the **Case details** screen with the **Reply** button highlighted.

[AWS Support](#) > [Your support cases](#) > Details

Case ID 12345678919 Info

[Resolve case](#)

Case details

Subject AWS Incident Detection and Response - Active Incident - [TEST CASE]	Status Unassigned
Case ID 12345678919	Severity Business-critical system down
Created YYYY-MM-DDTHH:MM:SS	Category Incident Detection and Response, Active Incident
Case type Technical	Language English
Opened by xyz@email.com	Additional contacts 123@email.com;abc@email.com

Correspondence

[Reply](#)

The following illustration shows the case Reply with the **Additional contacts** field and **Submit** button highlighted.

Reply

Do not share any sensitive information in case correspondences, such as credentials, credit cards, signed URLs, or personally identifiable information. Find more information [here](#).

Adding additional email IDs

Maximum 8000 characters (8000 remaining)

Attachments

Up to 3 attachments, each less than 5MB.

Contact methods [Info](#)

Web

We'll respond by email and Support Center.

Additional contacts - optional [Info](#)

When we contact you via email, we will copy the correspondence to the following email addresses

Use commas or semicolons to separate email addresses - Maximum 10 email addresses (10 remaining) or 254 characters (254 remaining)

10AWS Incident Detection and Response acknowledges your case within five minutes and engages you on a conference bridge with the appropriate AWS experts.

Request an Incident Response using the AWS Support API

Support cases can be programmatically created by using the [AWS Support API](#).

AWS Incident Detection and Response reporting

Incident Detection and Response provides operational and performance data to help you understand how the service is configured, the history of your incidents, and the performance of the Incident Detection and Response service.

Configuration data

- All accounts onboarded
- Names of all applications
- The alarms, runbooks, and support profiles associated with each application

Incident data

- The dates, number and duration of incidents for each application
- The dates, number and duration of incidents associated with a specific alarm
- Post Incident Report

Performance data

- Service Level Objective (SLO) performance

Reach out to your technical account manager for operational and performance data you may need.

Incident Detection and Response security and resiliency

The AWS [Share Responsibility Model](#) applies to data protection in AWS Support. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that you use.

For more information about data privacy, see the [Data Privacy FAQ](#).

For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the AWS Security Blog.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates to communicate with AWS resources. We recommend TLS 1.2 or later. For information, see [What Is An SSL/TLS Certificate?](#).
- Set up API and user activity logging with AWS CloudTrail. For information, see [AWS CloudTrail](#).
- Use AWS encryption solutions, along with all default security controls within AWS services. For information, see [AWS cryptographic services and tools](#).
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3. For information about Amazon Macie, see [Amazon Macie](#).
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form fields such as a **Name** field. This includes when you work with AWS Support or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form fields used for names may be used for billing or

diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

AWS Incident Detection and Response access to your accounts

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources.

AWS Incident Detection and Response and your alarm data

By default, Incident Detection and Response receives the Amazon resource name (ARN) and state of every CloudWatch alarm in your account and then starts the incident detection and response process when your onboarded alarm changes into the ALARM state. If you would like to customize what information incident detection and response receives about alarms from your account, contact your Technical Account Manager.

Document history

The following table describes the important changes to the documentation since the last release of the IDR guide.

- **Latest documentation update:** June 12, 2024

Change	Description	Date
Updated Account subscription	Updated the Account subscription section to include details on where to open a support case when you request to subscribe an account. Updated section: Account subscription	June 12, 2024
Post Incident Report for service events now available	Updated the Incident management for service events section to include information about the Post Incident Report for service events. Updated section: Incident management for service events	May 8, 2024
Added a new section: Offboard a workload	Added the Offload a workload section in Getting started to include information about offboarding workloads For more information, see Offboard a workload .	March 28, 2024
Updated Account subscription	Updated the Account subscription section to include information about offboarding workloads For more information, see Account subscription	March 28, 2024

Change	Description	Date
Updated Testing	<p>Updated the Testing section to include information on gameday testing as the last step in the onboarding process.</p> <p>Updated section: Test onboarded workloads</p>	February 29, 2024
Updated What is AWS Incident Detection and Response	<p>Updated the What is AWS Incident Detection and Response section.</p> <p>Updated section: What is AWS Incident Detection and Response?</p>	February 19, 2024
Updated Questionnaire section	<p>Updated the Workload onboarding questionnaire and added Alarm ingestion questionnaire. Renamed the section from Onboarding questionnaire to Workload onboarding and Alarm ingestion questionnaires.</p> <p>Updated section: Workload onboarding and alarm ingestion questionnaires</p>	February 2, 2024
Updated AWS Service Event and onboarding information	<p>Updated several sections with new information for onboarding.</p> <p>Updated sections:</p> <ul style="list-style-type: none"> • Incident management for service events • Workload discovery • Onboarding • Account subscription <p>New sections</p> <ul style="list-style-type: none"> • Provision access for application teams 	January 31, 2024

Change	Description	Date
Added a Related information section	Added a Related information section in Access provisioning . Updated section: Provision access for alert ingestion to Incident Detection and Response	January 17, 2024
Updated example steps	Updated the procedure for steps 2,3, and 4 in Example: Integrating notifications from Datadog and Splunk . Updated section: Example: Integrate notifications from Datadog and Splunk	December 21, 2023
Updated introduction graphic and text	Updated graphic in Ingest alarms from APMs that have direct integration with Amazon EventBridge . Updated section: Develop runbooks for AWS Incident Detection and Response	December 21, 2023
Updated runbook template	Updated the runbook template in Developing runbooks for AWS Incident Detection and Response . Updated section: Develop runbooks for AWS Incident Detection and Response	December 4, 2023

Change	Description	Date
Updated Alarm Configurations	<p>Updated Alarm Configurations with detailed information on CloudWatch alarm configuration.</p> <p>New section: Create CloudWatch alarms that fit your business needs in Incident Detection and Response</p> <p>New section: Use AWS CloudFormation templates to build CloudWatch alarms in Incident Detection and Response</p> <p>New section: Example uses cases for CloudWatch alarms in Incident Detection and Response</p>	September 28, 2023
Updated Getting Started	<p>Updated Getting Started with information on Workload change requests.</p> <p>New section: Request changes to an onboarded workload</p> <p>Updated section: Account subscription</p>	September 05, 2023
New section in Getting Started	Added Ingest alerts into AWS Incident Detection and Response Ingesting alerts into AWS Incident Detection and Response.	June 30 2023
Original document	AWS Incident Detection and Response first published	March 15 2023

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.