

User Guide

Incident Manager



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

Incident Manager: User Guide

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 Systems Manager Incident Manager?	. 1
Primary components and features	1
Benefits of using Incident Manager	3
Related services	5
Accessing Incident Manager	5
Incident Manager Regions and quotas	5
Pricing for Incident Manager	5
Incident lifecycle	6
Alerting and engagement	7
Triage	8
Investigation and mitigation	9
Post-incident analysis	10
Setting up	11
Sign up for an AWS account	11
Create a user with administrative access	12
Grant programmatic access	13
Required role for Incident Manager setup	14
Getting started	15
Prerequisites	15
Get prepared wizard	15
Managing incidents across AWS accounts and Regions	22
Cross-Region incident management	22
Cross-account incident management	. 23
Best practices	23
Set up and configure cross-account incident management	23
Limitations	25
Preparing for incidents	27
Monitoring	29
Configuring replication sets and Findings	29
Replication set	30
Managing tags for a replication set	. 31
Managing the Findings feature	
Creating and configuring contacts	
Contact channels	

Engagement plans	
Create a contact	
Import contact details to your address book	36
Managing responder rotations with on-call schedules	36
Creating an on-call schedule and rotation	
Managing an existing on-call schedule	41
Creating an escalation plan for responder engagement	47
Stages	47
Create an escalation plan	47
Creating and integrating chat channels for responders	48
Task 1: Create or update Amazon SNS topics for your chat channel	49
Task 2: Create a chat channel in AWS Chatbot	50
Task 3: Add the chat channel to a response plan in Incident Manager	53
Interacting through the chat channel	53
Integrating Systems Manager Automation runbooks for incident remediation	54
IAM permissions required to start and run runbook workflows	55
Working with runbook parameters	58
Define a runbook	60
Incident Manager runbook template	61
Creating and configuring response plans	62
Creating a response plan	63
Identifying potential causes of incidents from other services	
Enable and create a service role for findings	70
Configure permissions for cross-account findings support	71
Creating incidents automatically or manually	72
Creating incidents automatically with CloudWatch alarms	
Creating incidents automatically with EventBridge events	
Creating incidents using SaaS partners events	
Creating incidents using AWS service events	76
Creating incidents manually	77
Viewing incident details in the console	
Viewing the incident list in the console	
Viewing incident details in the console	
Top banner	
Incident notes	
Tabs	

Overview	
Diagnosis	81
Timeline	83
Runbooks	83
Engagements	84
Related items	
Properties	85
Performing a post-incident analysis	
Analysis details	86
Overview	
Metrics	
Timeline	87
Questions	
Actions	88
Checklist	
Analysis templates	88
AWS standard template	89
Create an analysis template	
Create an analysis	
Print a formatted incident analysis	90
Tutorials	
Using runbooks with Incident Manager	91
Task 1: Creating the runbook	92
Task 2: Creating an IAM role	
Task 3: Connecting the runbook to your response plan	97
Task 4: Assigning a CloudWatch alarm to your response plan	98
Task 5: Verifying the results	98
Managing security incidents	99
Tagging resources	102
Security	104
Data protection	105
Data encryption	106
Identity and Access Management	107
Audience	108
Authenticating with identities	109
Managing access using policies	112

How AWS Systems Manager Incident Manager works with IAM	114
Identity-based policy examples	122
Resource-based policy examples	126
Cross-service confused deputy prevention	128
Using service-linked roles	129
AWS managed policies for Incident Manager	131
Troubleshooting	138
Working with shared contacts and response plans in Incident Manager	140
Prerequisites for sharing contacts and response plans	141
Related services	141
Sharing a contact or response plan	
Stop sharing a shared contact or response plan	142
Identifying a shared contact or response plan	143
Shared contact and response plan permissions	143
Billing and metering	143
Instance limits	143
Compliance validation	144
Resilience	145
Infrastructure security	
Working with VPC endpoints (AWS PrivateLink)	146
Considerations for Incident Manager VPC endpoints	147
Creating an interface VPC endpoint for Incident Manager	147
Creating a VPC endpoint policy for Incident Manager	147
Configuration and vulnerability analysis	148
Security best practices	148
Preventative security best practices for Incident Manager	149
Detective security best practices for Incident Manager	150
Monitoring	
Monitoring metrics with Amazon CloudWatch	152
Viewing Incident Manager metrics on the CloudWatch console	154
Dimensions for Metrics	
Logging API calls using AWS CloudTrail	
Incident Manager management events in CloudTrail	
Incident Manager event examples	
Product and service integrations	
Integration with AWS services	160

Integration with other products and services	165
Storing PagerDuty access credentials in an AWS Secrets Manager secret	170
Troubleshooting	176
Error message: ValidationException – We were unable to validate the AWS	
Secrets Manager secret	176
Other troubleshooting issues	178
Document history	179

What Is AWS Systems Manager Incident Manager?

Incident Manager, a capability of AWS Systems Manager, is designed to help you mitigate and recover from *incidents* affecting your applications hosted on AWS.

In the context of AWS, an incident is any unplanned interruption or reduction in the quality of services that can have a significant impact on business operations. Therefore, it's crucial for organizations to establish a response strategy to efficiently mitigate and recover from incidents, and implement actions to prevent future incidents.

Incident Manager helps reduce the time to resolve incidents by:

- Providing automated plans for efficiently engaging the people responsible for responding to the incidents.
- Providing relevant troubleshooting data.
- Enabling automated response actions by using predefined Automation runbooks.
- Providing methods to collaborate and communicate with all stakeholders.

The features and workflows built into Incident Manager are based on the best practices for incident response that Amazon has been developing almost since its inception. Incident Manager integrates with such AWS services as Amazon CloudWatch, AWS CloudTrail, AWS Systems Manager, and Amazon EventBridge.

Primary components and features

This section describes the features in Incident Manager that you use to set up your incident response plans.

Response plan

A response plan functions as a template that defines what must be in place when an incident occurs. It includes such information as:

- Who is required to respond when an incident occurs.
- The established automated response to mitigate the incident.
- The collaboration tool that responders must use to communicate and receive automatic notifications about the incident.

Incident detection

You can configure Amazon CloudWatch alarms and Amazon EventBridge events to create incidents when conditions or changes that affect your AWS resources are detected.

Runbook automation support

You can initiate Automation runbooks from within Incident Manager to automate your critical response to incidents and provide detailed steps to first responders.

Engagement and escalation

An *engagement plan* specifies everyone to notify for each unique incident. You can specify individual contacts that you have added to Incident Manager or specify an on-call schedule that you created in Incident Manager. Engagement plans also specify an escalation path to help ensure visibility among stakeholders and active participation during the incident response process.

On-call schedules

An *on-call schedule* in Incident Manager consists of one or more rotations that you create for the schedule. For each rotation, you can include up to 30 contacts. When added to an escalation plan or response plan, the on-call schedule defines who is notified when an incident occurs that requires responder intervention. On-call schedules help ensure that you have full, redundant, 24/7 coverage as needed for your incident response.

Active collaboration

Incident responders actively respond to incidents through integration with the AWS Chatbot client. AWS Chatbot supports creating chat channels for Incident Manager that use Slack, Microsoft Teams, or Amazon Chime. Responders can communicate directly with one another, receive automated notifications about incidents, and—in Slack and Microsoft Teams—directly run some Incident Manager command line interface (CLI) operations.

Incident diagnosis

Responders can view up-to-date information in the Incident Manager console during an incident. Based on the changes in information, responders can then create follow-up items and remediate them by using Automation runbooks.

Findings from other services

To support responders' incident diagnosis, you can enable the Findings feature in Incident Manager. Findings are information about AWS CodeDeploy deployments and AWS CloudFormation stack updates that occurred around the time of an incident, and that involved one or more resources likely related to the incident. Having this information reduces the time needed to evaluate potential causes, which can reduce the mean time to recover (MTTR) from an incident.

Post-incident analysis

After an incident is resolved, you use a post-incident analysis to identify improvements to your incident response, including time to detection and mitigation. An analysis can also help you understand the root cause of the incidents. Incident Manager creates recommended follow-up action items that you can use to improve your incident response.

Benefits of using Incident Manager

Learn about the benefits of using Incident Manager in your incident detection and response operations.

This section describes the advantages that your organization can gain when you implement an Incident Manager response plan.

Diagnose issues efficiently and immediately

Amazon CloudWatch alarms and Amazon EventBridge events that you configure can create incidents automatically when there is any unplanned interruption or reduction in the quality of your services.

CloudWatch alarms detect and report when there are changes to the value of the metric or expression that is relative to a threshold over a number of time periods. EventBridge events are created as the result of change in an environment, application, or service that you have specified in an EventBridge rule. When you create an alarm or event, you can specify an action for an incident to be created in Incident Manager and the appropriate response plan to facilitate the engagement, escalation, and mitigation of the incident.

Incident Manager provides the ability to automatically collect and track the metrics related to an incident, through the use of CloudWatch metrics. In addition to the automated metrics generated for the incident when it is created through a CloudWatch alarm, you can add metrics manually in real time, to provide additional context and data to the responders in an incident.

Use the Incident Manager incident timeline to display points of interest in chronological order. Responders can also use the timeline to add custom events to describe what they did or what happened. Automated points of interest include:

- A CloudWatch alarm or EventBridge rule creates an incident.
- Incident metrics are reported to Incident Manager.
- Responders are engaged.
- Runbook steps complete successfully.

Engage effectively

Incident Manager brings incident responders together through the use of contacts, on-call schedules, escalation plans, and chat channels. You define individual contacts directly in Incident Manager and specify contact preferences (email, SMS, or voice). You add contacts to on-call schedule rotations to determine who is engaged to deal with incidents during a given period. Using your defined contacts and on-call schedules, you create escalation plans to engage the necessary responders at the right time during an incident.

Collaborate in real time

Communication during an incident is the key to faster resolution. Using an AWS Chatbot client set up to use Slack, Microsoft Teams, or Amazon Chime, you can bring together responders in their preferred connected chat channel where they directly interact with the incident and with one another. Incident Manager also displays the real-time actions of incident responders in the chat channel, providing context to others.

Automate service restoration

Incident Manager enables your responders to focus on the key tasks required to resolve an incident through the use of Automation *runbooks*. In Incident Manager, runbooks are a predefined series of actions taken to resolve an incident. They combine the power of automated tasks with manual steps as needed, leaving responders more available to analyze and respond to impact.

Prevent future incidents

Using Incident Manager post incident analysis, your team can develop more robust response plans and effect change across your applications to prevent future incidents and downtime. Post-incident analysis also provides for iterative learning and improvement of runbooks, response plans, and metrics.

Related services

Incident Manager integrates with several other AWS services and third-party services and tools to help you detect and resolve incidents, and to interact with its API operations indirectly and manage infrastructure. For information, see Product and service integrations with Incident Manager.

Accessing Incident Manager

You can access Incident Manager in any of the following ways:

- The Incident Manager console
- AWS CLI For general information, see <u>Getting started with the AWS CLI</u> in the AWS Command Line Interface User Guide. For information about CLI commands for Incident Manager, see <u>ssm-</u> incidents and <u>ssm-contacts</u> in the AWS CLI Command Reference.
- Incident Manager API For more information, see the <u>AWS Systems Manager Incident Manager</u> <u>API Reference</u>.
- AWS SDKs For more information, see Tools to Build on AWS.

Incident Manager Regions and quotas

Incident Manager isn't supported in all AWS Regions supported by Systems Manager.

To view information about Incident Manager Regions and quotas, see <u>AWS Systems Manager</u> <u>Incident Manager endpoints and quotas</u> in the *Amazon Web Services General Reference*.

Pricing for Incident Manager

There is a charge to use Incident Manager. For more information, see <u>AWS Systems Manager</u> pricing.

🚯 Note

Other AWS services, AWS content, and third-party content made available in connection with this service may be subject to separate charges and governed by additional terms.


For an overview of Trusted Advisor, a service that helps you optimize the costs, security, and performance of your AWS environment, see AWS Trusted Advisor in the AWS Support User Guide.

Incident lifecycle in Incident Manager

AWS Systems Manager Incident Manager provides a step-by-step framework based on best practices to identify and react to incidents, such as service outages or security threats. The primary focus of Incident Manager is to help restore affected services or applications to normal as quickly as possible through a complete incident lifecycle management solution.

As depicted in the following illustration, Incident Manager provides tools and best practices for every phase of the incident lifecycle:

- Alerting and engagement
- Triage
- Investigation and mitigation
- Post-incident analysis



Alerting and engagement

The alerting and engagement phase of the incident lifecycle focuses on bringing awareness to incidents within your applications and services. This phase begins before an incident is ever detected and requires a deep understanding of your applications. You can use <u>Amazon CloudWatch</u> <u>metrics</u> to monitor data about the performance of your applications, or use <u>Amazon EventBridge</u> to aggregate alerts from different sources, applications and services. After you've set up monitoring for your applications, you can begin alerting on metrics that stray outside the historical norm. To learn more about monitoring best practices, see <u>Monitoring</u>.

To support responders' incident diagnosis, you can enable the Findings feature in Incident Manager. Findings are information about AWS CodeDeploy deployments and AWS CloudFormation stack updates that occurred around the time of an incident. Having this information reduces the time needed to evaluate potential causes, which can reduce the mean time to recover (MTTR) from an incident.

Now that you are monitoring for incidents in your applications, you can define an incident *response plan* to use during an incident. To learn more about creating response plans, see <u>Creating and</u> <u>configuring response plans in Incident Manager</u>. Amazon EventBridge events or CloudWatch Alarms can automatically create an incident using with response plans as the template. To learn more about incident creation, see <u>Creating incidents automatically or manually in Incident</u> Manager.

Response plans launch related *escalation plans* and *engagement plans* to bring first responders into the incident. For more information about setting up escalation plans, see <u>Create an escalation plan</u>. Simultaneously, AWS Chatbot notifies responders using a *chat channel* directing them to the incident detail page. Using the chat channel and *incident details*, the team can communicate and triage an incident. For more information about setting up chat channels in Incident Manager, see <u>Task 2: Create a chat channel in AWS Chatbot</u>.

Triage

Triage is when first responders attempt to determine the impact to customers. The incident details view in the Incident Manager console provides the responders with timelines and metrics to help them assess the incident. Assessing the impact of an incident also lays the groundwork for response time, resolution, and communication for the incident. Responders prioritize incidents by using impact ratings from 1 (Critical) to 5 (No Impact).

Your organization can define the exact scope of each impact rating however you choose. The following table provides examples of how each impact level might typically be defined.

Impact code	Impact name	Sample defined scope
1	Critical	Full application failure that impacts most customers.
2	High	Full application failure that impacts a subset of customers .
3	Medium	Partial application failure that is customer-impacting.

Impact code	Impact name	Sample defined scope
4	Low	Intermittent failures that have limited impact on customers.
5	No Impact	Customers aren't currently impacted but urgent action is needed to avoid impact.

Investigation and mitigation

The *incident* details view provides your team with runbooks, timelines, and metrics. To see how you can work with an incident, see the Viewing incident details in the console.

Runbooks commonly provide investigation steps and can automatically pull data or attempt commonly used solutions. Runbooks also provide clear, repeatable steps that your team has found to be useful in mitigating incidents. The runbook tab focuses on the current runbook step and shows past and future steps.

Incident Manager integrates with Systems Manager Automation to build runbooks. Use runbooks to do any of the following:

- Manage instances and AWS resources
- Automatically run scripts
- Manage AWS CloudFormation resources

For more information about the supported action types, see <u>Systems Manager Automation actions</u> reference in the AWS Systems Manager User Guide.

The **Timeline** tab shows what actions have been taken. The timeline records each with a timestamp and automatically created details. To add custom events to the timeline, see the <u>Timeline</u> section in the *Incident details* page of this user guide.

The **Diagnosis** tab shows automatically populated metrics and manually added metrics. This view provides valuable information into the activities of your application during an incident.

The **Engagements** tab allows you to add additional contacts to the incident and helps provide the resources for the engaged contact to get up to speed quickly once involved in the incident. Contacts are engaged through defined escalation plans or personal engagement plans.

Using a *chat channel*, you can directly interact with your incident and other responders on your team. Using AWS Chatbot, you can configure chat channels in. Slack, Microsoft Teams, and Amazon Chime. In Slack and Microsoft Teams channels, responders can interact with incidents directly from the chat channel using a number of ssm-incidents commands. For more information, see Interacting through the chat channel.

Post-incident analysis

Incident Manager provides a framework for reflecting on an incident, taking steps needed to prevent the incident from occurring again in the future, and to improve incident response activities overall. Improvements can include:

- Changes to the applications involved in an incident. Your team can use this time to improve the system and make it more fault tolerant.
- Changes to an incident response plan. Take the time to incorporate learned lessons.
- Changes to runbooks. Your team can dive deep into steps needed for resolution and the steps that you can automate.
- Changes to alerting. After an incident, your team might have noticed critical points in the metrics you can use to alert the team sooner about an incident.

Incident Manager facilitates these potential improvements by using a set of post-incident analysis questions and action items alongside the incident timeline. To learn more about improvement through analysis, see <u>Performing a post-incident analysis in Incident Manager</u>.

Setting up AWS Systems Manager Incident Manager

We recommend setting up AWS Systems Manager Incident Manager in the account that you use to manage your operations. Before you use Incident Manager for the first time, complete the following tasks:

Topics

- Sign up for an AWS account
- <u>Create a user with administrative access</u>
- Grant programmatic access
- <u>Required role for Incident Manager setup</u>

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

- 1. Open https://portal.aws.amazon.com/billing/signup.
- 2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

When you sign up for an AWS account, an AWS account root user is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform tasks that require root user access.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <u>https://aws.amazon.com/</u> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the <u>AWS Management Console</u> as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see <u>Signing in as the root user</u> in the AWS Sign-In User Guide.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see Enable a virtual MFA device for your AWS account root user (console) in the IAM User Guide.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see Enabling AWS IAM Identity Center in the AWS IAM Identity Center User *Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see <u>Configure user access with the default IAM Identity Center directory</u> in the AWS IAM Identity Center User Guide.

Sign in as the user with administrative access

• To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see <u>Signing in to the AWS access portal</u> in the AWS Sign-In User Guide.

1. In IAM Identity Center, create a permission set that follows the best practice of applying leastprivilege permissions.

For instructions, see Create a permission set in the AWS IAM Identity Center User Guide.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see <u>Add groups</u> in the AWS IAM Identity Center User Guide.

Grant programmatic access

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	То	Ву
Workforce identity (Users managed in IAM Identity Center)	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	 Following the instructions for the interface that you want to use. For the AWS CLI, see <u>Configuring the AWS</u> <u>CLI to use AWS IAM</u> <u>Identity Center</u> in the AWS <u>Command Line Interface</u> <u>User Guide</u>. For AWS SDKs, tools, and AWS APIs, see <u>IAM Identity</u> <u>Center authentication</u> in the AWS SDKs and Tools <u>Reference Guide</u>.

Which user needs programmatic access?	То	Ву
IAM	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentia Is with AWS resources in the IAM User Guide.
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	 Following the instructions for the interface that you want to use. For the AWS CLI, see <u>Authenticating using IAM</u> <u>user credentials</u> in the AWS Command Line Interface User Guide. For AWS SDKs and tools, see <u>Authenticate using</u> <u>long-term credentials</u> in the AWS SDKs and Tools Reference Guide. For AWS APIs, see <u>Managing access keys for</u> <u>IAM users</u> in the IAM User Guide.

Required role for Incident Manager setup

Before you begin, your account must have the IAM permission iam:CreateServiceLinkedRole. Incident Manager uses this permission to create the AWSServiceRoleforIncidentManager in your account. For more information, see <u>Using service-linked roles for Incident Manager</u>.

Getting started with Incident Manager

This section walks through **Get prepared** in the Incident Manager console. You're required to complete **Get prepared** in the console before you can use it for incident management. The wizard walks you through setting up your replication set, at least one contact and one escalation plan, and your first response plan. The following guides will help you understand Incident Manager and the incident lifecycle:

- What Is AWS Systems Manager Incident Manager?
- Incident lifecycle in Incident Manager

Prerequisites

If you're using Incident Manager for the first time, see the <u>Setting up AWS Systems Manager</u> <u>Incident Manager</u>. We recommend setting up Incident Manager in the account that you use to manage your operations.

We recommend that you complete the Systems Manager quick setup before beginning the Incident Manager **Get prepared** wizard. Use Systems Manager <u>Quick Setup</u> to configure frequently used AWS services and features with recommended best practices. Incident Manager uses Systems Manager features to manage incidents associated with your AWS accounts and benefits from having Systems Manager configured first.

Get prepared wizard

The first time you use Incident Manager, you can access the **Get prepared** wizard from the Incident Manager service homepage. To access the **Get prepared** wizard after you first complete setup, choose **Prepare** on the **Incidents** list page.

- 1. Open the Incident Manager console.
- 2. On the Incident Manager service homepage, choose Get prepared.

General settings

1. Under General settings, choose Set up.

- Read the terms and conditions. If you agree to Incident Manager's terms and conditions, select
 I have read and agree to the Incident Manager terms and conditions, then choose Next.
- 3. In the **Regions** area, your current AWS Region appears as the first Region in your replication set. To add more Regions to your replication set, choose them from the list of Regions.

We recommend including at least two Regions. In case one Region is temporarily unavailable, incident-related activities can still be routed to the other Region.

🚯 Note

Creating the replication set creates the AWSServiceRoleforIncidentManager service-linked role in your account. To learn more about this role, see <u>Using service-linked roles for Incident Manager</u>.

4. To set up encryption for your replication set, do one of the following:

🚯 Note

All Incident Manager resources are encrypted. To learn more about how your data is encrypted, see <u>Data protection in Incident Manager</u>. For more information about your Incident Manager replication set, see <u>Configuring the Incident Manager replication set</u>.

- To use an AWS owned key, choose Use AWS owned key.
- To use your own AWS KMS key, choose **Choose an existing AWS KMS key**. For each Region you selected in step 3, choose an AWS KMS key, or enter an AWS KMS Amazon Resource Name (ARN).

🚺 Tip

If you don't have an available AWS KMS key, choose Create an AWS KMS key.

5. (Optional) In the **Tags** area, add one or more tags to the replication set. A tag includes a key and, optionally, a value.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For more information, see Tagging resources in Incident Manager.

Incident Manager

6. (Optional) In the **Service Access** area, to activate the Findings feature, choose the **Create service role for findings in this account** check box.

A *finding* is information about a code deployment or infrastructure change that occurred around the same time that an incident was created. A finding can be examined as a potential cause of the incident. Information about these potential causes is added to the **Incident details** page for the incident. With information about these deployments and changes readily at hand, responders don't need to manually search for this information.

🚺 Tip

To view information about the role to be created, choose View permissions.

7. Choose Create.

To learn more about replication sets and resiliency, see <u>Resilience in AWS Systems Manager</u> Incident Manager.

Contacts (optional)

1. Choose **Create contact**.

Incident Manager engages contacts during an incident. For more information about contacts, see Creating and configuring contacts in Incident Manager.

- 2. For **Name**, enter the contact's name.
- 3. For **Unique alias**, enter an alias to identify this contact.
- 4. In the **Contact channel** section., do the following to define how the contact is engaged during incidents:
 - a. For **Type**, choose **Email**, **SMS**, or **Voice**.
 - b. For **Channel name**, enter a unique name to help you identify the channel.
 - c. For **Detail**, enter the email address or phone number for the contact.

Phone numbers must have 9–15 characters and start with + followed by the country code and subscriber number.

d. To create another contact channel, choose **Add a new contact channel**. We recommend defining at least two channels for each contact.

5. In the **Engagement plan** area, do the following to define which channels to notify the contact through, and how long to wait for an acknowledgement through each channel. Select the contact channels to use to engage the contact during incidents.

🚺 Note

We recommend defining at least two devices in the engagement plan.

- a. For **Contact channel name**, choose a channel you specified in the **Contact channel** area.
- b. For **Engagement time (min)**, enter the number of minutes to wait before engaging the contact channel.

We recommend that you select at least one device to engage at the beginning of an engagement, specifying **0** (zero) minutes waiting time.

- c. To add more contact channels to the engagement plan, choose **Add engagement**.
- 6. (Optional) In the **Tags** area, add one or more tags to the contact. A tag includes a key and, optionally, a value.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For more information, see <u>Tagging resources in Incident Manager</u>.

- 7. To create the contact record and send activation codes to the defined contact channels, choose **Next**.
- 8. (Optional) In the **Contact channel activation** page, enter the activation code sent to each channel.

You can generate new activation codes later if you're not able to enter the codes now.

- 9. Repeat step four until you have added all of your contacts to Incident Manager.
- 10. After all contacts are entered, choose **Finish**.

(Optional) Escalation plans

1. Choose Create escalation plan.

An escalation plan escalates through your contacts during an incident, ensuring that Incident Manager engages the correct responders during an incident. For more information about escalation plans, see <u>Creating an escalation plan for responder engagement in Incident</u> Manager.

- 2. For Name, enter a unique name for the escalation plan.
- 3. For **Alias**, enter a unique alias to help you identify the escalation plan.
- 4. In the **Stage 1** area, do the following:
 - a. For **Escalation channel**, choose contact channels to engage.
 - b. If you want a contact to be able to halt the progression of escalation plan stages, select **Acknowledgment stops plan progression**.
 - c. To add more channels to a stage, choose **Add escalation channel**.
- 5. To create a new stage in the escalation plan, choose **Add stage** and add its stage details.
- 6. (Optional) In the **Tags** area, add one or more tags to the escalation plan. A tag includes a key and, optionally, a value.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For more information, see <u>Tagging resources in Incident Manager</u>.

7. Choose Create escalation plan.

Response plan

- 1. Choose **Create response plan**. Use the response plan to put together contacts and escalation plans you created. During this **Getting started** wizard, the following sections are optional, especially if this is your first time setting up a response plan:
 - Chat channel
 - Runbooks
 - Engagements
 - Third-party integrations

For information about adding these elements to response plans later, see <u>Preparing for</u> incidents in Incident Manager.

2. For **Name**, enter a unique, identifiable name for the response plan. The name is used to create the response plan ARN or in response plans with no display name.

- 3. (Optional) For **Display name**, enter a name to help you identify this response plan when creating incidents.
- 4. For **Title**, enter a title to help identify the type of incident that relates to this response plan. The value you specify is included in each incident's title. The alarm or event that started the incident is also added to the title.
- 5. For **Impact**, select the impact level you expect for incidents related to this response plan, such as **Critical** or **Low**.
- 6. (Optional) For **Summary**, enter a brief description that is used to provide an overview of the incident. Incident Manager automatically populates relevant information into the summary during an incident.
- 7. (Optional) For **Dedupe string**, enter a dedupe string. Incident Manager uses this string to prevent the same root cause from creating multiple incidents in the same account.

A deduplication string is a term or phrase the system uses to check for duplicate incidents. If you specify a deduplication string, Incident Manager searches for open incidents that contain the same string in the dedupeString field when it creates the incident. If a duplicate is detected, Incident Manager deduplicates the newer incident into the existing incident.

🚺 Note

By default, Incident Manager automatically deduplicates multiple incidents created by the same Amazon CloudWatch alarm or Amazon EventBridge event. You don't have to enter your own deduplication string to prevent duplication for these resource types.

8. (Optional) In the **Tags** area, add one or more tags to the response plan. A tag includes a key and, optionally, a value.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For more information, see Tagging resources in Incident Manager.

- 9. Select the contacts and escalation plans to apply to the incident from the **Engagements** dropdown.
- 10. Choose Create response plan.

After you've created a response plan, you can associate Amazon CloudWatch alarms or Amazon EventBridge events with the response plan. This will automatically create an incident based on an alarm or event. For more information, see <u>Creating incidents automatically or manually in Incident</u> Manager.

Managing incidents across AWS accounts and Regions in Incident Manager

You can configure Incident Manager, a capability of AWS Systems Manager, to work with multiple AWS Regions and accounts. This section describes cross-Region and cross-account best practices, set up steps, and known limitations.

Topics

- Cross-Region incident management
- Cross-account incident management

Cross-Region incident management

Incident Manager supports automated and manual incident creation in <u>several AWS Regions</u>. When you initially onboard with Incident Manager by using the **Get prepared** wizard, you can specify up to three AWS Regions for your *replication set*. For incidents automatically created by Amazon CloudWatch alarms or Amazon EventBridge events, Incident Manager attempts to create an incident in the same AWS Region as the event rule or alarm. If Incident Manager is experiencing an outage in first AWS Region where it is set up, CloudWatch or EventBridge automatically creates the incident in another of the available Regions specified for your replication set.

🛕 Important

Note the following important details.

- We recommend that you specify at least two AWS Regions in your replication set. If you don't specify at least two Regions, the system will fail to create incidents during the period when Incident Manager is unavailable.
- Incidents created by a cross-Region failover don't invoke runbooks specified in response plans.

For more information about on-boarding with Incident Manager and specifying additional Regions, see <u>Getting started with Incident Manager</u>.

Cross-account incident management

Incident Manager uses AWS Resource Access Manager (AWS RAM) to share Incident Manager resources across management and application accounts. This section describes cross-account best practices, how to set up cross-account functionality for Incident Manager, and known limitations of cross-account functionality in Incident Manager.

A management account is the account that you perform operations management from. In an organization setup, the management account owns the response plans, contacts, escalation plans, runbooks, and other AWS Systems Manager resources.

A application account is the account that owns the resources that make up your applications. These resources can be Amazon EC2 instances, Amazon DynamoDB tables, or any of the other resources that you use to build applications in the AWS Cloud. Application accounts also own the Amazon CloudWatch alarms and Amazon EventBridge events that create incidents in Incident Manager.

AWS RAM uses resource shares to share resources between accounts. You can share the response plan and contact resources between accounts in AWS RAM. By sharing these resources, application accounts and management accounts can interact with engagements and incidents. Sharing a response plan shares all past and future incidents created using that response plan. Sharing a contact shares all past and future engagements of the contact or response plan.

Best practices

Follow these best practices when sharing your Incident Manager resources across accounts:

- Regularly update the resource share with response plans and contacts.
- Regularly review resource share principals.
- Set up Incident Manager, runbooks, and chat channels in your management account.

Set up and configure cross-account incident management

The following steps describe how to set up and configure Incident Manager resources and use them for cross-account functionality. You may have configured some services and resources for cross-account functionality in the past. Use these steps as a checklist of requirements before starting your first incident using cross-account resources.

- 1. (Optional) Create organizations and organizational units using AWS Organizations. Follow the steps in the <u>Tutorial: Creating and configuring an organization</u> in the AWS Organizations User *Guide*.
- (Optional) Use the Systems Manager Quick Setup capability to set up the correct AWS Identity and Access Management roles for you to use when configuring your cross-account runbooks. For more information, see <u>Quick Setup</u> in the AWS Systems Manager User Guide.
- 3. Follow the steps listed in <u>Running automations in multiple AWS Regions and accounts</u> in the *AWS Systems Manager User Guide* to create runbooks in your Systems Manager automation documents. A runbook can be run by either a management account, or by one of your application accounts. Depending on your use case, you will need to install the appropriate AWS CloudFormation template for the roles necessary to create and view runbooks during an incident.
 - Running a runbook in the management account. The management account must download and install the <u>AWS-SystemsManager-AutomationReadOnlyRole</u> CloudFormation template. When installing AWS-SystemsManager-AutomationReadOnlyRole, specify the account IDs of all application accounts. This role will let your application accounts read the status of the runbook from the incident details page. The application account must install the <u>AWS-SystemsManager-AutomationReadOnlyRole</u> CloudFormation template. The incident details page uses this role to get the automation status from the management account.
 - Running a runbook in a application account. The management account must download and install the <u>AWS-SystemsManager-AutomationAdministrationReadOnlyRole</u> CloudFormation template. This role allows the management account to read the status of the runbook in the application account. The application account must download and install the <u>AWS-SystemsManager-AutomationReadOnlyRole</u> CloudFormation template. When installing AWS-SystemsManager-AutomationReadOnlyRole, specify the account ID of the management account and other application accounts. The management account and other application accounts assume this role to read the status of the runbook.
- 4. (Optional) In each application account in the organization, download and install the <u>AWS-SystemsManager-IncidentManagerIncidentAccessServiceRole</u> CloudFormation template. When installing AWS-SystemsManager-IncidentManagerIncidentAccessServiceRole, specify the account ID of the management account. This role provides the permissions that Incident Manager needs to access information about AWS CodeDeploy deployments and AWS CloudFormation stack updates. This information is reported as *findings* for an incident if the

Findings feature is enabled. For more information, see <u>Identifying potential causes of incidents</u> from other services as "findings" in Incident Manager.

- 5. To set up and create contacts, escalation plans, chat channels, and response plans, follow the steps detailed in <u>Preparing for incidents in Incident Manager</u>.
- 6. Add your contacts and response plan resources to either your existing resource share or a new resource share in AWS RAM. For more information, see <u>Getting started with AWS RAM</u> in the *AWS RAM User Guide*. Adding response plans to AWS RAM enables application accounts to access incidents and incident dashboards created using the response plans. Application accounts also gain the ability to associate CloudWatch alarms and EventBridge events to a response plan. Adding the contacts and escalation plans to AWS RAM enables application accounts to view engagements and engage contacts from the incident dashboard.
- 7. Add cross-account cross-Region functionality to your CloudWatch console. For steps and information, see <u>Cross-account cross-Region CloudWatch console</u> in the *Amazon CloudWatch User Guide*. Adding this functionality ensures that the application accounts and management account you've created can view and edit metrics from the incident and analysis dashboards.
- 8. Create a cross-account Amazon EventBridge event bus. For steps and information, see <u>Sending</u> <u>and receiving Amazon EventBridge events between AWS accounts</u>. You can then use this event bus to create event rules that detect incidents in application accounts and create incidents in the management account.

Limitations


The following are known limitations of Incident Manager's cross-account functionality:

- The account that creates a post-incident analysis is the only account that can view and change it. If you use a application account to create a post-incident analysis, only members of that account can view and change it. The same is true if you use a management account to create a postincident analysis.
- Timeline events aren't populated for automation documents run in application accounts.
 Updates of automation documents run in application accounts are visible in the Runbook tab of the incident.
- Amazon Simple Notification Service topics can't be used cross-account. Amazon SNS topics must be created in the same Region and account as the response plan it's used in. We recommend using the management account to create all SNS topics and response plans.

- Escalation plans can only be created using contacts in the same account. A contact that has been shared with you can't be added to an escalation plan in your account.
- Tags applied to response plans, incident records, and contacts can only be viewed and modified from the resource owner account.

Preparing for incidents in Incident Manager

Planning for an incident begins long before the incident lifecycle. As the following illustration shows, before starting to respond to incidents, you get prepared by setting up chat channels, creating escalation plans, specifying contacts, and determining the Automation runbooks to use in incident response. Then, use a response plan that specifies how monitoring occurs and whether responses are automated. After remediation is complete, you can analyze the incident and incident response to further refine your response plan for future incidents.



Topics

- Monitoring
- <u>Configuring replication sets and Findings in Incident Manager</u>
- <u>Creating and configuring contacts in Incident Manager</u>

- Managing responder rotations with on-call schedules in Incident Manager
- Creating an escalation plan for responder engagement in Incident Manager
- <u>Creating and integrating chat channels for responders in Incident Manager</u>
- Integrating Systems Manager Automation runbooks in Incident Manager for incident remediation
- Creating and configuring response plans in Incident Manager
- Identifying potential causes of incidents from other services as "findings" in Incident Manager

Monitoring

Monitoring the health of your AWS hosted applications is key to ensuring application up time and performance. When determining monitoring solutions, consider the following:

- Criticality of feature If the system were to fail, how critical would the impact to downstream users be.
- **Commonality of failure** How commonly does a system fail; systems that require frequent intervention should be closely monitored.
- Increased latency How much the time to complete a task has increased or decreased.
- Client-side versus server-side metrics If there is a discrepancy between related metrics on the client and server.
- **Dependency failures** Failures that your team can and should prepare for.

After creating response plans, you can use your monitoring solutions to automatically track incidents the moment they happen in your environment. For more information about incident tracking and creation, see <u>Viewing incident details in the Incident Manager console</u>.

For more information about architecting secure, high-performing, resilient, and efficient infrastructure applications and workloads, see the <u>AWS Well-Architected</u>.

Configuring replication sets and Findings in Incident Manager

After you have completed the Incident Manager Get prepared wizard, you can manage certain options on the **Settings** page. These options include your replication set, tags applied to the replication set, and the Findings feature.

Topics

- Configuring the Incident Manager replication set
- Managing tags for a replication set
- Managing the Findings feature

Configuring the Incident Manager replication set

The Incident Manager replication set replicates your data to many AWS Regions in order to do the following:

- Increase cross-Region redundancy
- Allow Incident Manager to access resources in different Regions and reduce latency for your users.
- Encrypt your data with either an AWS managed key or your own customer managed key.

All Incident Manager resources are encrypted by default. To learn more about how your resources are encrypted, see Data protection in Incident Manager.

To get started with Incident Manager, first create your replication set using the **Get prepared** wizard. To learn more about getting prepared in Incident Manager, see the <u>Get prepared wizard</u>.

Editing your replication set

By using the Incident Manager **Settings** page, you can edit your replication set. You can add Regions, delete Regions, and enable or disable replication set deletion protection. You can't edit the key used to encrypt your data. To change the key, delete and recreate the replication set.

Add a Region

- 1. Open the Incident Manager console, and then choose **Settings** in the left navigation pane.
- 2. Choose Add Region.
- 3. Select the **Region**.
- 4. Choose Add.

Delete a Region

1. Open the Incident Manager console, and then choose **Settings** in the left navigation pane.

- 2. Select the Region that you want to delete.
- 3. Choose **Delete**.
- 4. Enter **delete** into the text box, and choose **Delete**.

Deleting your replication set

Deleting the last Region in your replication set deletes the entire replication set. Before you can delete the last Region, disable the deletion protection by turning off **Deletion protection** on the **Settings** page. After you delete your replication set, you can create a new replication set by using the **Get prepared** wizard.

To delete a Region from your replication set, wait 24 hours after creating it. Attempting to delete a Region from your replication set sooner than 24 hours after creation causes the deletion to fail.

Deleting your replication set deletes all Incident Manager data.

Delete the replication set

- 1. Open the Incident Manager console, and then choose **Settings** in the left navigation pane.
- 2. Select the last Region in your replication set.
- 3. Choose Delete.
- 4. Enter **delete** into the text box, and choose **Delete**.

Managing tags for a replication set

Tags are optional metadata that you assign to a resource. Use tags to categorize a resource in different ways, such as by purpose, owner, or environment.

To manage tags for a replication set

- 1. Open the Incident Manager console, and then choose Settings in the left navigation pane.
- 2. In the Tags area, choose Edit.
- 3. To add a tag, do the following:
 - a. Choose **Add new tag**.
 - b. Enter a key and optional value for the tag.

- c. Choose Save.
- 4. To delete a tag, do the following:
 - a. Under the tag you want to delete, choose **Remove**.
 - b. Choose **Save**.

Managing the Findings feature

The Findings feature helps responders in your organization identify potential root causes of incidents soon after the incidents begin. Currently, Incident Manager provides findings for AWS CodeDeploy deployments and AWS CloudFormation stack updates.

For cross-account support for findings, after you enable the feature, you must complete an additional setup step in each application account in the organization.

To use the feature, you let Incident Manager create a service role that includes the required permissions to access data on your behalf.

To enable the Findings feature

- 1. Open the Incident Manager console, and then choose **Settings** in the left navigation pane.
- 2. In the **Findings** area, choose **Create service role**.
- 3.

Review information about the service role to be created, and then choose Create.

To disable the Findings feature

To stop using the Findings feature, delete the IncidentManagerIncidentAccessServiceRole role from each account where it has been created.

- 1. Open the IAM console at https://console.aws.amazon.com/iam/.
- 2. In the left navigation pane, choose **Roles**.
- 3. In the search box, enter **IncidentManagerIncidentAccessServiceRole**.
- 4. Choose the name of the role, and then choose **Delete**.
- 5. Enter the role name in the dialog box to confirm that you want to delete the role, and then choose **Delete**.

Creating and configuring contacts in Incident Manager

AWS Systems Manager Incident Manager contacts are responders to incidents. A contact can have multiple channels that Incident Manager can engage during an incident. You can define a contact's engagement plan to describe how and when Incident Manager engages the contact.

Topics

- Contact channels
- Engagement plans
- Create a contact
- Import contact details to your address book

Contact channels

Contact channels are the various methods Incident Manager uses to engage a contact.

Incident Manager supports the following contact channels:

- Email
- Short Message Service (SMS)
- Voice

Contact channel activation

To protect your privacy and security, Incident Manager sends a device activation code to you when you create contacts. To engage your devices during an incident, you must first activate them. To do so, enter the device activation code on the create contact page.

Certain features of Incident Manager include functionality that send notifications to a contact channel. By using these features, you consent to this service sending notifications about service disruptions or other events to the contact channels included in the specified workflow. This includes notifications sent to a contact as part of an on-call schedule rotation. Notifications may be sent by email, SMS message, or voice call as specified in a contact's details. You confirm by using these features that you're authorized to add the contact channels you provide to Incident Manager.

Opting out

You can cancel these notifications at any time by removing a mobile device as a contact channel. Individual notification recipients may also cancel notifications at any time by removing the device from their contact.

To remove a contact channel from a contact

- 1. Navigate to the Incident Manager console and choose **Contacts** from the left navigation.
- 2. Select the contact with the contact channel that you are removing and choose **Edit**.
- 3. Choose **Remove** next to the contact channel that you would like to remove.
- 4. Choose Update.

Contact channel deactivation

To deactivate a device, reply **UNSUBSCRIBE**. Replying **UNSUBSCRIBE** stops Incident Manager from engaging your device.

Contact channel reactivation

- 1. Reply **START** to the message from Incident Manager.
- 2. Navigate to the Incident Manager console and choose **Contacts** from the left navigation.
- 3. Select the contact with the contact channel that you are removing and choose **Edit**.
- 4. Choose Activate devices.
- 5. Enter the **Activation code** sent to the device by Incident Manager.
- 6. Choose Activate.

Engagement plans

Engagement plans define when Incident Manager engages the contact channels. You can engage contact channels multiple times at different stages from the start of an engagement. You can use engagement plans in an escalation plan or response plan. To learn more about escalation plans, see <u>Creating an escalation plan for responder engagement in Incident Manager</u>.

Create a contact

To create a contact, use the following steps.

1. Open the Incident Manager console and choose **Contacts** from the left navigation.

- 2. Choose **Create contact**.
- 3. Type the full name of the contact and provide a unique and identifiable alias.
- 4. Define a **Contact channel**. We recommend having two or more different types of contact channels.
 - a. Choose the type: email, SMS, or voice.
 - b. Enter an identifiable name for the contact channel.
 - c. Provide the contact channel details, such as email: arosalez@example.com
- 5. To define more than one contact channel, choose **Add contact channel**. Repeat step 4 for each new contact channel added.
- 6. Define an engagement plan.

🔥 Important

To engage a contact, you must define an engagement plan.

a. Choose a **Contact channel name**.

- b. Define how many minutes from the start of the engagement to wait until Incident Manager engages this contact channel.
- c. To add another contact channel, choose **Add engagement**.
- 7. After defining your engagement plan, choose **Create**. Incident Manager sends an activation code to each of the defined contact channels.
- 8. (Optional) To activate the contact channels, enter the activation code that Incident Manager sent to each defined contact channel.
- 9. (Optional) To send a new activation code, choose **Send new code**.
- 10. Choose Finish.

After you define a contact and activate its contact channels, you can add contacts to escalation plans to form a chain of escalation. To learn more about escalation plans, see <u>Creating an</u> <u>escalation plan for responder engagement in Incident Manager</u>. You can add contacts to a response plan for direct engagement. To learn more about creating response plans, see <u>Creating and</u> configuring response plans in Incident Manager.

Import contact details to your address book

When an incident is created, Incident Manager can notify responders by using voice or SMS notifications. To ensure that responders see that the call or SMS notification is from Incident Manager, we recommend that all responders download the Incident Manager <u>virtual card format</u> (.vcf) file to the address book on their mobile devices. The file is hosted in Amazon CloudFront and is available in the AWS commercial partition.

To download the Incident Manager .vcf file

- On your mobile device, either choose or enter the following URL: <u>https://</u> d26vhuvd5b89k2.cloudfront.net/aws-incident-manager.vcf.
- 2. Save or import the file to the address book on your mobile device.

Managing responder rotations with on-call schedules in Incident Manager

An on-call schedule in Incident Manager defines who is notified when an incident occurs that requires operator intervention. An on-call schedule consists of one or more rotations you create for the schedule. Each rotation can include up to 30 contacts.

After you create an on-call schedule, you can include it as an escalation in your escalation plan. When an incident associated with that escalation plan occurs, Incident Manager notifies the operator (or operators) who are on call according to the schedule. This contact can then acknowledge the engagement. In your escalation plan, you can designate one or more on-call schedules, as well as one or more individual contacts, across multiple stages of escalation. For more information, see Creating an escalation plan for responder engagement in Incident Manager.

🚺 Tip

As a best practice, we recommend adding contacts and on-call schedules as the escalation channels in an escalation plan. You should then choose an escalation plan as the engagement in a response plan. This approach provides the fullest coverage for incident response in your organization. Each on-call schedule supports up to eight rotations. Rotations can overlap or run concurrently. This increases the number of operators notified to respond when an incident occurs. You can also create rotations that run consecutively. This supports scenarios like "follow the sun" incident management where you have groups around the world that support the same service.

Use the topics in this section to help you create and manage on-call schedules for your incident response operations.

Topics

- Creating an on-call schedule and rotation in Incident Manager
- Managing an existing on-call schedule in Incident Manager

Creating an on-call schedule and rotation in Incident Manager

Create an on-call schedule with one or more rotations of contacts to engage to respond to incidents during their shifts.

Before you begin

Before you create an on-call schedule, ensure that you previously created the contacts you want to add to the rotations in the schedule. For information, see <u>Creating and configuring contacts in</u> Incident Manager.

Accounting for Daylight Savings Time (DST) changes

When you create a rotation, you specify the global time zone that serves as the basis for shift coverage times and dates you specify for this rotation. You can use any time zone defined by the <u>Internet Assigned Numbers Authority (IANA)</u>. For example: America/Los_Angeles, UTC, and Asia/Seoul. You can add more than one rotation to an on-call schedule. However, when the responders for each rotation are geographically located in different time zones, keep in mind any DST changes each rotation might be subject to.

For instance, America/Los_Angeles and Europe/Dublin observe different DST schedules. As a result, the time difference between the two zones can vary from 6 to 8 hours, depending on the time of the year. For example, a follow-the-sun on-call schedule has one rotation in the America/Los_Angeles time zone and one rotation in Europe/Dublin. In this example, the schedule can contain a one-hour shift gap or a one-hour shift overlap because of DST changes.

To avoid these situations, we recommend the following approach:

- 1. Use a single time zone for all rotations in an on-call schedule.
- 2. Calculate local times when you assign responders outside that particular time zone.

If you do decide to assign each rotation to its local time zone, review the schedule before any DST. Then, adjust the rotation shift times as needed to make sure that you avoid any unintended gaps or overlaps in your on-call coverage before any DST changes take effect.

To create on on-call schedule

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. Choose **Create on-call schedule**.
- For Schedule name, enter a name to help you identify the schedule, such as MyApp Primary
 On-call Schedule.
- 5. For **Schedule alias**, enter an alias for this schedule that is unique in the current AWS Region, such as **my-app-primary-on-call-schedule**.
- 6. (Optional) In the **Tags** area, apply one or more tag key name and value pairs to the on-call schedule.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For example, you can tag a schedule to identify the period of time in which it runs, the types of operators it contains, or the escalation plan it supports. For more information about tagging Incident Manager resources, see <u>Tagging resources in Incident Manager</u>.

7. Continue by adding one or more rotations to the on-call schedule.

Creating a rotation for an on-call schedule in Incident Manager

A rotation in an on-call schedule specifies when the shift is in effect. It also specifies the contacts that shifts rotate through. You can include up to eight rotations in a single on-call schedule.

You can add any individuals you created as a contact in Incident Manager to a rotation. For information about managing your contacts, see <u>Creating and configuring contacts in Incident</u> Manager.

As you configure your rotation, you can see how the overall schedule looks in a **Preview** calendar on the right side of the page.

- In the Rotation 1 section of the Create on-call schedule page, for Rotation name, enter a name that identifies the rotation, such as 00:00 - 7:59 Support, or Dublin Support Group.
- 2. For **Start date**, enter the date when this rotation becomes active in YYYY/MM/DD format, such as 2023/07/14.
- 3. For **Time zone**, select the global time zone that serves as the basis for shift coverage times and dates you specify for this rotation.

You can use any time zone defined by the Internet Assigned Numbers Authority (IANA). For example: "America/Los_Angeles", "UTC", "Asia/Seoul". For more information, see the <u>Time Zone</u> <u>Database</u> on the IANA website.

🔥 Warning

You can base each rotation on its own time zone. However, any Daylight Savings Time changes in the time zones you select can impact your intended coverage windows. For more information, see <u>Accounting for Daylight Savings Time (DST) changes</u> earlier in this topic.

4. For **Rotation start time**, enter the time when this rotation's shift begins in 24-hour hh:mm format, such as 16:00.

Note the differences in local time for contacts in time zones different from the one you specified. For example, if you choose America/Los_Angeles as the time zone and 00:00 as the rotation start time, this equals 08:00 in Dublin, Ireland, and 13:30 in Mumbai, India.

5. For **Rotation end time**, enter the time when this rotation's shift ends in 24-hour hh:mm format, such as 23:59.

🚯 Note

The length of time between the start and end of a rotation must be at least 30 minutes.

6. (Optional) To set the rotation length to 24 hours, select **24-hour coverage** and enter the start time for this rotation in the **Rotation start time** field. The **Rotation end time** value updates automatically.

For example, if you want your on-call to have 24-hour coverage with the shift change at 11 AM, choose **24-hour coverage** and enter **11:00** as the start time.

- 7. For **Active days**, select the days of the week that this rotation is active. If your on-call plan excludes weekend coverage for example, select all the days except **Sunday** and **Saturday**.
- 8. Continue by adding contacts to the rotation.

Adding contacts to a rotation in an on-call schedule in Incident Manager

For each rotation in your on-call schedule, you can add one or more contacts, up to a total of 30. You choose from contacts who are set up in your Incident Manager configuration.

When you add a contact to a rotation, the contact may receive notifications as part of their on-call duties. Notifications may be sent by email, SMS, or voice call as specified in a contact's details.

For information about managing your contacts and contacts notification options, see <u>Creating and</u> <u>configuring contacts in Incident Manager</u>.

To add contacts to a rotation in an on-call schedule

- On the Create on-call schedule page, in the Contacts section for the rotation, choose Add or remove contacts.
- 2. In the **Add or remove contacts** dialog box, select the aliases of the contacts to include in the rotation.

The order that you select the contacts in is the order that they are first listed in the rotation schedule. You can change the order after you add contacts.

- 3. Choose **Confirm**.
- To change a contact's position in the order, select the radio button for that user and use the Up
 (*)

and Down

•

buttons to update the contact order.

5. Continue by <u>specifying individual shift recurrence and length</u> for the rotation.

)

Specifying shift recurrence and length and adding tags to a rotation in Incident Manager

Shift recurrence specifies how frequently the contacts in a rotation rotate in and out of being on call. Recurrence lengths can be specified in a number of days, weeks, or months.

To specify shift recurrence and length and add tags to a rotation

- 1. On the **Create on-call schedule** page, in the **Recurrence settings** section for the rotation, do the following:
 - For **Shift recurrence type**, specify whether each on-call's shift lasts a number of days, weeks, or months by choosing from Daily, Weekly, and Monthly.
 - For **Shift length**, enter how many days, weeks, or months a shift lasts.

For example, if you chose Daily and enter **1**, each contact's on-call shift lasts one day. If you chose Weekly and enter **3**, each contact's on-call shift lasts three weeks.

2. (Optional) In the **Tags** area, apply one or more tag key name and value pairs to the rotation.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For example, you can tag a rotation to identify the location of the contacts assigned to it, the type of coverage it's meant to provide, or the escalation plan it will support. For more information about tagging Incident Manager resources, see <u>Tagging resources in Incident Manager</u>.

- 3. (Recommended) Use the calendar preview to ensure there are no unintended gaps in coverage for your on-call schedule.
- 4. Choose Create.

You can now add the on-call schedule as an escalation channel in an escalation plan. For information, see <u>Create an escalation plan</u>.

Managing an existing on-call schedule in Incident Manager

Use the content in this section to help you work with on-call schedules you have already created.

Topics

• Viewing on-call schedule details

- Editing an on-call schedule
- Copying an on-call schedule
- Creating an override for an on-call schedule rotation
- Deleting an on-call schedule

Viewing on-call schedule details

You can access an at-a-glance summary of an on-call schedule on the **View on-call schedule details** page. This page also contains information about who is currently on call and who is on call next. The page includes a calendar view that shows which contacts are on call at any specific time.

To view on-call schedule details

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. In the row for the on-call schedule to view, do one of the following:
 - To open a summary view of the calendar, choose the schedule alias.


-or-

Select the radio button for the row, and then choose View.

To open a calendar view of the schedule, choose View calendar

In calendar view, choose the name of a contact on a specific date in the schedule to see details about the assigned shift or create an override,.

• To turn on or turn off the display of a specific rotation in the calendar, choose the toggle next to the rotation's name.



Editing an on-call schedule

You can update the configuration for an on-call schedule and its rotations, except the following details:

- The schedule alias
- Rotation names
- Rotation start dates

To use an existing calendar as the basis for a new calendar with the ability to change these values, you can copy the calendar instead. For information, see <u>Copying an on-call schedule</u>.

To edit an on-call schedule

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. Do one of the following:
 - Select the radio button in the row for the on-call schedule to edit, and then choose Edit.
 - Choose the schedule alias for the on-call schedule to open the View on-call schedule details page, and then choose Edit.
- 4. Make any modifications needed to the on-call schedule and its rotations. You can change rotation configuration options such as the start and end times, contacts, and recurrence. You can add or remove rotations from the schedule as needed. The calendar preview reflects your changes as you make them.

For information about working with the options on the page, see <u>Creating an on-call schedule</u> and rotation in Incident Manager.

5. Choose Update.

🔥 Important

If you edit a schedule that contains overrides, your changes can affect the overrides. To ensure that your overrides remain configured as expected, we recommend reviewing your shift overrides closely after you update the schedule.

Copying an on-call schedule

To use the configuration of an existing on-call schedule as the starting point for a new schedule, you can create a copy of the calendar and modify it as needed.

To copy an on-call schedule

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. Select the radio button in the row for the on-call schedule to copy.
- 4. Choose **Copy**.
- 5. Make any modifications you need to the calendar and its rotations. You can change, add, or remove rotations as needed.

🚯 Note

When you copy an existing schedule, you must specify new start dates for each rotation. Copied schedules don't support rotations with start dates in the past.

For information about working with the options on the page, see <u>Creating an on-call schedule</u> and rotation in Incident Manager.

6. Choose Create copy.

Creating an override for an on-call schedule rotation

If you need to make one-off changes to an existing rotation schedule, you can create an override. An override lets you replace all or part of a contact's shift with another contact. You can also create an override that spans across multiple shifts.

You can only assign contacts to an override that are already assigned to the rotation.

In the calendar preview, overridden shifts are shown with a striped background instead of a solid background. The following image demonstrates that the contact named Zhang Wei is on call in an override. The override include parts of the shifts for John Doe and Martha Rivera, starting May 5th and ending May 11th.

Image: Schedule details Image: Schedule details Edit Deletails						
May 202 America/Los	3 Angeles (local timezone)	C	Create override	T A	Foday 🕨
Sun 30	Mon May 01 00:00 - 23:59 zhang-wei	Tue 02 00:00 - 23:59 zhang-wei	Wed 03 00:00 - 23:59 john-doe	Thu 04 00:00 - 23:59 john-doe	Fri 05 00:00 - 23:59 zhang-wei	Sat 06
07	08 00:00 - 23:59 zhang-wei	09 00:00 - 23:59 zhang-wei	10 00:00 - 23:59 zhang-wei	11 00:00 - 23:59 zhang-wei	12 00:00 - 23:59 martha-rivera	13
	15	16	17	18	19	20

To create an override for an on-call schedule

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. In the row for the on-call schedule to view, do one of the following:
 - Choose the schedule alias, then choose the **Schedule calendar** tab.
 - Choose View calendar
- 4. Do one of the following:
 - Choose Create override.
 - Choose the name of a contact in the calendar preview, and then choose **Override shift**.

5. In the **Create shift override** dialog box, do the following:

🚯 Note

An override must be at least 30 minutes in length. You can only specify an override for shifts that occur no more than six months in the future.

- a. For **Select rotation**, select the name of the rotation to create an override in.
- b. For **Start date**, select or enter the date when the override begins.
- c. For **Start time**, enter the time when the override begins in hh:mm format.
- d. For **End date**, select or enter the date when the override ends.
- e. For **End time**, enter the time when the override ends, in hh:mm format.
- f. For **Select override contact**, select the name of the rotation contact who is on call during the override period.
- 6. Choose **Create override**.

After you create an override, you can identify it by its striped background. When you choose the contact name for an overridden shift, an information box identifies it as an overridden shift. You can choose **Delete override** to remove it and restore the original on-call assignment.

Deleting an on-call schedule

When you no longer need a particular on-call schedule, you can delete it from Incident Manager.

If any escalation plans or response plans currently use the on-call schedule as an escalation channel, you should remove it from those plans before you delete the schedule.

To delete an on-call schedule

- 1. Open the Incident Manager console.
- 2. In the left navigation, choose **On-call schedules**.
- 3. Select the radio button in the row for the on-call schedule to delete.
- 4. Choose Delete.
- 5. In the **Delete on-call schedule?** dialog box, enter **confirm** in the text box.
- 6. Choose Delete.

Creating an escalation plan for responder engagement in Incident Manager

AWS Systems Manager Incident Manager provides escalation paths through your defined contacts or on-call schedules, collectively known as *escalation channels*. You can pull multiple escalation channels into an incident at the same time. If the designated contacts in the escalation channel don't respond, Incident Manager escalates to the next set of contacts. You can also choose if a plan stops escalating once a user acknowledges the engagement. You can add escalation plans to a response plan so escalation automatically starts at the beginning of an incident. You can also add escalation plans to an active incident.

Topics

- Stages
- Create an escalation plan

Stages

Escalation plans use stages where each stage lasts a defined number of minutes. Each stage has the following information:

- **Duration** The amount of time the plan waits until beginning the next stage. The first stage of the escalation plan begins once the engagement starts.
- Escalation channel An escalation channel is either a single contact or an on-call schedule composed of multiple contacts who rotate responsibilities on a defined schedule. The escalation plan engages each channel using its defined engagement plan. You can set up each escalation channel to stop the progression of the escalation plan before it continues to the next stage. Each stage can have multiple escalation channels.

For information about setting up individual contacts, see <u>Creating and configuring contacts in</u> <u>Incident Manager</u>. For information about creating on-call schedules, see <u>Managing responder</u> <u>rotations with on-call schedules in Incident Manager</u>.

Create an escalation plan

- 1. Open the Incident Manager console and choose **Escalation plans** from the left navigation.
- 2. Choose Create escalation plan.

- 3. For **Name**, enter a unique name for the escalation plan, such as **My Escalation Plan**.
- 4. For Alias, enter an alias to help you identify the plan, such as my-escalation-plan.
- 5. For **Stage duration**, enter the number of minutes for Incident Manager to wait until it continues to the next stage.
- 6. For **Escalation channel**, choose one or more contacts or on-call schedules to engage during this stage.
- 7. (Optional) To let a contact stop the escalation plan once they acknowledge the engagement, select **Acknowledgment stops plan progression**.
- 8. To add another channel to this stage, choose **Add escalation channel**.
- 9. To add another stage to the escalation plan, choose **Add stage**.
- 10. Repeat steps 5 through 9 until you finish adding the escalation channels and stages you want for this escalation plan.
- 11. (Optional) In the **Tags** area, apply one or more tag key name and value pairs to the escalation plan.

Tags are optional metadata that you assign to a resource. Tags allow you to categorize a resource in different ways, such as by purpose, owner, or environment. For example, you can tag an escalation plan to identify the type of incidents to use it for, the types of escalation channels it contains, or the escalation plan it supports. For more information about tagging Incident Manager resources, see <u>Tagging resources in Incident Manager</u>.

12. Choose Create escalation plan.

Creating and integrating chat channels for responders in Incident Manager

Incident Manager, a capability of AWS Systems Manager, gives incident responders the ability to communicate directly through *chat channels* during an incident. A *chat channel* is a chat room that you set up in <u>AWS Chatbot</u>. You then connect this channel to a response plan in Incident Manager.

During an incident, responders use the chat channel to communicate with one another about the incident. Incident Manager also pushes any updates and notifications about the incident directly to the chat channel. It sends these notifications using one or more Amazon Simple Notification Service (Amazon SNS) topics that you specify in your chat room configuration.

AWS Chatbot and Incident Manager support chat channels in the following applications:

- Slack
- Microsoft Teams
- Amazon Chime

The process for setting up a chat channel for use in your incidents consists of tasks in three different Amazon Web Services services.

Tasks

- Task 1: Create or update Amazon SNS topics for your chat channel
- Task 2: Create a chat channel in AWS Chatbot
- Task 3: Add the chat channel to a response plan in Incident Manager
- Interacting through the chat channel

Task 1: Create or update Amazon SNS topics for your chat channel

Amazon SNS is a managed service that provides message delivery from publishers to subscribers (also known as *producers* and *consumers*). Publishers communicate asynchronously with subscribers by sending messages to a *topic*, which is a logical access point and communication channel. Incident Manager uses one or more topics that you associate with a response plan to send notifications about an incident to the incident responders.

In a response plan, you can include one or more Amazon SNS topics to incident notifications. As a best practice, you should create an SNS topic in each AWS Region you have added to your replication set.

🚺 Tip

For a more linear setup workflow, we recommend that you configure your Amazon SNS topics for use with Incident Manager first. Once configured, you can create the chat channel.

To create or update Amazon SNS topics for your chat channel

1. Follow the steps in the <u>Creating an Amazon SNS topic</u> in the *Amazon Simple Notification* Service Developer Guide.

Note

After you create the topic, you edit it to update its access policy.

- 2. Select the topic that you created, and note or copy the Amazon Resource Name (ARN) of the topic, in a format such as arn:aws:sns:us-east-2:111122223333:My_SNS_topic.
- 3. Choose **Edit**, and then expand the **Access policy** section to configure additional access permissions beyond the defaults.
- 4. Add the following statement to the policy's **Statement** array:

```
{
    "Sid": "IncidentManagerSNSPublishingPermissions",
    "Effect": "Allow",
    "Principal": {
        "Service": "ssm-incidents.amazonaws.com"
    },
    "Action": "SNS:Publish",
    "Resource": "sns-topic-arn",
    "Condition": {
        "StringEqualsIfExists": {
            "AWS:SourceAccount": "account-id"
        }
    }
}
```

Replace the *placeholder values* as follows:

- sns-topic-arn is the Amazon Resource Name (ARN) of the topic that you created for this Region, in the format arn:aws:sns:us-east-2:111122223333:My_SNS_topic.
- *account-id* is the ID of the AWS account that you are working in, such as 111122223333.
- 5. Choose Save changes.
- 6. Repeat the process in each Region included in your replication set.

Task 2: Create a chat channel in AWS Chatbot

You can create a chat channel in Slack, Microsoft Teams, or Amazon Chime. You need only one chat channel for each response plan.

For your chat channels, we recommend following the principal of least privilege (not providing users with more permissions than needed to complete their tasks). You should also regularly review the membership of your AWS Chatbot chat channels. Reviews help check that only the appropriate responders and other stakeholders have access to your chat channels.

In Slack channels and Microsoft Teams channels created in AWS Chatbot, incident responders can run a number of Incident Manager CLI commands directly from the Slack or Microsoft Teams application. For more information, see Interacting through the chat channel.

🔥 Important

The users you add to your chat channel must be the same contacts listed on your escalation or response plan. You can also add additional users to chat channels, such as stakeholders and incident observers.

For general information about AWS Chatbot, see <u>What is AWS Chatbot</u> in the AWS Chatbot Administrator Guide.

Choose from the following applications to create your channel in:

Slack

The steps in this procedure provide the recommended permission settings to allow all channel users to use chat commands with Incident Manager. Using supported chat commands, your incident responders can update and interact with the incident directly from the Slack chat channel. For information, see Interacting through the chat channel.

To create a chat channel in Slack

- Follow the steps in <u>Tutorial: Get started with Slack</u> in the AWS Chatbot Administrator Guide and include the following in your configuration.
 - In step 10, for Role settings, choose Channel role.
 - In step 10d, for **Policy templates**, select **Incident Manager permissions**.
 - In step 11, for Channel guardrail policies, for Policy name, choose <u>AWSIncidentManagerResolverAccess</u>.
 - In step 12, in the **SNS topics** section, do the following:
 - For **Region 1**, select an AWS Region that is included in your replication set.

- For **Topics 1**, select the SNS topic you created in that Region to use to send incident notifications to the chat channel.
- For each additional Region in your replication set, choose **Add another Region** and add the additional Regions and SNS topics.

Microsoft Teams

The steps in this procedure provide the recommended permission settings to allow all channel users to use chat commands with Incident Manager. Using supported chat commands, your incident responders can update and interact with the incident directly from the Microsoft Teams chat channel. For information, see Interacting through the chat channel.

To create a chat channel in Microsoft Teams

- Follow the steps in <u>Tutorial: Get started with Microsoft Teams</u> in the AWS Chatbot Administrator Guide and include the following in your configuration:
 - In step 10, for **Role settings**, choose **Channel role**.
 - In step 10d, for **Policy templates**, select **Incident Manager permissions**.
 - In step 11, for **Channel guardrail policies**, for **Policy name**, choose <u>AWSIncidentManagerResolverAccess</u>.
 - In step 12, in the **SNS topics** section, do the following:
 - For **Region 1**, select an AWS Region that is included in your replication set.
 - For **Topics 1**, select the SNS topic you created in that Region to use to send incident notifications to the chat channel.
 - For each additional Region in your replication set, choose **Add another Region** and add the additional Regions and SNS topics.

Amazon Chime

To create a chat channel in Amazon Chime

- Follow the steps in <u>Tutorial: Get started with Amazon Chime</u> in the AWS Chatbot Administrator Guide and include the following in your configuration:
 - In step 11, for **Policy templates**, select **Incident Manager permissions**.

- In step 12, in the **SNS topics** section, select the SNS topics that will send notifications to the Amazon Chime webhook:
 - For Region 1, select an AWS Region that is included in your replication set.
 - For **Topics 1**, select the SNS topic you created in that Region to use to send incident notifications to the chat channel.
 - For each additional Region in your replication set, choose **Add another Region** and add the additional Regions and SNS topics.

🚯 Note

Chat commands, which incident responders can use in Slack and Microsoft Teams chat channels, are not supported in Amazon Chime.

Task 3: Add the chat channel to a response plan in Incident Manager

When you create or update a response plan, you can add chat channels for responders to communicate and receive updates through.

When following the steps in <u>Creating a response plan</u>, for the section <u>(Optional) Specifying an</u> <u>incident response chat channel</u>, select the channel you want to use for incidents related to this response plan.

Interacting through the chat channel

For channels in Slack and Microsoft Teams, Incident Manager enables responders to interact with incidents directly from the chat channel using the following ssm-incidents commands:

- start-incident
- list-response-plan
- get-response-plan
- create-timeline-event
- <u>delete-timeline-event</u>
- get-incident-record
- get-timeline-event

- list-incident-records
- list-timeline-events
- list-related-items
- update-related-items
- update-incident-record
- update-timeline-event

To run commands in an active incident's chat channel, use the following format. Replace *cli-options* with any options to be included for a command.

@aws ssm-incidents cli-options

For example:

```
@aws ssm-incidents start-incident --response-plan-arn arn:aws:ssm-
incidents::111122223333:response-plan/test-response-plan-chat --region us-east-2
```

@aws ssm-incidents create-timeline-event --event-data "\"example timeline event"\" -event-time 2023-03-31 T20:30:00.000 --event-type Custom Event --incident-record-arn
arn:aws:ssm-incidents::111122223333:incident-record/MyResponsePlanChat/98c397e6-7c10aa10-9b86-f199aEXAMPLE

@aws ssm-incidents list-incident-records

Integrating Systems Manager Automation runbooks in Incident Manager for incident remediation

You can use runbooks from <u>AWS Systems Manager Automation</u>, a capability of AWS Systems Manager, to automate common application and infrastructure tasks in your AWS Cloud environment.

Each runbook defines a *runbook workflow*, which is composed of the actions that Systems Manager performs on your managed nodes or other AWS resource types. You can use runbooks to automate the maintenance, deployment, and remediation of your AWS resources.

In Incident Manager, a runbook drives incident response and mitigation, and you specify a runbook to use as part of a response plan.

In your response plans, you can choose from dozens of pre-configured runbooks for commonly automated tasks, or you can create custom runbooks. When you specify a runbook in a response plan definition, the system can automatically start the runbook when an incident starts.

🛕 Important

Incidents created by a cross-Region failover don't invoke runbooks specified in response plans.

For more information about Systems Manager Automation, runbooks, and using runbooks with Incident Manager, see the following topics:

- To add a runbook to a response plan, see <u>Creating and configuring response plans in Incident</u> <u>Manager</u>.
- To learn more about runbooks, see <u>AWS Systems Manager Automation</u> in the AWS Systems Manager User Guide and the <u>AWS Systems Manager Automation runbook reference</u>.
- For information about the cost of using runbooks, see Systems Manager pricing.
- For information about automatically invoking runbooks when an incident is created by a Amazon CloudWatch alarm or an Amazon EventBridge event, see <u>Tutorial: Using Systems Manager</u> <u>Automation runbooks with Incident Manager</u>.

Topics

- IAM permissions required to start and run runbook workflows
- Working with runbook parameters
- Define a runbook
- Incident Manager runbook template

IAM permissions required to start and run runbook workflows

Incident Manager requires permissions to run runbooks as part of your incident response. To provide these permissions, you use AWS Identity and Access Management (IAM) roles, the *Runbook service role*, and the *Automation AssumeRole*.

The Runbook service role is a required service role. This role provides Incident Manager with the permissions it needs to access and start the workflow for the runbook.

The Automation AssumeRole provides the permissions needed to run the individual commands specified within the runbook.

Note

If no AssumeRole is specified, Systems Manager Automation attempts to use the Runbook service role for individual commands. If you don't specify an AssumeRole, you must add the necessary permissions to the Runbook service role. If you don't, the runbook fails to run those commands.

However, as a security best practice, we recommend using a separate AssumeRole. With a separate AssumeRole, you can limit the necessary permissions you must add to each role.

For more information about the Automation AssumeRole, see <u>Configuring a service role (assume</u> role) access for automations ' in the AWS Systems Manager User Guide.

You can create either type of role manually yourself in the IAM console.- You can also let Incident Manager create either one for you when you create or update a response plan.

Runbook service role permissions

Runbook service role permissions are provided through a policy similar to the following.

The first statement allows Incident Manager to start the Systems Manager StartAutomationExecution operation. This operation then runs on resources represented by the three Amazon Resource Name (ARN) formats.

The second statement allows the Runbook service role to assume a role in another account when that runbook runs in the impacted account. For more information, see <u>Running automations in</u> <u>multiple AWS Regions and accounts</u> in the *AWS Systems Manager User Guide*.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "
```

```
"Action": "ssm:StartAutomationExecution",
      "Resource": [
        "arn:aws:ssm:*:{{DocumentAccountId}}:automation-definition/{{DocumentName}}:*",
        "arn:aws:ssm:*:{{DocumentAccountId}}:document/{{DocumentName}}:*",
        "arn:aws:ssm:*::automation-definition/{{DocumentName}}:*"
      1
    },
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::*:role/AWS-SystemsManager-AutomationExecutionRole",
      "Condition": {
        "StringEquals": {
          "aws:CalledViaLast": "ssm.amazonaws.com"
        }
      }
    }
  ]
}
```

Automation AssumeRole permissions

When you create or update a response plan, you can choose from several AWS managed policies to attach to the AssumeRole that Incident Manager creates. These policies provide permissions to run a number of common operations used in Incident Manager runbook scenarios. You can choose one or more of these managed policies to provide permissions for your AssumeRole policy. The following table describes the policies that you can choose from when you create an AssumeRole from the Incident Manager console.

AWS managed policy name	Policy description
AmazonSSMAutomationRole	Grants permissions for the Systems Manager Automation service to run activities defined within runbooks. Assign this policy to administrators and trusted power users.
AWSIncidentManagerResolverAccess	Grants permission for users to start, view, and update incidents. You can also use them to create customer timeline events and related items in the incident dashboard.

You can use these managed policies to grant permissions for many common incident response scenarios. However, the permissions required for the specific tasks you need can vary. In these cases, you need to provide additional policy permissions for your AssumeRole. For information, see the AWS Systems Manager Automation runbook reference.

Working with runbook parameters

When you add a runbook to a response plan, you can specify the parameters the runbook should use at runtime. Response plans support parameters with both static and dynamic values. For static values, you enter the value when you define the parameter in the response plan. For dynamic values, the system determines the correct parameter value by collecting information from the incident. Incident Manager supports the following dynamic parameters:

Incident ARN

When Incident Manager creates an incident, the system captures the Amazon Resource Name (ARN) of the corresponding incident record and enters it for this parameter in the runbook.

í) Note

This value can only be assigned to parameters of type String. If assigned to a parameter of any other type, the runbook fails to run.

Involved resources

When Incident Manager creates an incident, the system captures the ARNs of the resources involved in the incident. These resource ARNs are then assigned to this parameter in the runbook.

About associated resources

Incident Manager can populate runbook parameter values with the ARNs of AWS resources specified in CloudWatch alarms, EventBridge events, and manually created incidents. This section describes the different types of resources for which Incident Manager can capture ARNs when populating this parameter.

CloudWatch alarms

When an incident is created from a CloudWatch alarm action, Incident Manager automatically extracts the following types of resources from the associated metrics. It then populates the chosen parameters with the following involved resources:

AWS service	Resource type
Amazon DynamoDB	Global secondary indexes
	Streams
	Tables
Amazon EC2	Images
	Instances
AWS Lambda	Function aliases
	Function versions
	Functions
Amazon Relational Database Service (Amazon	Clusters
RDS)	Database instances
Amazon Simple Storage Service (Amazon S3)	Buckets

EventBridge rules

When the system creates an incident from an EventBridge event, Incident Manager populates the chosen parameters with the Resources property in the event. For more information, see <u>Amazon</u> EventBridge events in the Amazon EventBridge User Guide.

Manually created incidents

When you create an incident by using the <u>StartIncident</u> API action, Incident Manager populates the chosen parameters by using information in the API call. Specifically, it populates parameters by using items of the type INVOLVED_RESOURCE that are passed in the relatedItems parameter.

🚯 Note

The INVOLVED_RESOURCES value can only be assigned to parameters of type StringList. If assigned to a parameter of any other type, the runbook fails to run.

Define a runbook

When creating a runbook, you can follow the steps provided here, or you can follow the more detailed guide provided in the <u>Working with runbooks</u> section in the *Systems Manager User Guide*. If you're creating a multi-account, multi-Region runbook, see <u>Running automations in multiple</u> <u>AWS Regions and accounts</u> in the *Systems Manager User Guide*.

Define a runbook

- 1. Open the Systems Manager console at <u>https://console.aws.amazon.com/systems-manager/</u>.
- 2. In the navigation pane, choose **Documents**.
- 3. Choose **Create automation**.
- 4. Enter a unique and identifiable runbook name.
- 5. Enter a description of the runbook.
- Provide an IAM role for the automation document to assume. This allows the runbook to run commands automatically. For more information, see <u>Configuring a service role access for</u> <u>Automation workflows</u>.
- 7. (Optional) Add any input parameters that the runbook starts with. You can use dynamic or static parameters when starting a runbook. Dynamic parameters use values from the incident that the runbook is started in. Static parameters use the value you provide.
- 8. (Optional) Add a **Target** type.
- 9. (Optional) Add tags.
- 10. Fill in the steps that the runbook will take when it runs. Each step requires:
 - A name.
 - A description of the purpose of the step.
 - The action to run during the step. Runbooks use the Pause action type to describe a manual step.
 - (Optional) Command properties.

11. After adding all required runbook steps, choose **Create Automation**.

To enable cross-account functionality, share the runbook in your management account with all application accounts that use the runbook during an incident.

Share a runbook

- 1. Open the Systems Manager console at https://console.aws.amazon.com/systems-manager/.
- 2. In the navigation pane, choose **Documents**.
- 3. In the documents list, choose the document you want to share and then choose View details. On the Permissions tab, verify that you're the document owner. Only a document owner can share a document.
- 4. Choose Edit.
- To share the command publicly, choose Public and then choose Save. To share the command privately, choose Private, enter the AWS account ID, choose Add permission, and then choose Save.

Incident Manager runbook template

Incident Manager provides the following runbook template to help your team start authoring runbooks in Systems Manager automation. You can use this template as is, or edit it to include details specific to your application and resources.

Find the Incident Manager runbook template

- 1. Open the Systems Manager console at https://console.aws.amazon.com/systems-manager/.
- 2. In the navigation pane, choose **Documents**.
- 3. In the **Documents** area, enter **AWSIncidents** in search field to display all Incident Manager runbooks.

🚺 Tip

Enter **AWSIncidents** - as free text instead of using the **Document name prefix** filter option.

Using a template

- 1. Open the Systems Manager console at https://console.aws.amazon.com/systems-manager/.
- 2. In the navigation pane, choose **Documents**.
- 3. Choose the template you want to update from the documents list.
- 4. Choose the **Content** tab, and then copy the content of the document.
- 5. In the navigation pane, choose **Documents**.
- 6. Choose **Create automation**.
- 7. Enter a unique and identifiable name.
- 8. Choose the **Editor** tab.
- 9. Choose Edit.
- 10. Paste or enter the copied details in the **Document editor** area.
- 11. Choose **Create automation**.

AWSIncidents-CriticalIncidentRunbookTemplate

The AWSIncidents-CriticalIncidentRunbookTemplate is a template that provides the Incident Manager incident lifecycle in manual steps. These steps are generic enough to use in most applications, but detailed enough for responders to get started with incident resolution.

Creating and configuring response plans in Incident Manager

Response plans let you plan for how to respond to an incident that impacts your users. A response plan works as a template that includes information about who to engage, the expected severity of the event, automatic runbooks to initiate, and metrics to monitor.

Best practices

You can reduce the impact on incidents on your teams when you plan for incidents ahead of time. Teams should consider the following best practices when you design a response plan.

 Streamlined engagement – Identify the most appropriate team for an incident. If you engage too wide a distribution list, or if you engage the wrong teams, you can cause confusion and waste responder time during an incident.

- Reliable escalation For your engagements in a response plan, we recommend selecting an engagement plan instead of contacts or on-call schedules. The engagement plan should specify the individual contacts or on-call schedules (which contain multiple rotating contacts) to engage during incidents. Because responders specified in your engagement plan can be unreachable at times, you should configure backup responders in your response plan to cover these scenarios. With backup contacts, if the primary and secondary contacts are unavailable or there are other unplanned gaps in coverage, Incident Manager still notifies a contact about the incident.
- **Runbooks** Use runbooks to provide repeatable, understandable steps that reduce the stress a responder experiences during an incident.
- Collaboration Use chat channels to streamline communication during incidents. Chat channels help responders stay up to date with information. They can also share information with other responders through these channels.

Creating a response plan

Use the following procedure to create a response plan and automate incident response.

To create a response plan

- 1. Open the Incident Manager console, and in the navigation pane, choose **Response plans**.
- 2. Choose Create response plan.
- 3. For **Name**, enter a unique and identifiable response plan name to use in the Amazon Resource Name (ARN) for the response plan.
- 4. (Optional) For **Display name**, enter a more human readable name to help identify the response plan when you create incidents.
- 5. Continue by specifying default values for incident records.

Specifying incident default values

To help you manage incidents more effectively, you can specify default values. Incident Manager applies these values to all incidents that are associated with a response plan.

To specify incident default values

1. For **Title**, enter a title for this incident to help you identify it on the Incident Manager home page.

- 2. For **Impact**, choose an impact level to indicate the potential scope of an incidents created from this response plan, such as **Critical** or **Low**. For information about impact ratings in Incident Manager, see Triage.
- 3. (Optional) For **Summary**, enter a brief summary the type of incidents created from this response plan.
- 4. (Optional) For **Dedupe string**, enter a dedupe string. Incident Manager uses this string to prevent the same root cause from creating multiple incidents in the same account.

A deduplication string is a term or phrase the system uses to check for duplicate incidents. If you specify a deduplication string, Incident Manager searches for open incidents that contain the same string in the dedupeString field when it creates the incident. If a duplicate is detected, Incident Manager deduplicates the newer incident into the existing incident.

🚺 Note

By default, Incident Manager automatically deduplicates multiple incidents created by the same Amazon CloudWatch alarm or Amazon EventBridge event. You don't have to enter your own deduplication string to prevent duplication for these resource types.

5. (Optional) Under **Incident Tags**, add tag keys and values to assign to incidents created from this response plan.

You must have the TagResource permission for the incident record resource to set incident tags within the response plan.

6. Continue by <u>specifying an optional chat channel</u> for resolvers to communicate with one another about incidents.

(Optional) Specifying an incident response chat channel

When you include a chat channel in a response plan, responders receive incident updates through the channel. They can interact with the incident directly from the chat channel by using chat commands.

Using AWS Chatbot, you can create a channel for Slack, for Microsoft Teams, or for Amazon Chime to use in your response plans. For information about creating a chat channel in AWS Chatbot, see the <u>AWS Chatbot Administrator Guide</u>.

<u> Important</u>

Incident Manager must have permissions to publish to a chat channel's Amazon Simple Notification Service (Amazon SNS) topic. Without permissions to publish to that SNS topic, you can't add it to the response plan. Incident Manager publishes a test notification to the SNS topic to verify permissions.

For more information about chat channels, see <u>Creating and integrating chat channels for</u> responders in Incident Manager.

To specify an incident response chat channel

1. For **Chat channel**, select an AWS Chatbot chat channel where responders can communicate during an incident.

🚺 Tip

To create a new chat channel in AWS Chatbot, choose **Configure new Chatbot client**.

- 2. For **Chat channel SNS topics**, choose additional SNS topics to publish to during the incident. Adding SNS topics in multiple AWS Regions increases redundancy in case a Region is down at the time of the incident.
- 3. Continue by <u>selecting the contacts, on-call schedules, and escalation plans</u> to be engaged during an incident.

(Optional) Select resources to engage in incident response

It's important to identify the most appropriate responders when an incident occurs. As a best practice, we recommend that you do the following:

- 1. Add contacts and on-call schedules as the escalation channels in an escalation plan.
- 2. Choose an escalation plan as the engagement in a response plan.

For more information about contacts and escalation plans, see <u>Creating and configuring contacts in</u> Incident Manager and <u>Creating an escalation plan for responder engagement in Incident Manager</u>.

To select resources to engage in incident response

- 1. For **Engagements**, choose any number of escalation plans, on-call schedules, and individual contacts.
- 2. Continue by optionally <u>specifying a runbook to run</u> as part of your incident mitigation.

(Optional) Specifying a runbook for incident mitigation

You can use runbooks from <u>AWS Systems Manager Automation</u>, a capability of AWS Systems Manager, to automate common application and infrastructure tasks in your AWS Cloud environment.

Each runbook defines an *runbook workflow*. A runbook workflow includes the actions that Systems Manager performs on your managed nodes or other AWS resource types. In Incident Manager, a runbook drives incident response and mitigation.

For more information about using runbooks in response plans, <u>Integrating Systems Manager</u> Automation runbooks in Incident Manager for incident remediation.

To specify a runbook for incident mitigation:

- 1. For **Runbook**, do one of the following:
 - Choose **Clone runbook from template** to make a copy of the default Incident Manager runbook. For **Runbook name**, enter a descriptive name for the new runbook.
 - Choose Select existing runbook. Select the Owner, Runbook, and Version to use.

🚺 Tip

To create a runbook from scratch, choose **Configure new runbook**. For information about creating runbooks, see <u>Integrating Systems Manager</u> Automation runbooks in Incident Manager for incident remediation.

2. In the **Parameters** area, supply any parameters requested for the runbook you selected.

The available parameters are those specified by the runbook. One runbook might require different parameters than another. Some parameters might be required and others optional.

In many cases, you can choose to manually enter a static value for a parameter, such as a list of Amazon EC2 instance IDs. You can also let Incident Manager provide the parameter values that were dynamically generated by an incident.

3. (Optional) For **AutomationAssumeRole**, specify the AWS Identity and Access Management (IAM) role to use. This role must have the permissions needed to run the individual commands specified within the runbook.

🚯 Note

If no AssumeRole is specified, Incident Manager attempts to use the Runbook service role to run the individual commands specified within the runbook.

Choose from the following:

- Enter ARN value Manually enter the Amazon Resource Name (ARN) of an AssumeRole, in the format arn: aws:iam::account-id:role/assume-role-name. For example, arn:aws:iam::123456789012:role/MyAssumeRole.
- Use existing service role Choose a role with the required permissions from a list of existing roles in your account.
- Create new service role Choose from among AWS managed policies to attach to your AssumeRole. After selecting this option, for AWS managed policies, choose one or more policies from the list.

You can accept the suggested default name for the new role, or enter a name that you choose.

🚺 Note

This new Runbook service role is associated with the specific runbook that you selected. It can't be used with different runbooks. This is because the Resource section of the policy won't support other runbooks.

4. For **Runbook service role**, specify the IAM role to use to provide the permissions needed to access and start the workflow for the runbook itself.

At minimum, the role must allow the ssm: StartAutomationExecution action for your specific runbook. For the runbook to work across accounts, the role must also allow the sts:AssumeRole action for the AWS-SystemsManager-AutomationExecutionRole role that you created during <u>Managing incidents across AWS accounts and Regions in Incident</u> Manager.

Choose from the following:

• **Create new service role** – Incident Manager creates a Runbook service role for you that includes the minimum required permissions to start the runbook workflow.

For **Role name**, you can accept the suggested default name, or enter a name that you choose. We recommend using the suggested name or keeping the name of the runbook in the name. This is because the new AssumeRole is associated with the specific runbook you selected and might not include the permissions required for other runbooks.

• Use existing service role – An IAM role that you or Incident Manager created previously grants the needed permissions.

For **Role name**, select the name of the existing role to use.

- 5. Expand **Additional options** and choose one of the following to specify the AWS account where the runbook workflow should run.
 - **Response plan owner's account** Start the runbook workflow in the AWS account that created it.
 - **Impacted account** Start the runbook workflow in the account that began or reported the incident.

Choose **Impacted account** when you use Incident Manager for cross-account scenarios and the runbook needs to access resources in the impacted account to remediate them.

6. Continue by optionally integrating a PagerDuty service into the response plan.

(Optional) Integrating a PagerDuty service into the response plan

To integrate a PagerDuty service into the response plan

When you integrate Incident Manager with PagerDuty, PagerDuty creates a corresponding incident whenever Incident Manager creates an incident. The incident in PagerDuty uses the paging workflow and escalation policies that you defined there in addition to those in Incident Manager. PagerDuty attaches timeline events from Incident Manager as notes on your incident.

- 1. Expand **Third-party integrations**, then choose the **Enable PagerDuty integration** check box.
- 2. For **Select secret**, select the secret in AWS Secrets Manager where you store the credentials to access your PagerDuty account.

For information about storing your PagerDuty credentials in a Secrets Manager secret, see Storing PagerDuty access credentials in an AWS Secrets Manager secret.

- 3. For **PagerDuty service**, select the service from your PagerDuty account where you want to create the PagerDuty incident.
- 4. Continue by adding optional tags and creating the response plan.

Adding tags and creating the response plan

To add tags and create the response plan

1. (Optional) In the **Tags** area, apply one or more tag key name/value pairs to the response plan.

Tags are optional metadata that you assign to a resource. With tags, you can categorize a resource in different ways, such as by purpose, owner, or environment. For example, you might want to tag a response plan to identify the type of incident it is meant to mitigation, the types of escalation channels it contains, or the escalation plan that will be associated with it. For more information about tagging Incident Manager resources, see <u>Tagging resources in Incident</u> <u>Manager</u>.

2. Choose Create response plan.

Identifying potential causes of incidents from other services as "findings" in Incident Manager

In Incident Manager, a *finding* is information about an AWS CodeDeploy deployments or AWS CloudFormation stack update that occurred around the time of an incident, and that involved one or more resources likely related to the incident. Each finding can be examined as a potential cause of the incident. Information about these potential causes is added to the **Incident details** page for an incident. With information about these deployments and changes readily at hand, responders don't need to manually search for this information. This lessens the time needed to evaluate potential causes, which can reduce the mean time to recover (MTTR) from an incident.

Currently, Incident Manager supports gathering findings from two AWS services: <u>AWS CodeDeploy</u> and <u>AWS CloudFormation</u>.

Findings is an opt-in feature. You can enable it in the <u>Get prepared wizard</u>, when you are first onboarding to Incident Manager, or later on the <u>Settings page</u>.

When you enable the Findings feature, Incident Manager creates a service role for you. This service role includes the permissions needed to retrieve findings from CodeDeploy and CloudFormation.

To work with findings in a cross-account scenario, enable the feature in the management account. After that, each application account in an AWS Resource Access Manager (AWS RAM) organization must create a corresponding service role.

Refer to the following topics to help you use the Findings feature.

Topics

- Enable and create a service role for findings
- Configure permissions for cross-account findings support

Enable and create a service role for findings

When you enable the Findings feature, Incident Manager creates a service role named IncidentManagerIncidentAccessServiceRole on your behalf. This service role provides the permissions Incident Manager needs to gather information about CodeDeploy deployments and CloudFormation stack updates that occurred around the time an incident was created.

🚯 Note

If you are using Incident Manager with an organization, the service role is created in the management account. To work with findings across other accounts in the organization, the service role must be created in each application account. For information about using a CloudFormation template to create this role in your application accounts, see step 4 in <u>Set</u> up and configure cross-account incident management.

This service role is associated with an AWS managed policy. For information about the permissions in this policy, see <u>AWS managed policy: AWSIncidentManagerIncidentAccessServiceRolePolicy</u>.

For information about enabling findings during the Incident Manager onboarding process, see <u>Getting started with Incident Manager</u>.

For information about enabling findings after you have completed the onboarding process, see <u>Managing the Findings feature</u>.

Configure permissions for cross-account findings support

To use the Findings feature across accounts with an organization set up in AWS RAM, each application account must configure permissions for Incident Manager to assume the management account's service role on its behalf.

These permissions can be configured in an application account by deploying an AWS CloudFormation template provided by AWS, which creates the role IncidentManagerIncidentAccessServiceRole.

For information about downloading and deploying this template in an application account, see step 4 in Managing incidents across AWS accounts and Regions in Incident Manager.

Creating incidents automatically or manually in Incident Manager

Incident Manager, a capability of AWS Systems Manager, helps you manage and quickly respond to incidents. You can configure Amazon CloudWatch and Amazon EventBridge to automatically create incidents based on CloudWatch alarms and EventBridge events. You can also create incidents manually on the incident list page or by using the <u>StartIncident</u> API action from the AWS CLI or the AWS SDK. Incident Manager deduplicates incidents created from the same CloudWatch alarm or EventBridge event into the same incident.

For incidents automatically created by CloudWatch alarms or EventBridge events, Incident Manager attempts to create an incident in the same AWS Region as the event rule or alarm. In the event that Incident Manager is not available in the AWS Region, CloudWatch or EventBridge automatically create the incident in one of the available Regions specified in your replication set. For more information, see <u>Managing incidents across AWS accounts and Regions in Incident Manager</u>.

When the system creates an incident, Incident Manager automatically collects information about the AWS resources involved in the incident and adds this information to the **Related items** tab. If you specified a runbook in your response plan, when the system creates an incident, Incident Manager can send the information about the AWS resources involved in the incident to the runbook. The system can then target those resources when it initiates the runbook and attempts to remediate the issue.

When the system creates an incident, it also creates a parent operational workitem (OpsItem) in OpsCenter, a component of Systems Manager, and links it to the incident as a related item. You can use this OpsItem to track related work and future incident analyses. Calls to OpsCenter incur costs. For more information about OpsCenter pricing, see <u>Systems Manager pricing</u>.

🛕 Important

Note the following important details.

 In the event that Incident Manager is not available, the system can only fail over and create incidents in other AWS Regions if you have specified at least two Regions in your replication set. For information about configuring a replication set, see <u>Getting started</u> with Incident Manager. Incidents created by a cross-Region failover don't invoke runbooks specified in response plans.

Creating incidents automatically with CloudWatch alarms

CloudWatch uses your CloudWatch metrics to alert you about changes in your environment and to automatically perform the start incident action. CloudWatch works with Systems Manager and Incident Manager to create an incident from a response plan template when an alarm goes into alarm state. This requires the following prerequisites:

- Incident Manager configured and replication set created. This step creates the Incident Manager service linked role in your account, providing the necessary permissions.
- A configured Incident Manager response plan. To learn how to configure Incident Manager response plans, see <u>Creating and configuring response plans in Incident Manager</u> in the *Incident preparation* section of this guide.
- Configured CloudWatch metrics monitoring your application. For monitoring best practices, see Monitoring in the Incident preparation section of this guide.

To create an alarm with a Start incident action

- 1. Create an alarm in CloudWatch. For more information, see <u>Using Amazon CloudWatch alarms</u> in the *Amazon CloudWatch User Guide*.
- 2. When choosing the action for the alarm to perform, select Add Systems Manager action.
- 3. Choose **Create incident** and select the **Response plan** for this incident.
- 4. Complete the remaining steps in your selected alarm type guide.

🚺 Tip

You can also add the create incident action to any existing alarm.

Creating incidents automatically with EventBridge events

EventBridge rules watch for event patterns. If the event matches the defined pattern, Incident Manager creates an incident using the chosen response plan.

Creating incidents using SaaS partners events

You can configure EventBridge to receive events from software as a service (SaaS) partner applications and services, allowing for third-party integration. After configuring EventBridge to receive events from third-party partners, you can create rules that match on partner events to create incidents. To see a list of third-party integrations, see <u>Receiving events from a SaaS partner</u>.

Configure EventBridge to receive events from a SaaS integration.

- 1. Open the Amazon EventBridge console at https://console.aws.amazon.com/events/.
- 2. In the navigation pane, choose **Partner event sources**.
- 3. Use the search bar to find the partner that you want and choose **Set up** for that partner.
- 4. Choose **Copy** to copy your account ID to the clipboard.

🚺 Note

To integrate with Salesforce use the steps described in the <u>Amazon AppFlow user</u> guide.

- 5. Go to the partner's website and follow the instructions to create a partner event source. Use your account ID for this. The event source that you create is available only on your account.
- 6. Go back to the EventBridge console and choose **Partner event sources** in the navigation pane.
- 7. Select the button next to the partner event source, and choose Associate with event bus.

Create a rule that triggers on events from a SaaS partner

- 1. Open the Amazon EventBridge console at <u>https://console.aws.amazon.com/events/</u>.
- 2. In the navigation pane, choose **Rules**.
- 3. Choose Create rule.
- 4. Enter a name and description for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

- 5. For **Event bus**, choose the event bus that corresponds to this partner.
- 6. For **Rule type**, choose **Rule with an event pattern**.
- 7. Choose Next.
- 8. For Event source, choose AWS events or EventBridge partner events.
- 9. For **Event pattern**, choose **Event pattern form**.
- 10. For Event source, choose EventBridge partners
- 11. For **Partners**, choose the name of the partner.
- 12. For **Event type**, choose **All Events** or choose the type of event to use for this rule. If you choose **All Events**, all events emitted by this partner event source will match the rule.

If you want to customize the event pattern, choose **Edit**, make your changes, and then choose **Save**.

- 13. Choose Next.
- 14. For **Select a target**, choose **Incident Manager response plan**, and then choose a **Response plan**.

🚺 Note

When selecting a response plan, all response plans that you own and have been shared with your account appear in the **Response plan** dropdown list.

- 15. EventBridge can create the IAM role needed for your rule to run:
 - To create an IAM role automatically, choose Create a new role for this specific resource.
 - To use an IAM role that you created before, choose **Use existing role**.
- 16. Choose Next.
- 17. (Optional) Enter one or more tags for the rule. For more information, see <u>Amazon EventBridge</u> <u>tags</u> in the *Amazon EventBridge User Guide*.
- 18. Choose Next.
- 19. Review your rule then choose **Create rule**.

Creating incidents using AWS service events

EventBridge also receives events from the AWS services listed in <u>Events from Supported AWS</u> <u>Services</u>. Similar to how you configure rules for SaaS partners, you can configure them for AWS services.

Create a rule that triggers on events from an AWS service

- 1. Open the Amazon EventBridge console at <u>https://console.aws.amazon.com/events/</u>.
- 2. In the navigation pane, choose **Rules**.
- 3. Choose Create rule.
- 4. Enter a name and description for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

- 5. For **Event bus**, choose **default**.
- 6. For **Rule type**, choose **Rule with an event pattern**.
- 7. Choose Next.
- 8. For **Event source**, choose **AWS events or EventBridge partner events**.
- 9. For **Event pattern**, choose **Event pattern form**.
- 10. For **Event source**, choose **AWS services**.
- 11. For **Service name**, choose the service that monitors for an incident.
- 12. For **Event type**, choose **All Events** or choose the type of event to use for this rule. If you choose **All Events**, all events emitted by this partner event source will match the rule.

If you want to customize the event pattern, choose **Edit**, make your changes, and then choose **Save**.

- 13. Choose Next.
- For Select a target, choose Incident Manager response plan, and then choose a Response plan.

1 Note

When selecting a response plan, all response plans that you own and have been shared with your account appear in the **Response plan** dropdown list.

15. EventBridge can create the IAM role needed for your rule to run:

- To create an IAM role automatically, choose **Create a new role for this specific resource**.
- To use an IAM role that you created before, choose **Use existing role**.
- 16. Choose Next.
- 17. (Optional) Enter one or more tags for the rule. For more information, see <u>Amazon EventBridge</u> <u>tags</u> in the *Amazon EventBridge User Guide*.
- 18. Choose Next.
- 19. Review your rule then choose **Create rule**.

Creating incidents manually

Responders can manually track an incident using the Incident Manager console by using a predefined response plan. Use the following steps to create an incident.

- 1. Open the Incident Manager console.
- 2. Choose Start incident.
- 3. For **Response plan**, choose a response plan from the list.
- 4. (Optional) To override the title provided by the defined response plan, enter an **Incident title**.
- 5. (Optional) To override the impact provided by the defined response plan, enter the **Impact** of the incident.

Viewing incident details in the Incident Manager console

AWS Systems Manager Incident Manager tracks your incidents from the moment they're detected to resolution and through post-incident analysis. You can find all incidents on the **Incident list** page in the Incident Manager console, with links directly to the **Incident details**.

Topics

- Viewing the incident list in the console
- Viewing incident details in the console

Viewing the incident list in the console

The **Incident list** page contains three sections: **Open incidents**, **Resolved incidents**, and **Analyses**. You can manually track new incidents and create analyses from this page. To learn more about manually tracking an incident, see <u>Creating incidents manually</u> in the *Incident creation* section of this guide. To learn about post-incident analysis, see the <u>Performing a post-incident analysis in</u> <u>Incident Manager</u> section of this guide.

The **Incident details** displays **Open incidents** in tiles with the title, impact, duration, and chat channel for that incident. After you resolve an incident, it moves to the **Resolved incidents** list. **Analyses** are in the second tab.

Viewing incident details in the console

The **Incident details** page provides detailed insights and tools you can use to manage an incident. From this page, you can start runbooks to mitigate an incident, add incident notes, engage other resolvers, and view incident details such as timelines, metrics, properties, and related resources.

As shown in the following image, the **Incident details** page includes several sections: Top banner, **Incident notes**, and seven tabs that contain additional information and resources. By default, the Top banner and **Incident notes** sections are displayed on all **Incident details** pages.

AWS Systems Manager > Incident Manager > I	Incident notes (2)	>			
Incident 1		C Refresh interval: 30 seconds	Edit properties Resolve incident	Use incident notes to highlight messages related to the incident	
Status Open	Impact	Chat channel 🛃	Duration 2m	Add incident note	
Tasks	Runbooks	Diagnosis	Engagements	November 8, 2023 at 12:31:50 (UTC-5:00) Work in progress to mitigate the impact and	:
- 12	1 waiting for input	-	-	runbook is in progress. November 8, 2023 at 12:30:38 (UTC-5:00)	:
Overview Diagnosis Timeline 10	On-call has been notified and impact is being assessed.				
Summary			Edit	п	
Pending tasks (0)					

This topic explains elements of the **Incident details** page and actions that you can perform from the page.

Top banner

The top banner on every incident details page includes the following information:

- Status The current status of an incident can be Open or Resolved.
- Impact The impact of the incident on your environment. It can be high, medium, and low. To change the impact of an incident, choose Edit properties.
- Chat channel A link to access the chat channel where you can view incident updates and notifications.
- Duration The amount of time lapsed before a responder resolves the incident.
- Runbooks The statuses for the runbooks associated with this incident. The status can be waiting for input, successful, or unsuccessful. If a runbook's status is waiting for input, you can select the runbook to view action details. You can select unsuccessful to view runbooks that are Timed out, Failed, or Canceled.
- Engagements The total number of engagements and the status of each engagement. When you create an engagement, its status is Engaged. Once you acknowledge the engagement, the status changes from Engaged to Acknowledged. Incident Manager doesn't support acknowledgement of third-party engagements. Such engagements remain in the Engaged status.

You can edit the incident title, impact, and chat channel by choosing **Edit** in the top-right corner of the banner.

Incident notes

The right side of the screen displays the **Incident notes** section. With notes, you can collaborate and communicate with other users that work on an incident. You can explain the mitigations that you applied, a potential root cause you identified, or the current status of the incident. As a best practice, use the **Incident notes** section to post status updates and actions you or others take on an incident. If you need to communicate with other resolvers in real time, use the chat channel available in Incident Manager.

To add a note, choose the **Add incident note** button, and then enter your note. Notes can contain updates about incident status or any other relevant information that provides visibility to other users. If required, you can also edit or delete incident notes.

🚯 Note

Any user with IAM permission to run the ssm-incidents:UpdateTimelineEvent and ssm-incidents:DeleteTimelineEvent actions can edit and delete notes. However, when you share an incident with another account, the resource policy doesn't include the ssm-incidents:DeleteTimelineEvent action. This prevents the user that you share the incident with from deleting the note. You can view the audit trail for a note from Incident Manager events in the AWS CloudTrail console.

Tabs

The incident details page has seven tabs, making it easier for responders to locate and view information during an incident. The tabs display a counter in the tab name, which indicates the number of updates to the tab. For more information about the contents of each tab as well as available actions, continue reading.

Overview

The **Overview** tab is the landing page for responders. It contains the incident summary, a list of recent timeline events, and the current runbook step.

Responders use the **Summary** to catch up on what actions have been taken, the results of any changes, possible next steps, and information about the impact of the incident. To update the summary, choose **Edit** in the top-right corner of the **Summary** section.

<u> Important</u>

If multiple responders are editing the summary field simultaneously, the responder who submits their edits last overwrites all other input.

The **Recent timeline events** section contains a timeline populated by Incident Manager with the five most recent events. Use this section to understand the status of the incident and what has recently occurred. To view a complete timeline, continue to the **Timeline** tab.

The overview page also displays the **Current runbook step**. This step might be an automatic step running in your AWS environment, or it may be a set of manual instructions for responders. To view the complete runbook, including prior and upcoming steps, choose the **Runbook** tab.

Diagnosis

The **Diagnosis** tab contains vital information about your AWS hosted applications and systems, including information about metrics and, if enabled, findings.

Working with metrics

Incident Manager uses Amazon CloudWatch to populate the metrics and alarm graphs found on this tab. To learn more about incident management best practices for defining alarms and metrics, see <u>Monitoring</u> in the *Incident planning* section of this user guide.

To add metrics

- Choose **Add** in the upper-right corner of this tab.
 - To add a metric from an existing CloudWatch dashboard, choose From existing
 CloudWatch dashboard.
 - a. Choose a **Dashboard**. This adds all metrics and alarms that are part of the chosen dashboard.
 - b. (Optional) You can also **Select metrics** from the dashboard to view specific metrics.

- Add a single metric by selecting From CloudWatch and pasting a metric source. To copy a metric source:
 - a. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
 - b. In the navigation pane, choose **Metrics**.
 - c. On the **All metrics** tab, enter a search term in the search field, such as a metric name or resource name, and choose **Enter**.

For example, if you search for the CPUUtilization metric, you will see the namespaces and dimensions associated with this metric.

- d. Choose one of the results from your search to view the metrics.
- e. Choose the **Source** tab and copy the source.

Metric alarm graphs can only be added to the incident details through the related response plan, or by selecting **From existing CloudWatch dashboard** when adding a metric.

To remove metrics, choose **Remove**, and then choose the metrics you want to remove from the provided **Metrics** dropdown.

Viewing findings from AWS CodeDeploy and AWS CloudFormation

After Findings is enabled and all required permissions configured, any findings that might be related to a specific incident are attached to the incident. Responders can view information about these findings on the **Incident details** page.

To view findings from CodeDeploy and CloudFormation

- 1. Open the Incident Manager console.
- 2. Choose the name of an incident to investigate.
- 3. On the **Diagnosis** tab, in the **Findings** area, compare the start times of any reported finding with the start time of the incident.
- 4. To view more details about a finding, in the **Reference** column, choose the link to the CodeDeploy or CloudFormation finding.

Timeline

Use the **Timeline** tab to track events that occur during an incident. Incident Manager automatically populates timeline events that identify significant occurrences during the incident. Responders can add custom events based on occurrences that are detected manually. During the post-incident analysis, the timeline tab provides valuable insights into how to better prepare and respond to incidents in the future. For more information about post-incident analysis, see <u>Performing a post-incident analysis in Incident Manager</u>.

To add a custom timeline event, choose **Add**. Select a date using the calendar, and then enter a time. All times are shown in your local time zone. Provide a brief description of the event that appears in the timeline.

To edit an existing custom event, select the event on the timeline and choose **Edit**. You can change the time, date, and description of *custom* events. You can only edit custom events.

Runbooks

The **Runbooks** tab of the incident details page is where responders can view runbook steps and start new runbooks.

To start a new runbook, choose **Start runbook** in the **Runbooks** section. Use the search field to find the runbook you want to start. Provide any required **Parameters** and the **Version** of the runbook you want to use when starting the runbook. Runbooks started during an incident from the **Runbooks** tab use the permissions of the currently signed-in account.

To navigate to a runbook definition in Systems Manager, choose the runbook's title under **Runbooks**. To navigate to the running instance of the runbook in Systems Manager, choose the execution details under **Execution details**. These pages display the template used to start the runbook and the specific details of the currently running instance of the automation document.

The **Runbook steps** section displays the list of steps that the selected runbook automatically takes or responders manually perform. The steps expand as they become the current step, displaying information required to complete the step, or details about what the step does. Automatic runbook steps resolve after the automation is complete. Manual steps require the responders to choose **Next step** at the bottom of each step. After a step is complete, the step output appears as a dropdown.

To cancel a runbook execution, choose **Cancel runbook**. This will stop the execution of the runbook and not complete any further steps in the runbook.

Engagements

The **Engagements** tab of the incident details drives the engagement of responders and teams. From this tab, you can see who has been engaged, who has responded, as well as which responders are going to be engaged as part of an escalation plan. Responders can engage other contacts directly from this tab. To learn more about creating contacts and escalation plans, see the <u>Creating</u> <u>and configuring contacts in Incident Manager</u> and <u>Creating an escalation plan for responder</u> <u>engagement in Incident Manager</u> sections of this guide.

You can configure response plans with contacts and escalation plans to automatically start engagement at the beginning of an incident. To learn more about configuring response plans, see the <u>Creating and configuring response plans in Incident Manager</u> section of this guide.

You can find information about each contact in the table. This table includes the following information:

- Name Links to the contact details page that displays their contact methods and engagement plan.
- Escalation plan Links to the escalation plan that engaged the contact.
- Contact source Identifies the service that engaged this contact, such as AWS Systems Manager or PagerDuty.
- **Engaged** Displays when the plan engaged a contact, or when to engage a contact as part of an escalation plan.
- Acknowledged Displays whether the contact acknowledged the engagement.

To acknowledge an engagement, the responder can do one of the following:

- Phone call Enter **1** when prompted.
- SMS Reply to the message with the provided code, or enter the provided code on the **Engagements** tab of the incident.
- Email Enter the provided code on the **Engagements** tab of the incident.

Related items

The **Related items** tab is used to collect resources related to incident mitigation. These resources can be ARNs, links to external resources, or files uploaded to Amazon S3 buckets. The table

displays a descriptive title and either an ARN, a link, or bucket details. Before using S3 buckets, review <u>Security Best Practices for Amazon S3</u> in the *Amazon S3 User Guide*.

When uploading files to an Amazon S3 bucket, versioning is either enabled or suspended on that bucket. When versioning is enabled on the bucket, files uploaded with the same name as an existing file are added as a new version of the file. If versioning is suspended, files uploaded with the same name as an existing file overwrite the existing file. To learn more about versioning, see Using versioning in S3 buckets in the Amazon S3 User Guide.

When removing a file-related item, the file is removed from the incident but is not removed from the Amazon S3 bucket. To learn more about removing objects from an Amazon S3 bucket, see <u>Deleting Amazon S3 objects</u> in the *Amazon S3 User Guide*.

Properties

The **Properties** tab provides the following details about the incident.

In the Incident properties section, you can view the following:

- Status Describes the current status of the incident. The incident can be Open or Resolved.
- **Start time** The time when the incident was created in Incident Manager.
- **Resolved time** The time that the incident was resolved in Incident Manager.
- Amazon Resource Name (ARN) The ARN of the incident. Use the ARN when referencing the incident from the chat or with AWS Command Line Interface (AWS CLI) commands.
- **Response Plan** Identifies the response plan for the selected incident. Choosing the response plan opens the response plan's details page.
- **Parent Opsitem** Identifies the Opsitem created as the parent of the incident. A parent Opsitem can have multiple related incidents and follow-up action items. Selecting the parent Opsitem opens the Opsitems details page in OpsCenter.
- **Analysis** Identifies the analysis created from this incident. Create an analysis from a resolved incident to improve your incident response process. Choose the analysis to open the analysis details page.
- Owner The account in which the incident was created.

In the **Tags** section, you can view and edit the tag keys and values associated with the incident record. For more information about tags in Incident Manager, see <u>Tagging resources in Incident</u> <u>Manager</u>.

Performing a post-incident analysis in Incident Manager

Post-incident analysis guides you through identifying improvements to your incident response, including time to detection and mitigation. An analysis can also help you understand the root cause of the incidents. Incident Manager creates recommended *action items* to improve your incident response.

Benefits of a post-incident analysis

- Improve incident response
- Understand the root cause of the problem
- Address root causes with deliverable action items
- Analyze the impact of incidents
- Capture and share learnings within an organization

What not to use an analysis for

An analysis is blameless and doesn't call out people by name.

"Regardless of what we discover, we understand and truly believe that everyone did the best job they could, given what they knew at the time, their skills and abilities, the resources available, and the situation at hand." - Norm Kerth, Project Retrospectives: A Handbook for Team Review

Analysis details

The analysis details page guides you through gathering information, assessing improvements, and creating action items. The analysis details page is similar to the incident details with some key differences such as historical metrics, editable timeline, and questions to improve future incidents.

Overview

The overview is a summary of the incident. This summary includes background, what occurred, why it happened, how it was mitigated, duration, and key action items to prevent the incident from happening again. The overview is high level. You'll explore more details in the **Questions** tab of the analysis.

Metrics

Use the metrics tab to visualize key metrics in your application over the duration of the incident. You can add metric graphs here that have one or more metrics depicted in the same graph. Metrics used during an incident are automatically populated on this tab. We recommend you adding a description, title, and annotations of key timepoints during the incident.

Some key time points you can consider when analyzing a metric graph:

- Deployment change
- Configuration change
- Incident start time
- Alarm time
- Time of engagement
- Mitigation start time
- Incident resolved time

Limitations

- CloudWatch alarms and metric expressions aren't imported from an incident.
- Metrics that are in a Region that Incident Manager doesn't support aren't imported from the incident.
- Metrics in application accounts require configuration of the CloudWatch-CrossAccountSharingRole prior to creating the analysis. For more information about the role, see Cross-Account Cross-Region CloudWatch console in the CloudWatch user guide.

Timeline

Describe key time points on the timeline as you dive deeper into understanding the incident. The incidents timeline is automatically populated on this tab. You can delete timepoints that aren't relevant to the analysis. You can also add and edit time points to more accurately describe the incident and its impact.

Use the timeline tab to answer questions you find on the **Questions** tab about the incident response.

Questions

Use Incident Manager questions to improve the time to resolution of incidents in your application and reduce the occurrence of incidents. As you answer questions, update the **Metrics** and **Timeline** tabs for accuracy. The questions focus on these key aspects of incident response:

- Detection Could you improve time to detection? Are there updates to metrics and alarms that would detect the incident sooner?
- Diagnosis Can you improve the time to diagnosis? Are there updates to your response plans or escalation plans that would engage the correct responders sooner?
- Mitigation Can you improve the time to mitigation? Are there runbook steps that you could add or improve?
- Prevention Can you prevent future incidents from occurring? To discover the root causes of an incident, Amazon uses the 5-Whys approach in problem investigation.

Actions

Incident Manager creates recommended action items for you to review as you complete the questions. You can choose to accept and complete these actions from this tab or you can dismiss these actions. You can review dismissed action items by choosing **Dismissed action items**. Action items are a type of OpsItem that are linked to the analysis and incident in OpsCenter.

Checklist

Before closing an analysis, use the checklist to review actions that a responder should take. As responders complete actions in the checklist, the icon next to the action changes from an ellipse to a check-mark, indicating that the action is complete. If you haven't completed checklist items, Incident Manager displays a message to confirm the responder wants to close the analysis without completing it.

Analysis templates

An analysis template provides a set of questions that dive deep into the root cause of incidents. You can use your answers to these questions to improve application performance and incident response.

AWS standard template

Incident Manager provides a standard template of questions based on AWS incident response and problem analysis best practices, titled AWSIncidents-PostIncidentAnalysisTemplate.

Create an analysis template

We encourage you to use the default AWSIncidents-PostIncidentAnalysisTemplate template and add additional questions or sections that are appropriate for your use cases. Create analysis templates based on the default template Use this template as a starting point to create analysis templates in your management account. You can then duplicate your analysis templates to each Region where you enabled Incident Manager.

Create an analysis template

- Call the GetDocument action and use its Name parameter to download AWSIncidents-PostIncidentAnalysisTemplate. For more information about the GetDocument syntax, see <u>Systems Manager API Reference</u>.
- 2. The content in the response contains the JSON building blocks for the analysis. Use the question building blocks to insert additional questions in the analysis. We recommend that you add questions or sections in the Incident questions section.
- 3. To create the new template, use the CreateDocument operation with the updated JSON from the previous step. You must include the following, where *Analysis_Template_Name* is the name of your template,
 - DocumentFormat: "JSON"
 - DocumentType: "ProblemAnalysisTemplate"
 - Name: "Analysis_Template_Name"

Create an analysis

- 1. To create an analysis, choose **Create analysis** from the incident details page of a closed incident.
- 2. Choose the analysis template to create this analysis from, and enter a descriptive name of the analysis.
- 3. Choose Create.

Print a formatted incident analysis

You can generate a copy of a complete or incomplete analysis that is formatted for printing. You can also save this copy as a PDF. You can print one analysis at a time. Batch printing of multiple analyses isn't currently supported.

To print a formatted analysis

- 1. Open the Incident Manager console.
- 2. Choose the **Analysis** tab.
- 3. Choose the title of the analysis that you want to print.
- 4. In the upper right corner of the analysis detail page, choose **Print**.
- 5. In the **Print incident analysis** dialog box, clear the sections of the analysis you don't want included in the printed version. By default, all sections are selected.
- 6. Choose **Print** to open the local print controls for your device.
- 7. Choose your printing destination or format. You can choose a local or network printer, or you can save the analysis to a PDF. Make any changes, if wanted, to the remaining printing options, and then choose **Print**.

🚺 Note

Local print controls refers to the user interface provided by your web browser and device.

Printing destinations are those configured for, and accessible from, your device.

Incident Manager tutorials

These AWS Systems Manager Incident Manager tutorials help you build a more robust incident management system. These tutorials cover common activities that occur during an incident or support incident response.

Topics

- Tutorial: Using Systems Manager Automation runbooks with Incident Manager
- Tutorial: Managing security incidents in Incident Manager

Tutorial: Using Systems Manager Automation runbooks with Incident Manager

You can use <u>AWS Systems Manager Automation</u> runbooks to simplify common maintenance, deployment, and remediation tasks for AWS services. In this tutorial, you'll create a custom runbook to automate an incident response in Incident Manager. The scenario for this tutorial involves an Amazon CloudWatch alarm assigned to an Amazon EC2 metric. When the instance enters a state that triggers the alarm, Incident Manager automatically performs the following tasks:

- 1. Creates an incident in Incident Manager.
- 2. Initiates a runbook that attempts to remediate the issue.
- 3. Publishes the runbook results to the incident details page in Incident Manager.

The process described in this tutorial can also be used with Amazon EventBridge events and other types of AWS resources. By automating your remediation response to alarms and events you can reduce the impact of an incident on your organization and its resources.

This tutorial describes how to edit a CloudWatch alarm assigned to an Amazon EC2 instance for an Incident Manager response plan. If you don't have an alarm, an instance, or a response plan configured, we recommend you configure those resources before you begin. For more information, see the following topics:

- Using Amazon CloudWatch alarms in the Amazon CloudWatch User Guide
- Amazon EC2 instances in the Amazon EC2 User Guide

- Amazon EC2 instances in the Amazon EC2 User Guide
- Creating and configuring response plans in Incident Manager

🛕 Important

You will incur costs by creating AWS resources and using runbook automation steps. For more information, see <u>AWS pricing</u>.

Topics

- Task 1: Creating the runbook
- Task 2: Creating an IAM role
- Task 3: Connecting the runbook to your response plan
- Task 4: Assigning a CloudWatch alarm to your response plan
- Task 5: Verifying the results

Task 1: Creating the runbook

Use the following procedure to create a runbook in the Systems Manager console. When invoked from an Incident Manager incident, the runbook restarts an Amazon EC2 instance and updates the incident with information about the runbook execution. Before you begin, verify that you have permission to create a runbook. For more information, see <u>Setting up Automation</u> in the AWS Systems Manager User Guide.

🔥 Important

Review the following important details about creating this tutorial's runbook:

- The runbook is intended for an incident created from a CloudWatch alarm source. If you use this runbook for other types of incidents, for example manually created incidents, then the timeline event in the first runbook step won't be found and the system returns an error.
- The runbook requires the CloudWatch alarm include a dimension called InstanceId. Alarms for Amazon EC2 instance metrics have this dimension. If you use this runbook

with other metrics (or with other incident sources, such as EventBridge), then you have to change the JsonDecode2 step to match the data captured in your scenario.

 The runbook attempts to remediate the issue that triggered the alarm by restarting the Amazon EC2 instance. For a real incident, you might not want to restart the instance. Update the runbook with the specific remediation actions that you want the system to take.

For more information about creating runbooks, see <u>Working with runbooks</u> in the AWS Systems Manager User Guide.

To create a runbook

- 1. Open the AWS Systems Manager console at <u>https://console.aws.amazon.com/systems-</u>manager/.
- 2. In the navigation pane, choose **Documents**.
- 3. Choose Automation.
- 4. For Name, enter a descriptive name for the runbook, such as **IncidentResponseRunbook**.
- 5. Choose the **Editor** tab, and then choose **Edit**.
- 6. Paste the following content into the editor:

```
description: This runbook attempts to restart an Amazon EC2 instance that caused an
 incident.
schemaVersion: '0.3'
parameters:
  IncidentRecordArn:
    type: String
    description: The incident
mainSteps:
  - name: ListTimelineEvents
    action: 'aws:executeAwsApi'
    outputs:
      - Selector: '$.eventSummaries[0].eventId'
        Name: eventId
        Type: String
    inputs:
      Service: ssm-incidents
      Api: ListTimelineEvents
      incidentRecordArn: '{{IncidentRecordArn}}'
```

```
filters:
       - key: eventType
         condition:
           equals:
             stringValues:
               - SSM Incident Trigger
   description: This step retrieves the ID of the first timeline event with the
CloudWatch alarm details.
 - name: GetTimelineEvent
   action: 'aws:executeAwsApi'
   inputs:
     Service: ssm-incidents
     Api: GetTimelineEvent
     incidentRecordArn: '{{IncidentRecordArn}}'
     eventId: '{{ListTimelineEvents.eventId}}'
   outputs:
     - Name: eventData
       Selector: $.event.eventData
       Type: String
   description: This step retrieves the timeline event itself.
 - name: JsonDecode
   action: 'aws:executeScript'
   inputs:
     Runtime: python3.8
     Handler: script_handler
     Script: |-
       import json
       def script_handler(events, context):
         data = json.loads(events["eventData"])
         return data
     InputPayload:
       eventData: '{{GetTimelineEvent.eventData}}'
   outputs:
     - Name: rawData
       Selector: $.Payload.rawData
       Type: String
   description: This step parses the timeline event data.
 - name: JsonDecode2
   action: 'aws:executeScript'
   inputs:
     Runtime: python3.8
     Handler: script_handler
     Script: |-
```

```
import json
      def script_handler(events, context):
        data = json.loads(events["rawData"])
        return data
    InputPayload:
      rawData: '{{JsonDecode.rawData}}'
  outputs:
     - Name: InstanceId
      Selector:
'$.Payload.detail.configuration.metrics[0].metricStat.metric.dimensions.InstanceId'
      Type: String
  description: This step parses the CloudWatch event data.
- name: RestartInstance
  action: 'aws:executeAutomation'
  inputs:
    DocumentName: AWS-RestartEC2Instance
    DocumentVersion: $DEFAULT
    RuntimeParameters:
      InstanceId: '{{JsonDecode2.InstanceId}}'
  description: This step restarts the Amazon EC2 instance
```

7. Choose Create automation.

Task 2: Creating an IAM role

Use the following tutorial to create an AWS Identity and Access Management (IAM) role that gives Incident Manager permission to intitiate a runbook specified in a response plan. The runbook in this tutorial restarts an Amazon EC2 instance. You will specify this IAM role in the next task when you connect the runbook to your response plan.

Create an IAM role that intitiates a runbook from a response plan

- 1. Open the IAM console at https://console.aws.amazon.com/iam/.
- 2. In the navigation pane, choose **Roles**, and then choose **Create role**.
- 3. Under **Trusted entity type**, verify that **AWS service** is selected.
- 4. Under Use case, in the Use cases for other AWS services field, enter Incident Manager.
- 5. Choose Incident Manager, and then choose Next.
- 6. On the **Add permissions** page, choose **Create policy**. The permissions editor will open in a new browser window or tab.

- 7. In the editor, choose the **JSON** tab.
- 8. Copy and paste the following permission policy into the JSON editor. Replace *account_ID* with your AWS account ID.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Resource": [
                "arn:aws:ssm:*:account_ID:automation-definition/
IncidentResponseRunbook:*",
                "arn:aws:ssm:*::automation-definition/AWS-RestartEC2Instance:*"
            ],
            "Action": "ssm:StartAutomationExecution"
        },
        {
            "Effect": "Allow",
            "Resource": "arn:aws:ssm:*:*:automation-execution/*",
            "Action": "ssm:GetAutomationExecution"
        },
        {
            "Effect": "Allow",
            "Resource": "arn:aws:ssm-incidents:*:*:*",
            "Action": "ssm-incidents:*"
        },
        {
            "Effect": "Allow",
            "Resource": "arn:aws:iam::*:role/AWS-SystemsManager-
AutomationExecutionRole",
            "Action": "sts:AssumeRole"
        },
        {
            "Effect": "Allow",
            "Resource": "*",
            "Action": [
                "ec2:StopInstances",
                "ec2:DescribeInstanceStatus",
                "ec2:StartInstances"
            ]
        }
    1
```

}

- 9. Choose Next: Tags.
- 10. (Optional) If needed, add tags to your policy.
- 11. Choose Next: Review.
- 12. In the **Name** field, enter a name that helps you identify this role as being used for this tutorial.
- 13. (Optional) Enter a description in the **Description** field.
- 14. Choose **Create policy**.
- 15. Navigate back to the browser window or tab for the role you are creating. The **Add permissions** page is displayed.
- 16. Choose the refresh button (located next to the **Create Policy** button), and then enter the name of the perimssions policy you created into the filter box.
- 17. Choose the permission policy you created, and then choose **Next**.
- 18. On the **Name, review, and create** page, for **Role name**, enter a name that helps you identify this role as being used for this tutorial.
- 19. (Optional) Enter a description in the **Description** field.
- 20. Review the role details, add tags if needed, and choose Create role.

Task 3: Connecting the runbook to your response plan

By connecting the runbook to your Incident Manager response plan, you ensure a consistent, repeatable, and timely mitigation process. The runbook also serves as a starting point for resolvers to determine their next course of action.

To assign the runbook to your response plan

- 1. Open the Incident Manager console.
- 2. Choose **Response plans**.
- 3. For **Response plan**, choose an existing response plan and choose **Edit**. If you do not have an existing response plan, choose **Create response plan** to create a new plan.

Complete the following fields:

- a. In the **Runbook** section, choose **Select existing runbook**.
- b. For **Owner**, verify that **Owned by me** is selected.

- c. For **Runbook**, choose the runbook you created in Task 1: Creating the runbook.
- d. For Version, choose Default at the time of execution.
- e. In the Inputs section, for the IncidentRecordArn parameter, choose Incident ARN.
- f. In the **Execution permissions** section, choose the IAM role you created in <u>Task 2: Creating</u> <u>an IAM role</u>.
- 4. Save your changes.

Task 4: Assigning a CloudWatch alarm to your response plan

Use the following procedure to assign a CloudWatch alarm for an Amazon EC2 instance to your response plan.

To assign a CloudWatch alarm to your response plan

- 1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
- 2. In the navigation pane, under Alarms, choose All alarms.
- 3. Choose an alarm for an Amazon EC2 instance that you want to connect to your response plan.
- 4. Choose **Actions**, and then choose **Edit**. Verify that the metric has a dimension called InstanceId.
- 5. Choose Next.
- 6. For **Configure actions wizard**, choose **Add Systems Manager action**.
- 7. Choose **Create incident**.
- 8. Choose the response plan you created in <u>Task 3: Connecting the runbook to your response</u> <u>plan</u>.
- 9. Choose Update alarm.

Task 5: Verifying the results

To verify that the CloudWatch alarm creates an incident and then processes the runbook specified in your response plan, you must trigger the alarm. After you trigger the alarm and the runbook finishes processing, you can verify the results of the runbook by using the following procedure. For information about triggering an alarm, see <u>set-alarm-state</u> in the *AWS CLI Command Reference*.

1. Open the Incident Manager console.

- 2. Choose the incident created by the CloudWatch alarm.
- 3. Choose the **Runbooks** tab.
- 4. View the actions performed on your Amazon EC2 instance in the **Runbook steps** section.

The following image demonstrates how the steps taken by the runbook you created in this tutorial are reported in the console. Each step is listed with a timestamp and a status message.

Runbook steps					
00:00:00 (UTC+11:00)					
15:01:44 – ElistTimelineEvents	⊘ Success				
15:01:44 – EtTimelineEvent	⊘ Success				
15:01:45 – 💌 JsonDecode	⊘ Success				
15:01:49 – 🕑 JsonDecode2	⊘ Success				
15:01:53 – Et RestartInstance	⊘ Success				
23:59:59					

To view all of the details in the CloudWatch alarm, expand the **JsonDecode2** step, and then expand **Output**.

<u> Important</u>

You must clean up any resource changes you implemented during this tutorial that you don't want to keep. This includes changes to Incident Manager resources such as resource plans and incidents, changes to CloudWatch alarms, and the IAM role you created for this tutorial.

Tutorial: Managing security incidents in Incident Manager

You can use AWS Security Hub, Amazon EventBridge, and Incident Manager together to identify and manage security incidents in your AWS hosted-applications. This tutorial walks you through configuring an EventBridge rule that creates an incident based on Security Hub automatically sent findings.

🚯 Note

This tutorial uses EventBridge Security Hub. You may incur costs from using these services.

Prerequisites

- Set up Security Hub. For more information, see <u>Setting up AWS Security Hub</u>.
- Create or update findings in Security Hub. For more information, see <u>Findings in AWS Security</u> <u>Hub</u>.
- Configure a response plan that Incident Manager will use as the template when creating your security incident. For more information, see <u>Preparing for incidents in Incident Manager</u>.

For this tutorial, we use a predefined pattern to create the EventBridge rule. To create the rule using a custom pattern, see <u>Using a custom pattern to create the rule</u> in the AWS Security Hub user guide.

Create an EventBridge rule

- 1. Open the Amazon EventBridge console at https://console.aws.amazon.com/events/.
- 2. In the navigation pane, choose **Rules**.
- 3. Choose **Create rule**.
- 4. Enter a **Name** and **Description** for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

- 5. For **Event bus**, choose **default**.
- 6. For **Rule type**, choose **Rule with an event pattern**.
- 7. Choose **Next**.
- 8. For **Event source**, choose **AWS events or EventBridge partner events**.
- 9. For **Event pattern**, choose **Event pattern form**.
- 10. For Event source, choose AWS services.
- 11. For **AWS service**, choose **Security Hub**.

- 12. For Event type, choose Security Hub Findings Imported.
- 13. By default, EventBridge configures the event pattern without any filter values. For each attribute, the **Any** *attribute name* option is selected. Update these filters to create incidents based on the security findings that most impact your environment.
- 14. Click Next.
- 15. For Target types, choose AWS service.
- 16. For Select a target, choose Incident Manager response plan.
- 17. For **Response plan**, choose a response plan to use as a template for created incidents.
- 18. EventBridge can create the IAM role needed for your rule to run.
 - To create an IAM role automatically, choose **Create a new role for the specific resource**.
 - To use an IAM role that already exists in your account, choose **Use existing role**.
- 19. (Optional) Enter one or more tags for the rule.
- 20. Choose Next.
- 21. Review the details of the rule and choose **Create rule**.

Now that you've created this EventBridge rule, security findings that match the attribute values you defined will create incidents in Incident Manager. You can triage, manage, monitor, and create post-incident analysis from these incidents.

Tagging resources in Incident Manager

Tags are optional metadata that you can assign to your Incident Manager resources in the AWS Regions specified in your replication set. You can assign tags to response plans, incident records, and contacts. You can also add tags to on-call schedules and rotations.You can also add tags to the replication set itself. Tags enable you to categorize and control access to these resources in different ways. Each tag consists of a key and an optional value, both of which you define. We recommend that you devise a set of tag keys that meets your needs for each Incident Manager resource type. Using a consistent set of tag keys makes it easier for you to manage these resources and manage access to them. You can search and filter resources based on tags. For more information about controlling access to resources by using tags, see <u>Controlling access to AWS</u> resources using tags in the *IAM User Guide*.

You can specify tags in the **Incident default** section when creating a response plan. These tags are applied to the incident record when an incident is created using the response plan.

Note

Tags don't have any semantic meaning. They are interpreted strictly as a string of characters.

You can add or remove tags by using the Incident Manager console. The following screenshot displays the **Tags** area of a console page, with fields for adding tag keys and values, and buttons for adding and removing tags.

 Tags - optional 					
Key		Value - optional			
Q Environment	×	Q Linux	×	Remove	
Add new tag					
You can add up to 49 more tags.					

To work with tags programmatically, use the following API actions:

- TagResource
- UntagResource
- ListTagsForResource

A Important

Tags applied to response plans, incident records, contacts, on-call schedules and rotations, and replication sets can be viewed and modified only from the resource owner account.

Security in AWS Systems Manager Incident Manager

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The <u>shared responsibility model</u> describes this as security *of* the cloud and security *in* the cloud:

- Security of the cloud AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the <u>AWS</u> <u>Compliance Programs</u>. To learn about the compliance programs that apply to AWS Systems Manager Incident Manager, see <u>AWS Services in Scope by Compliance Program</u>.
- Security in the cloud Your responsibility is determined by the AWS service that you use. You
 are also responsible for other factors including the sensitivity of your data, your company's
 requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Incident Manager. The following topics show you how to configure Incident Manager to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Incident Manager resources.

Topics

- Data protection in Incident Manager
- Identity and Access Management for AWS Systems Manager Incident Manager
- Working with shared contacts and response plans in Incident Manager
- <u>Compliance validation for AWS Systems Manager Incident Manager</u>
- Resilience in AWS Systems Manager Incident Manager
- Infrastructure security in AWS Systems Manager Incident Manager
- Working with AWS Systems Manager Incident Manager and interface VPC endpoints (AWS PrivateLink)
- Configuration and vulnerability analysis in Incident Manager
- Security best practices in AWS Systems Manager Incident Manager

Data protection in Incident Manager

The AWS <u>shared responsibility model</u> applies to data protection in AWS Systems Manager Incident Manager. 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. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the <u>Data Privacy FAQ</u>. For information about data protection in Europe, see the <u>AWS Shared</u> <u>Responsibility Model and GDPR</u> blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or 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 SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see <u>Working with CloudTrail trails</u> in the AWS CloudTrail User Guide.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

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

By default, Incident Manager encrypts data in transit using SSL/TLS.

Data encryption

Incident Manager uses AWS Key Management Service (AWS KMS) keys to encrypt your Incident Manager resources. For more information about AWS KMS, see the <u>AWS KMS Developer Guide</u>. AWS KMS combines secure, highly available hardware and software to provide a key management system scaled for the cloud. Incident Manager encrypts your data using your specified key and encrypts metadata using an AWS owned key. To use Incident Manager, you must set up your replication set, which includes setting up encryption. Incident Manager requires data encryption for use.

You can use an AWS owned key to encrypt your replication set or you can use your own customer managed key that you created in AWS KMS to encrypt the Regions in your replication set. Incident Manager only supports symmetric encryption AWS KMS keys to encrypt your data created within AWS KMS. Incident Manager doesn't support AWS KMS keys with imported key material, custom key stores, Hash-based Message Authentication Code (HMAC), or other types of keys. If you use customer managed keys, you use the <u>AWS KMS console</u> or AWS KMS APIs to centrally create the customer managed keys and define the key policies that control how Incident Manager can use the customer managed keys. When you use a customer managed key for encryption with Incident Manager, the AWS KMS customer managed key must be in the same Region as the resources. To learn more about setting up data encryption in Incident Manager, see <u>Get prepared wizard</u>.

There are additional charges for using AWS KMS customer managed keys. For more information, see <u>AWS KMS concepts - KMS keys</u> in the *AWS Key Management Service Developer Guide* and <u>AWS KMS pricing</u>.

🔥 Important

If you use a AWS KMS key (KMS key) to encrypt your replication set and Incident Manager data, but later decide to delete the replication set, make sure to delete the replication set *before* disabling or deleting the KMS key.

To allow Incident Manager to use your customer managed key to encrypt your data, you must add the following policy statements to the key policy of your customer managed key. To learn more about setting up and changing the key policy in your account, see <u>Using key policies in AWS KMS</u> in the AWS Key Management Service Developer Guide. The policy provides the following permissions:

- Allows Incident Manager to perform read-only operations to find the AWS KMS key for Incident Manager in your account.
- Allows Incident Manager to use the KMS key to create grants and describe the key, but only when it's acting on behalf of principals in the account who have permission to use Incident Manager. If the principals specified in the policy statement don't have permission to use the KMS keys and to use Incident Manager, the call fails, even when it comes from the Incident Manager service.

```
{
 "Sid": "Allow CreateGrant through AWS Systems Manager Incident Manager",
 "Effect": "Allow",
 "Principal": {
   "AWS": "arn:aws:iam::111122223333:user/ssm-lead"
 },
 "Action": [
   "kms:CreateGrant",
   "kms:DescribeKey"
 ],
 "Resource": "*",
 "Condition": {
   "StringLike": {
     "kms:ViaService": [
       "ssm-incidents.amazonaws.com",
       "ssm-contacts.amazonaws.com"
     ]
   }
 }
}
```

Replace the Principal value with the IAM principal that created your replication set.

Incident Manager uses an <u>encryption context</u> in all requests to AWS KMS for cryptographic operations. You can use this encryption context to identify CloudTrail log events where Incident Manager uses your KMS keys. Incident Manager uses the following encryption context:

```
• contactArn=ARN of the contact or escalation plan
```

Identity and Access Management for AWS Systems Manager Incident Manager

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Incident Manager resources. IAM is an AWS service that you can use with no additional charge.

Topics

- Audience
- Authenticating with identities
- Managing access using policies
- How AWS Systems Manager Incident Manager works with IAM
- Identity-based policy examples for AWS Systems Manager Incident Manager
- Resource-based policy examples for AWS Systems Manager Incident Manager
- <u>Cross-service confused deputy prevention in Incident Manager</u>
- Using service-linked roles for Incident Manager
- <u>AWS managed policies for AWS Systems Manager Incident Manager</u>
- Troubleshooting AWS Systems Manager Incident Manager identity and access

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Incident Manager.

Service user – If you use the Incident Manager service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Incident Manager features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Incident Manager, see <u>Troubleshooting AWS Systems Manager Incident</u> Manager identity and access.

Service administrator – If you're in charge of Incident Manager resources at your company, you probably have full access to Incident Manager. It's your job to determine which Incident Manager features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Incident Manager, see How AWS Systems Manager Incident Manager works with IAM.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Incident Manager. To view example Incident Manager identity-based policies that you can use in IAM, see <u>Identity-based policy examples for AWS</u> Systems Manager Incident Manager.

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see <u>How to sign in to your AWS</u> <u>account</u> in the AWS Sign-In User Guide.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <u>Signing AWS API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Multi-factor authentication</u> in the *AWS IAM Identity Center User Guide* and <u>Using multi-factor authentication (MFA) in AWS</u> in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your

root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see <u>Tasks that require root</u> user credentials in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see <u>What is IAM Identity Center?</u> in the AWS IAM Identity Center User Guide.

IAM users and groups

An <u>IAM user</u> is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-</u> term credentials in the *IAM User Guide*.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>When to create an IAM user</u> (instead of a role) in the *IAM User Guide*.

IAM roles

An <u>IAM role</u> is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by <u>switching roles</u>. You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Methods to assume a role in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- Federated user access To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see <u>Creating a role for a third-party Identity Provider</u> in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see <u>Permission sets</u> in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant crossaccount access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see <u>Cross account resource access in IAM</u> in the *IAM User Guide*.
- **Cross-service access** Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - Forward access sessions (FAS) When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

- Service role A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see <u>Creating a role to delegate permissions to an AWS service</u> in the *IAM User Guide*.
- Service-linked role A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Using an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see <u>When to create an IAM role (instead of a user)</u> in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see <u>Overview of JSON policies</u> in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam:GetRole action. A

user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Creating IAM policies</u> in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <u>Choosing between managed policies and inline policies</u> in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see <u>Permissions boundaries for IAM entities</u> in the *IAM User Guide*.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see <u>Service</u> <u>control policies</u> in the AWS Organizations User Guide.
- Session policies Session policies are advanced policies that you pass as a parameter when you
 programmatically create a temporary session for a role or federated user. The resulting session's
 permissions are the intersection of the user or role's identity-based policies and the session
 policies. Permissions can also come from a resource-based policy. An explicit deny in any of these
 policies overrides the allow. For more information, see <u>Session policies</u> in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

How AWS Systems Manager Incident Manager works with IAM

Before you use IAM to manage access to Incident Manager, learn what IAM features are available to use with Incident Manager.

IAM features you can use with AWS Systems Manager Incident Manager

IAM feature	Incident Manager support
Identity-based policies	Yes
Resource-based policies	Yes
Policy actions	Yes
Policy resources	Yes
Policy condition keys	No
ACLs	No
ABAC (tags in policies)	No
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how Incident Manager and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Incident Manager doesn't support policies that deny access to resources shared using AWS RAM.

Identity-based policies for Incident Manager

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Creating IAM policies</u> in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an

identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see <u>IAM JSON policy elements reference</u> in the *IAM User Guide*.

Identity-based policy examples for Incident Manager

To view examples of Incident Manager identity-based policies, see <u>Identity-based policy examples</u> for AWS Systems Manager Incident Manager.

Resource-based policies within Incident Manager

Supports resource-based policies: Yes

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see <u>Cross account resource access in IAM</u> in the *IAM User Guide*.

The Incident Manager service supports only two types of resource-based policies called using either the AWS RAM console or the PutResourcePolicy action, which is attached to a response plan or contact. This policy defines which principals can perform actions on the response plans, contacts, escalation plans, and incidents. Incident Manager uses resource based policies to share resources across accounts.

Incident Manager doesn't support policies that deny access to resources shared using AWS RAM.

To learn how to attach a resource-based policy to a response plan or contact, see <u>Managing</u> incidents across AWS accounts and Regions in Incident Manager.

Resource-based policy examples within Incident Manager

To view examples of Incident Manager resource-based policies, see <u>Resource-based policy examples</u> for AWS Systems Manager Incident Manager.

Policy actions for Incident Manager

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Incident Manager actions, see <u>Actions defined by AWS Systems Manager Incident</u> Manager in the *Service Authorization Reference*.

Policy actions in Incident Manager use the following prefixes before the action:

```
ssm-incidents
ssm-contacts
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
    "ssm-incidents:GetResponsePlan",
    "ssm-contacts:GetContact"
]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word Get, include the following action:

```
"Action": "ssm-incidents:Get*"
```

To view examples of Incident Manager identity-based policies, see <u>Identity-based policy examples</u> for AWS Systems Manager Incident Manager.

Incident Manager uses actions in two different namespaces, ssm-incidents and ssm-contacts. When creating policies for Incident Manager make sure to use the namespace correct for the action. SSM-Incidents is used for response plan and incident related action. SSM-Contacts is used for actions related to contacts and contact engagement. For example:

- ssm-contacts:GetContact
- ssm-incidents:GetResponsePlan

Policy resources for Incident Manager

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <u>Amazon Resource Name (ARN)</u>. You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

"Resource": "*"

To see a list of Incident Manager resource types and their ARNs, see <u>Resources defined by AWS</u> <u>Systems Manager Incident Manager</u> in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see <u>Actions defined by AWS Systems Manager</u> Incident Manager.

To view examples of Incident Manager identity-based policies, see <u>Identity-based policy examples</u> for AWS Systems Manager Incident Manager.

Incident Manager resources are used to create incidents, collaborate in chat channels, resolve incidents, and engage responders. If a user has access to a response plan they have access to all

incidents created from it. If a user has access to a contact or escalation plan they can engage the contact or contacts in the escalation plan.

Policy condition keys for Incident Manager

Supports service-specific policy condition keys: No

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see <u>IAM policy elements: variables and tags</u> in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see <u>AWS global condition context keys</u> in the *IAM User Guide*.

Access control lists (ACLs) in Incident Manager

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Attribute-based access control (ABAC) with Incident Manager

Supports ABAC (tags in policies): No

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or

roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the <u>condition element</u> of a policy using the aws:ResourceTag/key-name, aws:RequestTag/key-name, or aws:TagKeys condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see <u>What is ABAC?</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control (ABAC)</u> in the *IAM User Guide*.

Using temporary credentials with Incident Manager

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see <u>AWS services that</u> work with IAM in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see <u>Switching to a role (console)</u> in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see <u>Temporary security credentials in IAM</u>.

Cross-service principal permissions for Incident Manager

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see <u>Forward access sessions</u>.

Service roles for Incident Manager

Supports service roles: Yes

A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see <u>Creating a role to delegate permissions to an AWS service</u> in the *IAM User Guide*.

<u> M</u>arning

Changing the permissions for a service role might break Incident Manager functionality. Edit service roles only when Incident Manager provides guidance to do so.

Choosing an IAM role in Incident Manager

When you create a response plan resource in Incident Manager, you must choose a role to allow Incident Manager to run a Systems Manager automation document on your behalf. If you have previously created a service role or service-linked role, then Incident Manager provides you with a list of roles to choose from. It's important to choose a role that allows access to run your automation document instances. For more information, see <u>Integrating Systems Manager</u> <u>Automation runbooks in Incident Manager for incident remediation</u>. When you create a AWS Chatbot chat channel to be used during an incident you can select a service role that allows you to use commands directly from chat. To learn more about creating chat channels for incident collaboration, see <u>Creating and integrating chat channels for responders in Incident Manager</u>. To learn more about IAM policies in AWS Chatbot, see <u>Managing permissions for running commands</u> <u>using AWS Chatbot</u> in the *AWS Chatbot Administrator guide*.

Service-linked roles for Incident Manager

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For information about creating or managing Incident Manager service-linked roles, see Using service-linked roles for Incident Manager.

Identity-based policy examples for AWS Systems Manager Incident Manager

By default, users and roles don't have permission to create or modify Incident Manager resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see <u>Creating IAM policies</u> in the *IAM User Guide*.

For details about actions and resource types defined by Incident Manager, including the format of the ARNs for each of the resource types, see <u>Actions, resources, and condition keys for AWS</u> <u>Systems Manager Incident Manager</u> in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Incident Manager console
- Allow users to view their own permissions
- Accessing a response plan

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Incident Manager resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

• Get started with AWS managed policies and move toward least-privilege permissions – To get started granting permissions to your users and workloads, use the AWS managed policies

that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see <u>AWS managed policies</u> or <u>AWS</u> managed policies for job functions in the *IAM User Guide*.

- **Apply least-privilege permissions** When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see <u>Policies and permissions in IAM</u> in the *IAM User Guide*.
- Use conditions in IAM policies to further restrict access You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see <u>IAM JSON policy elements: Condition</u> in the *IAM User Guide*.
- Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see <u>IAM Access Analyzer policy validation</u> in the *IAM User Guide*.
- Require multi-factor authentication (MFA) If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see <u>Configuring MFA-protected API access</u> in the *IAM User Guide*.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> in the *IAM User Guide*.

Using the Incident Manager console

To access the AWS Systems Manager Incident Manager console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Incident Manager resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can resolve incident using the Incident Manager console, also attach the Incident Manager IncidentManagerResolverAccess AWS managed policy to the entities. For more information, see Adding permissions to a user in the *IAM User Guide*.

IncidentManagerResolverAccess

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
```

```
"iam:ListUsers"
],
"Resource": "*"
}
]
}
```

Accessing a response plan

In this example, you want to grant an IAM user in your Amazon Web Services account access to one of your Incident Manager response plans, exampleplan. You also want to allow the user to add, update, and delete the response plan.

The policy grants the ssm-incidents:ListResponsePlans, ssm-

incidents:GetResponsePlan, ssm-incidents:UpdateResponsePlan and ssmincident:ListResponsePlan permissions to the user.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"ListResponsePlans",
         "Effect":"Allow",
         "Action":[
            "ssm-incidents:ListResponsePlans"
         ],
         "Resource":"arn:aws:ssm-incidents:::*"
      },
      {
         "Sid": "ViewSpecificResponsePlanInfo",
         "Effect":"Allow",
         "Action":[
            "ssm-incidents:GetResponsePlan"
         ],
         "Resource":"arn:aws:ssm-incidents:*:111122223333:response-plan/exampleplan"
      },
      {
         "Sid": "ManageResponsePlan",
         "Effect":"Allow",
         "Action":[
            "ssm-incidents:UpdateResponsePlan"
         ],
         "Resource":"arn:aws:ssm-incidents:*:111122223333:response-plan/exampleplan/*"
```

}

Resource-based policy examples for AWS Systems Manager Incident Manager

AWS Systems Manager Incident Manager supports resource-based permissions policies for Incident Manager response plans and contacts.

Incident Manager doesn't support resource-based policies that deny access to resources shared using AWS RAM.


To learn how to create a response plan or contact, see <u>Creating and configuring response plans in</u> Incident Manager and <u>Creating and configuring contacts in Incident Manager</u>.

Restricting Incident Manager response plan access by organization

The following example grants permissions to users in the organization with the organization ID: oabc123def45 to respond to incidents created using the response plan myplan.

The Condition block uses the StringEquals conditions and the aws:PrincipalOrgID condition key, which is an AWS Organizations specific condition key. For more information about these condition keys, see <u>Specifying conditions in a policy</u>.

```
{
    "Version": "2012-10-17",
    "Statement": [
    {
        "Sid": "OrganizationAccess",
        "Effect": "Allow",
        "Principal": "*",
        "Condition": {
            "StringEquals": {"aws:PrincipalOrgID":"o-abc123def45"}
        },
        "Action": [
            "ssm-incidents:GetResponsePlan",
            "ssm-incidents:StartIncident",
            "ssm-incidents:UpdateIncidentRecord",
        "
}
```



Providing Incident Manager contact access to a principal

The following example grants permission to the principal with the ARN arn:aws:iam::999988887777:root to create engagements to the contact mycontact.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "PrincipalAccess",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::999988887777:root"
            },
            "Action": [
                "ssm-contacts:GetContact",
                "ssm-contacts:StartEngagement",
                "ssm-contacts:DescribeEngagement",
                "ssm-contacts:ListPagesByContact"
            ],
            "Resource": [
                "arn:aws:ssm-contacts:*:111122223333:contact/mycontact"
                "arn:aws:ssm-contacts:*:111122223333:engagement/mycontact/*"
            ]
        }
    ]
}
```

Cross-service confused deputy prevention in Incident Manager

The confused deputy problem is an information security issue that occurs when an entity without permission to perform an action calls a more-privileged entity to perform the action. This can allow malicious actors to run commands or modify resources they otherwise would not have permission to run or access.

In AWS, cross-service impersonation can lead to a confused deputy scenario. Cross-service impersonation is when one service (the *calling service*) calls another service (the *called service*). A malicious actor can use the calling service to alter resources in another service using permissions that they normally would not have.

AWS provides service principals with managed access to resources on your account to help you protect your resources' security. We recommend using the <u>aws:SourceArn</u> and <u>aws:SourceAccount</u> global condition context keys in your resource policies. These keys limit the permissions that AWS Systems Manager Incident Manager gives another service to that resource. If you use both global condition context keys, the aws:SourceAccount value and the account referenced in the aws:SourceArn value must use the same account ID when used in the same policy statement.

The value of aws:SourceArn must be the ARN of the affected incident record. If you don't know the full ARN of the resource, or if you are specifying multiple resources, use the aws:SourceArn global context condition key with the * wildcard for the unknown portions of the ARN. For example, you can set aws:SourceArn to arn:aws:ssm-incidents::111122223333:*.

In the following trust policy example, we use the aws:SourceArn condition key to restrict access to the service role based on the incident record's ARN. Only incident records created from the response plan myresponseplan are able to use this role.

```
{
   "Version": "2012-10-17",
   "Statement": {
     "Effect": "Allow",
     "Principal": { "Service": "ssm-incidents.amazonaws.com" },
     "Action": "sts:AssumeRole",
     "Condition": {
        "ArnLike": {
            "aws:SourceArn": "arn:aws:ssm-incidents:*:111122223333:incident-record/
myresponseplan/*"
        }
}
```

}

} }

Using service-linked roles for Incident Manager

AWS Systems Manager Incident Manager uses AWS Identity and Access Management (IAM) <u>service-linked roles</u>. A service-linked role is a unique type of IAM role that is linked directly to Incident Manager. Service-linked roles are predefined by Incident Manager and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Incident Manager easier because you don't have to manually add the necessary permissions. Incident Manager defines the permissions of its service-linked roles, and unless defined otherwise, only Incident Manager can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your Incident Manager resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see <u>AWS Services That Work</u> <u>with IAM</u> and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Incident Manager

Incident Manager uses the service-linked role named **AWSServiceRoleforIncidentManager** – allows Incident Manager to manage Incident Manager incident records and related resources on your behalf.

The AWSServiceRoleforIncidentManager service-linked role trusts the following services to assume the role:

ssm-incidents.amazonaws.com

The role permissions policy <u>AWSIncidentManagerServiceRolePolicy</u> allows Incident Manager to complete the following actions on the specified resources:

• Action: ssm-incidents:ListIncidentRecords on all resources related to the action.

- Action: ssm-incidents:CreateTimelineEvent on all resources related to the action.
- Action: ssm:CreateOpsItem on all resources related to the action.
- Action: ssm:AssociateOpsItemRelatedItem on all resources related to the action.
- Action: ssm-contacts: StartEngagement on all resources related to the action.
- Action: cloudwatch:PutMetricData on CloudWatch metrics inside the AWS/ IncidentManager namespace

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see <u>Service-Linked Role Permissions</u> in the *IAM User Guide*.

Creating a service-linked role for Incident Manager

You don't need to manually create a service-linked role. When you create a replication set in the AWS Management Console, the AWS CLI, or the AWS API, Incident Manager creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a replication set, Incident Manager creates the service-linked role for you again.

Editing a service-linked role for Incident Manager

Incident Manager does not allow you to edit the AWSServiceRoleforIncidentManager service-linked role. After you create a service-linked role, you can't change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Editing a Service-Linked Role in the IAM User Guide.

Deleting a service-linked role for Incident Manager

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that isn't actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

To delete the service-linked role you must first delete the replication set. Deleting the replication set deletes all data created and stored in Incident Manager, including response plans, contacts, and

escalation plans. You will also lose all previously created incidents. Any alarms and EventBridge rules pointing to deleted response plans will no longer create an incident on alarm or rule match. To delete the replication set you must delete every Region in the set.

🚯 Note

If the Incident Manager service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To delete the Regions in the replication set used by the AWSServiceRoleforIncidentManager

- 1. Open the Incident Manager console and choose **Settings** from the left navigation.
- 2. Select a Region in the **Replication set**.
- 3. Choose **Delete**.
- 4. To confirm deletion of the Region, enter the Region name and choose **Delete**.
- 5. Repeat these steps until you have deleted all Regions in your replication set. When deleting the final Region, the console informs you that it deletes the replication set with it.

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleforIncidentManager service-linked role. For more information, see <u>Deleting a</u> <u>Service-Linked Role</u> in the *IAM User Guide*.

Supported Regions for Incident Manager service-linked roles

Incident Manager supports using service-linked roles in all of the Regions where the service is available. For more information, see <u>AWS Regions and Endpoints</u>.

AWS managed policies for AWS Systems Manager Incident Manager

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you

reduce permissions further by defining <u>customer managed policies</u> that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see <u>AWS managed policies</u> in the *IAM User Guide*.

AWS managed policy: AWSIncidentManagerIncidentAccessServiceRolePolicy

You can attach AWSIncidentManagerIncidentAccessServiceRolePolicy to your IAM entities. Incident Manager also attaches this policy to an Incident Manager role that allows Incident Manager to perform actions on your behalf.

This policy grants read-only permissions that allow Incident Manager to read resources in certain other AWS services to identify findings related to incidents in those services.

Permissions details

This policy includes the following permissions.

- cloudformation Allows principals to describe AWS CloudFormation stacks. This is required for Incident Manager to identify CloudFormation events and resources related to an incident.
- codedeploy Allows principals to read AWS CodeDeploy deployments. This is required for Incident Manager to identify CodeDeploy deployments and targets related to an incident.
- autoscaling Allows principals to determine if an Amazon Elastic Compute Cloud (EC2) instance is part of an Auto Scaling group. This is needed so Incident Manager can provide findings for EC2 instances that are part of Auto Scaling groups.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "IncidentAccessPermissions",
            "Effect": "Allow",
            "Action": [
                "cloudformation:DescribeStackEvents",
                "cloudformation:DescribeStackResources",
                "codedeploy:BatchGetDeployments",
                "codedeploy:ListDeployments",
                "codedeploy:ListDeploymentTargets",
                "autoscaling:DescribeAutoScalingInstances"
            ],
            "Resource": "*"
        }
    ]
}
```

To view more details about the policy, including the latest version of the JSON policy document, see <u>AWSIncidentManagerIncidentAccessServiceRolePolicy</u> in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSIncidentManagerServiceRolePolicy

You can't attach AWSIncidentManagerServiceRolePolicy to your IAM entities. This policy is attached to a service-linked role that allows Incident Manager to perform actions on your behalf. For more information, see Using service-linked roles for Incident Manager.

This policy grants Incident Manager permissions to list incidents, create timeline events, create OpsItems, associate related items to OpsItems, start engagements, and publish CloudWatch metrics related to an incident.

Permissions details

This policy includes the following permissions.

- ssm-incidents Allows principals to list incidents and create timeline events. This is required so responders can collaborate during an incident on the incident dashboard.
- ssm Allows principals to create OpsItems and associate related items. This is required to create
 a parent OpsItem when an incident starts.
- ssm-contacts Allows principals to start engagements. This is required for Incident Manager to engage contacts during an incident.
- cloudwatch Allows principals to publish CloudWatch metrics. This is required for Incident Manager to publish metrics related to an incident.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "UpdateIncidentRecordPermissions",
            "Effect": "Allow",
            "Action": [
                 "ssm-incidents:ListIncidentRecords",
                "ssm-incidents:CreateTimelineEvent"
            ],
            "Resource": "*"
        },
        {
            "Sid": "RelatedOpsItemPermissions",
            "Effect": "Allow",
            "Action": [
                "ssm:CreateOpsItem",
                "ssm:AssociateOpsItemRelatedItem"
            ],
            "Resource": "*"
        },
        {
            "Sid": "IncidentEngagementPermissions",
            "Effect": "Allow",
            "Action": "ssm-contacts:StartEngagement",
            "Resource": "*"
        },
        {
            "Sid": "PutCloudWatchMetricPermission",
            "Effect": "Allow",
            "Action": [
```

```
"cloudwatch:PutMetricData"
],
"Resource": "*",
"Condition": {
    "StringEquals": {
        "cloudwatch:namespace": "AWS/IncidentManager"
        }
    }
}
```

To view more details about the policy, including the latest version of the JSON policy document, see <u>AWSIncidentManagerServiceRolePolicy</u> in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSIncidentManagerResolverAccess

You can attach AWSIncidentManagerResolverAccess to your IAM entities to allow them to start, view, and update incidents. This also allows them to create customer timeline events and related items in the incident dashboard. You can also attach this policy to the AWS Chatbot service role or directly to your customer managed role associated with any chat channel used for incident collaboration. To learn more about IAM policies in AWS Chatbot, see <u>Managing permissions for running commands using AWS Chatbot</u> in the *AWS Chatbot Administrator Guide*.

Permissions details

This policy includes the following permissions.

 ssm-incidents – Allows you to start incidents, list response plans, list incidents, update incidents, list timeline events, create custom timeline events, update custom timeline events, delete custom timeline events, list related items, create related items, and update related items.

```
"ssm-incidents:StartIncident"
            ],
            "Resource": "*"
        },
        {
            "Sid": "ResponsePlanReadOnlyPermissions",
            "Effect": "Allow",
            "Action": [
                "ssm-incidents:ListResponsePlans",
                "ssm-incidents:GetResponsePlan"
            ],
            "Resource": "*"
        },
        {
            "Sid": "IncidentRecordResolverPermissions",
            "Effect": "Allow",
            "Action": [
                "ssm-incidents:ListIncidentRecords",
                "ssm-incidents:GetIncidentRecord",
                "ssm-incidents:UpdateIncidentRecord",
                "ssm-incidents:ListTimelineEvents",
                "ssm-incidents:CreateTimelineEvent",
                "ssm-incidents:GetTimelineEvent",
                "ssm-incidents:UpdateTimelineEvent",
                "ssm-incidents:DeleteTimelineEvent",
                "ssm-incidents:ListRelatedItems",
                "ssm-incidents:UpdateRelatedItems"
            ],
            "Resource": "*"
        }
    ]
}
```

To view more details about the policy, including the latest version of the JSON policy document, see <u>AWSIncidentManagerResolverAccess</u> in the AWS Managed Policy Reference Guide.

Incident Manager updates to AWS managed policies

View details about updates to AWS managed policies for Incident Manager since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Incident Manager Document history page.

Change	Description	Date
AWSIncidentManager IncidentAccessServiceRolePo Licy – Policy update	Incident Manager has added a new permission to AWSIncidentManager IncidentAccessServ iceRolePolicy , in support of the Findings feature, that allows it to check whether an EC2 instance is part of an Auto Scaling group.	February 20, 2024
AWSIncidentManager IncidentAccessServ iceRolePolicy – New policy	Incident Manager added a new policy that grants Incident Manager permissions to call other AWS services as a part of managing an incident.	November 17, 2023
AWSIncidentManager ServiceRolePolicy Policy update	Incident Manager added a new permission that allows Incident Manager to publish metrics into your account.	Dec 16, 2022
AWSIncidentManager ResolverAccess – New policy	Incident Manager added a new policy to allow you to start incidents, list response plans, list incidents, update incidents, list timeline events, create custom timeline events, update custom timeline events, delete custom timeline events, list related items, create related items, and update related items.	April 26, 2021

Change	Description	Date
AWSIncidentManager ServiceRolePolicy – New policy	Incident Manager added a new policy to grant Incident Manager permissions to list incidents, create timeline events, create OpsItems, associate related items to OpsItems, and start engagements related to an incident.	April 26, 2021
Incident Manager started tracking changes	Incident Manager started tracking changes for its AWS managed policies.	April 26, 2021

Troubleshooting AWS Systems Manager Incident Manager identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Incident Manager and IAM.

Topics

- I am not authorized to perform an action in Incident Manager
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my Amazon Web Services account to access my Incident
 Manager resources

I am not authorized to perform an action in Incident Manager

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional ssm-incidents: *GetWidget* permissions.

User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform: ssmincidents:GetWidget on resource: my-example-widget

In this case, the policy for the mateojackson user must be updated to allow access to the *myexample-widget* resource by using the ssm-incidents: *GetWidget* action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam: PassRole action, your policies must be updated to allow you to pass a role to Incident Manager.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Incident Manager. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform: iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the iam: PassRole action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon Web Services account to access my Incident Manager resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Incident Manager supports these features, see <u>How AWS Systems Manager</u> <u>Incident Manager works with IAM</u>.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the IAM User Guide.
- To learn how to provide access to your resources to third-party AWS accounts, see <u>Providing</u> access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see <u>Cross account resource access in IAM</u> in the *IAM User Guide*.

Working with shared contacts and response plans in Incident Manager

With contact sharing, as a contact owner, you can share contact information, escalation plans, and engagements with other AWS accounts or within an AWS organization. You can create and manage contacts and escalation plans centrally, and ensure that others can engage the correct contacts during an incident.

With response plan sharing, as a response plan owner, you can share a response plan and the related incidents with other AWS accounts or within an AWS organization. You can create and manage response plans centrally so that responders in consumer accounts can interact with incidents as they happen.

A contact or response plan owner can share contacts and response plans with:

- Specific AWS accounts inside or outside of its organization in AWS Organizations
- An organizational unit inside its organization in AWS Organizations
- Its entire organization in AWS Organizations

Contents

- Prerequisites for sharing contacts and response plans
- Related services

- Sharing a contact or response plan
- Stop sharing a shared contact or response plan
- Identifying a shared contact or response plan
- Shared contact and response plan permissions
- Billing and metering
- Instance limits

Prerequisites for sharing contacts and response plans

To share a contact or response plan with your organization or organizational unit in AWS Organizations:

- You must own the resource in your AWS account. You can't share a contact or response plan that has been shared with you.
- You must enable sharing with AWS Organizations. For more information, see <u>Enable Sharing</u> with AWS Organizations in the AWS RAM User Guide.

Related services

Contact and response plan sharing integrates with AWS Resource Access Manager (AWS RAM). With AWS RAM, you can share your AWS resources with any AWS account or through AWS Organizations. You share resources that you own by creating a *resource share*. A resource share specifies the resources to share, and the consumers with whom to share them. Consumers can be individual AWS accounts, organizational units, or an entire organization in AWS Organizations.

For more information about AWS RAM, see the AWS RAM User Guide.

Sharing a contact or response plan

After you share a response plan, the consumers have access to all past, current, and future incidents created using that response plan.

After you share a contact, the consumers have access to the contact information, engagement plan, escalation plans, and engagements that occur during an incident. Consumers can also engage a contact or escalation plan during an incident.

If you're part of an organization in AWS Organizations and sharing within your organization is enabled, consumers in your organization are automatically granted access to the shared contact or response plan. Otherwise, consumers receive an invitation to join the resource share and are granted access to the shared contact or response plan after accepting the invitation.

You can share a contact or response plan that you own by using the AWS RAM console or the AWS CLI.

To share a contact or response plan that you own by using the AWS RAM console

See <u>Creating a Resource Share</u> in the AWS RAM User Guide.

To share a contact or response plan that you own by using the AWS CLI

Use the create-resource-share command.

Stop sharing a shared contact or response plan

When a resource owner stops sharing a contact or response plan with a consumer, the contacts, response plans, escalation plans, engagements, and incidents no longer appear in the consumer's console.

Note

The consumer continues to see the contacts, response plans, escalation plans, engagements, or incidents without updates, if they're viewing them in the console, until they refresh the page or navigate away from the page.

To stop sharing a shared contact or response plan that you own, you must remove it from the resource share. You can do this by using the AWS RAM console or the AWS CLI.

To stop sharing a shared contact or response plan that you own by using the AWS RAM console

See <u>Updating a Resource Share</u> in the AWS RAM User Guide.

To stop sharing a shared contact or response plan that you own by using the AWS CLI

Use the disassociate-resource-share command.

Identifying a shared contact or response plan

Owners and consumers can identify shared contacts and response plans by using the Incident Manager console and AWS CLI.

To identify a shared contact or response plan by using the Incident Manager console

i Note

Contacts, response plans, escalation plans, engagements, and incidents are generally not identifiable as a shared resource in the Incident Manager console. In places where the Amazon Resource Name (ARN) is visible, the ARN contains the owner's account ID.

To identify a shared contact or response plan by using the AWS CLI

Use the <u>ListResponsePlans</u> or <u>ListContacts</u> commands. The command returns the contacts and response plans that you own and contacts and response plans that are shared with you. The ARN shows the AWS account ID of the contact or response plan owner.

Shared contact and response plan permissions

Permissions for owners

Owners can update, view, share, stop sharing, and use contacts and response plans. Contacts and response plans include related engagements and incidents.

Permissions for consumers

Consumers can use and view only response plans and contacts. Contacts and response plans include related engagements and incidents.

Billing and metering

The owner of the resource is billed for the resource. Consumers aren't billed for resources shared with them. There aren't extra costs associated with sharing a resource.

Instance limits

Sharing a resource doesn't affect the limits of the resource in the owner's or consumer's account. Only the owner's account is used to calculate the limits of the resource.

Compliance validation for AWS Systems Manager Incident Manager

Third-party auditors assess the security and compliance of AWS Systems Manager Incident Manager as part of multiple AWS compliance programs. These include SOC, PCI, FedRAMP, HIPAA, and others.

To learn whether an AWS service is within the scope of specific compliance programs, see <u>AWS</u> <u>services in Scope by Compliance Program</u> and choose the compliance program that you are interested in. For general information, see <u>AWS Compliance Programs</u>.

You can download third-party audit reports using AWS Artifact. For more information, see <u>Downloading Reports in AWS Artifact</u>.

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- <u>Security and Compliance Quick Start Guides</u> These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on AWS that are security and compliance focused.
- <u>Architecting for HIPAA Security and Compliance on Amazon Web Services</u> This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

🚯 Note

Not all AWS services are HIPAA eligible. For more information, see the <u>HIPAA Eligible</u> <u>Services Reference</u>.

- <u>AWS Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>AWS Customer Compliance Guides</u> Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).

- <u>Evaluating Resources with Rules</u> in the *AWS Config Developer Guide* The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>AWS Security Hub</u> This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see <u>Security Hub controls reference</u>.
- <u>Amazon GuardDuty</u> This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- <u>AWS Audit Manager</u> This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in AWS Systems Manager Incident Manager

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see <u>AWS Global Infrastructure</u>.

Incident Manager is a global-regional service and does not currently support Availability Zones.

In addition to the AWS global infrastructure, Incident Manager offers several features to help support your data resiliency and backup needs. During the Getting prepared wizard you're asked to set up a replication set. This regional replication set ensures that your data and resources are accessible from multiple Regions, making incident management across a cloud-network more manageable. This replication also ensures that your data is safe and accessible in the event that one of your Regions goes down.

For more information about using the Incident Manager replication set, see <u>Configuring the</u> <u>Incident Manager replication set</u>.

Infrastructure security in AWS Systems Manager Incident Manager

As a managed service, AWS Systems Manager Incident Manager is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see <u>AWS Cloud Security</u>. To design your AWS environment using the best practices for infrastructure security, see <u>Infrastructure Protection</u> in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Incident Manager through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the <u>AWS Security Token Service</u> (AWS STS) to generate temporary security credentials to sign requests.

Working with AWS Systems Manager Incident Manager and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and AWS Systems Manager Incident Manager by creating an *interface VPC endpoint*. Interface endpoints are powered by AWS PrivateLink. With AWS PrivateLink, you can privately access Incident Manager API operations without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.. Instances in your VPC don't need public IP addresses to communicate with Incident Manager API operations. Traffic between your VPC and Incident Manager stays within the Amazon network.

Each interface endpoint is represented by one or more <u>Elastic Network Interfaces</u> in your subnets.

For more information, see Interface VPC endpoints (AWS PrivateLink) in the Amazon VPC User Guide.

Considerations for Incident Manager VPC endpoints

Before you set up an interface VPC endpoint for Incident Manager, ensure that you review Interface endpoint properties and limitations and AWS PrivateLink quotas in the Amazon VPC User Guide.

Incident Manager supports making calls to all of its API actions from your VPC. To use all of Incident Manager, you must create two VPC endpoints: one for ssm-incidents and one for ssm-contacts.

Creating an interface VPC endpoint for Incident Manager

You can create a VPC endpoint for Incident Manager using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see <u>Creating an interface endpoint</u> in the *Amazon VPC User Guide*.

Create a VPC endpoint for Incident Manager using the following service names:

- com.amazonaws.region.ssm-incidents
- com.amazonaws.*region*.ssm-contacts

If you use private DNS for the endpoint, you can make API requests to Incident Manager using its default DNS name for the Region. For example, you can use the names ssm-incidents.us-east-1.amazonaws.com or ssm-contacts.us-east-1.amazonaws.com.

For more information, see <u>Accessing a service through an interface endpoint</u> in the Amazon VPC User Guide.

Creating a VPC endpoint policy for Incident Manager

You can attach an endpoint policy to your VPC endpoint that controls access to Incident Manager. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which these actions can be performed.

For more information, see <u>Controlling access to services with VPC endpoints</u> in the *Amazon VPC User Guide*.

Example: VPC endpoint policy for Incident Manager actions

The following is an example of an endpoint policy for Incident Manager. When attached to an endpoint, this policy grants access to the listed Incident Manager actions for all principals on all resources.

```
{
    "Statement":[
        {
          "Principal":"*",
          "Effect":"Allow",
          "Action":[
          "ssm-contacts:ListContacts",
          "ssm-incidents:ListResponsePlans",
          "ssm-incidents:StartIncident"
        ],
        "Resource":"*"
      }
   ]
}
```

Configuration and vulnerability analysis in Incident Manager

Configuration and IT controls are a shared responsibility between AWS and you, our customer. For more information, see the AWS shared responsibility model.

Security best practices in AWS Systems Manager Incident Manager

AWS Systems Manager Incident Manager provides many security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

Topics

- Preventative security best practices for Incident Manager
- Detective security best practices for Incident Manager

Preventative security best practices for Incident Manager

Implement least privilege access

When granting permissions, you decide who is getting what permissions to which Incident Manager resources. You enable specific actions that you want to allow on those resources. Therefore, grant only the permissions that are required to perform a task. Implementing least privilege access is fundamental in reducing security risk and the impact that could result from errors or malicious intent.

The following tools are available to implement least privilege access:

- <u>Controlling access to AWS resources using policies and Permissions boundaries for IAM entities</u>
- Service Control Policies

Creating and managing contacts

When activating contacts, Incident Manager reaches out to the device to confirm the activation. Ensure the device information is correct before activating the device. This reduces the possibility that Incident Manager contacts the wrong device or person during activation.

Regularly review your contacts and escalation plans to ensure that only contacts that need to be contacted during an incident are being contacted. Regularly review the contacts to remove outdated or incorrect information. If a contact should no longer be informed when an incident occurs, remove them from the related escalation plans or remove them from Incident Manager.

Make chat channels private

You can make your incident chat channels private to implement least privilege access. Consider using a different chat channel with a scoped down user list for each response plan template. This ensures only the correct responders are pulled into a chat channel that may contain sensitive information.

Slack channels created in AWS Chatbot inherit the permissions of the IAM role used to configure AWS Chatbot. This enables responders in an AWS Chatbot enabled Slack channel to call any allow-listed action, such as Incident Manager APIs and retrieving metrics graphs.

Keep AWS tools up to date

AWS regularly releases updated versions of tools and plugins that you can use in your AWS operations. Keeping these resources up to date ensures that users and instances in your account have access to the latest functionality and security features in these tools.

- AWS CLI The AWS Command Line Interface (AWS CLI) is an open source tool that enables you to interact with AWS services using commands in your command-line shell. To update the AWS CLI, you run the same command used to install the AWS CLI. We recommend creating a scheduled task on your local machine to run the command appropriate to your operating system at least once every two weeks. For information about installation commands, see <u>Installing the</u> <u>AWS Command Line Interface</u> in the AWS Command Line Interface User Guide.
- AWS Tools for Windows PowerShell The Tools for Windows PowerShell are a set of PowerShell
 modules that are built on the functionality exposed by the AWS SDK for .NET. The Tools for
 Windows PowerShell enable you to script operations on your AWS resources from the PowerShell
 command line. Periodically, as updated versions of the Tools for Windows PowerShell are
 released, you should update the version that you're running locally. For information, see
 Updating the AWS Tools for Windows PowerShell on Windows or Updating the AWS Tools for
 Windows PowerShell on Linux or macOS.

Related content

Security best practices for Systems Manager

Detective security best practices for Incident Manager

Identify and audit all your Incident Manager resources

Identification of your IT assets is a crucial aspect of governance and security. Identify your Systems Manager resources to assess their security posture and take action on potential areas of weakness. Create resource groups for your Incident Manager resources. For more information, see <u>What are</u> resource groups? in the AWS Resource Groups User Guide.

Use AWS CloudTrail

AWS CloudTrail provides a record of actions taken by a user, role, or an AWS service in Incident Manager. Using the information collected by AWS CloudTrail, you can determine the request that was made to Incident Manager, the IP address from which the request was made, who made the request, when it was made, and additional details. For more information, see <u>Logging AWS Systems</u> Manager Incident Manager API calls using AWS CloudTrail.

Monitor AWS security advisories

Regularly check security advisories posted in Trusted Advisor for your AWS account. You can do this programmatically using <u>describe-trusted-advisor-checks</u>.

Further, actively monitor the primary email address registered to each of your AWS accounts. AWS will contact you, using this email address, about emerging security issues that might affect you.

AWS operational issues with broad impact are posted on the <u>AWS Service Health Dashboard</u>. Operational issues are also posted to individual accounts through the AWS Health Dashboard. For more information, see the <u>AWS Health documentation</u>.

Related content

Amazon Web Services: Overview of Security Processes (whitepaper)

<u>Getting Started: Follow Security Best Practices as You Configure Your AWS Resources</u> (AWS Security Blog)

IAM Best Practices

Security Best Practices in AWS CloudTrail

Monitoring in Incident Manager

AWS Systems Manager Incident Manager integrates with the following services that offer monitoring and logging capabilities:

CloudWatch metrics

Use CloudWatch metrics to retrieve statistics about data points for your AWS Systems Manager Incident Manager operations as an ordered set of time series data, known as *metrics*. You can use these metrics to verify that your system is performing as expected. For more information, see Monitoring metrics in Incident Manager with Amazon CloudWatch.

CloudTrail logs

Use AWS CloudTrail to capture detailed information about the calls made to AWS APIs. You can store these calls as log files in Amazon Simple Storage Service.. You can use these CloudTrail logs to determine such information as which call was made, the source IP address where the call came from, who made the call, and when the call was made. The CloudTrail logs contain information about the calls to API actions for Incident Manager. IFor more information, see Logging AWS Systems Manager Incident Manager API calls using AWS CloudTrail.

Trusted Advisor

AWS Trusted Advisor can help you monitor your AWS resources to improve performance, reliability, security, and cost effectiveness. Four Trusted Advisor checks are available to all users; more than 50 checks are available to users with a Business or Enterprise support plan. For Incident Manager, Trusted Advisor checks that a replication set's configuration uses more than one AWS Region to support regional failover and response. For more information, see <u>AWS</u> <u>Trusted Advisor</u> in the *AWS Support User Guide*.

Monitoring metrics in Incident Manager with Amazon CloudWatch

Incident Manager provides aggregate metrics that you can monitor in Amazon CloudWatch. You can use these metrics to identify incident and response plan trends.

These metrics include:

• Number of incidents created over a given period of time

- The time to respond to and resolve those incidents
- Number of incidents resolved

You can monitor Incident Manager metrics to better understand your operational health, and take meaningful actions to drive the operational excellence of your incident response. Incident Manager metrics are available in all Incident Manager Regions. Your metrics will be available to view in Amazon CloudWatch for all the Regions you specified in your replication set when on-boarding to Incident Manager. You can view the published metrics in the Region that actions for the incident were taken. There is no additional charge for these metrics.

On the CloudWatch console, you can build dashboards with these metrics to:

- Measure and review your existing incident load
- Track whether your incident load is increasing, decreasing, or remaining the same
- More effectively use Incident Manager to reduce the frequency, duration, and impact of your incidents

This page describes the Incident Manager metrics available on the CloudWatch console.

🔥 Important

For a customer-generated event, if the <u>source</u> value in TriggerDetailsis named using non-ASCII characters, then metrics for the event won't be reported in Amazon CloudWatch metrics, which doesn't support non-ASCII text. source can provided programatically only, such as by using an SDK or the AWS CLI.

Incident Manager sends the following metrics to CloudWatch.

Metric	Description
NumberOfCreateIncidents	Number of incidents created. Valid Dimensions: [](Empty dimension), [ResponsePlan], [Impact], [Source], [ResponsePlan , Impact], [ResponsePlan , Source]

Metric	Description
	Unit: Count
NumberOfResolveIncidents	Number of incidents resolved. Valid Dimensions: [](Empty dimension), [ResponsePlan], [Impact], [Source], [ResponsePlan , Impact], [ResponsePlan , Source] Unit: Count
TimeToFirstAcknowledgement	Time difference between the incident create time and the time the first acknowledgment was made to the incident. Valid Dimensions: [](Empty dimension), [ResponsePlan], [Impact], [Source], [ResponsePlan , Impact], [ResponsePlan , Source] Unit: Seconds
TimeToResolveIncident	Time difference between when the incident was created and when it was resolved. Valid Dimensions:](Empty dimension), [ResponseP lan], [Impact], [Source], [ResponsePlan , Impact], [ResponsePlan , Source] Unit: Seconds

Viewing Incident Manager metrics on the CloudWatch console

To view Incident Manager metrics in the CloudWatch console

- 1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
- 2. In the navigation pane, choose Metrics.

- 3. Select the IncidentManager namespace.
- 4. On the **Metrics** tab, choose a dimension, and then choose a metric.

For more information about working with CloudWatch metrics, see the following topics in the *Amazon CloudWatch User Guide*:

- Metrics
- Using Amazon CloudWatch metrics

Dimensions for Metrics

Incident Manager metrics use the IncidentManager namespace and provide metrics for the following dimension(s):

Dimension	Description
By Response Plan	View aggregate metrics by response plan.
By Impact Level	View aggregate metrics by the level of severity.
By Source	View metrics for incidents created manually, by CloudWatch alarm, or EventBridge event.
Across All Incidents	View aggregate metrics for all incidents in the current AWS Region.
Response Plan name and Source	View aggregate metrics for each combination of response plan and source.
Response Plan Name and Impact Level	View aggregate metrics for each combination of response plan and level of severity.

Logging AWS Systems Manager Incident Manager API calls using AWS CloudTrail

AWS Systems Manager Incident Manager is integrated with <u>AWS CloudTrail</u>, a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures all API calls for Incident Manager as events. The calls captured include calls from the Incident Manager console and code calls to the Incident Manager API operations. Using the information collected by CloudTrail, you can determine the request that was made to Incident Manager, the IP address from which the request was made, when it was made, and additional details.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see <u>Working with CloudTrail Event history</u> in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a <u>CloudTrail</u> <u>Lake</u> event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see <u>Creating a trail for your AWS account</u> and <u>Creating a trail for an organization</u> in the AWS CloudTrail User *Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For more information about CloudTrail pricing, see <u>AWS CloudTrail Pricing</u>. For information about Amazon S3 pricing, see <u>Amazon S3 Pricing</u>.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to <u>Apache ORC</u> format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying <u>advanced</u> <u>event selectors</u>. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see <u>Working</u> with AWS CloudTrail Lake in the AWS CloudTrail User Guide.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the <u>pricing option</u> you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see <u>AWS CloudTrail Pricing</u>.

Incident Manager management events in CloudTrail

<u>Management events</u> provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations. By default, CloudTrail logs management events.

AWS Systems Manager Incident Manager logs all Incident Manager control plane operations as management events. For a list of the AWS Systems Manager Incident Manager control plane operations that Incident Manager logs to CloudTrail, see the <u>AWS Systems Manager Incident</u> <u>Manager API Reference</u>.

Incident Manager event examples

An event represents a single request from any source and includes information about the requested API operation, the date and time of the operation, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so events don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the StartIncident action.

```
{
   "eventVersion": "1.08",
   "userIdentity": {
       "type": "IAMUser",
       "principalId": "1234567890abcdef0",
       "arn": "arn:aws:iam::246873129580111122223333:user/nikki_wolf",
       "accountId": "abcdef01234567890",
       "accessKeyId": "021345abcdef6789",
       "userName": "nikki_wolf"
   },
   "eventTime": "2024-04-22T23:20:10Z",
   "eventSource": "ssm-incidents.amazonaws.com",
   "eventName": "StartIncident",
   "awsRegion": "us-east-2",
   "sourceIPAddress": "192.0.2.0",
   "userAgent": "aws-cli/2.0.58 Python/3.7.4 Darwin/19.6.0 exe/x86_64 command/
ssmincidents.start-incident",
   "requestParameters": {
       test-response-plan-non-dedupe-v1",
       "clientToken": "12345678-1111-2222-3333-abcdefghijkl"
   },
   "responseElements": {
       "incidentRecordArn": "arn:aws:ssm-incidents::444455556666:incident-record/
security-test-response-plan-non-dedupe-v1/abcdefgh-abcd-1234-1234-1234567890"
   },
   "requestID": "abcdefgh-1234-abcd-1234-1234567abcdef",
   "eventID": "12345678-1234-1234-abcd-abcdef1234567",
   "readOnly": false,
   "eventType": "AwsApiCall",
   "managementEvent": true,
   "eventCategory": "Management",
   "recipientAccountId": "12345678901234567"
}
```

The following example shows a CloudTrail log entry that demonstrates the DeleteContactChannel action.

```
"eventVersion": "1.08",
```

{

```
"userIdentity": {
        "type": "IAMUser",
        "principalId": "1234567890abcdef0",
        "arn": "arn:aws:iam::246873129580111122223333:user/nikki_wolf",
        "accountId": "abcdef01234567890",
        "accessKeyId": "021345abcdef6789",
        "userName": "nikki_wolf"
    },
    "eventTime": "2024-04-08T02:27:21Z",
    "eventSource": "ssm-contacts.amazonaws.com",
    "eventName": "DeleteContactChannel",
    "awsRegion": "us-east-2",
    "sourceIPAddress": "192.0.2.0",
    "userAgent": "Apache-HttpClient/UNAVAILABLE (Java/1.8.0_282)",
    "requestParameters": {
        "contactChannelId": "arn:aws:ssm-contacts:us-west-2:5555555555555device/
bnuomysohc/abcdefgh-abcd-1234-1234-1234567890"
    },
    "responseElements": null,
    "requestID": "abcdefgh-1234-abcd-1234-1234567abcdef",
    "eventID": "12345678-1234-1234-abcd-abcdef1234567",
    "readOnly": true,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "12345678901234567"
}
```

For information about CloudTrail record contents, see <u>CloudTrail record contents</u> in the AWS *CloudTrail User Guide*.

Product and service integrations with Incident Manager

Incident Manager, a capability of AWS Systems Manager, integrates with the following products, services, and tools.

Integration with AWS services

Incident Manager integrates with the AWS services and tools described in the following table.

AWS CDK	The AWS CDK is a development framework for using code to define your cloud infrastru cture and using AWS CloudFormation for provisioning. The AWS CDK supports multiple programming languages including TypeScript, JavaScript, Python, Java, and C#/.Net.
	For information about using the AWS CDK with Incident Manager, see the following sections in the AWS CDK API Reference:
	 <u>@aws-cdk/aws-ssmincidents module</u> <u>@aws-cdk/aws-ssmcontacts module</u>
AWS Chatbot	<u>AWS Chatbot</u> enables DevOps and software development teams to use messaging program chat rooms to monitor and respond to operational events in their AWS Cloud.
	Using AWS Chatbot with Incident Manager, you can create <i>chat channels</i> that responders can use to monitor and respond to incidents . AWS Chatbot supports Slack chat rooms, Microsoft Teams channels, and Amazon Chime chat rooms as chat channels.
	As part of creating a chat channel, you also create a <i>topic</i> in Amazon Simple Notificat

	 ion Service (Amazon SNS). <u>Amazon SNS</u> is a managed service that provides message delivery from publishers to subscribers. In incident response plans, when you associate a chat channel you have created with the plan, you also choose one or more topics that you associated with the chat channel. These SNS topics are used to send notifications about an incident to the incident responders. For more information, see <u>Creating and</u> <u>integrating chat channels for responders in</u> <u>Incident Manager</u>.
AWS CloudFormation	AWS CloudFormation is a service that you can use to create a template with all the resources you need for your application, and then configure and provision the resources for you. It will also configure all the dependencies, so you can focus more on your application and less on managing resources. For information about using AWS CloudForm ation with Incident Manager, see the following topics in the <u>AWS CloudFormation User Guide</u> : <u>Incident Manager resource type reference</u> <u>Contacts resource type reference resource</u> type reference

Amazon CloudWatch	CloudWatch monitors your AWS resources and the applications you run on AWS in real time. You can use CloudWatch to collect and track metrics, which are variables you can measure for your resources and applications. You can configure CloudWatch alarms to create incidents in Incident Manager. CloudWatch works with Systems Manager and Incident Manager to create an incident from a response plan template when an alarm goes
	into alarm state. For more information, see <u>Creating incidents</u> <u>automatically with CloudWatch alarms</u> .
Amazon Chime	Amazon Chime is an online workplace that combines meetings, chat, and business calls. You can meet, chat, and place business calls inside and outside your organization using Amazon Chime.
	You can integrate an Amazon Chime room into your Incident Manager operations by creating a chat channel for Amazon Chime in <u>AWS</u> <u>Chatbot</u> , and then adding that channel to a response plan.
	For more information, see <u>Creating and</u> integrating chat channels for responders in Incident Manager.

Amazon EventBridge	 EventBridge is a serverless service that uses events to connect application component s, making it easier for you to build scalable event-driven applications. You can configure EventBridge rules to watch for event patterns in your AWS resources and create an incident in Incident Manager when an event matches a pattern that you have defined. Your rules can monitor for event patterns in dozens of AWS services and thirdparty applications and services. For more information, see <u>Creating incidents automatically with EventBridge events</u>.
AWS Secrets Manager	 <u>Secrets Manager</u> helps you manage, retrieve, and rotate database credentials, application credentials, OAuth tokens, API keys, and other secrets throughout their lifecycles. When you integrate Incident Manager with the PagerDuty service, you create a secret in Secrets Manager that contains your PagerDuty credentials. For more information, see <u>Storing PagerDuty access credentials in an AWS Secrets Manager secret.</u>

AWS Systems Manager

<u>Systems Manager</u> is an operations hub that you can use to view and control your applicati on infrastructure and a secure end-to-end management solution for cloud environments. The following Systems Manager capabilities integrate directly with Incident Manager:

 <u>Automation</u> – An Automation runbook defines the actions that Systems Manager performs on your AWS resources. In Incident Manager, a runbook defines a series of automated and manual steps to use to resolve your incidents.

For information about creating Automatio n runbooks for use with Incident Manager, see <u>Integrating Systems Manager Automatio</u> <u>n runbooks in Incident Manager for incident</u> remediation.

 OpsCenter – OpsCenter provides a central location where operations engineers and IT professionals can manage operational work items, called *OpsItems*, related to AWS resources. You can create OpsItems directly from a post-incident analysis to follow up on related work.

For more information, see <u>Performing a</u> post-incident analysis in Incident Manager.

AWS Trusted Advisor

<u>Trusted Advisor</u> is a tool available to AWS customers with a Basic or Developer support plan. Trusted Advisor inspects your AWS environment, and then makes recommend ations when opportunities exist to save money, improve system availability and performance, or help close security gaps.

For Incident Manager, Trusted Advisor checks that a replication set's configuration uses more than one AWS Region to support Regional failover and response.

Integration with other products and services

You can integrate or use Incident Manager with the third-party services described in the following table.

Jira Cloud	
------------	--

Using the AWS Service Management Connector , you can integrate Incident Manager with <u>Jira</u> <u>Cloud</u> (Atlassian), a third-party cloud-based workflow platform.

After you configure integration with Jira Cloud, when you create a new incident in Incident Manager, the integration creates the incident in Jira Cloud as well. If you update an incident in Incident Manager, it makes these updates to the corresponding incident in Jira Cloud. If you resolve an incident in either Incident Manager or Jira Cloud, the integration resolves the incident in both services based on which preferences you configure.

For more information, see <u>Integrating AWS</u> Systems Manager Incident Manager (Jira

	<u>Cloud</u>) in the AWS Service Management Connector Administrator Guide.
Jira Service Management	Using the AWS Service Management Connector , you can integrate Incident Manager with <u>Jira</u> <u>Service Management</u> , a third-party cloud-bas ed workflow platform.
	After you configure integration with Jira Service Management, when you create a new incident in Incident Manager, the integration creates the incident in Jira Service Management as well. If you update an incident in Incident Manager, it makes these updates to the corresponding incident in Jira Service Management. If you resolve an incident in either Incident Manager or Jira Service Management, the integration resolves the incident in both services based on which preferences you configure. For more information, see <u>Configuring Jira</u> <u>Service Management</u> in the AWS Service Management Connector Administrator Guide.
Microsoft Teams	 <u>Microsoft Teams</u> provides collaborative cloud- based tools for team messaging, audio and video conferencing, and file sharing. You can integrate a Microsoft Teams channel into your Incident Manager operations by creating a chat channel for Microsoft Team in <u>AWS Chatbot</u>, and then adding that channel to a response plan.
	For more information, see <u>Creating and</u> integrating chat channels for responders in Incident Manager.

PagerDuty

<u>PagerDuty</u> is an incident response tool that supports paging workflows and escalation policies.

When you integrate Incident Manager with PagerDuty, you can add a PagerDuty service to your response plan. After that, a correspon ding incident is created in PagerDuty whenever an incident in created in Incident Manager. The incident in PagerDuty uses the paging workflow and escalation policies that you defined there in addition to those in Incident Manager. PagerDuty attaches timeline events from Incident Manager as notes on your incident.

To integrate Incident Manager with PagerDuty , you must first create a secret in AWS Secrets Manager that contains your PagerDuty credentials.

For information about adding a PagerDuty REST API Key and other required details to a secret in AWS Secrets Manager, see <u>Storing</u> <u>PagerDuty access credentials in an AWS Secrets</u> <u>Manager secret</u>.

For information about adding a PagerDuty service from your PagerDuty account to a response plan in Incident Manager, see the steps for <u>Integrate a PagerDuty service into the</u> <u>response plan</u> in the topic <u>Creating a response</u> <u>plan</u>.

ServiceNow

Using the AWS Service Management Connector , you can integrate Incident Manager with <u>ServiceNow</u>, a third-party cloud-based workflow platform.

After you configure integration with ServiceNo w, when you create a new incident in Incident Manager, the integration creates the incident in ServiceNow as well. If you update an incident in Incident Manager, it makes these updates to the corresponding incident in ServiceNow. If you resolve an incident in either Incident Manager or ServiceNow, the integrati on resolves the incident in both services based on which preferences you configure.

For more information, see <u>Integrating AWS</u> <u>Systems Manager Incident Manager in</u> <u>ServiceNow</u> in the AWS Service Management Connector Administrator Guide.

<u>Slack</u> provides collaborative cloud-based tools for team messaging, audio and video conferencing, and file sharing.

You can integrate a Slack channel into your Incident Manager operations by creating a chat channel for Slack in <u>AWS Chatbot</u>, and then adding that channel to a response plan.

For more information, see <u>Creating and</u> integrating chat channels for responders in Incident Manager.

Slack

Terraform

HashiCorp <u>Terraform</u> is an open-source infrastructure as code (IaC) software tool that provides a command line interface (CLI) workflow to manage various cloud services. For Incident Manager, you can use Terraform to manage or provision the following:

SSM Incident Manager Contacts resources

- <u>aws_ssmcontacts_contact</u>
- <u>aws_ssmcontacts_contact_channel</u>
- <u>aws_ssmcontacts_plan</u>
- <u>aws_ssmcontacts_rotation</u>

SSM Contacts data sources

- <u>aws_ssmcontacts_contact</u>
- aws_ssmcontacts_contact_channel
- <u>aws_ssmcontacts_plan</u>
- aws_ssmcontacts_rotation

SSM Incident Manager resources

- aws_ssmincidents_replication_set
- <u>aws_ssmincidents_response_plan</u>

SSM Incident Manager data sources

- aws_ssmincidents_replication_set
- <u>aws_ssmincidents_response_plan</u>

Storing PagerDuty access credentials in an AWS Secrets Manager secret

After you turn on integration with PagerDuty for a response plan, Incident Manager works with PagerDuty in the following ways:

- Incident Manager creates a corresponding incident in PagerDuty when your create a new incident in Incident Manager.
- The paging workflow and escalation policies you created in PagerDuty are used in the PagerDuty environment. However, Incident Manager doesn't import your PagerDuty configuration.
- Incident Manager publishes timeline events as notes to the incident in PagerDuty, up to a maximum of 2,000 notes.
- You can choose to automatically resolve PagerDuty incidents when you resolve the related incident in Incident Manager.

To integrate Incident Manager with PagerDuty, you must first create a secret in AWS Secrets Manager that contains your PagerDuty credentials. These allow Incident Manager to communicate with your PagerDuty service. You can then include a PagerDuty service in response plans that you create in Incident Manager.

This secret you create in Secrets Manager must contain, in the proper JSON format, the following:

- An API key from your PagerDuty account. You can use either a General Access REST API Key or a User Token REST API Key.
- A valid user email address from your PagerDuty subdomain.
- The PagerDuty service region where you deployed your subdomain.

🚯 Note

All services in a PagerDuty subdomain are deployed to the same service region.

Prerequisites

Before creating the secret in Secrets Manager, ensure that you meet the following requirements.

KMS key

You must encrypt the secret you create with a *customer managed key* that you have created in AWS Key Management Service (AWS KMS). You specify this key when you create the secret that stores you PagerDuty credentials.

🔥 Important

Secrets Manager provides the option of encrypting the secret with an AWS managed key, but this encryption mode is not supported.

The customer managed key must meet the following requirements:

- Key type: Choose Symmetric.
- Key usage: Choose Encrypt and decrypt.
- **Regionality**: If you want to replicate your response plan to multiple AWS Regions, ensure that you select **Multi-Region key**.

Key policy

The user that is configuring the response plan must have permission for kms:GenerateDataKey and kms:Decrypt in the key's resource-based policy. The ssm-incidents.amazonaws.com service principal must have permission for kms:GenerateDataKey and kms:Decrypt in the key's resource based policy.

The following policy demonstrates these permissions. Replace each *user input placeholder* with your own information.

```
"Resource": "*"
    },
    {
        "Sid": "Allow creator of response plan to use the key",
        "Effect": "Allow",
        "Principal": {
            "AWS": "IAM_ARN_of_principal_creating_response_plan"
        },
        "Action": [
            "kms:Decrypt",
            "kms:GenerateDataKey*"
        ],
        "Resource": "*"
    },
    {
        "Sid": "Allow Incident Manager to use the key",
        "Effect": "Allow",
        "Principal": {
            "Service": "ssm-incidents.amazonaws.com"
        },
        "Action": [
            "kms:Decrypt",
            "kms:GenerateDataKey*"
        ],
        "Resource": "*"
    }
]
```

For information about creating a new customer managed key, see <u>Creating symmetric</u> <u>encryption KMS keys</u> in the AWS Key Management Service Developer Guide. For more information about AWS KMS keys, see AWS KMS concepts.

If an existing customer managed key meets all the previous requirements, you can edit its policy to add these permissions. For information about updating the policy in a customer managed key, see <u>Changing a key policy</u> in the AWS Key Management Service Developer Guide.

🚺 Tip

}

You can specify a condition key to limit access even further. For example, the following policy allows access through Secrets Manager in the US East (Ohio) Region (us-east-2) only:

```
{
    "Sid": "Enable IM Permissions",
    "Effect": "Allow",
    "Principal": {
        "Service": "ssm-incidents.amazonaws.com"
    },
    "Action": ["kms:Decrypt", "kms:GenerateDataKey*"],
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "kms:ViaService": "secretsmanager.us-east-2.amazonaws.com"
        }
    }
}
```

GetSecretValue permission

The IAM identity (user, role, or group) that creates the response plan must have the IAM permission secretsmanager:GetSecretValue.

To store PagerDuty access credentials in an AWS Secrets Manager secret

- Follow the steps through Step 3a in <u>Create an AWS Secrets Manager secret</u> in the AWS Secrets Manager User Guide.
- 2. For Step 3b, for Key/value pairs, do the following:
 - Choose the **Plaintext** tab.
 - Replace the default contents of the box with the following JSON structure:



- In the JSON sample you pasted, replace the *placeholder values* as follows:
 - pagerduty-token: The value of a General Access REST API Key or a User Token REST API Key from your PagerDuty account.

For related information, see <u>API Access Keys</u> in the *PagerDuty Knowledge Base*.

 pagerduty-region: The service region of the PagerDuty data center that hosts your PagerDuty subdomain.

For related information, see <u>Service Regions</u> in the *PagerDuty Knowledge Base*.

pagerduty-email: The valid email address for a user that belongs to your PagerDuty subdomain.

For related information, see <u>Manage Users</u> in the *PagerDuty Knowledge Base*.

The following example shows a completed JSON secret containing the required PagerDuty credentials:

```
{
    "pagerDutyToken": "y_NbAkKc66ryYEXAMPLE",
    "pagerDutyServiceRegion": "US",
    "pagerDutyFromEmail": "JohnDoe@example.com"
}
```

- 3. On Step 3c, for **Encryption key**, choose a customer managed key you created that meets the requirements listed under the previous **Prerequisites** section.
- 4. On Step 4c, for **Resource permissions**, do the following:
 - Expand Resource permissions.
 - Choose Edit permissions.
 - Replace the default contents of the policy box with the following JSON structure:

```
{
    "Effect": "Allow",
    "Principal": {
        "Service": "ssm-incidents.amazonaws.com"
    },
    "Action": "secretsmanager:GetSecretValue",
    "Resource": "*"
}
```

- Choose Save.
- 5. On Step 4d, for **Replicate secret**, do the following if you replicated your response plan to more than one AWS Region:

- Expand Replicate secret.
- For **AWS Region**, select the Region where you replicated your response plan to.
- For **Encryption key**, choose a customer managed key you created in, or replicated to, this Region that meets the requirements listed under the **Prerequisites** section.
- For each additional AWS Region, choose **Add Region** and select the Region name and customer managed key.
- 6. Complete the remaining steps in <u>Create an AWS Secrets Manager secret</u> in the AWS Secrets Manager User Guide.

For information about how to add a PagerDuty service to a Incident Manager incident workflow, see Integrate a PagerDuty service into the response plan in the topic Creating a response plan.

Related information

How to Automate Incident Response with PagerDuty and AWS Systems Manager Incident Manager (AWS Cloud Operations and Migrations Blog)

Secret encryption in AWS Secrets Manager in the AWS Secrets Manager User Guide

Troubleshooting AWS Systems Manager Incident Manager

If you encounter issues while using AWS Systems Manager Incident Manager, you can use the following information to resolve them according to our best practices. If the issues you encounter are outside the scope of the following information, or if they persist after you've tried to resolve them, contact <u>AWS Support</u>.

Topics

- Error message: ValidationException We were unable to validate the AWS Secrets Manager secret
- Other troubleshooting issues

Error message: ValidationException – We were unable to validate the AWS Secrets Manager secret

Problem 1: The AWS Identity and Access Management (IAM) identity (user, role, or group) that creates the response plan doesn't have the secretsmanager:GetSecretValue IAM permission. IAM identities must have this permission to validate Secrets Manager secrets.

• **Solution**: Add the missing secretsmanager: GetSecretValue permission to the IAM policy for the IAM identity that creates the response plan. For information, see <u>Adding IAM identity</u> permissions (console) or <u>Adding IAM policies (AWS CLI)</u> in the *IAM User Guide*.

Problem 2: The secret doesn't have a resource-based policy attached that allows the IAM identity to run the <u>GetSecretValue</u> action, or the resource-based policy denies permission to the identity.

• Solution: Create or add an Allow statement to the secret's resource-based policy that grants permission for secrets: GetSecretValue to the IAM identity. Or, if you use a Deny statement that includes the IAM identity, update the policy so the identity can run the action. For information, see <u>Attach a permissions policy to an AWS Secrets Manager secret</u> in the AWS Secrets Manager User Guide.

Problem 3: The secrets doesn't have a resource-based policy attached that allows access to the Incident Manager service principal, ssm-incidents.amazonaws.com.

• **Solution**: Create or update the resource-based policy for the secret and include the following permission:

```
{
    "Effect": "Allow",
    "Principal": {
        "Service": ["ssm-incidents.amazonaws.com"]
    },
    "Action": "secretsmanager:GetSecretValue",
    "Resource": "*"
}
```

Problem 4: The AWS KMS key selected to encrypt the secret isn't a customer managed key, or the selected customer managed key doesn't provide the IAM permissions kms:Decrypt and kms:GenerateDataKey* to the Incident Manager service principal. Alternately, the IAM identity that creates the response plan may not have the IAM permission <u>GetSecretValue</u>.

 Solution: Ensure that you meet the requirements described under Prerequisites in the topic <u>Storing PagerDuty access credentials in an AWS Secrets Manager secret</u>.

Problem 5: The ID of the secret that contains the General Access REST API Key or User Token REST API Key isn't valid.

• **Solution**: Ensure that you entered the ID of the Secrets Manager secret accurately, with no trailing space. You must work in the same AWS Region that stores the secret you want to use. You can't use a deleted secret.

Problem 6: In rare cases, the Secrets Manager service may experience an issue, or Incident Manager might have trouble communicating with it.

• **Solution**: Wait a few minutes, then try again. Check the <u>AWS Health Dashboard</u> for any issues that might affect either service.

Other troubleshooting issues

If the previous steps didn't resolve your issue, you can find additional help from the following resources:

- For IAM issues specific to Incident Manager when you access the <u>Incident Manager console</u>, see <u>Troubleshooting AWS Systems Manager Incident Manager identity and access</u>.
- For general authentication and authorization issues when you access the AWS Management Console, see Troubleshooting IAM in the *IAM User Guide*

Document history for Incident Manager

Change

Update to managed policy AWSIncidentManager IncidentAccessServ iceRolePolicy

Additional HashiCorp Terraform support: On-call rotations

New feature: Findings from other AWS services

Description

Incident Manager has added a new permission to AWSIncidentManager IncidentAccessServ iceRolePolicy , in support of the Findings feature, that allows it to check whether an EC2 instance is part of an Auto Scaling group. For more information, see <u>Incident</u> <u>Manager updates to AWS</u> <u>managed policies</u>.

Terraform has added to its support for Incident Manager. You can now provision or manage Incident Manager oncall resources using Terraform . For information about this and other third-party integrations with Incident Manager, see <u>Integration with</u> <u>other products and services</u>.

Findings provide you with information about changes related to AWS CloudForm ation stacks and AWS CodeDeploy deployments that occurred around the same time that an incident Date

February 20, 2024

February 2, 2024

November 15, 2023

was created in Incident Manager. In the Incident Manager console, you can view summary information about those changes and, in many cases, access links to the CloudFormation or CodeDeploy consoles for complete details about the change. Findings reduce the time required to evaluate potential causes of incidents. They also reduce the chances of responders accessing the wrong account or console to investigate the cause of an incident. This feature also introduces a new managed policy, AWSIncide ntManagerIncidentA ccessServiceRolePo licy, which allows Incident Manager to read resources in other AWS services to identify findings related to incidents. For more information, see the following topics:

- Working with findings
- <u>AWS managed policy:</u> <u>AWSIncidentManager</u> <u>IncidentAccessServ</u> <u>iceRolePolicy</u>

Updated lists of integrations with Incident Manager

The topic Product and service integrations with Incident Manager has been expanded to list and describe all AWS services and third-party tools that you can integrate with Incident Manager into your incident detection and response operations.

Integration with AWS Trusted Advisor

Trusted Advisor now checks that a replication set's configuration uses more than one AWS Region to support regional failover and response. For incidents created by CloudWatch alarms or EventBridge events, Incident Manager creates an incident in the same AWS Region as the alarm or event rule. If Incident Manager is temporarily unavailable in that Region, the system attempts to create an incident in another Region in the replication set. If the replicati on set includes only one Region, the system fails to create an incident record while Incident Manager is unavailable. To help avoid this situation, Trusted Advisor reports when a replication set is configured for only one Region. For information about working with Trusted Advisor, see AWS Trusted Advisor in the AWS Support User Guide.

April 28, 2023

Use Microsoft Teams as a chat channel in response plans

Through integration with Microsoft Teams and AWS Chatbot, you can now use Microsoft Teams for the chat channel in your response plans. This is in addition to support for Slack and Amazon Chime chat channels. During an incident, Incident Manager sends status notifications directly to a chat channel to keep all responders informed. Responders can also communicate with one another and incident-related AWS CLI commands in the Microsoft Teams applicati on to update and interact with the incidents. For more information, see Working with chat channels in Incident Manager.

April 4, 2023

New feature: On-call schedules

Print a formatted incident analysis or save as PDF An on-call schedule in **Incident Manager defines** who is notified when an incident occurs that requires operator intervention. An on-call schedule consists of one or more rotations you create for the schedule. Each rotation can include up to 30 contacts. After you create an on-call schedule, you can include it as an escalation in your escalation plan. When an incident associated with that escalation plan occurs, **Incident Manager notifies** the operator (or operators) who are on call according to the schedule. For more information, see Working with on-call schedules in Incident Manager.

January 17, 2023

The incident analysis page now includes a **Print** button to generate a version of the analysis that's formatted for printing. Using the printer destinations configured for your device, you can save the incident analysis as a PDF or send it to a local or network printer. For more information, see <u>Print a formatted incident</u> <u>analysis</u>. March 28, 2023

PagerDuty integration: Incident Manager now copies incident timeline events to PagerDuty incidents When you turn on integration with PagerDuty in a response plan, Incident Manager adds timeline events created from that plan to the correspon ding incident record in PagerDuty. PagerDuty adds timeline events as notes on the incident, up to a maximum of 2,000 notes. To learn more about these changes, see the following topics:

- Store PagerDuty access credentials in an AWS Secrets Manager secret
- Integrate a PagerDuty service into the response plan

Incident Manager integration with CloudWatch metrics. You can now have incidentrelated metrics published in CloudWatch. For more information, see <u>CloudWatc</u> <u>h metrics</u>. The <u>AWSIncide</u> <u>ntManagerServiceRolePolicy</u> has included an additiona l permission to allow our service to publish metrics on your behalf. December 15, 2022

December 15, 2022

Launched Incident notes and updated the Incident Details screen You can collaborate and communicate with other users that work on an incident using **Incident notes**. Additionally, you can view runbooks and engagements statuses from the **Incident Details** screen. For more information, see <u>Incident</u> <u>Details</u>. November 16, 2022

Integrate PagerDuty escalatio n plans and paging workflows into Incident Manager response plans You can now integrate Incident Manager with PagerDuty and add a PagerDuty service to a response plan. After you configure integration, Incident Manager can create a corresponding incident in PagerDuty for each new incident created in Incident Manager. PagerDuty uses the paging workflow and escalation policies you define in the PagerDuty environme nt.

For more information, see the following topics:

- Product and service integrations with Incident Manager
- Store PagerDuty access credentials in an AWS Secrets Manager secret
- Integrate a PagerDuty service into the response plan in the topic Creating a response plan
- <u>Troubleshooting</u>

November 16, 2022

November 16, 2022

Launched Incident notes and updated the Incident Details screen.

Tagging support for replicati on sets You can collaborate and communicate with other users that work on an incident using **Incident notes**. Additionally, you can view runbooks and engagements statuses from the **Incident Details** screen. For more information, see <u>Incident</u> <u>Details</u>.

You can now assign tags to your replication set in AWS Systems Manager Incident Manager. This adds to existing support for assigning tags to response plans, incident records, and contacts in the AWS Regions specified in your replication set. For informati on, see the following topics:

- Get prepared wizard
- <u>Tagging Incident Manager</u> resources

November 2, 2022

Incident Manager integrati on with Atlassian Jira Service Management You can integrate Incident Manager with Jira Service Management by using the **AWS Service Management** Connector for Jira Service Management. After you configure integration, new incidents created in Incident Manager create a correspon ding incident in Jira. If you update an incident in Incident Manager, the updates are added to the correspon ding incident in Jira. If you resolve an incident in either Incident Manager or Jira, the corresponding incident is also resolved, based on configure d preferences. For more information, see Configuring Jira Service Management in the AWS Service Managemen t Connector Administrator Guide.

October 6, 2022

Enhanced tagging support	Incident Manager supports assigning tags to response plans, incident records, and contacts in the AWS Regions specified in your replicati on set. Incident Manager also supports automatically assigning tags to incidents created from response plans. For more information, see <u>Tagging Incident Manager</u> <u>resources</u> .	June 28, 2022
Incident Manager integration with ServiceNow	You can integrate Incident Manager with <u>ServiceNow</u> by using the AWS Service Management Connector for ServiceNow. After you configure integration, new incidents created in Incident Manager create a correspon ding incident in ServiceNo w. If you update an incident in Incident Manager, the updates are added to the corresponding incident in ServiceNow. If you resolve an incident in either Incident Manager or ServiceNow, the corresponding incident is also resolved, based on configured preferences. For more information, see Integrating AWS Systems	June 9, 2022
	<u>Manager Incident Manager in</u> ServiceNow.	

Import contact details

When an incident is created, Incident Manager can notify responders by using voice or SMS notifications. To ensure that responders see that the call or SMS notification is from Incident Manager, we recommend that all responders download the Incident Manager virtual card format (.vcf) file to the address book on their mobile devices. For more informati on, see Import contact details to your address book. May 18, 2022

Multiple feature improveme nts to enhance incident creation and remediation Incident Manager launched the following feature improvements to enhance incident creation and remediation:

- Automatically create incidents in other AWS Regions: In the event that Incident Manager is not available in an AWS Region when Amazon CloudWatc h or Amazon EventBridge create an incident, these services now automatically create the incident in one of the available Regions specified in your replication set. For more information, see Cross-Region incident management.
- Automatically populate runbook parameters with incident metadata: You can now configure Incident Manager to collect information about AWS resources from incidents
 . Incident Manager can then populate runbook parameters with the collected information. For more information, see <u>Tutorial: Using Systems</u> Manager Automation

May 17, 2022

	Manager. • Automatically collect AWS resource information: When the system creates an incident, Incident Manager now automatically collects information about the AWS resources involved in the incident. Incident Manager then adds this information to the Related items tab.	
<u>Multi-runbook support</u>	Incident Manager now supports running multiple runbooks during an incident for the incident details page.	January 14, 2022
Incident Manager launched in new AWS Regions	Incident Manager is now available in these new Regions: us-west-1, sa-east-1 , ap-northeast-2, ap-south-1, ca-central-1, eu-west-2, and eu-west-3. For more informati on about Incident Manager Regions and quotas, see the <u>AWS General Reference</u> reference guide.	November 8, 2021
Console engagement acknowledgement	You can now acknowledge engagements directly from the Incident Manager console.	August 5, 2021

runbooks with Incident

<u>Properties tab</u>	Incident Manager introduce d a properties tab to the incident details page, providing more informati on about the incidents, the parent OpsItem, and the related post-incident analysis.	August 3, 2021
Incident Manager launch	Incident Manager is an incident management console designed to help users mitigate and recover from incidents affecting their AWS hosted applications.	May 10, 2021