



Admin and Other Tasks

Data Infrastructure Insights

NetApp
January 17, 2025

Table of Contents

- Admin and Other Tasks 1
- Data Infrastructure Insights API 1
- Monitoring your Environment 12

Admin and Other Tasks

Data Infrastructure Insights API

The Data Infrastructure Insights API enables NetApp customers and independent software vendors (ISVs) to integrate Data Infrastructure Insights with other applications, such as CMDB's or other ticketing systems.


Your Data Infrastructure Insights [feature set role](#) will determine which APIs you can access. User and Guest roles have fewer privileges than Administrator role. For example, if you have Administrator role in Monitor and Optimize, but User role in Reporting, you can manage all API types except Data Warehouse.

Requirements for API Access

- An API Access Token model is used to grant access.
- API Token management is performed by Data Infrastructure Insights users with the Administrator role.

API Documentation (Swagger)

The latest API information is found by logging in to Data Infrastructure Insights and navigating to **Admin > API Access**. Click the **API Documentation** link.



The API Documentation is Swagger-based, which provides a brief description and usage information for the API, and allows you to try it out in your tenant. Depending on your user role and/or Data Infrastructure Insights edition, the API types available to you may vary.

POST

/assets/annotations Create annotation definition



Parameters

Try it out

No parameters

Request body

application/json



Request body should include required name, type, optional description and enumValues (if enum type). Enums should contain name and label. Example:

```
{
  "name": "StorageLocation",
  "type": "FIXED_ENUM",
  "description": "Storage Location",
  "enumValues": [
    {
      "name": "PT_LISBON",
      "label": "Lisbon (Portugal)"
    },
    {
      "name": "US_WALTHAM",
      "label": "Waltham (USA)"
    }
  ]
}
```

Example Value | Schema

```
{}
```

API Access Tokens

Before using the Data Infrastructure Insights API, you must create one or more **API Access Tokens**. Access tokens are used for specified API types, and can grant read and/or write permissions. You can also set the expiration for each access token. All APIs under the specified types are valid for the access token. Each token is void of a username or password.

To create an Access Token:

- Click **Admin > API Access**
- Click **+API Access Token**
 - Enter Token Name
 - Select API Types
 - Specify the permissions granted for this API access
 - Specify Token Expiration



Your token will only be available for copying to the clipboard and saving during the creation process. Tokens can not be retrieved after they are created, so it is highly recommended to copy the token and save it in a secure location. You will be prompted to click the **Copy API Access Token** button before you can close the token creation screen.

You can disable, enable, and revoke tokens. Tokens that are disabled can be enabled.

Tokens grant general purpose access to APIs from a customer perspective; managing access to APIs in the scope of their own tenant. Customer administrators may grant and revoke these tokens without direct involvement from Data Infrastructure Insights back end personnel.

The application receives an Access Token after a user successfully authenticates and authorizes access, then passes the Access Token as a credential when it calls the target API. The passed token informs the API that the bearer of the token has been authorized to access the API and perform specific actions specified by the scope that was granted during authorization.

The HTTP header where the Access Token is passed is **X-CloudInsights-ApiKey:**.

For example, use the following to retrieve storages assets:

```
curl https://<tenant_host_name>/rest/v1/assets/storages -H 'X-CloudInsights-ApiKey:<API_Access-Token>'
```

Where *<API_Access-Token>* is the token you saved during API access creation.

See the swagger pages for examples specific to the API you wish to use.

API Type

The Data Infrastructure Insights API is category-based, and currently contains the following types:

- ASSETS type contains asset, query, and search APIs.
 - Assets: Enumerate top-level objects and retrieve a specific object or an object hierarchy.
 - Query: Retrieve and manage Data Infrastructure Insights queries.
 - Import: Import annotations or applications and assign them to objects
 - Search: Locate a specific object without knowing the object's unique ID or full name.
- DATA COLLECTION type is used to retrieve and manage data collectors.
- DATA INGESTION type is used to retrieve and manage ingestion data and custom metrics, such as from Telegraf agents
- LOG INGESTION is used to retrieve and manage log data

Additional types and/or APIs may become available over time. You can find the latest API information in the [API Swagger documentation](#).

Note that the API types to which a user has access depend also on the [User Role](#) they have in each Data Infrastructure Insights feature set (Monitoring, Workload Security, Reporting).

Inventory Traversal

This section describes how to traverse a hierarchy of Data Infrastructure Insights objects.

Top Level Objects

Individual objects are identified in requests by unique URL (called "self" in JSON) and require knowledge of

object type and internal id. For some of the top level objects (Hosts, Storages, and so on), REST API provides access to the complete collection.

The general format of an API URL is:

```
https://<tenant>/rest/v1/<type>/<object>
```

For example, to retrieve all storages from a tenant named *mysite.c01.cloudinsights.netapp.com*, the request URL is:

```
https://mysite.c01.cloudinsights.netapp.com/rest/v1/assets/storages
```

Children and Related Objects

Top level objects, such as Storage, can be used to traverse to other children and related objects. For example, to retrieve all disks for a specific storage, concatenate the storage “self” URL with “/disks”, for example:

```
https://<tenant>/rest/v1/assets/storages/4537/disks
```

Expands

Many API commands support the **expand** parameter, which provides additional details about the object or URLs for related objects.

The one common expand parameter is *expands*. The response contains a list of all available specific expands for the object.

For example, when you request the following:

```
https://<tenant>/rest/v1/assets/storages/2782?expand=_expands
```

The API returns all available expands for the object as follows:

```

{
  "id": "1247936",
  "self": "/rest/v1/assets/storages/1247936",
  "name": "amsprdclu01",
  "simpleName": "amsprdclu01",
  "naturalKey": "5DF483F0-1729-11DC-9A79-123478563412",
  "ip": "10.64.0.132",
  "serialNumber": "1-80-000011",
  "model": "FAS3270,FAS6290",
  "vendor": "NetApp",
  "microcodeVersion": "8.1.3 clustered Data ONTAP",
  "capacity": {
    "description": "Storage Capacity",
    "unitType": "MB",
    "total": {
      "value": 8.23185105E8
    }
  },
  "storagePools": {
    "value": 5.43220974E8
  }
},
  "isActive": true,
  "createTime": "2013-05-07T16:52:21-0700",
  "family": "FAS3200,FAS6200",
  "managementUrl": null,
  "virtualizedType": "STANDARD",
  "protocols":
  [
    "NAS",
    "NFS",
    "CIFS",
    "FC",
    "ISCSI"
  ],
  "expands": {
    "performance": {
      "url": "/rest/v1/assets/storages/1247936/performance",
      "name": "Performance Data"
    },
    "storageNodes": {
      "url": "/rest/v1/assets/storages/1247936/storageNodes",
      "name": "Storage Storage Nodes"
    },
    "storagePools": {
      "url": "/rest/v1/assets/storages/1247936/storagePools",
      "name": "Storage Storage Pools"
    },
    "storageResources": {
      "url": "/rest/v1/assets/storages/1247936/storageResources",
      "name": "Storage Storage Resources"
    },
    "internalVolumes": {
      "url": "/rest/v1/assets/storages/1247936/internalVolumes",
      "name": "Storage Internal Volumes"
    },
    "volumes": {
      "url": "/rest/v1/assets/storages/1247936/volumes",
      "name": "Storage Volumes"
    },
    "disks": {
      "url": "/rest/v1/assets/storages/1247936/disks",
      "name": "Disks"
    },
    "datasources": {
      "url": "/rest/v1/assets/storages/1247936/datasources",
      "name": "Storage Datasources"
    },
    "ports": {
      "url": "/rest/v1/assets/storages/1247936/ports",
      "name": "Storage Ports"
    },
    "annotations": {
      "url": "/rest/v1/assets/storages/1247936/annotations",
      "name": "Storage Annotations"
    },
    "qtrees": {
      "url": "/rest/v1/assets/storages/1247936/qtrees",
      "name": "Qtrees"
    }
  },
  ".....":

```

Each expand contains data, a URL, or both. The expand parameter supports multiple and nested attributes, for example:

```
https://<tenant>/rest/v1/assets/storages/2782?expand=performance,storageRe  
sources.storage
```

Expand allows you to bring in a lot of related data in one response. NetApp advises that you do not request too much information at one time; this can cause performance degradation.

To discourage this, requests for top-level collections cannot be expanded. For example, you cannot request expand data for all storage objects at once. Clients are required to retrieve the list of objects and then choose specific objects to expand.

Performance Data

Performance data is gathered across many devices as separate samples. Every hour (the default), Data Infrastructure Insights aggregates and summarizes performance samples.

The API allows access to both the samples and the summarized data. For an object with performance data, a performance summary is available as *expand=performance*. Performance history time series are available through nested *expand=performance.history*.

Examples of Performance Data objects include:

- StoragePerformance
- StoragePoolPerformance
- PortPerformance
- DiskPerformance

A Performance Metric has a description and type and contains a collection of performance summaries. For example, Latency, Traffic, and Rate.

A Performance Summary has a description, unit, sample start time, sample end time, and a collection of summarized values (current, min, max, avg, etc.) calculated from a single performance counter over a time range (1 hour, 24 hours, 3 days, and so on).

<https://tenant.cloudinsights.netapp.com/rest/v1/assets/storages/1/performance?expand=history>

Details

Response body

```
{
  "self": "/rest/v1/assets/storages/1/performance",
  "cacheHitRatio": {
    "read": {
      "description": "Cache Hit Ratio - Read",
      "unitType": "%",
      "start": null,
      "end": null,
      "current": null,
      "min": null,
      "max": null,
      "avg": null,
      "sum": null,
      "isDownsampled": false
    },
    "write": {
      "description": "Cache Hit Ratio - Write",
      "unitType": "%",
      "start": null,
      "end": null,
      "current": null,
      "min": null,
      "max": null,
      "avg": null,
      "sum": null,
      "isDownsampled": false
    }
  }
}
```

Self

Performance Metric

Response body

```
}
},
"history": [
  [
    1578418848140,
    {
      "latency.total": 1.30578,
      "latency.read": 3.64681,
      "ioDensity.read": 9.62065,
      "iops.write": 686.35502,
      "ioDensity.total": 31.36259,
      "capacity.raw": 80024.92772,
      "throughput.read": 7.32371,
      "iops.total": 1488.7974,
      "latency.write": 0.39495,
      "ioDensity.write": 14.45856,
      "iops.read": 456.69703,
      "capacity.storagePools": 56058.1041,
      "throughput.write": 14.59581,
      "throughput.total": 21.91953
    }
  ],
  [
    1578419748198,
    {

```

History

Timestamp

Counter Values

The resulting Performance Data dictionary has the following keys:

- "self" is the object's unique URL

- “history” is the list of pairs of timestamp and map of counters values
- Every other dictionary key (“diskThroughput” and so on) is the name of a performance metric.

Each performance data object type has a unique set of performance metrics. For example, the Virtual Machine performance object supports “diskThroughput” as a performance metric. Each supported performance metric is of a certain “performanceCategory” presented in the metric dictionary. Data Infrastructure Insights supports several performance metric type listed later in this document. Each performance metric dictionary will also have the “description” field that is a human-readable description of this performance metric and a set of performance summary counter entries.

The Performance Summary counter is the summarization of performance counters. It presents typical aggregated values like min, max, and avg for a counter and also the latest observed value, time range for summarized data, unit type for counter and thresholds for data. Only thresholds are optional; the rest of attributes are mandatory.

Performance summaries are available for these types of counters:

- Read – Summary for read operations
- Write – Summary for write operations
- Total – Summary for all operations. It may be higher than the simple sum of read and write; it may include other operations.
- Total Max – Summary for all operations. This is the maximum total value in the specified time range.

Object Performance Metrics

The API can return detailed metrics for objects on your tenant, for example:

- Storage Performance Metrics such as IOPS (Number of input/output requests per second), Latency, or Throughput.
- Switch Performance Metrics, such as Traffic Utilization, BB Credit Zero data, or Port Errors.

See the [API Swagger documentation](#) for information on metrics for each object type.

Performance History Data

History data is presented in performance data as a list of timestamp and counter maps pairs.

History counters are named based on the performance metric object name. For example, the virtual machine performance object supports “diskThroughput” so the history map will contain keys named “diskThroughput.read”, “diskThroughput.write” and “diskThroughput.total”.



Timestamp is in UNIX time format.

The following is an example of a performance data JSON for a disk:

```

"performance": {
  "self": "/rest/v1/assets/disks/4013931/performance",
  "iops": {
    "performanceCategory": "IOPS",
    "description": "Disk IOPS",
    "read": {
      "description": "Disk Read Iops",
      "unitType": "IO/s",
      "start": 1399305599999,
      "end": 1402604368055,
      "current": 1,
      "min": 0,
      "max": 6,
      "avg": 0.5532
    },
    [...]
  },
  "total": {
    "description": "Disk Total Throughput",
    "unitType": "MB/s",
    "start": 1399305599999,
    "end": 1402604368055,
    "current": 0,
    "min": 0,
    "max": 2,
    "avg": 0.1702
  }
},
"history":
[
  [
    1399300412690,
    {
      "utilization.total": 12,
      "iops.total": 26,
      "iops.write": 22,
      "iops.read": 4,
      "throughput.read": 0,
      "utilization.read": 2.12,
      "throughput.total": 5,
      "utilization.write": 10.24,
      "throughput.write": 5
    }
  ]
]

```

Objects with Capacity Attributes

Objects with capacity attributes use basic data types and the `CapacityItem` for representation.

CapacityItem

`CapacityItem` is a single logical unit of capacity. It has “value” and “highThreshold” in units defined by its parent object. It also supports an optional breakdown map that explains how the capacity value is constructed. For example, the total capacity of a 100 TB storagePool would be a `CapacityItem` with a value of 100. The breakdown may show 60 TB allocated for “data” and 40 TB for “snapshots”.

Note

“highThreshold” represents system defined thresholds for the corresponding metrics, which a client can use to generate alerts or visual cues on values that are out of acceptable configured ranges.

The following shows the capacity for StoragePools with multiple capacity counters:

StoragePoolCapacity

Model properties:

```
{
  description: string
  unitType: 'MB' or 'GB' or 'TB' or 'KiB' or 'MiB' or 'TiB'
  total: CapacityItem
  used: CapacityItem
  provisioned: CapacityItem
  reservedCapacity: CapacityItem
  softLimit: Double
  rawToUsableRatio: Double
  isDedupeEnabled: boolean
  dedupeSavings: NumericValueWithUnit
  isCompressionEnabled: boolean
  compressionSavings: NumericValueWithUnit
  isThinProvisioningSupported: boolean
}
```

close

Using Search to Look Up Objects

The search API is a simple entry point to the system. The only input parameter to the API is a free-form string and the resulting JSON contains a categorized list of results. Types are different asset types from the Inventory, such as storages, hosts, dataStores, and so on. Each type would contain a list of objects of the type that match the search criteria.

Data Infrastructure Insights is an extensible (wide open) solution that allows integrations with third party orchestration, business management, change control and ticketing systems as well as custom CMDB integrations.

Cloud Insight's RESTful API is a primary point of integration that allows simple and effective movement of data as well as allows users to gain seamless access to their data.

Disabling or Revoking an API Token

To temporarily disable an API token, on the API token list page, click the "three dots" menu for the API, and select *Disable*. You can Re-enable the token at any time using the same menu and selecting *Enable*.

To permanently remove an API token, from the menu, select "Revoke". You cannot re-enable a revoked token; you must create a new token.


API Access Tokens (252) 

+ API Access Token

Bulk Actions

Filter...



<input type="checkbox"/>	Name ↑	Description	Token	API Type	Permission	Expires On	Status	
<input type="checkbox"/>	10.197.120.70		...RpTMJ4	Data Ingestion	Write Only	11/06/2021	Expired	
	22		...nUBDhe	Data Ingestion	Write Only	06/17/2022	Enabled	Disable
	22TOKEN2010560		...8gXq7K	All Categories	Read Only	06/17/2022	Enabled	Edit Description
	ActiveIQ_POC_token		...scmES6	Data Ingestion	Read/Write	11/12/2021	Expired	Revoke

Rotating Expired API Access Tokens

API access tokens have an expiration date. When an API access token expires, users need to generate a new token (of type *Data Ingestion* with Read/Write permissions) and reconfigure Telegraf to use the newly-generated token instead of the expired token. The steps below detail how to do this.

Kubernetes

Note that these commands are using the default namespace "netapp-monitoring". If you have set your own namespace, substitute that namespace in these and all subsequent commands and files.

Note: If you have the latest NetApp Kubernetes Monitoring Operator installed and using an API access token that is renewable, expiring tokens will automatically be replaced by new/refreshed API access tokens. There is no need to perform the manual steps listed below.

- Edit the the NetApp Kubernetes Monitoring Operator.

```
kubectl -n netapp-monitoring edit agent agent-monitoring-netapp
```

- Modify the `spec.output-sink.api-key` value, replacing the old API token with the new API token.

```
spec:
...
  output-sink:
    - api-key:<NEW_API_TOKEN>
```

RHEL/CentOS and Debian/Ubuntu

- Edit the Telegraf configuration files, and replace all instances of the old API token with the new API token.

```
sudo sed -i.bkup 's/<OLD_API_TOKEN>/<NEW_API_TOKEN>/g'
/etc/telegraf/telegraf.d/*.conf
```

- Restart Telegraf.

```
sudo systemctl restart telegraf
```

Windows

- For each Telegraf configuration file in `C:\Program Files\telegraf\telegraf.d`, replace all instances of the old API token with the new API token.

```
cp <plugin>.conf <plugin>.conf.bkup
(Get-Content <plugin>.conf).Replace('<OLD_API_TOKEN>',
'<NEW_API_TOKEN>') | Set-Content <plugin>.conf
```

- Restart Telegraf.

```
Stop-Service telegraf
Start-Service telegraf
```

Monitoring your Environment

Auditing

To identify changes both expected (for tracking) or unexpected (for troubleshooting), you can view an audit trail of the Data Infrastructure Insights system events and user activities.

Viewing Audited Events

To View the Audit page, click **Admin > Audit** in the menu. The Audit page is displayed, providing the following details for each audit entry:

- **Time** - Date and time of the event or activity
- **User** - The User who initiated the activity
- **Role** - The user's role in Data Infrastructure Insights (guest, user, administrator)
- **IP** - The IP address associated with the event
- **Action** - Type of activity, for example Login, Create, Update
- **Category** - The category of activity
- **Details** - Details of the activity

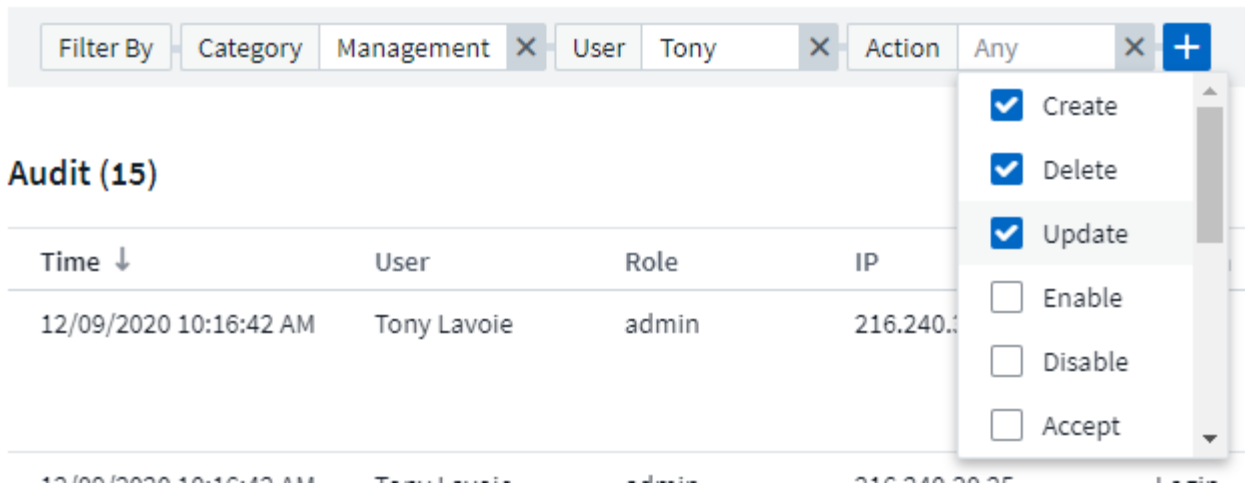
Displaying audit entries

There are a number of different ways to view audit entries:

- You can display audit entries by choosing a particular time period (1 hour, 24 hours, 3 days, etc.).
- You can change the sort order of entries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.

By default, the table displays the entries in descending time order.

- You can use the filter fields to show only the entries you want in the table. Click the [+] button to add additional filters.



More on Filtering

You can use any of the following to refine your filter:

Filter	What it does	Example	Result
* (Asterisk)	enables you to search for everything	vol*rhel	returns all resources that start with "vol" and end with "rhel"
? (question mark)	enables you to search for a specific number of characters	BOS-PRD??-S12	returns BOS-PRD 12 -S12, BOS-PRD 23 -S12, and so on
OR	enables you to specify multiple entities	FAS2240 OR CX600 OR FAS3270	returns any of FAS2440, CX600, or FAS3270
NOT	allows you to exclude text from the search results	NOT EMC*	returns everything that does not start with "EMC"
None	searches for blank/NULL/None in any field where selected	None	returns results where the target field is not empty
Not *	as with <i>None</i> above, but you can also use this form to search for NULL values in <i>text-only</i> fields	Not *	returns results where the target field is not empty.
""	searches for an exact match	"NetApp*"	returns results containing the exact literal string <i>NetApp*</i>

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators OR and NOT will also be treated as literal strings when enclosed in double quotes.

Audited Events and Actions

The events and actions audited by Data Infrastructure Insights can be categorized in the following broad areas:

- **User account:** Log in, log out, role change, etc.

Example: *User **Tony Lavoie** logged in from **10.1.120.15**, user agent **Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/85.0.4183.121 Safari/537.36**, login method(s) **BlueXP Portal Login***

- **Acquisition Unit:** create, delete, etc.

Example: *Acquisition unit **AU-Boston-1** removed.*

- **Data Collector:** add, remove, modify, postpone/resume, change acquisition unit, start/stop, etc.

Example: *Datasource **FlexPod Lab** removed, vendor **NetApp**, model **ONTAP Data Management Software**, ip **192.168.106.5**.*

- **Application:** add, assign to object, remove, etc.

Example: *Internal Volume **ocisedev:t1appSVM01:t1appFlexVol01** added to application **Test App**.*

- **Annotation:** add, assign, remove, annotation rule actions, annotation value changes, etc.

Example: *Annotation value **Boston** added to annotation type **SalesOffice**.*

- **Query:** add, remove, etc.

Example: *Query **TL Sales Query** is added.*

- **Monitor:** add, remove, etc.

Example: *Monitor **Aggr Size - CI Alerts Notifications Dev** updated*


- **Notification:** change email, etc.

Example: *Recipient **ci-alerts-notifications-dl** created*

Exporting Audit Events

You can export the results of your Audit display to a .CSV file, which will allow you to analyze the data or import it into another application.

Steps

1. On the Audit page, set the desired time range and any filters you want. Data Infrastructure Insights will export only the Audit entries that match the filtering and time range you have set.
2. Click the *Export* button  in the upper right of the table.

The displayed Audit events will be exported to a .CSV file, up to a maximum of 10,000 rows.

Retention of Audit Data

The amount of time Data Infrastructure Insights retains Audit data is based on your Edition:

- Basic Edition: Audit data is retained for 30 days
- Standard and Premium Editions: Audit data is retained for 1 year plus 1 day

Audit entries older than the retention time are automatically purged. No user interaction is needed.

Troubleshooting

Here you will find suggestions for troubleshooting problems with Audit.

Problem:	Try this:
I see Audit messages telling me that a monitor has been exported.	Export of a custom monitor configuration is typically used by NetApp engineers during development and testing of new features. If you did not expect to see this message, please consider exploring the actions of the user named in the audited action or contact support.

Active IQ

NetApp [Active IQ](#) provides a series of visualizations, analytics, and other support-related services to NetApp customers for their hardware / software systems. The data reported by Active IQ can enhance troubleshooting of system problems and also provide insight into optimization and predictive analysis related to your devices.



Data Infrastructure Insights collects the **Risks** for any NetApp Clustered Data ONTAP storage system that is monitored and reported by Active IQ. Risks reported for the storage systems are collected automatically by Data Infrastructure Insights as part of its data collection from those devices. You must add the appropriate data collector to Data Infrastructure Insights to collect Active IQ risk information.

Data Infrastructure Insights will not show risk data for ONTAP systems that are not monitored and reported by Active IQ.

The risks reported are shown in Data Infrastructure Insights on the *storage* and *storage node* asset landing pages, in the "Risks" table. The table shows Risk Detail, Category of risk, and Potential Impact of the risk, and also provides a link to the Active IQ page summarizing all risks for the storage node (NetApp Support account sign-in required).

Object ↑	Risk Detail	Category	Potential Impact	Source
tawny01	The following certificates have expired or are expiring within 30 days: Expired: 53CF9553, 53C504D4, 53D671B4, Expiring within 30 days: None	System Configuration	Clients may not be able to connect to the cluster over secure (SSL based) protocols.	Active IQ ↗
tawny01	None of the NIS servers configured for SVM(s) tawny_svm_oci_markc can be contacted.	CIFS Protocol	Potential CIFS and NFS outages may occur.	Active IQ ↗
tawny01	ONTAP version 8.3.2 has entered the Self-Service Support period.	ONTAP	Self-Service Support is the time period where NetApp does not provide support for a version of a software product, but related documentation is still available on the NetApp Support Site.	Active IQ ↗

A count of reported risks is also shown in the landing page's Summary widget, with a link to the appropriate Active IQ page. On a *storage* landing page, the count is a sum of risks from all underlying storage nodes.

Storage Summary		
Model: FAS6210	Microcode Version: 8.3.2 clustered Data ONTAP	Management: HTTPS://10.197.143.25:443
Vendor: NetApp	Raw Capacity: 80,024.3 GB	FC Fabrics Connected: 0
Family: FAS6200	Latency - Total: 0.77 ms	Performance Policies:
Serial Number: 1-80-000013	IOPS - Total: 1,819.19 IO/s	Risks: 108 risks detected by  Active IQ 
IP: 10.197.143.25	Throughput - Total: 41.69 MB/s	

Opening the Active IQ page

When clicking on the link to an Active IQ page, if you are not currently signed in to your Active IQ account, you must perform the following steps to view the Active IQ page for the storage node.

1. In the Data Infrastructure Insights Summary widget or Risks table, click the "Active IQ" link.
2. Sign in to your NetApp Support account. You are taken directly to the storage node page in Active IQ.

Querying for Risks

In Data Infrastructure Insights, you can add the **monitoring.count** column to a storage or storage node query. If the returned result includes Active IQ-Monitored storage systems, the monitoring.count column will display the number of risks for the storage system or node.

Dashboards

You can build widgets (e.g. pie chart, table widget, bar, column, scatter plot, and single value widgets) in order to visualize object risks for storage and storage nodes for NetApp Clustered Data ONTAP systems monitored by Active IQ. "Object Risks" can be selected as a column or metric in these widgets where Storage or Storage Node is the object of focus.

Object Risks by Array



12 items found

Name	Object Risks ↓
tawny	110
rtp-sa-select01	23
rtp-sa-cl02	N/A
rtp-sa-cl07	N/A
oci-phonehome	N/A
durlabprdfs01	N/A
oci-3070-01	N/A

Object Risks by Node

25 items found

Name	Object Risks ↓
tawny02	55
tawny01	55
rtp-sa-select01-1	17
rtp-sa-select01-2	6
tnode02	N/A
rtp-sa-cl06-01	N/A
rtp-sa-cl06-02	N/A



Total Storage Risks

133.00
Object Risks

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.