AWS Incident Detection and Response User Guide

AWS Incident Detection and Response Concepts and Procedures

Version March 15, 2023

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

Table of Contents

What is AWS Incident Detection and Response?

AWS Incident Detection and Response is an add-on to AWS Enterprise Support that offers you proactive monitoring and incident management for your selected workloads. AWS Incident Detection and Response is designed to improve your operations, increase workload resiliency, and accelerate your recovery from critical incidents. With AWS Incident Detection and Response, AWS incident managers monitor your workloads and engage you on a call bridge within minutes of a critical alarm to guide you through recovery. AWS Incident Detection and Response offers the following benefits:

- **Improved visibility**: AWS works with you to define critical metrics and alarms to provide improved visibility into the application and infrastructure layers of your workloads.
- **Early issue detection**: AWS Incident Managers monitor your onboarded workloads 24x7 to detect critical incidents.
- **Faster resolution**: Recover faster from disruptions through rapid engagement with AWS experts using jointly defined response plans and runbooks.
- **Reduced potential for failure**: Proactively mitigate issues by improving the architecture and operations of your workloads with best practice guidance from AWS.

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.

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:

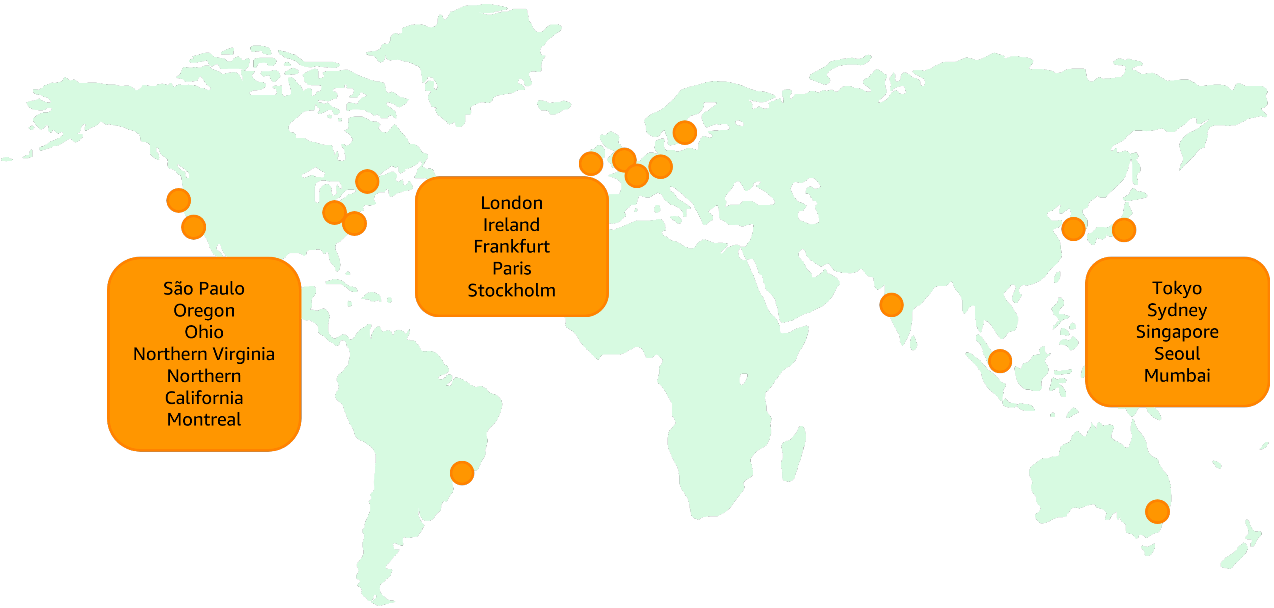

AWS Incident Detection and Response RACI

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
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		

Activity	Customer	Incident Detection and Response
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		
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.

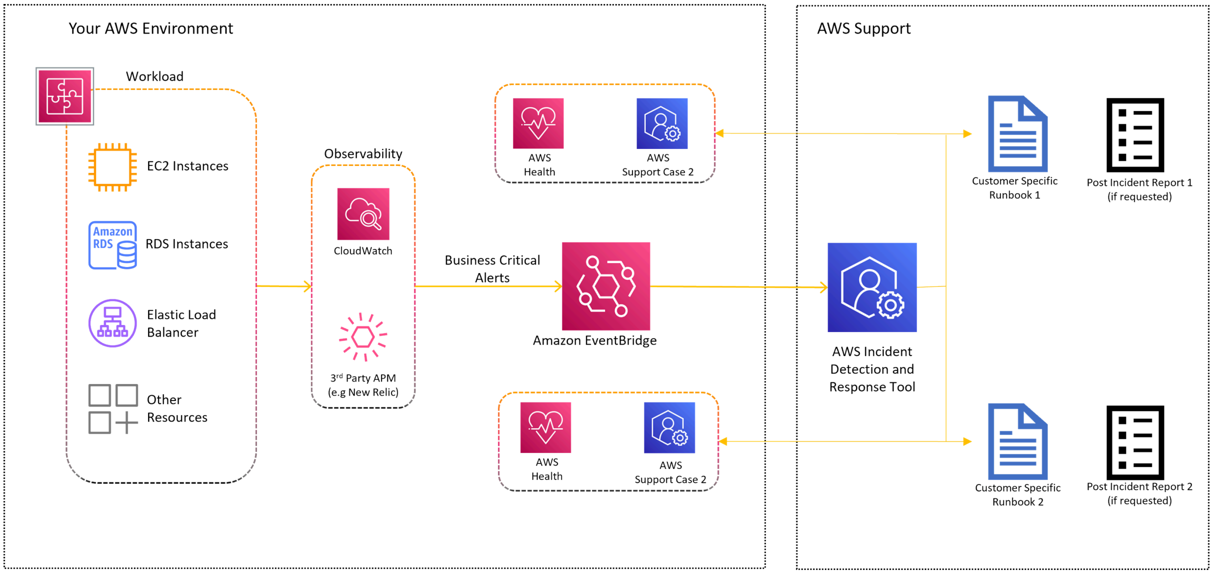

Getting Started with AWS Incident Detection and Response

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

For example, a monolithic application can be hosted in a single account (for example, Employee Performance App in Fig.1) whereas another application (for example, Storefront Webapp in Fig. 1) broken into microservices may stretch across different accounts. A workload may also share resources such as a database with other applications/workloads as shown in Fig. 1.

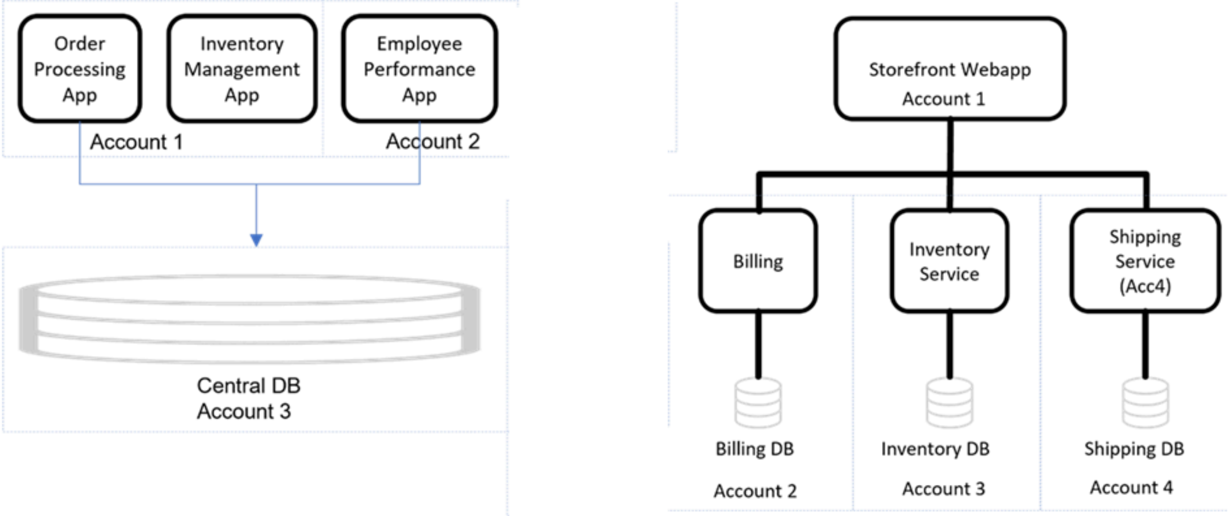

Note
You can make changes to your runbooks, workload information or the alarms monitored on AWS Incident Detection and Response by working with you technical account manager (TAM). Your TAM collects all the required information and works with the incident management team to make the required updates on AWS Incident Detection and Response.

AWS Incident Detection and Response onboarding a workload process

To onboard a workload, the constituent accounts of the workload must be subscribed to the service before the workload can be onboarded for monitoring. The process for onboarding a workload to Incident Detection and Response is shown in the following graphic.

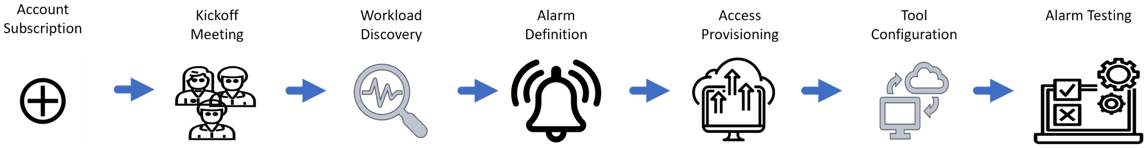

The steps required for onboarding a workload to Incident Detection and Response are shown in the following table. While the table shows the duration of each task, the actual dates for each task will be defined based on the availability of your resources and schedule.

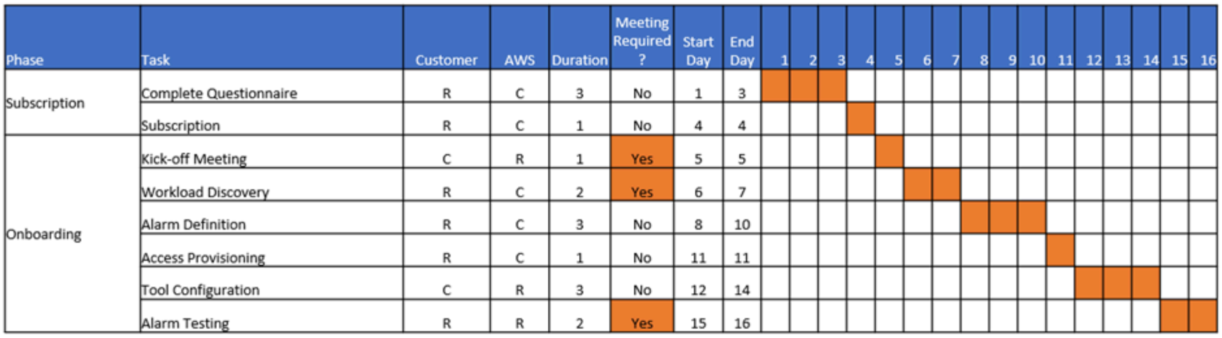

Phase	Task	Customer	AWS	Duration	Meeting Required ?	Start Day	End Day	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Subscription	Complete Questionnaire	R	C	3	No	1	3	🟧	🟧	🟧													
	Subscription	R	C	1	No	4	4				🟧												
Onboarding	Kick-off Meeting	C	R	1	Yes	5	5					🟧											
	Workload Discovery	R	C	2	Yes	6	7						🟧	🟧									
	Alarm Definition	R	C	3	No	8	10								🟧	🟧	🟧						
	Access Provisioning	R	C	1	No	11	11											🟧					
	Tool Configuration	C	R	3	No	12	14												🟧	🟧	🟧		
	Alarm Testing	R	R	2	Yes	15	16															🟧	🟧

AWS Incident Detection and Response Account subscription

You subscribe an account to the service by raising a billing support case from a payer account and listing the individual linked accounts you wish to subscribe. To subscribe an account, go to the AWS Support Center and click **Create case** as shown. Only accounts on Enterprise Support can be subscribed to the service.

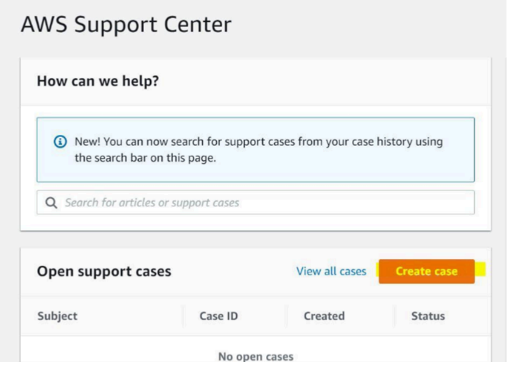

AWS Support Center

How can we help?

New! You can now search for support cases from your case history using the search bar on this page.

Open support cases

Subject	Case ID	Created	Status

No open cases

1. Complete the support case form:
   a. Select Service = Billing
   b. Category = Other Billing Questions and
   c. Severity= General question
2. Provide the list of the accounts you wish to enroll to AWS Incident Detection and Response. The request should specify if this is meant to "add" or "remove" accounts and should include the account IDs to be added, or removed. You can provide a desired start date. Workload onboarding starts after your accounts are subscribed to the service.

AWS Incident Detection and Response kick-off meeting

Responsible: Technical account manager

Consulted: Customer

The kick off meeting is used to outline the program, entitlements, and align expectations when working with the AWS Incident Detection and Response service. AWS technical account management (TAM) teams explore your current technical architectures, monitoring and alarming configurations, and your production AWS Regions. The meeting is also used to address any concerns you may have on the workload you want onboard.

Key Outputs:

- An onboarding questionnaire is delivered to you by the TAM team.
- An onboarding schedule defining the dates and time for future engagements during the onboarding phase.

AWS Incident Detection and Response workload discovery

Responsible: Technical account manager

Consulted: Customer

During workload discovery, AWS seeks to gain as much context about your workloads as you are willing to share, including a description of the workload outcomes, workload architecture diagrams, and details of the AWS services employed in your workload. The specific details required are captured in the onboarding questionnaire.

Key Outputs:

- AWS understands your services and AWS Regions where your workloads operate.

AWS Incident Detection and Response alarm configuration

Responsible: Customer

Consulted: Technical account manager

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 should only enter "Alarm" state upon 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 details 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) # of errors >=3 for 10 datapoints (different clients) in a 5 minute window	arn:aws:cloudwatch:us-west-2:xxxxx:alarm:httperrorcode400	Ticket cut to Service team	EC2, Amazon Aurora

For more details, see AWS Incident Detection and Response monitoring and observability (p. 19).

Key Outputs:

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

AWS Incident Detection and Response access provisioning

Responsible: Customer

Consulted: Technical account manager

Consulted: AWS incident manager

In order to create the managed EventBridge rule for AWS Incident Detection and Response to receive alerts from your account, you must install the `AWSServiceRoleForHealth_EventProcessor` service-linked role (SLR). More information about this SLR, including the associated AWS managed policy, is available in the *AWS Health User Guide* Using service-linked roles.

You can install this service-linked role in your account by following the following instructions from the *IAM User Guide* Create service-linked role. You can also use the following 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.

Developing runbooks for AWS Incident Detection and Response

You can also download an example Incident Detection and Response runbook: aws-idr-runbook-example.zip.

Responsible: AWS Incident Manager

Consulted: Technical account manager

Consulted: Customer

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 as a template for an incident response runbook in
 Incident Manager using the AWS Incident Detection and Response services."
schemaVersion: "0.3"
```

```
mainSteps:
 - name: Priority
 action: aws:pause
 inputs: {}
 nextStep: Information
 description: |-
Priority actions
 This section is for defining actions which should be actioned immediately upon incident
 engagement before delving into triage and remediation actions.
 These actions are extremely time sensitive and a must action regardless of following
 steps. Remove this section if not needed.
Action Title -
  1 - Step by step...
  2 - etc
 - name: Information
 action: aws:pause
 inputs: {}
 nextStep: Triage
 description: |-
 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 (either reference an AWS
 Incident Manager escalation plan, a list of AMS Incident Manager contacts, or detailed
 instructions.
 Initial engagement - example of detailed instruction plan
  1 - Step by step
  2 - etc.
 Incident call setup - example of detailed instruction plan
  1 - Step by step
  2 - etc.
 Joining existing incident call - example of detailed instruction plan
  1 - Step by step
  2 - etc.
 Engagement Escalation - General
  Via AWS Incident Manager, engage escalation plan alpha
 Engagement Escalation - AppTeam Beta
  Via AWS Incident Manager, engage contact beta on-call
 Engagement Escalation - Security
 Important - This escalation is not for use with a security incident, see Incident
 Detection and Response security and resiliency.
 This engagement plan is for when security is team is participating in a regular incident
 event.
 Via AWS Incident Manager, engage contact security review
 ---
 Escalation instructions
 Describe actions Incident Manager engineer may take if engagement plans are unresponsive
 or insufficient.
 Use the general engagement escalation plan
 ---
 Communication plans
 Describe how Incident Manager engineer communicates with designated stakeholders outside
 the incident call and communication channels.
 Initiating communication plan - example of detailed instruction plan
  1 - Step by step
  2 - etc.
 Updates - example of detailed instruction
  1 - Step by step
  2 - etc.
 ---
```

```
Application architecture overview
This section provides an overview of the application/workload architecture for Incident
Manager 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
- name: Triage
action: aws:pause
inputs: {}
nextStep: Investigate
description: |-

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 - example
 1 - This runbook is for desc. If this incident is in regards to descdesc, in the Runbooks
tab,
   stop this runbook and start runbook title of other runbook.
 2 - Review Metrics tab and Related items tab in Incident Manager, populate with
additional reporting metrics and/or
   item references as needed.
 3 - Edit the title of the incident to be of the form desc formatting.
 4 - Evaluate desc how to review impact, change incident impact if
Impact - example description of impact requirements
 1 – Critical impact, full application failure that impacts many to all customers.
 2 – High impact, partial application failure with impact to many customers.
 3 – Medium impact, the application is providing reduced service to many customers.
 4 – Low impact, the application is providing reduced service to few customers.
 5 – No impact, customers are not currently impacted but urgent action is needed to avoid
impact.
---
Incident engagement - example
This section describes the conditions to begin an incident engagement call.
Upon conditions, execute the Initial engagement plan.
Incident communications - example
This section describes the conditions to begin the communications plan. Upon conditions,
execute the Initiating communication plan.
- name: Investigate
action: aws:pause
inputs: {}
nextStep: Mitigation
description: |-
Investigation
This section describes performing investigation of known and unknown symptoms.
Known issues
A known issue - example
 1 - Step by step instructions to determine if known issue
 2 - etc
 3 - includes which runbook to reference for a known issue in the Mitigation step
Some other known issue - example
 1 - Step by step instructions to determine if known issue
 2 - etc
 3 - includes which runbook to reference for a known issue in the Mitigation step
---
All other issues
This section provides general guidance for investigating. May include instructions as to
which additional escalation and/or engagement plans to execute based on findings.
For example, when errors are increasing for metrics alpha execute Engagement Escalation -
AppTeam Beta to add Beta team to the incident engagement.
- name: Mitigation
action: aws:pause
inputs: {}
nextStep: Recovery
description: |-
Collaborate
 Communicate any changes or important information from the previous step to the members of
the incident call.
 Engage additional contacts or teams using their escalation plan from the Contacts tab.
Implement mitigation
 Update the Timeline tab of the incident when a possible mitigation is identified.
If needed, review the mitigation with others in the associated chat channel before
proceeding.
 etc
- name: Recovery
action: aws:pause
inputs: {}
isEnd: true
description: |-
Monitor customer impact
 View the Metrics tab of the incident to monitor for recovery of your key performance
indicators (KPIs).
 Update the Impact field in the incident when customer incident has been reduced or
resolved.
Identify action items
 Add entries in the Timeline tab of the incident to record key decisions and actions
taken, including temporary mitigation that might have been implemented.
 etc.
```

AWS Incident Detection and Response testing

Consulted: AWS incident manager

Responsible: Customer

The last step in workload onboarding is alarm testing. We test to confirm that alarms generated on your workloads can be detected by AWS Incident Detection and Response and that they are linked to the appropriate runbooks and response plans. You test creating an incident by manually changing the alarm to the "Alarm" state and observing that all of the expected steps defined in your runbook happen. To learn more about CloudWatch alarms, see set-alarm-state. Testing should be done in as close to a real-world scenario as possible.

Key Outputs:

- You confirm that your runbooks are configured to your requirements.
- Incident management team confirms that a test incident is created successfully.
- Incident management team sends a Go Live confirmation email to you.
- Production monitoring begins for the workload.

AWS Incident Detection and Response questionnaire

You can also download the AWS Incident Detection and Response Questionnaire: aws-idr-questionnaire.zip.

AWS Incident Detection and Response general questions

General Questions

Question	Example Response
Enterprise Name	Amazon Inc.
What is the abbreviated and expanded name of the workload to be on-boarded?	*Workload Name*
Please describe the primary end user and the function of this workload.	This workload is an e-commerce website that allows Amazon users to purchase various types of inventory. This workload is the primary revenue generator for our business.
What is the Geographic location (Country) and Timezone of the primary customer contact?	Australia, Australian Eastern Standard Time(AEST)
Which organization or business units service this workload?	Consumer Organization.
Are there any compliance and regulatory requirements for this workload? If so, are there any actions required from AWS Incident Detection and Response during an incident to assist you in complying?	The workload operates in Europe and is subject to GDPR

AWS Incident Detection and Response architecture questions

Architecture Questions

Question	Example Response
Are there any architectural diagrams you can share outlining the various components of this workload? Please attach on delivery.	See attached
What AWS accounts support this workload and what resources does each account own (if multiple)?	123456789101 owns US-EAST-1 resources 123456789102 owns EU-WEST-1 resources
If multiple AWS accounts support this workload, are all of the accounts cleared for cross account support?	Yes

Question	Example Response
What AWS Regions does this workload operate in?	US-EAST-1, US-WEST-2, EU-WEST1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	Route 53 - Routes internet traffic using a A record from our registered domain name www.MyIncidentDetectionAndResponseWorkload.com to our ELB. Account#: 123456789012 Region: US-EAST-1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	Application Load Balancer - Routes incoming traffic to a target group of ECS containers. Account#: 123456789012 Region: US-EAST-1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	ECS - Compute infrastructure for main business logic fleet. Responsible for handling incoming user requests and making queries to persistence layer. Account#: 123456789012 Region: US-EAST-1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	RDS - Amazon Aurora cluster stores user data accessed by ECS business logic layer. Account#: 123456789012 Region: US-EAST-1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	S3 - Stores website assets. Account#: 123456789012 Region: US-EAST-1
Please describe how each AWS service is utilized by this workload in sentence or two. If the same service is utilized for multiple, different functions please include a new row for each function.	...
Are there any on-premise (non-AWS) components for this workload? If so, what are they and what functions are performed?	No
Is this workload currently architected for high availability at AZ / Regional Level. Please provide a detailed answer.	Yes. Warm standby. Replication to backup DB in Region X...
If yes to the above question, are there any automated failover mechanisms in place?	Yes. Auto Scaling / DNS

Question	Example Response
How is business continuity, backup and disaster recovery currently handled for this workload?	Each Region is capable of handling all traffic on minimal notice. Persistence layer data is backed up regularly to S3. No expectations for Incident Detection and Response

AWS Incident Detection and Response Runbook questions

Runbook Questions

Question	Example Response
Which communication method would you like AWS' Incident Management team to use when engaging your resolver team with Amazon Chime Conference Bridge details when an incident occurs on your workload? 1) Email Only, 2) Email and Phone Call. Note: AWS Incident Detection and Response can join a meeting on your preferred conferencing platform (for example: Microsoft Teams, Zoom, WebEx), but requires you to provide the conference bridge details to join, when an incident occurs.	Email Only. Contact details (Email/Phone) available in column I of Part-2
What is the bottom up escalation matrix (upwards through the leadership) in place for this workload? Provide Name, Role, Email Address and Mobile Number. At what interval should AWS Incident Detection and Response escalate if no response is received from the initial engagement?	Follow the Engagement Escalation - Plan if the SRE does not join the bridge, 10 minutes after the Initial Engagement. 1 - After 10 minutes, if no one has joined the Chime Bridge then Reply All to the previous email by addressing Jane Deer (Technical Lead) jane.deer@xyz.com and call Jane Deer (+1 234-567-8910) 2 - If Jane Deer cannot be reached then Reply All to the previous email by addressing Jim Buck (Director) jum.buck@xyz.com and call Jim Buck (+1 234-567-8910) ... n - If Jim Buck cannot be reached then Reply All to the previous email by addressing John Doe (General Manager) john.doe@xyz.com and call John Doe (+1 234-567-8910) *Please specify explicitly if you do not wish to be contacted via Phone Call for escalations.*
Once impact is confirmed, do you want AWS Incident Detection and Response to communicate regular status updates through email throughout the lifecycle of the incident?	Yes/No Send status updates to these recipients: sre@xyz.com; jane.deer@xyz.com;…."

Question	Example Response
If yes, please provide the email addresses that updates should be sent to. At what interval should status updates be sent through email?	Every 1 hr / 2 hr / .. / n hour.
[Optional] If you have any CloudWatch Dashboards that could be used by Incident Detection and Response to monitor key metrics during an incident, provide their Region and ARN.	No Dashboard available / Refer to these CloudWatch dashboards in US-EAST-1 region: arn:aws:cloudwatch::123456789012:dashboard/5xxErrors arn:aws:cloudwatch::123456789012:dashboard/latency …

AWS Incident Detection and Response 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 flapping 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 / `CW Alarm ARN` / 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 network connectivity, or issues with our DNS implementation that result in users not being able to access the workload.	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. Resolvers: Site Reliability Engineers	Engage the Site Reliability Engineering team by sending an email to `SRE@xyz.com` Create an AWS Premimum Support case for our ELB, and Route 53 services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the *XYZ* Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)
Workload Request Latency / `CW Alarm ARN` / 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 `SRE@xyz.com` Create an AWS Premimum Support case for our ECW, and RDS services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the *XYZ* Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)
Workload Request Availability / `CW Alarm ARN` / Availability < 95% for 5 datapoints within 5 minutes , treat missing data as missing.	This metric represents the availability for HTTP requests to be fulfilled by the workload. (# of HTTP 200 / # of Requests) per period. This alarm represents the availability of the workload.	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 `SRE@xyz.com` Create an AWS Premimum Support case for our ELB, and Route 53 services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the *XYZ* Team

Metric name / ARN / Threshold	Description	Notes	Actions requested
			through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)
New Relic Alarm Example			
End to End Integration test / *CW Alarm ARN* / 3% failure rate for 1 minute metrics over 3 minutes duration , treat missing data as missing Workload Identifier: End to End Test Workflow, AWS Region: US-EAST-1, AWS Account ID: 012345678910	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. This alarm represents the ability to process business transactions for the workload.	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 *SRE@xyz.com* Create an AWS Premimum Support case for our ECS, and DynamoDB services. If IMMEDIATE action is needed: Check EC2 Free memory/disk space and inform the *XYZ* Team through email to restart the instance, or run a log flush. (if immediate action is not needed, leave blank)

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.

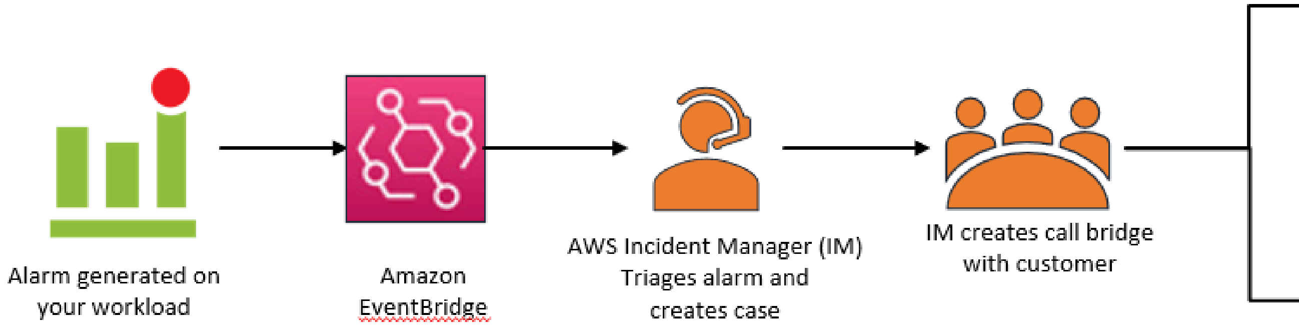
Alarm generated on your workload
Amazon EventBridge
AWS Incident Manager (IM) Triages alarm and creates case
IM creates call bridge with customer

1. **Alarm generation**: Alarms triggered on your workloads are pushed through Amazon EventBridge to the AWS Incident Detection and Response service. The AWS Incident Detection and Response service automatically pulls up the runbook associated with your alarm and notifies an incident manager.
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 we monitor AWS services specifically for your applications we may 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 affords you the opportunity to implement your recovery plans or workarounds early to mitigate the impact of the AWS Service Event.
3. **Incident Resolution**: The incident manager co-ordinates the incident across the required AWS teams and ensures 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/actions taken to mitigate or resolve the incident. The Post Incident Report may inform learnings for reducing a reccurence of the incident, or for improving the management of a future occurrence of a similar incident. The Post Incident Report is not a Root Cause Analysis (RCA) and customers may request for RCA in addition to the Post Incident Report. An example of a Post Incident Report is provided next.

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.
```

AWS Incident Detection and Response incident management for service events

AWS Incident Detection and Response engagement during a service event is focused on providing you with the relevant information you need to implement your recovery plans. AWS notifies you of an ongoing service event in your AWS Regions, whether or not your workload is impacted. If your workload is impacted by the service event, AWS Incident Detection and Response creates an AWS Support case to engage you, receive feedback on impact and sentiment, and provide predetermined guidance to invoke your recovery plans during the event. You also receive a notification through AWS Health containing details of the service event. Customers who are not affected by the AWS-owned service event (for example, operating in a different region, do not use the AWS service that is impaired, etc.) will continue to be supported by our standard engagement. For more information about AWS Health, see What is AWS Health?.

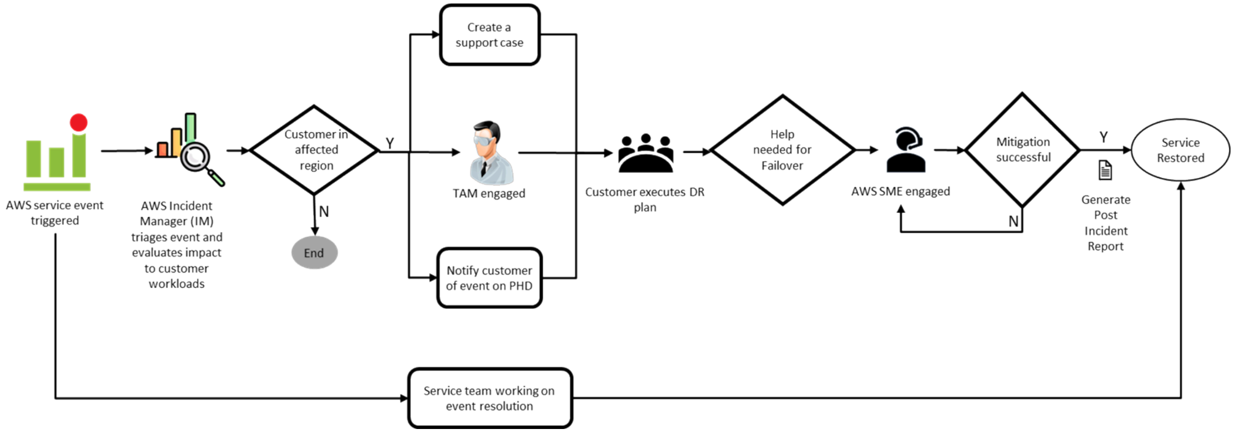
AWS service event triggered
AWS Incident Manager (IM) triages event and evaluates impact to customer workloads
Customer in affected region
Y
N
End
Create a support case
TAM engaged
Notify customer of event on PHD
Customer executes DR plan
Help needed for Failover
AWS SME engaged
Mitigation successful
Y
N
Generate Post Incident Report
Service Restored
Service team working on event resolution

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 our 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:** March 15, 2023

Change	Description	Date
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 General Reference*.