



Infrastructure Performance User Guide

AWS Network Manager



AWS Network Manager: Infrastructure Performance 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 Infrastructure Performance? | 1 |
| How Infrastructure Performance works | 3 |
| How latency is calculated | 3 |
| Region availability | 4 |
| Monitor network performance using the console | 5 |
| Health status timeline | 8 |
| Network latency | 9 |
| Monitor network performance using the CLI | 10 |
| CloudWatch subscriptions | 13 |
| Subscription metrics | 13 |
| Manage subscriptions using the console | 14 |
| Manage subscriptions using the CLI | 14 |
| DescribeAwsNetworkPerformanceMetricSubscriptions | 14 |
| DisableAwsNetworkPerformanceMetricSubscription | 16 |
| EnableAwsNetworkPerformanceMetricSubscription | 16 |
| Identity and access management | 18 |
| How Infrastructure Performance works with IAM | 18 |
| Infrastructure Performance identity-based policies | 18 |
| Infrastructure Performance IAM roles | 20 |
| Allow IAM users to access Infrastructure Performance | 21 |
| Create an IAM policy | 21 |
| Document history | 23 |

What is Infrastructure Performance for AWS Network Manager?

Infrastructure Performance allows you to obtain near real-time and historical network latency across AWS Regions and across or within Availability Zones for a specified time period. However, network performance data is not available prior to Oct 26 2022 00:00:00 GMT+000. Performance metrics are aggregated according to a time period that you choose. Network performance results are returned for that time period, showing whether the performance was healthy or if there was any performance degradation. This can help you to more easily evaluate whether network performance might affect your applications. Infrastructure Performance also supports publishing metrics as subscriptions to Amazon CloudWatch, allowing you to view performance metrics in CloudWatch. There is no cost associated with using Infrastructure Performance. However, you might incur charges when using subscriptions to view performance in CloudWatch. For more information about CloudWatch pricing guidelines, see [Amazon CloudWatch pricing](#).

Note

Infrastructure Performance is designed to give you an overview showing whether network performance is normal or degraded. It does not provide details about the specific causes of degradation, such as a hardware failure.

Infrastructure Performance works in the following geographical areas:

- **Availability zones** — Generate network performance metrics between two Availability Zones or within a single Availability Zone.
- **Regions** — Generate network performance metrics to see performance between two Regions.

You can use the Region metrics to evaluate a Region expansion strategy. Infrastructure Performance can help you gather information needed to understand network performance between Regions where you currently have a presence, and Regions that you're exploring expanding into. For example, you might have a presence in `us-west-2` and are looking to expand into `eu-central-1`. Use Infrastructure Performance to check the latency between these two Regions to help you make a more informed choice for Region expansion. For a list of the Regions that support Infrastructure Performance, see [the section called "Region availability"](#).

Note

Infrastructure Performance does not incorporate performance metrics for paths through VPC networking resources, such as transit gateways, NAT gateways, VPC endpoints, Elastic Load Balancing, or Amazon EC2 network interfaces.

How Infrastructure Performance for AWS Network Manager works

Network latency measurements are an aggregate of latency performance information captured by active probing between AWS managed probes within the AWS global network. Using Infrastructure Performance, AWS Network Manager allows you to view network performance based on the following metrics:

- Inter-Region latency metrics, generated by aggregating latency measurements from probes located across AWS Regions, filtered on your chosen source and destination AWS Region pairs.
- Inter-Availability Zone latency metrics, generated by aggregating latency measurements from probes located across Availability Zones, filtered on your chosen source and destination Availability Zones.
- Intra-Availability Zone latency metrics, generated by aggregating latency measurements between all probes within a single Availability Zone. This includes all probes that are across AWS data centers, and within the same data center, for the chosen Availability Zone.

The placement of the probes is not related to EC2 instances or AWS services running in your account. For these reasons, the same latency metric data is presented in all AWS accounts.

How latency is calculated

Metric data is generated by computing the median (P50) of all latency measurements from AWS managed probes for every five-minute (5M) interval.

Latency metrics represent round-trip traffic latency of synthetic traffic streams between selected sources and destinations. This is a combination of the latency of the underlying physical network and the latency between the host and the underlying physical network. The Infrastructure Performance dashboard also provides a network health status, indicating whether the aggregate network latency is normal or degraded. The health status is determined based on expected AWS network performance between AWS Regions and Availability Zones. The health status is available only for inter-Region and inter-Availability Zone metrics.

Note

In case of a degraded health status, AWS is automatically notified of the incident.

Region availability

Infrastructure Performance is available in the following AWS Regions:

| AWS Region | Description |
|----------------|--------------------------|
| us-east-1 | US East (N. Virginia) |
| us-east-2 | US East (Ohio) |
| us-west-1 | US West (N. California) |
| us-west-2 | US West (Oregon) |
| ap-east-1 | Asia Pacific (Hong Kong) |
| ap-south-1 | Asia Pacific (Mumbai) |
| ap-northeast-3 | Asia Pacific (Osaka) |
| ap-northeast-2 | Asia Pacific (Seoul) |
| ap-southeast-1 | Asia Pacific (Singapore) |
| ap-southeast-2 | Asia Pacific (Sydney) |
| ap-northeast-1 | Asia Pacific (Tokyo) |
| ap-southeast-3 | Asia Pacific (Jakarta) |
| ca-central-1 | Canada (Central) |
| eu-central-1 | Europe (Frankfurt) |
| eu-west-1 | Europe (Ireland) |

| AWS Region | Description |
|------------|---------------------------|
| eu-west-2 | Europe (London) |
| eu-west-3 | Europe (Paris) |
| eu-north-1 | Europe (Stockholm) |
| eu-south-1 | Europe (Milan) |
| sa-east-1 | South America (São Paulo) |
| af-south-1 | Africa (Cape Town) |
| me-south-1 | Middle East (Bahrain) |

Monitor Infrastructure Performance using the AWS Network Manager console

View the network performance between two Regions, between two Availability Zones, or in the two connections within an Availability Zone. After you choose a starting and ending date, and a period by which to aggregate performance, the network performance metrics display in the **Health status timeline** and **Network latency timeline** for that time period. You can also subscribe or unsubscribe performance metrics from Amazon CloudWatch. When you subscribe, you're able to view performance metrics in CloudWatch.

Note

Enabling subscriptions with CloudWatch might incur a charge. For more information about CloudWatch pricing guidelines, see [Amazon CloudWatch pricing](#).

Choose between two Regions, two Availability Zones, or connections within an Availability Zone. You can choose multiple Region, Availability Zones or intra-Availability Zones, but you can't have a combination of all three. Combinations of Regions, Inter-Availability Zones, or Intra-Availability Zones are not supported.

To choose a Region or Availability Zone

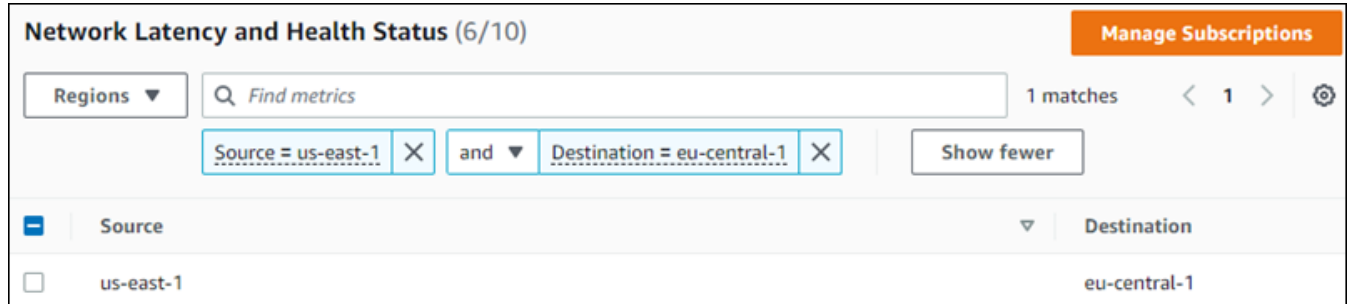
1. Access the Network Manager console at <https://console.aws.amazon.com/networkmanager/home/>.
2. Under **Monitoring and troubleshooting**, choose **Infrastructure Performance**.
3. From the Regions drop-down list, choose whether you want to see **Inter-Region**, **Inter-Availability Zone**, or **Intra-Availability Zone** performance data.

By default the table populates all possible **Source** and **Destination** pairs of Regions or Availability Zones, or all possible pairs within an intra-Availability Zone. This section displays whether that pair is subscribed to Amazon CloudWatch. Sort on **Source**, **Destination**, or **Subscription** to more easily select which Regions or Availability Zones you want to check network performance for.

4. Scroll through the pages and choose one or more pairs.
5. (Optional). To choose specific Regions or Availability Zone, choose **Find metrics**.
 1. In the **Find metrics** field, choose one of the following **Properties**:
 - **Source** — The Region or Inter/Intra Availability Zone that's the source of the network traffic.
 - **Destination** — The Region or Inter/Intra Availability Zone that's the destination for the network source.
 - **Subscription** — Indicates whether the Region or Availability Zone is **Subscribed** or **Unsubscribed** for CloudWatch notifications. For more information about managing subscriptions with CloudWatch, see [CloudWatch subscriptions](#).
 2. For the property operator, choose one of the following operators:
 - **Source =** — Returns network performance results only for the chosen Region, Availability Zone, or an **unsubscribed** subscription status.
 - **Source !=** — Does not equal. Excludes the chosen Region, Availability Zone, or **unsubscribed** subscription status from the network performance results.
 3. From the **Source Values** drop-down list, choose the value to apply to the operator.
 4. Repeat the previous step if you want to include multiple metrics in the results.
 5. By default, each metric that you add is separated by **and** connecting logic, restricting results to the criteria that meet all of the requirements. To return results that match any of your chosen metrics, choose **or** from the drop-down list.

All metrics use the same connecting logic. Combinations of **and/or** logic are not supported. If you change the connecting logic between two metrics, all other metrics will use the same connecting logic.

6. The following example shows a filter for a single Source Region, `us-east-1` **and** a single Destination Region, `eu-central-1`. The results return one combination:



6. Choose a time range, and then choose **Apply**. You can choose a **Relative range** or an **Absolute range**.

Relative range options use a date range based on your current date and time:

- **Last 30 minutes**
- **Last 1 hour**
- **Last 6 hour**
- **Last 1 day**
- **Last 1 week**
- **Custom range** — Set a custom date range. Enter a **Duration**, and then choose one of the following **Units of time**:
 - **seconds**
 - **minutes** (default)
 - **hours**
 - **days**
 - **weeks**
 - **years**

Absolute range allows you to set a specific start date and time and end date and time. The date range cannot exceed 30 days.

Note

Performance metrics are not available prior to Oct 26, 2022 00:00:00 GMT+000.

- On the calendar control, choose the start and end dates. Or,
 - Enter a **Start date** and **End date** using a **YYYY/MM/DD** format.
 - Enter a **Start time** and **End time** using an **hh:mm** format.

If you choose dates using the calendar control, the default start and end times are set to **00:00**. You can change those times by entering your preferred times in the fields.

When using an absolute range, both start and end dates are required.

7. Choose the time period by which you want the performance aggregated. Options are:
 - **5M** — Five minutes (default)
 - **15M** — Fifteen minutes
 - **1H** — One hour
 - **1D** — One day
 - **1W** — One week
8. The **Health status timeline** and **Network latency sections** update.

Health status timeline

The **Health status timeline** shows a single consolidated view of network performance for the chosen Region, Inter-Availability Zone, or Intra-Availability Zone pairs over the starting and end times, consolidated by the aggregation period. The timeline displays the following network latency statuses:


- **Healthy** (green) — Network performance fell within the expected range.
- **Degraded** (yellow) — Network latency fell below the expected range.

Hover over the status bar at a particular date and time to display the health status at that time for the pairs you've chosen.

In the following example, network performance is checked for two pairs of Regions: eu-north-1 and me-south-1, as well as eu-north-1 and ap-south-1. The time period uses a **Relative time** of **Last 1 hour** on November 11, and the aggregation period is **5M** (five minutes). From **11:55 AM** to **12:05 PM** during this hour there was **Degraded** performance. Hovering over the status bar shows that degraded performance occurred only between the eu-north-1 and me-south-1 Regions, while network performance for the other Region pair was healthy.

Note

This example is provided for illustrative purposes only and does not represent the actual network health status between these Regions during this time.

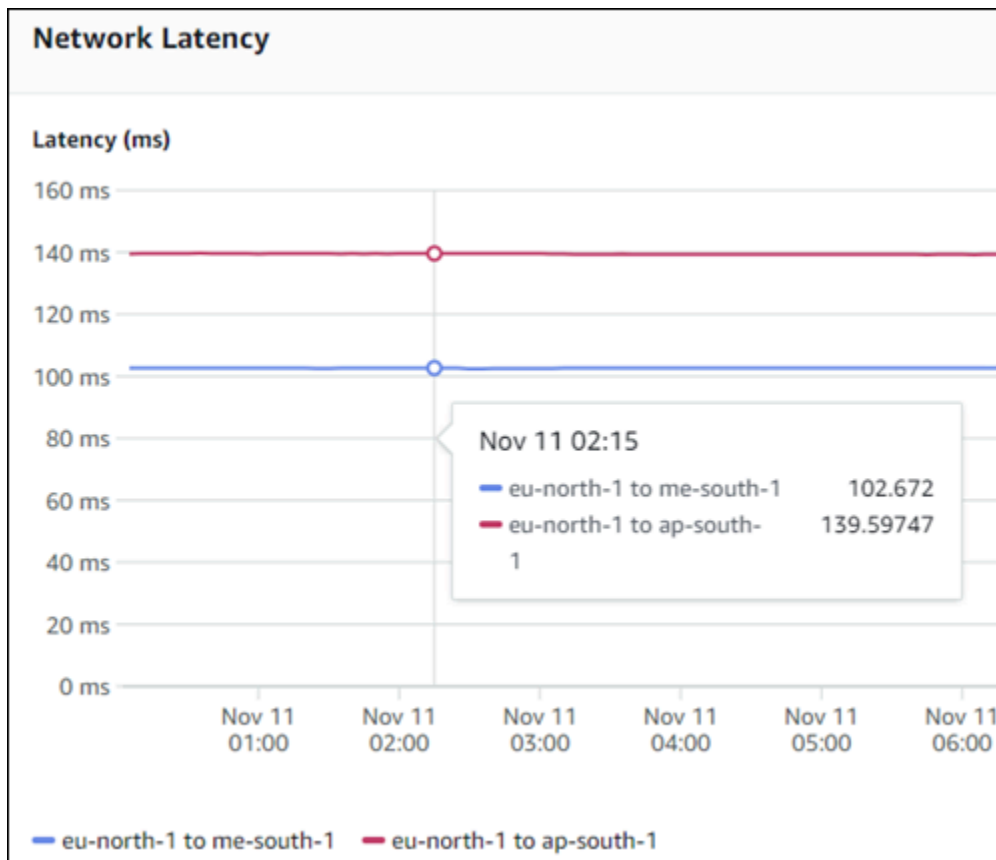


Network latency

The Network latency section shows the actual network performance speed between the chosen Region, Availability Zone, or within an Availability Zone pairs over the starting and end times, consolidated by the aggregation period. Latency can run from a minimum of **0 ms** to a maximum of **160 ms**. Each chosen pair is represented by individual lines in the section, with a legend at the bottom of the page mapping a line to a pair. Hover anywhere over the graph at a particular time to view a comprehensive list of all chosen pairs and the actual latency at that time. Or, if you want to see the latency for a specific pair, choose that line. The following example shows the latency of two Region pairs : eu-north-1 and me-south-1, as well as eu-north-1 and ap-south-1 at **02:15** on **Nov 11**.

Note

This example is provided for illustrative purposes only and does not represent the actual network latency between these Regions during this time.



Monitor Infrastructure Performance using the AWS CLI

[GetAwsNetworkPerformanceData](#) returns network performance metrics. For more information, see [GetAwsNetworkPerformanceData](#).

In the following example, network performance is queried between two Regions, us-east-1 and us-west-2 on 2022-10-26 between a start-time of 12:00:00 and an end time of 12:30:00. The request also uses the following parameters:

- **Metric** indicates what type of metric is being requested. In this example, `aggregate-latency`, indicates that the performance metrics are aggregated and returned for latency.
- **Statistic** is the median value of the metric. In this example, `p50` is the statistic of all the data points gathered within those five minutes.

Note

`p50` is the only supported statistic.

- Period indicates the interval at which performance the metric is returned. In this example, aggregated-latency metrics are returned for every five-minutes.

```
aws ec2 --region ap-southeast-3 get-aws-network-performance-data --start-time
2022-10-26T12:00:00.000Z --end-time 2022-10-26T12:30:00.000Z --data-queries
Id=id1,Source=us-east-1,Destination=us-west-2,Metric=aggregate-
latency,Statistic=p50,Period=five-minutes
```

The results return the following:

```
{
  "DataResponses": [
    {
      "Id": "id1",
      "Source": "us-east-1",
      "Destination": "us-west-2",
      "Metric": "aggregate-latency",
      "Statistic": "p50",
      "Period": "five-minutes",
      "MetricPoints": [
        {
          "StartDate": "2022-10-26T12:00:00+00:00",
          "EndDate": "2022-10-26T12:05:00+00:00",
          "Value": 62.44349,
          "Status": "OK"
        },
        {
          "StartDate": "2022-10-26T12:05:00+00:00",
          "EndDate": "2022-10-26T12:10:00+00:00",
          "Value": 62.483498,
          "Status": "OK"
        },
        {
          "StartDate": "2022-10-26T12:10:00+00:00",
          "EndDate": "2022-10-26T12:15:00+00:00",
          "Value": 62.51248,
          "Status": "OK"
        },
        {
          "StartDate": "2022-10-26T12:15:00+00:00",
          "EndDate": "2022-10-26T12:20:00+00:00",
          "Value": 62.635475,

```

```
        "Status": "OK"
      },
      {
        "StartDate": "2022-10-26T12:20:00+00:00",
        "EndDate": "2022-10-26T12:25:00+00:00",
        "Value": 62.733974,
        "Status": "OK"
      },
      {
        "StartDate": "2022-10-26T12:25:00+00:00",
        "EndDate": "2022-10-26T12:30:00+00:00",
        "Value": 62.773975,
        "Status": "OK"
      },
      {
        "StartDate": "2022-10-26T12:30:00+00:00",
        "EndDate": "2022-10-26T12:35:00+00:00",
        "Value": 62.75349,
        "Status": "OK"
      }
    ]
  }
}
```

Amazon CloudWatch subscriptions for Infrastructure Performance

Infrastructure Performance subscriptions publish network performance metrics in five-minute periods to Amazon CloudWatch for any pair of inter Regions, inter-Availability Zones, or intra-Availability Zones that you optionally subscribe to. Once a subscription is enabled, performance metrics continue to publish to CloudWatch for those pairs until you unsubscribe any pair that you no longer want to publish performance metrics for.

Note

There's a separate CloudWatch metrics subscription charge for each inter-Region, inter-Availability Zone, or intra-Availability Zone pair you subscribe to. You won't be charged for any CloudWatch performance metrics that you unsubscribe from, or for any pair that you've not enabled subscriptions for. For more information about pricing guidelines, see [Amazon CloudWatch pricing](#).

Subscription metrics for CloudWatch

When subscriptions are enabled, you can use CloudWatch to view metrics. For more information on using CloudWatch for your Infrastructure Performance subscriptions, see [What is CloudWatch Management](#) in the *Amazon CloudWatch User Guide*.

The following EC2 namespace metric for Infrastructure Performance is tracked in CloudWatch:

| Metric | Description |
|--------------------------------|---|
| AggregateAWSNetworkPerformance | The latency between Regions, inter-Availability Zones, or intra-Availability Zones. |

- [the section called “Manage subscriptions using the console”](#)
- [the section called “Manage subscriptions using the CLI”](#)

Manage CloudWatch subscriptions using the AWS Network Manager console

Subscriptions allow you to send and view Infrastructure Performance metrics in Amazon CloudWatch. Using the console you can either subscribe to, or unsubscribe from, CloudWatch performance metrics.

To manage CloudWatch subscriptions

1. Open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. In the navigation pane, choose **Network Manager**.
3. Under **Monitoring and troubleshooting**, choose **Infrastructure Performance**.
4. Choose the **Regions**, **Inter-Availability Zone**, or **Intra-Availability Zone** that you want to manage subscriptions for. For detailed steps on how to choose them, see [the section called “Monitor network performance using the console”](#).
5. Choose **Manage Subscriptions**.
6. From the **Update Subscriptions** drop-down list, choose whether you want to **Subscribe** or **Unsubscribe** the pair.
7. Choose **Confirm**.

Manage CloudWatch subscriptions using the AWS CLI

Use the AWS CLI or APIs to get Infrastructure Performance metrics in AWS Network Manager and to enable or disable Amazon CloudWatch subscriptions.

CloudWatch subscription APIs

- [DescribeAwsNetworkPerformanceMetricSubscriptions](#)
- [DisableAwsNetworkPerformanceMetricSubscription](#)
- [EnableAwsNetworkPerformanceMetricSubscription](#)

DescribeAwsNetworkPerformanceMetricSubscriptions

This provides details about any Infrastructure Performance metric subscriptions for Amazon CloudWatch. For more information, see [DescribeAwsNetworkPerformanceMetricSubscriptions](#).

In this example, a filter is used where:

- The filter Name is Source.
- The Value of the filter Name is us-east-1.

In this example, CloudWatch subscriptions are being described for the Region us-east-1.

```
aws ec2 --region us-east-1 describe-aws-network-performance-metric-subscriptions --
filters Name=Source,Values=us-east-1
```

The results describe the us-east-1 subscriptions, including the following:

- Source is the originating location of the subscription. In this example, the Source is the Region us-east-1.
- Destination is the target location of the subscription. In this example, the Destination is the Region us-east-2.
- Metric indicates what type of metric is being requested. In this example, aggregate-latency, indicates that the performance metrics are aggregated and returned for latency.
- Statistic is the median value of the metric. In this example, p50 is the statistic of all the data points gathered within those five minutes.

 **Note**

p50 is the only supported statistic.

- Period indicates the interval at which performance the metric is returned. In this example, aggregated-latency metrics are returned for every five-minutes.

```
{
  "Subscriptions": [
    {
      "Source": "us-east-1",
      "Destination": "us-east-2",
      "Metric": "aggregate-latency",
      "Statistic": "p50",
      "Period": "five-minutes"
    }
  ]
}
```

```
]
}
```

DisableAwsNetworkPerformanceMetricSubscription

This disables Amazon CloudWatch subscriptions for Infrastructure Performance. For more information, see [DisableAwsNetworkPerformanceMetricSubscription](#).

The following example disables subscriptions between a source Region, `us-east-1`, and a destination Region, `us-east-2`. In addition to the source and destination parameters, the request uses the following parameters:

- `metric` indicates what type of metric is being requested. In this example, `aggregate-latency`, indicates that the performance metrics are aggregated and returned for latency.
- `statistic` is the median value of the metric. In this example, `p50` is the statistic of all the data points gathered within those five minutes.

Note

`p50` is the only supported statistic.

```
aws ec2 --region us-east-1 disable-aws-network-performance-metric-subscription --source
us-east-1 --destination us-east-2 --metric aggregate-latency --statistic p50
```

The results return the following Boolean, `true`, indicating that the Amazon CloudWatch subscription between the two Regions was disabled. CloudWatch will no longer receive these metrics.

```
{
  "Output": true
}
```

EnableAwsNetworkPerformanceMetricSubscription

This enables Amazon CloudWatch subscriptions for Infrastructure Performance. For more information, see [EnableAwsNetworkPerformanceMetricSubscription](#).

In the following example, subscriptions are enabled between a source Region, `us-east-1`, and a destination Region, `us-east-2`. In addition to the source and destination parameters, the request also uses the following parameters:

- `metric` indicates what type of metric is being requested. In this example, `aggregate-latency`, indicates that the performance metrics are aggregated and returned for latency.
- `statistic` is the median value of the metric. In this example, `p50` is the statistic of all the data points gathered within those five minutes.

 **Note**

`p50` is the only supported statistic.

```
aws ec2 --region us-east-1 enable-aws-network-performance-metric-subscription --source
us-east-1 --destination us-east-2 --metric aggregate-latency --statistic p50
```

The results return the following Boolean, `true`, indicating that an Amazon CloudWatch subscription has been enabled between the two Regions. CloudWatch will begin receiving metrics.

```
{
  "Output": true
}
```

Identity and access management for Infrastructure Performance

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 Infrastructure Performance resources. IAM is an AWS service that you can use with no additional charge.

To use Infrastructure Performance, you'll need an AWS account and AWS credentials. To increase the security of your AWS account, we recommend that you use an IAM user to provide access credentials instead of using your AWS account credentials. For more information, see [AWS account root user credentials and IAM identities](#) in the *Amazon Web Services General Reference* and [IAM best practices](#) in the *IAM User Guide*.

The following sections provide details on how an IAM administrator can use IAM to help secure your AWS resources, by controlling who can perform Infrastructure Performance actions.

Contents

- [How Infrastructure Performance works with IAM](#)
- [Allow IAM users or groups to access Infrastructure Performance](#)

How Infrastructure Performance works with IAM

The following sections describe how Infrastructure Performance works with IAM.

Contents

- [Infrastructure Performance identity-based policies](#)
- [Infrastructure Performance IAM roles](#)

Infrastructure Performance identity-based policies

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. Infrastructure Performance supports specific actions and resources. There are no Infrastructure Performance service-specific condition keys that can be used in the `Condition` element of policy statements. To learn about all of the

elements that you use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Actions

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.

Infrastructure Performance shares its API namespace with Amazon EC2. Policy actions in Infrastructure Performance use the following prefix before the action: `ec2:`. For example, to grant someone permission to create a path with the `GetAwsNetworkPerformanceData` API operation, you include the `ec2:GetAwsNetworkPerformanceData` action in their policy. Policy statements must include either an `Action` or `NotAction` element.

To specify multiple actions in a single statement, separate them with commas as shown in the following example.

```
"Action": [  
    "ec2:action1",  
    "ec2:action2"  
]
```

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

```
"Action": "ec2:Describe*"
```

The following actions are supported by Infrastructure Performance:

- `DescribeAwsNetworkPerformanceMetricSubscriptions`
- `DisableAwsNetworkPerformanceMetricSubscription`

- `EnableAwsNetworkPerformanceMetricSubscription`
- `GetAwsNetworkPerformanceData`

Resources

Infrastructure Performance does not support 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": "*" 
```

Condition keys

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. For example, you might want a policy to be applied only after a specific date. To express conditions, use predefined condition keys.

Infrastructure Performance does not provide any service-specific condition keys, but it does support using some global condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

All Amazon EC2 actions support the `aws:RequestedRegion` and `ec2:Region` condition keys. For more information, see [Example: Restricting Access to a Specific Region](#).

The `Condition` element is optional.

Infrastructure Performance IAM roles

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with Infrastructure Performance

You can use temporary credentials to sign in with federation, to assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

Infrastructure Performance supports using temporary credentials.

Service-linked roles

Infrastructure Performance has no service-linked roles.

Service roles

Infrastructure Performance has no service roles.

Allow IAM users or groups to access Infrastructure Performance

Any user that signs in to the AWS Management Console or AWS Command Line Interface (AWS CLI) must have permissions to access specific resources. You provide those permissions by using AWS Identity and Access Management (IAM), through policies.

The following procedure shows you how to attach an IAM policy to a user or group that allows full access to Infrastructure Performance.

Note

We recommend creating a new IAM policy that grants only the permissions necessary to use Infrastructure Performance.

Create an IAM policy

Create a policy that provides users full access to Infrastructure Performance. Then attach the policy to a user or group.

To create and attach an IAM policy using the console

1. Sign in to the IAM console at <https://console.aws.amazon.com/iam/> with administrator credentials.
2. In the navigation pane, choose **Policies**.
3. In the content pane, choose **Create policy**.
4. Choose the **JSON** tab.
5. Paste the following JSON policy document in the text field.

```
{
```



```
"Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DisableAwsNetworkPerformanceMetricSubscription",
        "ec2:DescribeAwsNetworkPerformanceMetricSubscriptions",
        "ec2:EnableAwsNetworkPerformanceMetricSubscription",
        "ec2:GetAwsNetworkPerformanceData"
      ],
      "Resource": "*"
    }
  ]
}
```

When you are finished, choose **Review policy**.

6. On the **Review** page, enter a name for the policy, for example, `InfrastructurePerformancePolicy`. Optionally, enter a description for **Description**.
7. In **Summary**, review the policy to see the permissions that it grants, and then choose **Create policy**.
8. Attach the new policy to your user or group.

For information on attaching a policy to a user, see [Changing permissions for an IAM user](#) in the *IAM User Guide*. For information on attaching a policy to a group, see [Attaching a policy to an IAM Group](#) in the *IAM User Guide*.

Document history for Infrastructure Performance

The following table describes the releases for Infrastructure Performance.

| Change | Description | Date |
|---------------------------------|---|-------------------|
| Initial release | This release introduces Infrastructure Performance. | November 28, 2022 |