



Working with Annotations

Data Infrastructure Insights

NetApp

February 19, 2026

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/task_defining_annotations.html on February 19, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Working with Annotations	1
Defining annotations	1
Default annotation types	1
Creating custom annotations	2
Using annotations	4
Defining annotations	4
Manually assigning annotations to assets	4
Assigning annotations using annotation rules	5
Creating annotation rules	6
Modifying annotation rules	7
Changing the Order of Rules	7
Deleting annotation rules	7
Importing Annotations	8
Importing	8
.CSV File Format	8
Import Behavior	9

Working with Annotations

Defining annotations

When customizing Data Infrastructure Insights to track data for your corporate requirements, you can define specialized notes, called annotations, and assign them to your assets.

You can assign annotations to assets with information such as asset end of life, data center, building location, storage tier, or volume service level.

Using annotations to help monitor your environment includes the following high-level tasks:

- Creating or editing definitions for all annotation types.
- Displaying asset pages and associating each asset with one or more annotations.

For example, if an asset is being leased and the lease expires within two months, you might want to apply an end-of-life annotation to the asset. This helps prevent others from using that asset for an extended time.

- Creating rules to automatically apply annotations to multiple assets of the same type.
- Filter assets by their annotations.

Default annotation types

Data Infrastructure Insights provides some default annotation types. These annotations can be used to filter or group data.

You can associate assets with default annotation types such as the following:

- Asset life cycle, such as birthday, sunset, or end of life
- Location information about a device, such as data center, building, or floor
- Classification of assets, such as by quality (tiers), by connected devices (switch level), or by service level
- Status, such as hot (high utilization)

The following table lists the Data Infrastructure Insights-provided annotation types.

Annotation types	Description	Type
Alias	User-friendly name for a resource	Text
Compute Resource Group	Group assignment used by the Host and VM Filesystems data collector	List
Data Center	Physical location	List
Hot	Devices under heavy use on a regular basis or at the threshold of capacity	Boolean
Note	Comments associated with a resource	Text

Service Level	A set of supported service levels that you can assign to resources. Provides an ordered options list for internal volumes, qtree, and volumes. Edit service levels to set performance policies for different levels.	List
Sunset	Threshold set after which no new allocations can be made to that device. Useful for planned migrations and other pending network changes.	Date
Switch Level	Predefined options for setting up categories for switches. Typically, these designations remain for the life of the device, although you can edit them. Available only for switches.	List
Tier	Can be used to define different levels of service within your environment. Tiers can define the type of level, such as speed needed (for example, gold or silver). This feature is available only on internal volumes, qtrees, storage arrays, storage pools, and volumes.	List
Violation Severity	Rank (for example, major) of a violation (for example, missing host ports or missing redundancy), in a hierarchy of highest to lowest importance.	List

 Alias, Data Center, Hot, Service Level, Sunset, Switch Level, Tier, and Violation Severity are system-level annotations, which you cannot delete or rename; you can change only their assigned values.

Creating custom annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While Data Infrastructure Insights provides a set of default annotations, you might find that you want to view data in other ways. The data in custom annotations supplements device data already collected, such as storage manufacturer, number volumes, and performance statistics. The data you add using annotations is not discovered by Data Infrastructure Insights.

Steps

1. In the Data Infrastructure Insights menu, click **Manage > Annotations**.

The Annotations page displays the list of annotations.

2. Click **+Add**

3. Enter a **Name** and **Description** of the annotation.

You can enter up to 255 characters in these fields.

4. Click **Type** and then select one of the following options that represents the type of data allowed in this annotation:

Annotation types

- Boolean: Creates a drop-down list with the choices of yes and no. For example, the "Direct Attached" annotation is Boolean.
- Date: This creates a field that holds a date. For example, if the annotation will be a date, select this.

- List: Creates either of the following:
 - A drop-down fixed list

When others are assigning this annotation type on a device, they cannot add more values to the list.

- A drop-down flexible list

If you select the Add new values on the fly option when you create this list, when others are assigning this annotation type on a device, they can add more values to the list.

- Number: Creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is "Floor", the user could select the Value Type of "number" and enter the floor number.
- Text: Creates a field that allows free-form text. For example, you might enter "Language" as the annotation type, select "Text" as the value type, and enter a language as a value.



After you set the type and save your changes, you cannot change the type of the annotation. If you need to change the type, you have to delete the annotation and create a new one.

1. If you select List as the annotation type, do the following:

- a. Select **Add new values on the fly** if you want the ability to add more values to the annotation when on an asset page, which creates a flexible list.

For example, suppose you are on an asset page and the asset has the City annotation with the values Detroit, Tampa, and Boston. If you selected the **Add new values on the fly** option, you can add additional values to City like San Francisco and Chicago directly on the asset page instead of having to go to the Annotations page to add them. If you do not choose this option, you cannot add new annotation values when applying the annotation; this creates a fixed list.

- b. Enter a value and description in **Value** and **Description** fields.
- c. Click **Add** to add additional values.
- d. Click the Trash icon to delete a value.

2. Click **Save**

Your annotations appear in the list on the Annotations page.

A Note About Boolean Annotations

When filtering on a Boolean annotation, you may be presented with the following values on which to filter:

- **Any**: This will return *all* results, including results set to "Yes", "No", or not set at all.
- **Yes**: Returns only "Yes" results. Note that DII shows "Yes" as a check mark in most tables. Values may be set to "True", "On", etc.; DII treats all of these as "Yes".
- **No**: Returns only "No" results. Note that DII shows "No" as an "X" in most tables. Values may be set to "False", "Off", etc.; DII treats all of these as "No".
- **None**: Returns only results where the annotation has not been set at all. Also referred to as "Null" values.

After you finish

In the UI, the annotation is available immediately for use.

Using annotations

You create annotations and assign them to assets you monitor. Annotations are notes that provide information about an asset, such as physical location, end of life, storage tier, or volume service levels.

Defining annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While Data Infrastructure Insights provides a set of default annotations, such as asset life cycle (birthday or end of life), building or data center location, and tier, you might find that you want to view data in other ways.

The data in custom annotations supplements device data already collected, such as switch manufacturer, number of ports, and performance statistics. The data you add using annotations is not discovered by Data Infrastructure Insights.

Before you begin

- List any industry terminology to which environment data must be associated.
- List corporate terminology to which environment data must be associated.
- Identify any default annotation types that you might be able to use.
- Identify which custom annotations you need to create. You need to create the annotation before it can be assigned to an asset.

Use the following steps to create an annotation.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Enrich > Annotations**
2. Click **+ Annotation** to create a new annotation.
3. Enter a Name, Description, and type for the new annotation.

For example, enter the following to create a text annotation that defines the physical location of an asset in Data Center 4:

- Enter a name for the annotation, such as "Location"
- Enter a description of what the annotation is describing, such as "Physical location is Data Center 4"
- Enter the 'type' of annotation it is, such as "Text".

Manually assigning annotations to assets

Assigning annotations to assets helps you sort, group, and report on assets in ways that are relevant to your business. Although you can assign annotations to assets of a particular type automatically using annotation rules, you can assign annotations to an individual asset by using its asset page.

Before you begin

- You must have created the annotation you want to assign.

Steps

1. Log in to your Data Infrastructure Insights environment.

2. Locate the asset to which you want to apply the annotation.
 - You can locate assets by querying, choosing from a dashboard widget, or search. When you have located the asset you want, click the link to open the asset's landing page.
3. On the asset page, in the User Data section, click **+ Annotation**.
4. The Add Annotation dialog box displays.
5. Select an annotation from the list.
6. Click **Value** and do either of the following, depending on type of annotation you selected:
 - If the annotation type is list, date, or Boolean, select a value from the list.
 - If the annotation type is text, type a value.
7. Click **Save**.

If you want to change the value of the annotation after you assign it, click the annotation field and select a different value.

If the annotation is of list type for which the *Add new values on the fly* option is selected, you can type a new value in addition to selecting an existing value.

Assigning annotations using annotation rules

To automatically assign annotations to assets based on criteria that you define, you configure annotation rules. Data Infrastructure Insights assigns the annotations to assets based on these rules. Data Infrastructure Insights also provides two default annotation rules, which you can modify to suit your needs or remove if you do not want to use them.

Creating annotation rules

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

Before you begin

You must have created a query for the annotation rule.

About this task

Although you can edit the annotation types while you are creating the rules, you should have defined the types ahead of time.

Steps

1. Click **Manage > Annotation rules**

The Annotation Rules page displays the list of existing annotation rules.

2. Click **+ Add**.

3. Do the following:

- a. In the **Name** box, enter a unique name that describes the rule.

This name will appear in the Annotation Rules page.

- b. Click **Query** and select the query that is used to apply the annotation to assets.
 - c. Click **Annotation** and select the annotation you want to apply.

- d. Click **Value** and select a value for the annotation.

For example, if you choose Birthday as the annotation, you specify a date for the value.

- e. Click **Save**

- f. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

Creating annotation rules

You can use annotation rules to automatically apply annotations to multiple assets based on criteria that you define. Data Infrastructure Insights assigns the annotations to assets based on these rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Cloud Insight evaluates the annotation rules.

Before you begin

You must have created a query for the annotation rule.

Steps

1. In the Data Infrastructure Insights menu click **Manage > Annotation rules**.
2. Click **+ Rule** to add a new annotation rule.

The Add Rule dialog is displayed.

3. Do the following:

- a. In the **Name** box, enter a unique name that describes the rule.

The name appears in the Annotation Rules page.

- b. Click **Query** and select the query that Data Infrastructure Insights uses to identify the assets the annotation applies to.
 - c. Click **Annotation** and select the annotation you want to apply.
 - d. Click **Value** and select a value for the annotation.

For example, if you choose Birthday as the annotation, you specify a date for the value.

- e. Click **Save**

- f. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.



In a large Data Infrastructure Insights environment, you may notice that running annotation rules seems to take a while to complete. This is because the indexer runs first and must complete prior to running the rules. The indexer is what gives Data Infrastructure Insights the ability to search or filter for new or updated objects and counters in your data. The rules engine waits until the indexer completes its update before applying the rules.

Modifying annotation rules

You can modify an annotation rule to change the rule's name, its annotation, the annotation's value, or the query associated with the rule.

Steps

1. In the Data Infrastructure Insights menu, Click **Manage > Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

2. Locate the Annotation Rule you want to modify.

You can filter the annotation rules by entering a value in the filter box or click a page number to browse through the annotation rules by page.

3. Click the menu icon for the rule that you want to modify.

4. Click **Edit**

The Edit Rule dialog is displayed.

5. Modify the annotation rule's name, annotation, value, or query.

Changing the Order of Rules

Annotation rules are processed from the top of the rules list to the bottom. To change the order in which a rule is processed, do the following:

Steps

1. Click on the menu icon for the rule you want to move.

2. Click **Move Up** or **Move Down** as needed until the rule appears in the location you want.

Note that when running multiple rules that update the same annotation on an asset, the first rule (as run from the top down) applies the annotation and updates the asset, then the second rule applies but doesn't change any annotation that was already set by the previous rule.

Deleting annotation rules

You might want to delete annotation rules that are no longer used.

Steps

1. In the Data Infrastructure Insights menu, Click **Manage > Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

2. Locate the Annotation Rule you want to delete.

You can filter the annotation rules by entering a value in the filter box or click a page number to browse through the annotation rules by page.

3. Click the menu icon for the rule that you want to delete.

4. Click **Delete**

A confirmation message is displayed, prompting whether you want to delete the rule.

5. Click **OK**

Importing Annotations

Data Infrastructure Insights includes an API for importing annotations or applications from a CSV file, and assigning them to objects you specify.



The Data Infrastructure Insights API is available in **Data Infrastructure Insights Premium Edition**.

Importing

The **Admin > API Access** links contain [documentation](#) for the **Assets/Import** API. This documentation contains information on the .CSV file format.

ASSETS.import

PUT /assets/import Import assets from a CSV file.

Import annotations and applications from the given CSV file. The format of the CSV file is following:

```
Project , <Annotation Type> [, <Annotation Type> ...] [, Application] [, Tenant] [, Line_Of_Business] [, Business_Unit] [, <Project>]
<Object Type Value 1>, <Object Name or Key 1>, <Annotation Value> [, <Annotation Value> ...] [, <Application>] [, <Tenant>] [, <Line_Of_Business>] [, <Business_Unit>] [, <Project>]
<Object Type Value 2>, <Object Name or Key 2>, <Annotation Value> [, <Annotation Value> ...] [, <Application>] [, <Tenant>] [, <Line_Of_Business>] [, <Business_Unit>] [, <Project>]
<Object Type Value 3>, <Object Name or Key 3>, <Annotation Value> [, <Annotation Value> ...] [, <Application>] [, <Tenant>] [, <Line_Of_Business>] [, <Business_Unit>] [, <Project>]
...
<Object Type Value N>, <Object Name or Key N>, <Annotation Value> [, <Annotation Value> ...] [, <Application>] [, <Tenant>] [, <Line_Of_Business>] [, <Business_Unit>] [, <Project>]
```

.CSV File Format

The general format of the CSV file is as follows. The first line of the file defines the import fields and specifies the order of the fields. This is followed by separate lines for each annotation or application. You do not need to define every field. However, the subsequent annotation lines must follow the same order as the definition line.

```
[Object Type] , [Object Name or ID] , Annotation Type [, Annotation Type, ...
...] [, Application] [, Tenant] [, Line_Of_Business] [, Business_Unit] [, Project]
```

See the API Documentation for examples of .CSV files.

You can import and assign annotations from a .CSV file from within the API swagger itself. Simply choose the file to use and click the *Execute* button:

The screenshot shows a user interface for importing data from a CSV file. At the top, there is a 'Parameters' section with a 'Cancel' button. Below it is a 'Request body' section with a dropdown menu set to 'multipart/form-data'. The main area contains a 'CSV file to import' field and a 'data' field of type 'string(\$binary)' with a 'Choose File' button and a 'No file chosen' message. At the bottom, there are 'Execute' and 'Clear' buttons.

Import Behavior

During the import operation, data is added, merged, or replaced, depending on the objects and object types that are being imported. While importing, keep in mind the following behaviors.

- Adds an annotation or application if none exists with the same name in the target system.
- Merges an annotation if the annotation type is a list, and an annotation with the same name exists in the target system.
- Replaces an annotation if the annotation type is anything other than a list, and an annotation with the same name exists in the target system.

Note: If an annotation with the same name but with a different type exists in the target system, the import fails. If objects depend on the failed annotation, those objects may show incorrect or unwanted information. You must check all annotation dependencies after the import operation is complete.

- If an annotation value is empty then that annotation is removed from the object. Inherited annotations are not affected.
- Date type annotation values must be passed in as unix time in milliseconds.
- When annotating volumes or internal volumes, the object name is a combination of storage name and volume name using the "->" separator. For example: <Storage Name>-><Volume Name>
- If an object name contains a comma, the whole name must be in double quotes. For example: "NetApp1,NetApp2"->023F
- When attaching annotating to storages, switches, and ports, the 'Application' column will be ignored.
- Tenant, Line_Of_Business, Business_Unit, and/or Project makes a business entity. As with all business entities, any of the values can be empty.

The following object types can be annotated.

OBJECT TYPE	NAME OR KEY
Host	id-><id> or <Name> or <IP>
VM	id-><id> or <Name>
StoragePool	id-><id> or <Storage Name>-><Storage Pool Name>
InternalVolume	id-><id> or <Storage Name>-><Internal Volume Name>

Volume	id-><id> or <Storage Name>-><Volume Name>
Storage	id-><id> or <Name> or <IP>
Switch	id-><id> or <Name> or <IP>
Port	id-><id> or <WWN>
Qtree	id-><id> or <Storage Name>-><Internal Volume Name>-><Qtree Name>
Share	id-><id> or <Storage Name>-><Internal Volume Name>-><Share Name>-><Protocol>[-><Qtree Name (optional in case of default Qtree)>]

Copyright information

Copyright © 2026 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—with 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.