



Tanium™ API Gateway User Guide

Version 1.2.53

April 08, 2022

The information in this document is subject to change without notice. Further, the information provided in this document is provided “as is” and is believed to be accurate, but is presented without any warranty of any kind, express or implied, except as provided in Tanium’s customer sales terms and conditions. Unless so otherwise provided, Tanium assumes no liability whatsoever, and in no event shall Tanium or its suppliers be liable for any indirect, special, consequential, or incidental damages, including without limitation, lost profits or loss or damage to data arising out of the use or inability to use this document, even if Tanium Inc. has been advised of the possibility of such damages.

Any IP addresses used in this document are not intended to be actual addresses. Any examples, command display output, network topology diagrams, and other figures included in this document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Please visit <https://docs.tanium.com> for the most current Tanium product documentation.

This documentation may provide access to or information about content, products (including hardware and software), and services provided by third parties (“Third Party Items”). With respect to such Third Party Items, Tanium Inc. and its affiliates (i) are not responsible for such items, and expressly disclaim all warranties and liability of any kind related to such Third Party Items and (ii) will not be responsible for any loss, costs, or damages incurred due to your access to or use of such Third Party Items unless expressly set forth otherwise in an applicable agreement between you and Tanium.

Further, this documentation does not require or contemplate the use of or combination with Tanium products with any particular Third Party Items and neither Tanium nor its affiliates shall have any responsibility for any infringement of intellectual property rights caused by any such combination. You, and not Tanium, are responsible for determining that any combination of Third Party Items with Tanium products is appropriate and will not cause infringement of any third party intellectual property rights.

Tanium is committed to the highest accessibility standards for our products. To date, Tanium has focused on compliance with U.S. Federal regulations - specifically Section 508 of the Rehabilitation Act of 1998. Tanium has conducted 3rd party accessibility assessments over the course of product development for many years and has most recently completed certification against the WCAG 2.1 / VPAT 2.3 standards for all major product modules in summer 2021. In the recent testing the Tanium Console UI achieved supports or partially supports for all applicable WCAG 2.1 criteria. Tanium can make available any VPAT reports on a module-by-module basis as part of a larger solution planning process for any customer or prospect.

As new products and features are continuously delivered, Tanium will conduct testing to identify potential gaps in compliance with accessibility guidelines. Tanium is committed to making best efforts to address any gaps quickly, as is feasible, given the severity of the issue and scope of the changes. These objectives are factored into the ongoing delivery schedule of features and releases with our existing resources.

Tanium welcomes customer input on making solutions accessible based on your Tanium modules and assistive technology requirements. Accessibility requirements are important to the Tanium customer community and we are committed to prioritizing these compliance efforts as part of our overall product roadmap. Tanium maintains transparency on our progress and milestones and welcomes any further questions or discussion around this work. Contact your sales representative, email Tanium Support at support@tanium.com, or email accessibility@tanium.com to make further inquiries.

Tanium is a trademark of Tanium, Inc. in the U.S. and other countries. Third-party trademarks mentioned are the property of their respective owners.

© 2022 Tanium Inc. All rights reserved.

Table of contents

- API Gateway overview** 7
 - API Gateway topology 7
 - Query explorer 8
 - Query variables 10
 - Schema reference 10
 - Authentication 11
 - Rate limits 11
 - Root endpoint 11
 - Example cURL syntax 11
 - Pagination 12
 - Cursors 12
 - Connection and edges 12
 - Arguments 13
 - Filters 13
 - Simple filters 14
 - Compound filters 15
 - Negated filters 15
 - Field filters 15
 - Integration with other Tanium products 17
- Getting started with API Gateway** **18**
 - Step 1: Review the requirements 18
 - Step 2: Install API Gateway 18
 - Step 3: Install any integrated solutions that use the API Gateway 18
 - Step 4: Grant API Gateway permissions 18
 - Step 5: Test queries through the Tanium™ Console 18
 - Step 6: (Optional) Test queries through cURL 18
 - Step 7: Explore sample queries and mutations 18

API Gateway requirements	19
Core platform dependencies	19
Solution dependencies	19
Tanium recommended installation	19
Import specific solutions	19
Required dependencies	19
Feature-specific dependencies	20
Tanium™ Module Server	20
Endpoints	20
Host and network security requirements	20
Ports	20
Security exclusions	21
User role requirements	21
Installing API Gateway	23
Before you begin	23
Import API Gateway	23
Manage solution dependencies	24
Upgrade API Gateway	24
Verify API Gateway version	24
Troubleshoot issues	24
Using API Gateway	25
Test a query in the Tanium Console	25
Troubleshooting API Gateway	27
Collect logs	27
Update platform setting for All-in-One deployment	27
Queries return unexpected results or errors	27
Request error handling	28
Syntax Validation	28
Structural Validation	28
Error Extensions	29

Error codes	35
Uninstall API Gateway	36
Contact Tanium Support	37
Reference: Filter syntax	38
Endpoint query filter syntax	38
General Fields	38
Computer group membership	38
Sensors	39
Parameterized sensors	39
Multi-column sensors	40
Sensor readings	42
Reference: API Gateway examples	43
General examples	43
Get server time	43
Get endpoints	43
Get endpoints IDs from Tanium Data Service	45
Get endpoints using aliases	46
Get endpoints using multiple sensors	49
Get rich endpoint data	52
Get a set of endpoints	55
Unregistered sensor query	57
Unregistered parameterized sensor query	59
Endpoint cursor query	63
Paginated query	63
Software characteristics query with filter	66
Action examples	68
Create action (subset of endpoints)	68
Get action details	69
Deploy examples	71
Deploy a package to all endpoints	71

Get package details	72
Get Deploy packages	77
Get software deployment status	79
Direct Connect examples	81
Open a connection to an endpoint	81
Ping the connection to an endpoint	82
Get data from an endpoint	82
Get process from an endpoint	83
Get alerts from an endpoint	85
Stop a process on an endpoint	87
Close connection to an endpoint	87

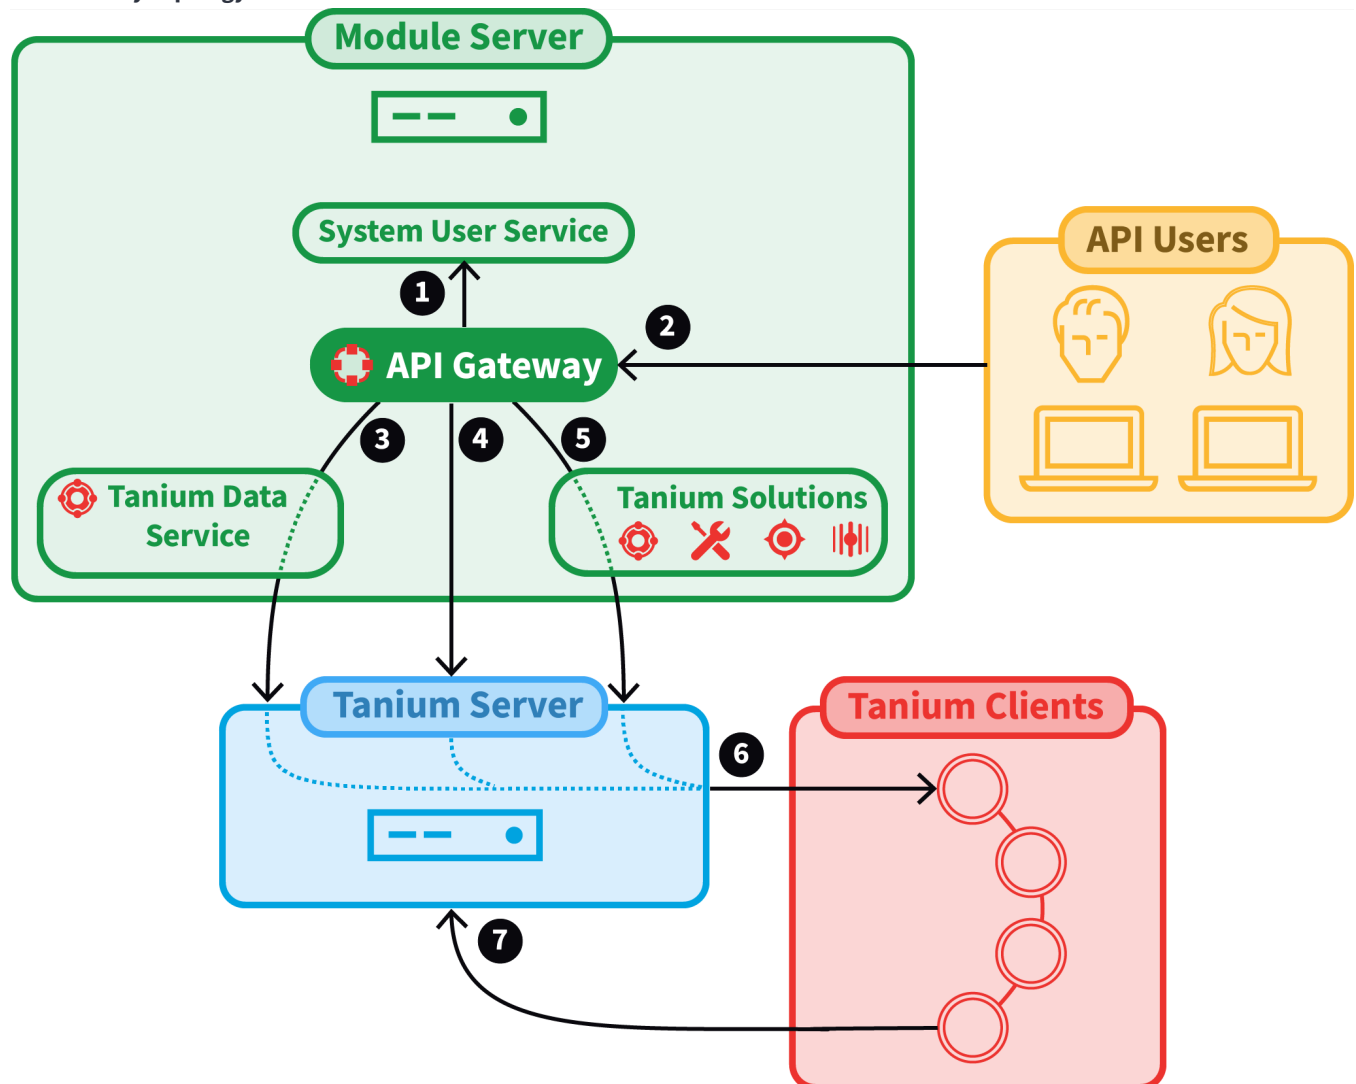
API Gateway overview

Tanium™ API Gateway provides a single and stable API integration point for various Tanium solutions. It is designed for Tanium partners and customers interested in building integrated solutions with the Tanium™ Core Platform.

API Gateway topology

The API Gateway service runs on the Tanium™ Module Server, and requires the System User service to run background processes. The API Gateway service passes user requests to Tanium™ Data Service or the Tanium™ Server; the Tanium Server collects data from Tanium™ Clients or updates data based on the request, and returns the results to the API Gateway service. The API Gateway service passes the results to the user that submitted the request.

API Gateway topology



With the API Gateway, information passes between users and the Tanium Server through the following workflow (matching the numbers in [API Gateway topology on page 7](#)):

1. The API Gateway requires the System User service to be running to perform background functions.
2. Users send requests to the API Gateway to be performed on endpoints, such as action deployment and queries for question results.
3. By default, the API Gateway relays user queries to the Tanium Data Service, which continuously collects results for all registered sensors through the Tanium Server. Continuous collection from endpoints ensures that results are available in the Tanium Data Service cache for both online and offline endpoints.
4. Users can optionally specify queries that go directly from the API Gateway to the Tanium Server for collecting live results from endpoints that are currently online.
5. Users can specify requests that go from the API Gateway to Tanium solutions to be performed on endpoints. For example, a user might request Tanium™ Direct Connect to establish a direct connection to a specific endpoint.
6. The Tanium Server issues questions and deploys actions (and associated content such as packages) to Tanium Clients.
7. Tanium Clients return the question results or action statuses to the Tanium Server. The server responds to the request with the results and statuses. If the user sent the request through the API Gateway query explorer, the results pane displays the response.

Query explorer

API Gateway includes an interactive query explorer that you can use to write and run queries and mutations in the Tanium Console. Use the query explorer to try new queries and discover what data is available.

You can find the query explorer on the API Gateway **Overview** page:

Query explorer

Documentation Explorer pane (expandable)

Query pane

Results pane

Query variables pane (expandable)

GRAPHQL

```

1 # Welcome to GraphQL
2 #
3 # GraphQL is an in-browser tool for writing, validating, and
4 # testing GraphQL queries.
5 #
6 # Type queries into this side of the screen, and you will see intelligent
7 # typeheads aware of the current GraphQL type schema and live syntax and
8 # validation errors highlighted within the text.
9 #
10 # GraphQL queries typically start with a "{" character. Lines that start
11 # with a # are ignored.
12 #
13 # An example GraphQL query might look like:
14 #
15 #   {
16 #     field(arg: "value") {
17 #       subField
18 #     }
19 #   }
20 #
21 # Keyboard shortcuts:
22 #
23 # Prettify Query: Shift-Ctrl-P (or press the prettify button above)
24 #
25 # Merge Query: Shift-Ctrl-M (or press the merge button above)
26 #
27 # Run Query: Ctrl-Enter (or press the play button above)
28 #
29 # Auto Complete: Ctrl-Space (or just start typing)
30 #
31 #
32 #

```


QUERY VARIABLES



If the query explorer does not appear on the API Gateway **Overview** page, click **Customize Page** and make sure the **Query Explorer** option is selected.

Query variables


If a query or mutation uses variables, expand the **QUERY VARIABLES** pane and include the variables in the pane that expands.



Schema reference

API Gateway contains a schema reference that documents all queries, mutations, and objects that are available in API Gateway. The schema reference is generated directly from the schema; refer to the schema reference in API Gateway for the most up-to-date documentation.

To view the schema reference in the query explorer, click **Docs** to expand the **Documentation Explorer** pane of the query explorer.





The query explorer uses the GraphQL interactive browser to send GraphQL queries and mutations to the Tanium Server. For more information about the options that are available in GraphQL, see <https://graphql.org/learn/>.

Authentication

Requests that are sent from the query explorer in the Tanium Console are authenticated and authorized with the session ID of the user who is signed in. The Tanium Server uses the role-based access control (RBAC) permissions of the user account to determine which content you can query and mutate.

Requests that are sent from outside the Tanium Console are authenticated and authorized with either session IDs or API tokens. You must include an API token or session ID in the authorization header of all requests that are sent to API Gateway. API Gateway uses the RBAC permissions of the requesting user to determine content access for all queries and mutations. For an example cURL query that shows the authorization header, see [Example cURL syntax on page 11](#).



BEST PRACTICE

Use API tokens to send requests through API Gateway instead of session IDs. While session IDs time out after five minutes of inactivity, you can set a longer timeout for API tokens. You can create API tokens in the Tanium Console or through the Tanium Core Platform REST API. For more information, see [Tanium Console User Guide: Managing API tokens](#).



IMPORTANT

When requests are sent outside the Tanium Console, make sure to use the correct URL to send requests. See [Root endpoint on page 11](#) and [Example cURL syntax on page 11](#) for examples.

Rate limits

API Gateway has no specific rate limits.

Root endpoint

To send queries and mutations outside the Tanium Console, use the following address:

```
https://<server>/plugin/products/gateway/graphql
```

Example cURL syntax

```
curl --request POST \  
  --url https://localhost/plugin/products/gateway/graphql \  
  --header 'Content-Type: application/json' \  
  --header 'session: token-356d5f5bbb3671f28e24f65be3bdd54d9d81001ca823efaabc5fbff251' \  
  --data '{"query": "{\n  now\n}"}'
```

Pagination

Queries that return many results are paginated to reduce resource utilization. API Gateway uses the standard [Relay GraphQL pagination specification](#) to provide users the option to explicitly control pagination.

Paginated queries return a connection type that is prefixed with the name of the data type, such as `EndpointConnection`. Queries accept standard arguments to control the pagination.

Example of a paginated query:

```
1  {
2    endpoints {
3      edges {
4        node {
5          id
6          serialNumber
7        }
8      }
9      pageInfo {
10     hasNextPage
11     endCursor
12   }
13 }
14 }
```

Cursors

Use cursors to control relay pagination. Cursors are opaque strings that point to records within queried collections, and can be used to request the records after the cursor. All collections support forward traversal, and some also support backward traversal.

Cursors are valid only for the query for which they were returned. Cursors are generally valid for five minutes after their most recent use. Any queries that deviate from this policy are documented in the query field.

Connection results are stable and consistent when traversed with cursors unless documented in the query field.

Connection and edges

The connection includes an `edge` field that returns a list of typed edges, such as `EndpointEdge`. Each edge contains at least two fields:

- A `node` field with the actual data type, such as `Endpoint`
- A `cursor` field with a cursor for the record

The connection type also includes a `pageInfo` field that contains at least two fields:

- `hasNextPage` indicates if there are more records
- `endCursor` is the cursor of the last record in the returned page, if any

Some connection types feature other metadata, such as `totalRecords`.

Arguments

Paginated queries support at least two arguments: `first` indicates the number of records to return, and `after` is the value of the record cursor that precedes the records in the requested page. When fully paginated, this value is the same as the `endCursor` value from the previous page. Both arguments have sensible defaults.

Paginated queries that support backward traversal allow two corresponding arguments: `last` and `before`.



A single query supports either forward or backward traversal, but not both. The server returns an error response for queries with arguments for both forward and backward traversals.

When a paginated request extends beyond the collection, the query returns only the available results.

Example of a request for a page of data within a collection:

```

1  {
2    endpoints(after: "the-cursor-value", first: 10) {
3      edges {
4        node {
5          id
6          serialNumber
7        }
8      }
9      pageInfo {
10     hasNextPage
11     endCursor
12   }
13 }
14 }
```

Filters

Most queries that return multiple results provide support to filter the results. Such queries provide a `filter` argument.

Simple filters

Simple filters are single filters that constrain the values of fields that participate in the query. You can specify simple filters in the `path` property with a period to separate levels in the graph starting at the record type. For example:

```
1 | {
2 |   endpoints(filter: {path: "primaryUser.email", value: "user@example.com"}) {
3 |     edges {
4 |       node {
5 |         id
6 |         primaryUser {
7 |           email
8 |         }
9 |       }
10 |     }
11 |   }
12 | }
```

The query does not need to return the filtered path. Not all field paths are filterable. Refer to the schema to see which paths cannot be filtered.

Simple filters must also contain a string `value` property.

You can specify an operator in the `op` property, which is an enumerated type and defaults to the equality operator. For example:

```
1 | {
2 |   endpoints(filter: {path: "processor.logicalProcessors", value: "4", op: GTE}) {
3 |     edges {
4 |       node {
5 |         id
6 |       }
7 |     }
8 |   }
9 | }
```



Not all operators are valid for all fields.

NOTE

Compound filters

Compound filters contain multiple simple or compound filters that appear in the `filters` property. By default, all child filters must pass for a record to be included. If the `or` argument is given with a `true` value, then a record is included if any child filter matches.

Example of a simple compound filter:

```
1 {
2   endpoints(filter: {filters: [{path: "serialNumber" value: "x"}, {path: "name", value:
3     "y"}]}) {
4     edges {
5       node {
6         id
7       }
8     }
9   }
```

Negated filters

You can negate both simple and compound filters with a `negated` property of `true`. For example, the following query returns endpoints whose serial number does not contain the letter `x`:

```
1 {
2   endpoints(filter: {path: "serialNumber", value: "x", negated: true}) {
3     edges {
4       node {
5         id
6       }
7     }
8   }
9 }
```

Field filters

Filters apply to the entire record. Some records contain fields that are collections; you can also filter these fields. When you filter a field, the filter applies to both the child collection and to the records. For example, if you search for endpoints with an installed application named `Tanium Client` with a filter on the field, API Gateway returns only those endpoints with such an application, as well as only the matching application:

```

1  {
2    endpoints {
3      edges {
4        node {
5          installedApplications(filter: {path: "name", value: "Tanium Client"}) {
6            name
7            version
8          }
9        }
10     }
11  }
12 }

```

Field filters can be simple or compound. Compound filters are limited to one level of children that must use the equality operator, and require all child filters. For example:

```

1  {
2    endpoints {
3      edges {
4        node {
5          installedApplications(
6            filter: {
7              filters: [
8                {path: "name", value: "Tanium Client"},
9                {path: "version": value: "7.5.0.0"}
10             ]
11          }
12        ) {
13          name
14          version
15        }
16      }
17    }
18  }
19 }

```


For more information on specific filter syntax, see [Reference: Filter syntax on page 38](#).

Integration with other Tanium products

The following solutions are supported by API Gateway:

- Tanium Core Platform
 - Actions
 - Tanium™ Data Service
 - Tanium™ Direct Connect
 - Packages
- Tanium™ Blob
- Tanium™ Deploy
- Tanium™ Performance

Getting started with API Gateway

Step 1: Review the requirements

Review the system, network, security, and user role requirements: see [API Gateway requirements on page 19](#).

Step 2: Install API Gateway

See [Installing API Gateway on page 23](#).

Step 3: Install any integrated solutions that use the API Gateway

Import any integrated solutions that you want to use. For information on which Tanium solutions use the API Gateway, see [Integration with other Tanium products on page 17](#).

Step 4: Grant API Gateway permissions

Grant permissions to users to use API Gateway. Users with the **Administrator** reserved role have access by default. See [User role requirements on page 21](#).

Step 5: Test queries through the Tanium™ Console

Use the interactive query explorer to test queries in the Tanium Console. See [Test a query in the Tanium Console on page 25](#).

Step 6: (Optional) Test queries through cURL

Test queries through cURL. See [Example cURL syntax on page 11](#).

Step 7: Explore sample queries and mutations

Explore sample queries and mutations to see what you can do with API Gateway. See [Reference: API Gateway examples on page 43](#).

API Gateway requirements

Review the requirements before you install and use API Gateway.

Core platform dependencies

Make sure that your environment meets the following requirements:

- **Tanium™ Core Platform servers:** 7.4.4 or later
- **Tanium™ Client:** No client requirements.
- **Tanium™ Console:** 2.0 or later
- **Tanium content:** API Gateway uses sensors that are included in the Core Content and Core AD Query content packs.

Solution dependencies

Other Tanium solutions are required for API Gateway to function (required dependencies) or for specific API Gateway features to work (feature-specific dependencies). The installation method that you select determines if the Tanium Server automatically imports dependencies or if you must manually import them.



NOTE

Some API Gateway dependencies have their own dependencies, which you can see by clicking the links in the lists of [Required dependencies on page 19](#) and [Feature-specific dependencies on page 20](#). Note that the links open the user guides for the latest version of each solution, not necessarily the minimum version that API Gateway requires.

Tanium recommended installation

If you select **Tanium Recommended Installation** when you import API Gateway, the Tanium Server automatically imports all your licensed solutions at the same time. See [Tanium Console User Guide: Import all modules and services](#).

Import specific solutions

If you select only API Gateway to import and are using Tanium Core Platform 7.5.2.3531 with Tanium Console 3.0.72 or later, the Tanium Server automatically imports the latest available versions of any required dependencies that are missing. If some required dependencies are already imported but their versions are earlier than the minimum required for API Gateway, the server automatically updates those dependencies to the latest available versions.

If you select only API Gateway to import and you are using Tanium Core Platform 7.5.2.3503 or earlier with Tanium Console 3.0.64 or earlier, you must manually import or update required dependencies. See [Tanium Console User Guide: Import, re-import, or update specific solutions](#).

Required dependencies

API Gateway has the following required dependencies at the specified minimum versions:

- Tanium [Interact](#) 2.9.83 or later
- Tanium System User 1.0.40 or later

Feature-specific dependencies

If you select only API Gateway to import, you must manually import or update its feature-specific dependencies regardless of the Tanium Console or Tanium Core Platform versions. API Gateway has the following feature-specific dependencies at the specified minimum versions:

- Tanium Blob 1.0.6 or later
- Tanium [Direct Connect](#) 1.10.39 or later
- Tanium [Deploy](#) 2.9.123 or later
- Tanium [Performance](#) 1.10.57 or later

Tanium™ Module Server

API Gateway is installed and runs as a service on the Module Server host computer. The impact on the Module Server is minimal and depends on usage.

For information about Module Server sizing in a Windows deployment, see [Tanium Core Platform Deployment Guide for Windows: Host system sizing guidelines](#).

Endpoints

API Gateway does not directly deploy packages to endpoints. However, you can use API Gateway to deploy packages through Tanium Deploy. For Tanium Deploy endpoint requirements, see [Tanium Deploy User Guide: Endpoints](#).

For Tanium Client operating system support, see [Tanium Client Management User Guide: Client version and host system requirements](#).

Host and network security requirements

Specific ports and processes are needed to run API Gateway.

Ports

The following ports are required for API Gateway communication.

Source	Destination	Port	Protocol	Purpose
Module Server	Module Server (loopback)	17600	TCP	Internal purposes, not externally accessible



Configure firewall policies to open ports for Tanium traffic with TCP-based rules instead of application identity-based rules. For example, on a Palo Alto Networks firewall, configure the rules with service objects or service groups instead of application objects or application groups.

Security exclusions

If security software is in use in the environment to monitor and block unknown host system processes, Tanium recommends that a security administrator create exclusions to allow the Tanium processes to run without interference. The configuration of these exclusions varies depending on AV software. For a list of all security exclusions to define across Tanium, see [Tanium Core Platform Deployment Reference Guide: Host system security exclusions](#).







API Gateway security exclusions

Target Device	Notes	Exclusion Type	Exclusion
Module Server		Process	<Module Server>\services\gateway-service\taniumgateway.exe

User role requirements

The following tables list the role permissions required to use API Gateway. For more information about role permissions and associated content sets, see [Tanium Console User Guide: Managing RBAC](#).

API Gateway user role permissions

Permission	API Gateway User ¹	Gateway Service Account	Gateway Service Account - All Content Sets
Gateway Api Access API Gateway	 EXECUTE		
Gateway Service Account Provides access for the API Gateway service.		 EXECUTE	

¹ This role provides module permissions for Tanium Interact. You can view which Interact permissions are granted to this role in the Tanium Console. For more information, see [Tanium Interact User Guide: User role requirements](#).

Provided API Gateway administration and platform content permissions

Permission	Permission Type	API Gateway User	Gateway Service Account	Gateway Service Account - All Content Sets
Action Group	Administration	✘	✔ READ WRITE	✘
Computer Group	Administration	✘	✔ READ	✘
Global Settings	Administration	✘	✔ READ	✘
Sensor	Platform Content	✘	✘	✔ READ ¹
Token - Use	Administration	✔ SPECIAL	✘	✘
Plugin	Platform Content	✔ EXECUTE ² READ ²	✘	✘
¹ This permission applies to all content sets. ² This permission applies to the Interact content set.				

Installing API Gateway

Use the Tanium Console **Solutions** page to install API Gateway and choose either automatic or manual configuration:

- **Automatic configuration** (Tanium Core Platform 7.4.2 or later only): API Gateway is installed with any required dependencies and other selected products. This option is the best practice for most deployments. For more information about the automatic configuration for API Gateway, see [Tanium Console User Guide: Import all modules and services](#).
- **Manual configuration:** Manually install API Gateway and the required dependencies. For more information, see [Import API Gateway on page 23](#).

Before you begin

- Read the [release notes](#).
- Review the [API Gateway requirements on page 19](#).
- If you are upgrading from a previous version, see [Upgrade API Gateway on page 24](#).
- Assign the correct roles to users for API Gateway. Review the [User role requirements on page 21](#).
 - To import the API Gateway solution, you must be assigned the **Administrator** reserved role.
- If you install API Gateway in an All-in-One deployment, update the **api_token_trusted_ip_address_list** platform setting. For more information, see [Update platform setting for All-in-One deployment on page 27](#).

Import API Gateway

Perform the following steps to install the API Gateway solution on the Tanium Server.



NOTE

If you have multiple Tanium Servers in an active-active configuration, you only need to perform these steps on one Tanium Server if you have Tanium Core Platform 7.4.3.1204 or later.

1. Sign in to the Tanium Console with an account that has the **Administrator** reserved role.
2. From the Main menu, go to **Administration > Configuration > Solutions**.
3. In the **Content** section, select the checkbox for **API Gateway** and click **Install**.



TIP

If you need to install any prerequisite Tanium solutions or content, select the corresponding checkboxes for those solutions as well.

4. Review the content to import and click **Begin Install**.

Manage solution dependencies


Other Tanium solutions are required for API Gateway to function (required dependencies) or for specific API Gateway features to work (feature-specific dependencies). See [Solution dependencies](#).

Upgrade API Gateway

For the steps to upgrade API Gateway, see [Tanium Console User Guide: Import, re-import, or update specific solutions](#). After the upgrade, verify that the correct version is installed: see [Verify API Gateway version on page 24](#).

Verify API Gateway version

After you import or upgrade API Gateway, verify that the correct version is installed:

1. Refresh your browser.
2. From the Main menu, go to **Administration > Shared Services > API Gateway** to open the API Gateway **Overview** page.
3. To display version information, click Info .

Troubleshoot issues

If you experience issues with installing API Gateway, see [Queries return unexpected results or errors on page 27](#).

Using API Gateway

Use API Gateway to build API-based integrations with the Tanium Core Platform. This service consolidates information from multiple Tanium modules into a unified view of information on the endpoints in the environment. API Gateway intelligently routes requests to the services and sources that provide the most recent information and the most reliable mutations.

API Gateway uses GraphQL to request data (queries) and to make changes (mutations). With GraphQL, you can compose queries in API Gateway to retrieve the exact data that you want as well as filter the results to a set of endpoints.

Use API Gateway to:

- Query endpoints through the Tanium Server, or access data through Tanium Data Service
- Create, delete, and query actions
- Query packages
- Open a connection to an endpoint through Tanium Direct Connect and retrieve data from the endpoint

For examples of available functions, see [Reference: API Gateway examples on page 43](#).

Test a query in the Tanium Console


To access the query explorer in the Tanium Console and run a query, perform the following steps:

1. From the Main menu, go to **Administration > Shared Services > API Gateway**.
2. Enter a query in the query pane. For example, paste the following query to get the time from the Tanium Server:

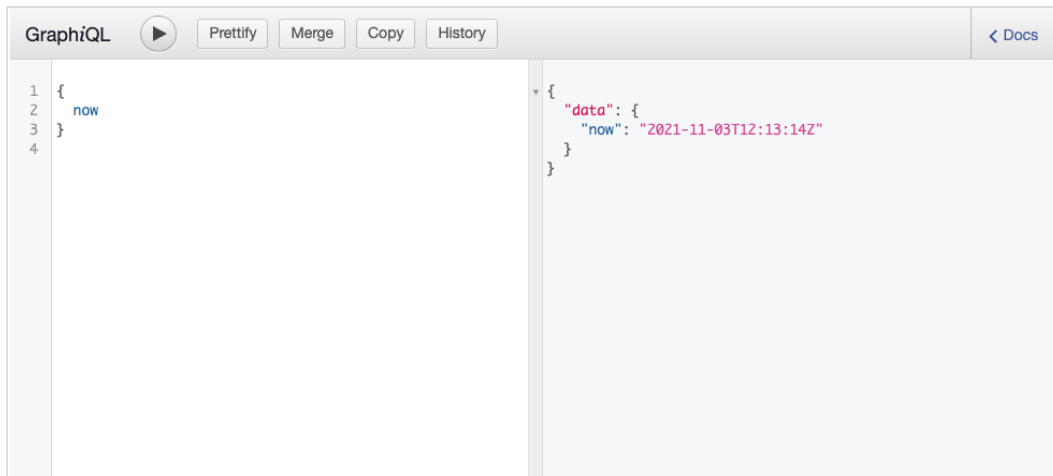
```
1 | {
2 |   now
3 | }
```



If the query explorer does not appear on the API Gateway **Overview** page, click **Customize Page** and make sure the **Query Explorer** option is selected.

3. (Optional) If a query or mutation uses variables, expand the **QUERY VARIABLES** pane and include the variables in the pane that expands.
4. Click Execute Query .

API Gateway sends the query to the server and returns the response in the results pane.



The screenshot shows the GraphQL query explorer interface. The top bar includes the title 'GraphQL', a play button, and buttons for 'Prettify', 'Merge', 'Copy', and 'History'. A '< Docs' link is also present. The left pane contains the query:

```
1 {  
2   now  
3 }  
4
```

The right pane shows the JSON response:

```
{  
  "data": {  
    "now": "2021-11-03T12:13:14Z"  
  }  
}
```


For more information on the query explorer, see [Query explorer on page 8](#).

Troubleshooting API Gateway

If API Gateway is not performing as expected, you might need to troubleshoot issues.

Collect logs

The information is saved as a ZIP file that you can download with your browser.

1. From the API Gateway **Overview** page, click Help , then the **Troubleshooting** tab.
2. Click **Download Support Package**.
A `tanium-api-gateway-support-[timestamp].zip` file downloads to the local download directory.
3. Contact Tanium Support to determine the best option to send the ZIP file. For more information, see [Contact Tanium Support on page 37](#).

Tanium API Gateway maintains logging information in the `gateway-service.log` file in the `\Program Files\Tanium\Tanium Module Server\services\gateway-files\logs\` directory.

Update platform setting for All-in-One deployment

For All-in-One deployments, before you install System User Service as a required dependency for API Gateway, you must first add the `127.0.0.1` IPv4 address if you enabled IPv4, or the `::1` IPv6 address if you enabled IPv6, to the `api_token_trusted_ip_address_list` platform Server setting. Otherwise, the installation process does not complete.



All-in-One deployments are supported only for proof-of-concept (POC) demonstrations.

1. From the Main menu, go to **Administration > Configuration > Platform Settings**.
2. In the **Name** column, click `api_token_trusted_ip_address_list`.
3. You have the following options:
 - If you enabled IPv4, enter `127.0.0.1` in the **Value** field.
 - If you enabled IPv6, enter `::1` in the **Value** field.
4. Click **Save**.

Queries return unexpected results or errors

- The API Gateway service redirects queries and mutations to other Tanium solutions. If API Gateway returns unexpected results or errors, make sure that all prerequisites are installed at the minimum recommended version. For information, see [API Gateway requirements on page 19](#).

- If all queries and mutations return a **502 Gateway Timeout** error, make sure the Tanium System User service and the Tanium API Gateway service are running on the Tanium Module Server.
- If you recently installed API Gateway or the System User service, restart the Tanium Module Server.
- Queries and mutations that use the **eid** element require Interact 2.9 or later.

Request error handling

When requests cannot be fully satisfied, the gateway service gives the caller accurate, practical error information. The gateway service honors the GraphQL error specification with some custom extensions for greater specificity.

Syntax Validation

Before processing any portion of a request, the service parses its syntax. If the request is not valid GraphQL, or if it does not conform to the schema, the service returns a response with the syntax errors, including line and column positions whenever possible.

Examples of syntax errors include:

- Unbalanced parentheses, braces, brackets, and quotation marks
- References to fields that do not exist

The GraphQL error specification requires reporting these errors to the caller by line and column in the request.



When possible, the Query pane in the Query Explorer identifies syntax errors in the request.

NOTE

Structural Validation

If the request passes syntax validation, the service evaluates the request for structural validity, which checks for errors not expressible with the GraphQL type system. Structural error examples include:

- Restrictions on type values, such as:
 - A string may not be empty
 - An integer must be positive
- Constraints on input objects, such as:
 - Exactly one of two or more entries must be present
- Existence of values in a known set, such as:
 - A named sensor must exist
- References to fields that cannot be satisfied in the current environment, such as:
 - A sensor does not exist

The GraphQL `extensions` error field for custom fields contains these errors.

Error Extensions

Errors that pertain to a specific field in the request are identified by a `fieldPath` entry in the `extensions` object. Similar to the `path` entry on the GraphQL error type, the `fieldPath` consists of a list of strings and integers which indicate keys in objects and offsets in lists, respectively. The root of the `fieldPath` is the request field that triggered the GraphQL error.



The numbering for list items starts at 0 for the first item and increments for each additional item. For example, if a `fieldPath` contains 5, this refers to the sixth item in the list.

The `extensions` object may also have these fields:

Extensions field	Description
<code>code</code>	an error code attributable to the field itself
<code>context</code>	a map of contextual information that refines or specifies the error
<code>message</code>	a short error explanation

For example, with the following syntactically valid request:

```
1  {
2    endpoints {
3      edges {
4        node {
5          name
6          serialNumber
7        }
8      }
9    }
10 }
```

If the serial number is not available in the underlying data store, the gateway service returns a response including an extended error that the value does not exist:

```
1  {
2    "errors": [
3      {
4        "message": "Invalid request",
```

```

5     "path": ["endpoints"],
6     "extensions": {
7         "fieldPath": ["edges", "node", "serialNumber"],
8         "code": "validation_exists",
9         "message": "must refer to an existing sensor",
10        "context": {
11            "source": "tds"
12        }
13    }
14 }
15 ]
16 }

```

Invalid argument errors contain an extended `argumentErrors` entry. The value of that entry is a list of objects that may have the following entries:

Extensions field	Description
code	an error code attributable to the field itself
context	a map of contextual information that refines or specifies the error
path	the path to the argument error, relative to the field
message	a short error explanation

For example, with the following request that does not define a `name` or `path` value:

```

1  {
2    endpoints(filter: {filters: [{memberOf: {name: ""}}, {path:""}]}) {
3      edges {
4        node {
5          name
6        }
7      }
8    }
9  }

```

The gateway service returns a response including both errors stating a value is a required:

```

1  {
2    "errors": [
3      {
4        "message": "Invalid request",
5        "path": ["endpoints"],
6        "extensions": {
7          "argumentErrors": [
8            {
9              "code": "validation_required",
10             "path": ["filter", "filters", 0, "memberOf", "name"]
11            },
12            {
13              "code": "validation_required",
14              "path": ["filter", "filters", 1, "path"]
15            }
16          ]
17        }
18      }
19    ]
20  }

```

Extended errors may identify argument errors within sub-fields. For example, the following request specifies a non-existent sensor name:

```

1  {
2    endpoints {
3      edges {
4        node {
5          sensorReadings(sensors:[{name: "A Non-Existent Sensor"}]) {
6            columns {
7              values
8            }
9          }
10         }

```

```
11 |     }
12 |   }
13 | }
```

The gateway service returns a response including the sub-field error that the value does not exist:

```
1 | {
2 |   "errors": [
3 |     {
4 |       "message": "Invalid request",
5 |       "path": ["endpoints"],
6 |       "extensions": {
7 |         "fieldPath": ["edges", "node", "sensorReadings"],
8 |         "argumentErrors": [
9 |           {
10 |            "code": "validation_exists",
11 |            "path": ["sensors", 0, "name"]
12 |           }
13 |         ]
14 |       }
15 |     }
16 |   ]
17 | }
```

Whenever possible, the response includes the valid portions of a query. For example, the following request references a valid sensor (Custom Tags) and a not valid sensor (A Non-existent Sensor):

```
1 | {
2 |   endpoints {
3 |     edges {
4 |       node {
5 |         name
6 |         ipAddress
7 |         os {
```



```

8         generation
9     }
10    sensorReadings(sensors: [{name: "Custom Tags"}, {name: "A Non-existent Sensor"}])
11    {
12        columns {
13            sensor {
14                name
15            }
16            name
17            values
18        }
19    }
20 }
21 }
22 }

```

The response includes the error from A Non-Existent Sensor not being found, and all information that can otherwise be returned from the Custom Tags sensor:

```

1  {
2    "errors": [
3      {
4        "message": "sensors: (1: (name: must refer to an existing sensor.)).",
5        "path": [
6          "endpoints"
7        ],
8        "extensions": {
9          "argumentErrors": [
10         {
11           "code": "validation_exists",
12           "context": {
13             "source": "tds"
14         },

```

```
15         "message": "must refer to an existing sensor",
16         "path": [
17             "sensors",
18             1,
19             "name"
20         ]
21     }
22 ],
23     "fieldPath": [
24         "edges",
25         "node",
26         "sensorReadings"
27     ]
28 }
29 }
30 ],
31 "data": {
32     "endpoints": {
33         "edges": [
34             {
35                 "node": {
36                     "name": "example-host",
37                     "ipAddress": "192.0.2.10",
38                     "os": {
39                         "generation": "Windows 10"
40                     },
41                     "sensorReadings": {
42                         "columns": [
43                             {
44                                 "sensor": {
45                                     "name": "Custom Tags"
46                                 },
47                                 "name": "Custom Tags",
48                                 "values": [
```

```

49         "custom-tag-1",
50         "custom-tag-2"
51     ]
52 }
53 ]
54 }
55 }
56 }
57 ]
58 }
59 }
60 }

```



Similar to lists in the `fieldPath` referring to offsets in lists, if an error occurs on an item in a list, the error message might refer to the list offset instead of the field name. In the above response, the `message` refers to item 1 in the list; because list numbering starts at 0, this points to the second sensor (A Non-existent Sensor).

Error codes

General error codes include:

Error code	Description
502 Bad Gateway	The gateway service received an invalid response from the server.
internal_server_error	The response field cannot be resolved due to a server error.

Syntax error codes include:

Error code	Description
GRAPHQL_PARSE_FAILED	The request contains unbalanced parentheses, braces, brackets, or quotation marks. See the <code>message</code> for more information.
GRAPHQL_VALIDATION_FAILED	The request contains an unrecognized argument, or a required argument is not included. See the <code>message</code> for more information.

Structural error codes include:

Error code	Description
validation_exists	The field must refer to an existing entity.
validation_in_invalid	The field must have a value from a specific set.
validation_invalid_use	The filter is not valid in its context.
validation_key_unexpected	The field is not allowed.
validation_min_greater_equal_than_required	The field must have a value greater than or equal to its documented minimum.
validation_nil	The field must be absent.
validation_nil_or_not_empty_required	The field must be absent or have a non-empty value.
validation_not_nil_required	The field must have a value.
validation_overdetermined	The field's object must not have inconsistent entries.
validation_required	The field must have a non-empty value.
validation_timestamp	The field's value must be a valid RFC 3339 string.
validation_unique	The field's value must be unique in its documented container.

Uninstall API Gateway

If you need to uninstall API Gateway, perform the following steps.



Consult with Tanium Support before you uninstall or reinstall API Gateway.

IMPORTANT

1. Sign in to the Tanium Console as a user with the Administrator role.
2. From the Main menu, go to **Administration > Configuration > Solutions**.
3. In the **Content** section, select the **API Gateway** row and click **Uninstall**.
4. Review the summary and click **Yes** to proceed with the uninstallation.
5. When prompted to confirm, enter your password.



NOTE

The uninstall does not remove the API Gateway log from the Tanium Module Server. To remove the log after the uninstall completes, manually delete the `\Program Files\Tanium\Tanium Module Server\services\gateway-files\` directory.

Contact Tanium Support

To contact Tanium Support for help, sign in to <https://support.tanium.com>.

Reference: Filter syntax

Endpoint query filter syntax

Filter endpoint query requests using the following filter syntax.

General Fields

You can filter on field values. Unless otherwise specified, define a filter object using the syntax:

```
{path: "field-path", value: "field-value"}
```

Use dot notation for sub-fields not at the root level.

For example, the following request filters and returns only endpoints whose primary user email address is `user@example.com`:

```
1  {
2    endpoints(filter: {path: "primaryUser.email", value: "user@example.com"}) {
3      edges {
4        node {
5          id
6          primaryUser {
7            email
8          }
9        }
10     }
11  }
12 }
```

Computer group membership

You can filter on membership in named computer groups. Define a filter object using the syntax:

```
{memberOf: {name: "computer-group-name"}}
```

For example, the following request filters and returns only endpoints that are members of the All Linux computer group:

```

1 | {
2 |   endpoints(filter: {memberOf: {name: "All Linux"}}) {
3 |     edges {
4 |       node {
5 |         id
6 |       }
7 |     }
8 |   }
9 | }

```

Sensors

Endpoint queries can filter on readings of arbitrary named sensors. Define a filter object using the syntax:

```
{sensor: {name: "sensor-name", value: "filter-value"}}
```

For example, the following request filters based on the Total Memory sensor returning a value of 16384 MB for endpoints, and returns only those endpoints:

```

1 | {
2 |   endpoints(filter: {sensor: {name: "Total Memory", value: "16384 MB"}}) {
3 |     edges {
4 |       node {
5 |         id
6 |       }
7 |     }
8 |   }
9 | }

```

Parameterized sensors

You can also filter on parameterized sensors. Define a filter object using the syntax:

```
{sensor:
  {name: "sensor-name",
  params: [
    {name: "parameter-1-name", value: "parameter-1-value"},

```

```

    {name: "parameter-2-name", value: "parameter-2-value"},
    {name: "parameter-n-name", value: "parameter-n-value"}
  ]
},
value: "filter-value"
}

```

Add a `params` `name/value` object for every parameter.

For example, the following request filters based on the Custom Tag Exists sensor, returning endpoints tagged with the custom tag `the-tag`, and returns only those endpoints:

```

1  {
2    endpoints(filter: {sensor: {name: "Custom Tag Exists", params:[{name: "tag", value: "the-
3      tag"}]}}, value: "true"}) {
4      edges {
5        node {
6          id
7        }
8      }
9    }

```

Multi-column sensors

You can filter multi-column sensors by specifying the column name. Define a filter object using the syntax:

```

{sensor: {
  name: "sensor-name",
  column: "sensor-column-name",
  value: "filter-value"}
}

```

For example, the following request filters based on the CPU Details sensor returning endpoints with an Intel CPU, and returns only those endpoints:

```

1  {
2    endpoints(filter: {sensor: {name: "CPU Details", column: "CPU", value: "Intel"}}) {

```



```

3 |     edges {
4 |         node {
5 |             id
6 |         }
7 |     }
8 | }
9 | }

```

If you want to select across multiple columns of each row of a multi-column sensor, define a `filter.filters` list, then define a list element for each column filter, using the following syntax:

```

filters: [
  {sensor: {column: "column-1-name"}, value: "column-1-value"},
  {sensor: {column: "column-2-name"}, value: "column-2-value"},
  {sensor: {column: "column-n-name"}, value: "column-n-value"}
]

```

You can only use a single level of child filters, and can only use the default `EQ` operator.

For example, the following request filters based on the CPU Details sensor returning endpoints with 4 total cores and an Intel CPU, and returns only those endpoints:

```

1 | {
2 |   endpoints(filter: {
3 |     sensor: {name: "CPU Details"},
4 |     filters: [
5 |       {sensor: {column: "Total Cores"}, value: "4"},
6 |       {sensor: {column: "CPU"}, value: "Intel"}
7 |     ]
8 |   }) {
9 |     edges {
10 |       node {
11 |         id
12 |       }
13 |     }
14 |   }
15 | }

```

In this form, the sensor is identified in the top filter and the columns are identified in the child filters. You can only use a single level of child filters, and can only use the default EQ operator.

Sensor readings

After you filter endpoints, you can filter your results to return data from additional sensors. As a child of the node object, define the sensor reading filter using the following syntax:

```
sensorReadings
  (sensors: [{name: "sensor-name"}]) {
    columns {
      column-1
      column-2
      column-n
    }
  }
```

For example, the following request first filters endpoints and only returns endpoints with Firefox as an installed application, then for each endpoint, returns the results of the Is Linux sensor:

```
1  {
2    endpoints(
3      filter: {sensor: {name: "Installed Applications" column:"Name"}, op: CONTAINS, value:
4        "firefox"}
5    ) {
6      edges {
7        node {
8          name
9          ipAddress
10         sensorReadings(sensors: [{name: "Is Linux"}]) {
11           columns {
12             sensor {name}
13             name
14             values
15           }
16         }
17       }
18     }
19 }
```

Reference: API Gateway examples

Use the following query examples to learn about the functionality and syntax of queries and mutations in API Gateway.

- [General examples on page 43](#)
- [Action examples on page 68](#)
- [Deploy examples on page 71](#)
- [Direct Connect examples on page 81](#)

General examples

The following queries retrieve data from the endpoints in your environment.

Get server time

The following query retrieves the local time from the Tanium Server.

```
1 | {  
2 |   now  
3 | }
```

Example response:

```
1 | {  
2 |   "data": {  
3 |     "now": "2021-11-08T19:22:03Z"  
4 |   }  
5 | }
```

Get endpoints

The following query retrieves known endpoints from the Tanium Server.

```
1 | {  
2 |   endpoints(source: {ts: {expectedCount: 1, stableWaitTime: 10}}) {
```

```
3     edges {
4       node {
5         computerID
6         name
7         serialNumber
8         ipAddress
9       }
10    }
11  }
12 }
```

Example response:

```
1  {
2    "data": {
3      "endpoints": {
4        "edges": [
5          {
6            "node": {
7              "computerID": "937672696",
8              "name": "ubuntu-test",
9              "serialNumber": "Not Specified",
10             "ipAddress": "10.168.20.30"
11           }
12         },
13         {
14           "node": {
15             "computerID": "1867570226",
16             "name": "CentOS-test-1",
17             "serialNumber": "Not Specified",
18             "ipAddress": "10.168.20.40"
19           }
20         },
21         {
```

```

22     "node": {
23         "computerID": "2711217959",
24         "name": "CentOS-test-2",
25         "serialNumber": "Not Specified",
26         "ipAddress": "10.168.20.50"
27     }
28 }
29 ]
30 }
31 }
32 }

```

Get endpoints IDs from Tanium Data Service

The following query retrieves all endpoint IDs from Tanium Data Service.

```

1  {
2  endpoints {
3      edges {
4          node {
5              id
6          }
7      }
8  }
9  }

```

Example response:

```

1  {
2  "data": {
3      "endpoints": {
4          "edges": [
5              {
6                  "node": {

```

```

7         "id": "12345"
8     }
9 },
10    {
11        "node": {
12            "id": "54321"
13        }
14    },
15    {
16        "node": {
17            "id": "21212"
18        }
19    }
20 ]
21 }
22 }
23 }

```

Get endpoints using aliases

The following query retrieves known endpoints using aliases to rename the information returned and prevent data collision due to overlapping field names. The `tds` alias retrieves the `id` and `name` of known endpoints from Tanium Data Service. The `ts` alias retrieves the IPv6 address `value` of known endpoints from the Tanium Server and the sensor `name`.

```

1  {
2    tds:endpoints {
3      edges {
4        node {
5          id
6          name
7        }
8      }
9    }
10   ts:endpoints(source: {ts: {stableWaitTime: 10}}) {

```

```

11     edges {
12       node {
13         sensorReadings(sensors: [{name: "IPv6 Address"}]) {
14           columns {
15             name
16             values
17           }
18         }
19       }
20     }
21   }
22 }

```

Example response:

```

1  {
2    "data": {
3      "tds": {
4        "edges": [
5          {
6            "node": {
7              "id": "12345",
8              "name": "example-name"
9            }
10         },
11         {
12           "node": {
13             "id": "54321",
14             "name": "example-name-2"
15           }
16         }
17       ]
18     },
19     "ts": {

```

```
20     "edges": [  
21     {  
22         "node": {  
23             "sensorReadings": {  
24                 "columns": [  
25                     {  
26                         "name": "IPv6 Address",  
27                         "values": [  
28                             "2001:db8::3"  
29                         ]  
30                     }  
31                 ]  
32             }  
33         }  
34     },  
35     {  
36         "node": {  
37             "sensorReadings": {  
38                 "columns": [  
39                     {  
40                         "name": "IPv6 Address",  
41                         "values": [  
42                             "2001:db8::2"  
43                         ]  
44                     }  
45                 ]  
46             }  
47         }  
48     }  
49 ]  
50 }  
51 }  
52 }
```


Get endpoints using multiple sensors

The following query retrieves known endpoints using multiple sensors. The request specifies the sensor name as one of the fields returned in the request, to identify the sensor that reported the information.

```
1 {
2   endpoints {
3     edges {
4       node {
5         computerID
6         sensorReadings(sensors: [{name:"Installed Applications"},{name:"Running
Applications"}]) {
7           columns {
8             name
9             values
10            sensor {name}
11          }
12        }
13      }
14    }
15  }
16 }
```

Example response:

```
1 {
2   "data": {
3     "endpoints": {
4       "edges": [
5         {
6           "node": {
7             "computerID": "987654321",
8             "sensorReadings": {
9               "columns": [
10              {
11                "name": "Name",
```

```
12         "values": [  
13             "Tanium Client 7.4.7.1094"  
14         ],  
15         "sensor": {  
16             "name": "Installed Applications"  
17         }  
18     },  
19     {  
20         "name": "Version",  
21         "values": [  
22             "7.4.7.1094"  
23         ],  
24         "sensor": {  
25             "name": "Installed Applications"  
26         }  
27     },  
28     {  
29         "name": "Silent Uninstall String",  
30         "values": [  
31             "\"C:\\Program Files (x86)\\Tanium\\Tanium Client\\uninst.exe\""br/>32         ],  
33         "sensor": {  
34             "name": "Installed Applications"  
35         }  
36     },  
37     {  
38         "name": "Uninstallable",  
39         "values": [  
40             "Is Uninstallable"  
41         ],  
42         "sensor": {  
43             "name": "Installed Applications"  
44         }  
45     },
```

```
46     {
47         "name": "Name",
48         "values": [
49             "Tanium Client"
50         ],
51         "sensor": {
52             "name": "Running Applications"
53         }
54     },
55     {
56         "name": "Version",
57         "values": [
58             "7.4.7.1094"
59         ],
60         "sensor": {
61             "name": "Running Applications"
62         }
63     },
64     {
65         "name": "Process Name",
66         "values": [
67             "TaniumClient.exe"
68         ],
69         "sensor": {
70             "name": "Running Applications"
71         }
72     }
73 ]
74 }
75 }
76 }
77 ]
78 }
79 }
80 }
```

Get rich endpoint data

The following query demonstrates using nested fields to retrieve categorized endpoint data.



The `first:2` argument retrieves two records; set this value higher to retrieve more records at a time. For more information on pagination arguments, see [Pagination on page 12](#).

```
1 {
2   endpoints (first:2) {
3     edges {
4       node {
5         name
6         computerID
7         ipAddress
8         isVirtual
9         chassisType
10        systemUUID
11        domainName
12        os {
13          name
14          platform
15          generation
16        }
17        processor {
18          architecture
19          cacheSize
20          consumption
21          cpu
22          family
23          manufacturer
24          speed
25        }
26        lastLoggedInUser
27      }
28    }
29    pageInfo {
```

```
30     startCursor
31     endCursor
32     hasNextPage
33   }
34 }
35 }
```

Example response:

```
1  {
2    "data": {
3      "endpoints": {
4        "edges": [
5          {
6            "node": {
7              "name": "Test-01",
8              "computerID": "1234567890",
9              "ipAddress": "192.0.2.10",
10             "isVirtual": false,
11             "chassisType": "TSE-Error: Unknown - dmidecode unavailable",
12             "systemUUID": "TSE-Error: Unknown - dmidecode unavailable",
13             "domainName": "(none)",
14             "os": {
15               "name": "Red Hat Enterprise Linux Server release 5.11 (Tikanga)",
16               "platform": "Linux",
17               "generation": "Red Hat Enterprise Linux 5"
18             },
19             "processor": {
20               "architecture": "x86_64",
21               "cacheSize": "16384 KB",
22               "consumption": "9.9 %",
23               "cpu": "Intel Core Processor (Haswell, no TSX, IBRS)",
24               "family": "6",
25               "manufacturer": "GenuineIntel",
```

```
26         "speed": "2600 Mhz"
27     },
28     "lastLoggedInUser": "reboot"
29 }
30 },
31 {
32     "node": {
33         "name": "Test-02",
34         "computerID": "3216549870",
35         "ipAddress": "192.0.2.20",
36         "isVirtual": true,
37         "chassisType": "Virtual",
38         "systemUUID": "[no results]",
39         "domainName": "(none)",
40         "os": {
41             "name": "CentOS Linux release 8.4.2105",
42             "platform": "Linux",
43             "generation": "CentOS 8"
44         },
45         "processor": {
46             "architecture": "x86_64",
47             "cacheSize": "35840 KB",
48             "consumption": "18.6 %",
49             "cpu": "Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz",
50             "family": "6",
51             "manufacturer": "GenuineIntel",
52             "speed": "2400 Mhz"
53         },
54         "lastLoggedInUser": "tester-5"
55     }
56 }
57 ],
58 "pageInfo": {
59     "startCursor": "4267468:0",
```

```

60     "endCursor": "4267468:1",
61     "hasNextPage": true
62   }
63 }
64 }
65 }

```

Get a set of endpoints

The following query retrieves a set of endpoints. The query demonstrates the use of the `sensorReadings` field and contains a filter argument to retrieve endpoints whose names contain the letter `a`. The results are paginated to 3 records.

```

1  {
2    endpoints(first: 3, filter: {op: MATCHES, path: "name", value: "a.*"}) {
3      edges {
4        node {
5          name
6          ipAddress
7          sensorReadings(sensors: [{name: "EID Last Seen"}]) {
8            columns {
9              name
10             values
11           }
12         }
13       }
14     }
15   }
16 }

```

Example response:

```

1  {
2    "data": {
3      "endpoints": {

```

```
4     "edges": [
5       {
6         "node": {
7           "name": "test-1",
8           "ipAddress": "192.0.2.10",
9           "sensorReadings": {
10            "columns": [
11              {
12                "name": "EID Last Seen",
13                "values": [
14                  "Mon, 08 Nov 2021 21:29:28 +0000"
15                ]
16              }
17            ]
18          }
19        }
20      },
21      {
22        "node": {
23          "name": "test-2",
24          "ipAddress": "192.0.2.20",
25          "sensorReadings": {
26            "columns": [
27              {
28                "name": "EID Last Seen",
29                "values": [
30                  "Mon, 08 Nov 2021 21:29:28 +0000"
31                ]
32              }
33            ]
34          }
35        }
36      },
37      {
```



```

38     "node": {
39         "name": "test-3",
40         "ipAddress": "192.0.2.30",
41         "sensorReadings": {
42             "columns": [
43                 {
44                     "name": "EID Last Seen",
45                     "values": [
46                         "Mon, 08 Nov 2021 21:17:17 +0000"
47                     ]
48                 }
49             ]
50         }
51     }
52 }
53 ]
54 }
55 }
56 }

```

Unregistered sensor query

The following query retrieves the operating system platform from all endpoints.



In API Gateway, a sensor is unregistered if the sensor is not represented by a named field in the API Gateway schema. This has no correlation to registering sensors in Tanium Data Service.

```

1  {
2    endpoints {
3      edges {
4        node {
5          id
6          name
7          sensorReadings(sensors: [{name: "OS Platform"}]) {

```

```
8         columns {
9             name
10            values
11        }
12    }
13 }
14 }
15 }
16 }
```

Example response:

```
1 {
2     "data": {
3         "endpoints": {
4             "edges": [
5                 {
6                     "node": {
7                         "id": "12345",
8                         "name": "Test-01",
9                         "sensorReadings": {
10                            "columns": [
11                                {
12                                    "name": "OS Platform",
13                                    "values": [
14                                        "Linux"
15                                    ]
16                                }
17                            ]
18                        }
19                    }
20                },
21                {
22                    "node": {
```

```

23         "id": "54321",
24         "name": "Test-03",
25         "sensorReadings": {
26             "columns": [
27                 {
28                     "name": "OS Platform",
29                     "values": [
30                         "Linux"
31                     ]
32                 }
33             ]
34         }
35     }
36 }
37 ]
38 }
39 }
40 }

```

Unregistered parameterized sensor query

The following query checks to see if each endpoint contains the `C:\Windows\py.exe` file.



In API Gateway, a sensor is unregistered if the sensor is not represented by a named field in the API Gateway schema. This has no correlation to registering sensors in Tanium Data Service.

```

1  {
2  endpoints(source: {ts: {}}) {
3  edges {
4  node {
5  name
6  id
7  sensorReadings(

```

```

8         sensors: [{name: "File Exists", params: [{name: "file", value:
"C:\\Windows\\py.exe"}]}]
9     ) {
10         columns {
11             sensor {
12                 name
13                 params {
14                     name
15                     value
16                 }
17             }
18             values
19         }
20     }
21 }
22 }
23 }
24 }

```

Example response:

```

1 {
2   "data": {
3     "endpoints": {
4       "edges": [
5         {
6           "node": {
7             "name": "Test-01",
8             "id": "12345",
9             "sensorReadings": {
10              "columns": [
11                {
12                  "sensor": {
13                    "name": "File Exists",

```

```
14         "params": [
15             {
16                 "name": "file",
17                 "value": "C:\\Windows\\py.exe"
18             }
19         ]
20     },
21     "values": [
22         "File does not exist"
23     ]
24 }
25 ]
26 }
27 }
28 },
29 {
30     "node": {
31         "name": "[no results]",
32         "id": "54225",
33         "sensorReadings": {
34             "columns": [
35                 {
36                     "sensor": {
37                         "name": "File Exists",
38                         "params": [
39                             {
40                                 "name": "file",
41                                 "value": "C:\\Windows\\py.exe"
42                             }
43                         ]
44                     },
45                     "values": [
46                         "[no results]"
47                     ]
48                 }
49             ]
50         }
51     }
52 }
```

```
48         }
49     ]
50 }
51 }
52 },
53 {
54     "node": {
55         "name": "[no results]",
56         "id": "65456",
57         "sensorReadings": {
58             "columns": [
59                 {
60                     "sensor": {
61                         "name": "File Exists",
62                         "params": [
63                             {
64                                 "name": "file",
65                                 "value": "C:\\Windows\\py.exe"
66                             }
67                         ]
68                     },
69                     "values": [
70                         "[no results]"
71                     ]
72                 }
73             ]
74         }
75     }
76 }
77 ]
78 }
79 }
80 }
```

Endpoint cursor query

The following query retrieves cursors for endpoint records. You can use these cursors to retrieve specific endpoint records on a page. For more information on cursors, see [Pagination on page 12](#).

```
1  {
2    endpoints {
3      pageInfo {
4        hasPreviousPage
5        hasNextPage
6        startCursor
7        endCursor
8      }
9    }
10 }
```

Example response:

```
1  {
2    "data": {
3      "endpoints": {
4        "pageInfo": {
5          "hasPreviousPage": false,
6          "hasNextPage": true,
7          "startCursor": "4277520:0",
8          "endCursor": "4277520:19"
9        }
10     }
11  }
12 }
```

Paginated query

The following query retrieves the first five endpoint records after the given cursor. For more information on cursors, see [Pagination on page 12](#).

```

1  {
2    endpoints(after: "4277520:4", first: 5) {
3      edges {
4        node {
5          name
6          id
7          ipAddress
8        }
9      }
10     pageInfo {
11       hasNextPage
12       startCursor
13       endCursor
14     }
15   }
16 }

```

Example response:

```

1  {
2    "data": {
3      "endpoints": {
4        "edges": [
5          {
6            "node": {
7              "name": "Test-06",
8              "id": "6172",
9              "ipAddress": "192.0.2.10"
10           }
11         },
12         {
13           "node": {
14             "name": "Test-07",
15             "id": "87654",

```



```
16         "ipAddress": "192.0.2.20"
17     }
18 },
19 {
20     "node": {
21         "name": "Test-14",
22         "id": "43584",
23         "ipAddress": "192.0.2.30"
24     }
25 },
26 {
27     "node": {
28         "name": "Test-03",
29         "id": "37233",
30         "ipAddress": "[no results]"
31     }
32 },
33 {
34     "node": {
35         "name": "Test-55",
36         "id": "12139",
37         "ipAddress": "[no results]"
38     }
39 }
40 ],
41 "pageInfo": {
42     "hasNextPage": true,
43     "startCursor": "4277520:5",
44     "endCursor": "4277520:9"
45 }
46 }
47 }
48 }
```

Software characteristics query with filter

The following query retrieves endpoints that contain software installed by Deploy, where the package ID is 123.

```
1  {
2    endpoints {
3      edges {
4        node {
5          ipAddress
6          isVirtual
7          domainName
8          os {
9            generation
10         }
11         lastLoggedInUser
12         deployedSoftwarePackages (
13           filter: {filters: [{op: EQ, path: "id", value: "123"}, {op: EQ, path:
14             "applicability", value: "Installed"}]}
15         ) {
16           id
17         }
18       }
19     }
20 }
```

Example response:

```
1  {
2    "data": {
3      "endpoints": {
4        "edges": [
5          {
6            "node": {
7              "ipAddress": "192.0.2.10",
```

```
8         "isVirtual": true,
9         "domainName": "(none)",
10        "os": {
11            "generation": "Debian 9"
12        },
13        "lastLoggedInUser": "[no results]",
14        "deployedSoftwarePackages": [
15            {
16                "id": "123"
17            }
18        ]
19    },
20 },
21 {
22     "node": {
23         "ipAddress": "192.0.2.20",
24         "isVirtual": true,
25         "domainName": "(none)",
26         "os": {
27             "generation": "CentOS 7"
28         },
29         "lastLoggedInUser": "admin",
30         "deployedSoftwarePackages": [
31             {
32                 "id": "123"
33             }
34         ]
35     }
36 },
37 {
38     "node": {
39         "ipAddress": "192.0.2.30",
40         "isVirtual": true,
41         "domainName": "(none)",
```

```

42     "os": {
43         "generation": "Ubuntu 20.04"
44     },
45     "lastLoggedInUser": "admin",
46     "deployedSoftwarePackages": [
47         {
48             "id": "123"
49         }
50     ]
51     }
52 }
53 ]
54 }
55 }
56 }

```

Action examples

The following query and mutation retrieve action details and create an action.

Create action (subset of endpoints)

The following mutation deploys an action to increase the verbosity of log levels on Debian endpoints.

```

1  mutation {
2    createAction(
3      action: {description: "Increasing log verbosity level on all debian endpoints for
troubleshooting", target: {targetGroup: "All Debian", platforms: [Linux]},
changeClientSetting: {name: LOG_VERBOSITY_LEVEL, value: "41"}}
4    ) {
5      id
6    }
7  }

```

Example response:

```
1 {
2   "data": {
3     "createAction": {
4       "id": "82"
5     }
6   }
7 }
```

Get action details

The following parameterized query retrieves details and any results for an action.

```
1 query ($id: ID!) {
2   lastActionDetails(id: $id) {
3     id
4     name
5     comment
6     expireSeconds
7     creationTime
8     startTime
9     expirationTime
10    distributeSeconds
11    status
12    stoppedFlag
13  }
14  lastActionResults(id: $id) {
15    id
16    waiting
17    downloading
18    running
19    waitingToRetry
20    completed
21    expired
22    failed
```

```
23     pendingVerification
24     verified
25     failedVerification
26   }
27 }
28
```

Include the endpoint ID in the **QUERY VARIABLES** panel:

```
1  {
2    "id": 12323
3  }
```

Example response:

```
1  {
2    "data": {
3      "lastActionDetails": {
4        "id": "219",
5        "name": "Distribute Tanium Standard Utilities",
6        "comment": "Distribute Tanium Standard Utilities",
7        "expireSeconds": 3900,
8        "creationTime": "2022-03-14T18:05:41Z",
9        "startTime": "2022-03-14T18:05:35Z",
10       "expirationTime": "2022-03-14T19:10:35Z",
11       "distributeSeconds": 3600,
12       "status": "OPEN",
13       "stoppedFlag": false
14     },
15     "lastActionResults": {
16       "id": "219",
17       "waiting": 0,
18       "downloading": 0,
```

```
19     "running": 0,  
20     "waitingToRetry": 0,  
21     "completed": 1,  
22     "expired": 0,  
23     "failed": 0,  
24     "pendingVerification": 0,  
25     "verified": 0,  
26     "failedVerification": 0  
27   }  
28 }  
29 }
```

Deploy examples

The following queries and mutation require Deploy, retrieve information about software packages deployed in your environment, and allow you to deploy a software package to endpoints.

Deploy a package to all endpoints

The following mutation deploys a package to `All Computers`.

```
1  mutation {  
2    manageSoftware(  
3      operation: INSTALL  
4      softwarePackageID: 2  
5      start: "2021-10-27T00:00:00Z"  
6      end: "2021-11-03T00:00:00Z"  
7      target: {targetGroup: "All Computers"}  
8    ) {  
9      ID  
10     name  
11   }  
12 }
```

Example response:

```
1 {
2   "data": {
3     "manageSoftware": {
4       "ID": "2",
5       "name": "Install Tanium Standard Utilities (Linux)"
6     }
7   }
8 }
```

Get package details

The following query retrieves multiple fields for all packages.

```
1 query PackagesQuery {
2   packages {
3     items {
4       id
5       name
6       displayName
7       command
8       commandTimeout
9       expireSeconds
10      contentSet {
11        id
12        name
13      }
14      processGroupFlag
15      skipLockFlag
16      metadata {
17        adminFlag
18        name
19        value
20      }
21      sourceHash
```



```
22     sourceHashChangedFlag
23     sourceID
24     sourceName
25     parameters {
26         key
27         value
28     }
29     rawParameterDefinition
30     parameterDefinition {
31         parameterType
32         model
33         parameters {
34             model
35             parameterType
36             key
37             label
38             helpString
39             defaultValue
40             validationExpressions {
41                 model
42                 parameterType
43                 expression
44                 helpString
45             }
46         promptText
47         heightInLines
48         maxChars
49         values
50         restrict
51         allowEmptyList
52         minimum
53         maximum
54         stepSize
55         snapInterval
```

```
56     dropdownOptions {
57         model
58         parameterType
59         name
60         value
61     }
62     componentType
63     startDateRestriction {
64         model
65         parameterType
66         type
67         interval
68         intervalCount
69         unixTimeStamp
70     }
71     endDateRestriction {
72         model
73         parameterType
74         type
75         interval
76         intervalCount
77         unixTimeStamp
78     }
79     startTimeRestriction {
80         model
81         parameterType
82         type
83         interval
84         intervalCount
85         unixTimeStamp
86     }
87     endTimeRestriction {
88         model
89         parameterType
```

```

90         type
91         interval
92         intervalCount
93         unixTimeStamp
94     }
95     allowDisableEnd
96     defaultRangeStart {
97         model
98         parameterType
99         type
100        interval
101        intervalCount
102        unixTimeStamp
103    }
104    defaultRangeEnd {
105        model
106        parameterType
107        type
108        interval
109        intervalCount
110        unixTimeStamp
111    }
112    separatorText
113    }
114 }
115 verifyExpireSeconds
116 }
117 }
118 }

```

Example response:

```

1 | {
2 |   "data": {

```

```

3     "packages": {
4         "items": [
5             {
6                 "id": "1",
7                 "name": "Distribute Tanium Standard Utilities",
8                 "displayName": "Distribute Tanium Standard Utilities",
9                 "command": "cmd.exe /c cscript.exe //E:VBScript install-standard-utils.vbs
10                \"Tools\\StdUtils\\",
11                "commandTimeout": 2700,
12                "expireSeconds": 3300,
13                "contentSet": {
14                    "id": "5",
15                    "name": "Client Management"
16                },
17                "processGroupFlag": true,
18                "skipLockFlag": false,
19                "metadata": [],
20                "sourceHash":
21                "60b3e906f92929da67341792db9675d5cd91686546f01b57857686c8c6d84fa8",
22                "sourceHashChangedFlag": false,
23                "sourceID": 0,
24                "sourceName": "",
25                "parameters": [],
26                "rawParameterDefinition": null,
27                "parameterDefinition": null,
28                "verifyExpireSeconds": 600
29            },
30            {
31                "id": "2",
32                "name": "Distribute Tanium Standard Utilities (Linux)",
33                "displayName": "Distribute Tanium Standard Utilities (Linux)",
34                "command": "/bin/bash distribute-tools.sh STRICT",
35                "commandTimeout": 120,
36                "expireSeconds": 720,
37                "contentSet": {

```

```

36         "id": "5",
37         "name": "Client Management"
38     },
39     "processGroupFlag": true,
40     "skipLockFlag": false,
41     "metadata": [],
42     "sourceHash":
43     "4ed7a30a1ca6c81be5a71b892dfadcdb489122cc641fa4b644f53255134215c9",
44     "sourceHashChangedFlag": false,
45     "sourceID": 0,
46     "sourceName": "",
47     "parameters": [],
48     "rawParameterDefinition": null,
49     "parameterDefinition": null,
50     "verifyExpireSeconds": 600
51 }
52 ]
53 }
54 }

```

Get Deploy packages

The following query retrieves all Deploy packages.

```

1  {
2    softwarePackages {
3      edges {
4        node {
5          id
6          productName
7          productVendor
8          productVersion
9        }

```

```
10 |     }
11 |   }
12 | }
```

Example response:

```
1 | {
2 |   "data": {
3 |     "softwarePackages": {
4 |       "edges": [
5 |         {
6 |           "node": {
7 |             "id": "19",
8 |             "productName": "Firefox (x64 en-US)",
9 |             "productVendor": "Mozilla",
10 |            "productVersion": "98.0"
11 |          }
12 |        },
13 |        {
14 |          "node": {
15 |            "id": "30",
16 |            "productName": "Power BI Desktop (x64)",
17 |            "productVendor": "Microsoft",
18 |            "productVersion": "2.102.845.0"
19 |          }
20 |        },
21 |        {
22 |          "node": {
23 |            "id": "43",
24 |            "productName": "VLC media player (64-bit)",
25 |            "productVendor": "VideoLAN",
26 |            "productVersion": "3.0.16.0"
27 |          }
28 |        },
```

```
29     {
30       "node": {
31         "id": "46",
32         "productName": "Visual Studio Code (x64 en-us)",
33         "productVendor": "Microsoft",
34         "productVersion": "1.65.2"
35       }
36     }
37   ]
38 }
39 }
40 }
```

Get software deployment status

The following query retrieves the deployment status of all Deploy packages.

```
1  {
2    softwareDeployment {
3      ID
4      name
5      status {
6        completeCount
7        downloadCompleteWaitingCount
8        downloadingCount
9        failedCount
10       notApplicableCount
11       runningCount
12       waitingCount
13     }
14     errors {
15       error
16       count
17     }
18   }
19 }
```

```
18 | }
19 | }
```

Example response:

```
1 | {
2 |   "data": {
3 |     "softwareDeployment": [
4 |       {
5 |         "ID": "5",
6 |         "name": "Install Tanium Standard Utilities",
7 |         "status": {
8 |           "completeCount": 1,
9 |           "downloadCompleteWaitingCount": 0,
10 |          "downloadingCount": 0,
11 |          "failedCount": 0,
12 |          "notApplicableCount": 0,
13 |          "runningCount": 0,
14 |          "waitingCount": 0
15 |        },
16 |        "errors": null
17 |      },
18 |      {
19 |        "ID": "6",
20 |        "name": "Install Tanium Standard Utilities (Linux)",
21 |        "status": {
22 |          "completeCount": 0,
23 |          "downloadCompleteWaitingCount": 0,
24 |          "downloadingCount": 1,
25 |          "failedCount": 0,
26 |          "notApplicableCount": 0,
27 |          "runningCount": 0,
28 |          "waitingCount": 0
29 |        },
```



```
30 |         "errors": null
31 |     }
32 | ]
33 | }
34 | }
```

Direct Connect examples

The following queries and mutations use Direct Connect to connect to a single endpoint, retrieve data, stop a process, and then close the connection. Queries that retrieve information from endpoints require Performance.

Open a connection to an endpoint

The following mutation uses Direct Connect to establish a connection to the endpoint with an ID of `12323`. You can retrieve IDs through the [Get endpoints IDs from Tanium Data Service on page 45](#) query.



Direct Connect connections close after two minutes of inactivity.

```
1 | mutation {
2 |   openDirectConnection(input: {endpointID: "12323"}) {
3 |     connectionID
4 |   }
5 | }
```

Example response:

```
1 | {
2 |   "data": {
3 |     "openDirectConnection": {
4 |       "connectionID": "86d9a9ac-0229-481b-9d88-5f1bcb1b177b"
5 |     }
6 |   }
7 | }
```

Ping the connection to an endpoint

The following mutation retrieves the status for a Direct Connect connection. Use this mutation to check connection details or to keep the connection active. You need the `connectionID` that is returned by the mutation to open the connection.



Direct Connect connections close after two minutes of inactivity.

```
1 mutation ($connectionID: ID!) {  
2   pingDirectConnection(input: {connectionID: $connectionID}) {  
3     result  
4   }  
5 }
```

Include the connection ID in the **QUERY VARIABLES** panel:

```
1 {  
2   "connectionID": "5fc564d6-5767-47fc-abb6-25cba65409d8"  
3 }
```

Example response:

```
1 {  
2   "data": {  
3     "pingDirectConnection": {  
4       "result": true  
5     }  
6   }  
7 }
```

Get data from an endpoint

After you establish a connection to an endpoint through Direct Connect, you can query the endpoint for specific information. The following query retrieves the CPU usage on the endpoint:

```
1 {
```

```

2 |   directEndpoint (input : {endpointID: "12323"}) {
3 |     performance {
4 |       cpuUsagePercent
5 |     }
6 |   }
7 | }

```

Example response:

```

1 | {
2 |   "data": {
3 |     "directEndpoint": {
4 |       "performance": {
5 |         "cpuUsagePercent": 28.751501243887798
6 |       }
7 |     }
8 |   }
9 | }

```

Get process from an endpoint

After you establish a connection to an endpoint through Direct Connect, you can query the endpoint for process information. The following query retrieves the state of all processes running on the endpoint.

```

1 | {
2 |   directEndpoint (input : {endpointID: "12323"}) {
3 |     processes {
4 |       all {
5 |         pid
6 |         ppid
7 |         name
8 |         commandLine
9 |         userName
10 |        groupName

```

```
11     memoryResidentBytes
12   }
13 }
14 }
15 }
```

Example response:

```
1  {
2    "data": {
3      "directEndpoint": {
4        "processes": {
5          "all": [
6            {
7              "pid": 2092,
8              "ppid": 496,
9              "name": "TaniumReceiver.exe",
10             "commandLine": "\"C:\\Program Files\\Tanium\\Tanium
11             Server\\TaniumReceiver.exe\" --service",
12             "userName": "admin",
13             "groupName": "test-group",
14             "memoryResidentBytes": 59842560
15           },
16           {
17             "pid": 5760,
18             "ppid": 1112,
19             "name": "TaniumClient.exe",
20             "commandLine": "\"C:\\Program Files (x86)\\Tanium\\Tanium
21             Client\\TaniumClient.exe\" -c",
22             "userName": "SYSTEM",
23             "groupName": "NT AUTHORITY",
24             "memoryResidentBytes": 17965056
```

```

25         "pid": 1036,
26         "ppid": 496,
27         "name": "TaniumBlobService.exe",
28         "commandLine": "\"C:\\Program Files\\Tanium\\Tanium Module
Server\\services\\blob-service\\TaniumBlobService.exe\"",
29         "userName": "SYSTEM",
30         "groupName": "NT AUTHORITY",
31         "memoryResidentBytes": 7426048
32     }
33 ]
34 }
35 }
36 }
37 }

```

Get alerts from an endpoint

After you establish a connection to an endpoint through Direct Connect, you can query the endpoint for alert information. The following query retrieves alerts from an endpoint.

```

1  {
2    directEndpoint (input : {endpointID: "12323"}) {
3      alerts {
4        all {
5          schema
6          key
7          type
8          ref
9          topProcessesExpr
10         labels
11         pendingAt
12         start
13         resolvedAt
14         leadup

```

```
15     value
16   }
17 }
18 }
19 }
```

Example response:

```
1  {
2    "data": {
3      "directEndpoint": {
4        "alerts": {
5          "all": [
6            {
7              "schema": 1,
8              "key": "available-mem{heuristic=\"available-mem\"}",
9              "type": "available-mem",
10             "ref": null,
11             "topProcessesExpr": null,
12             "labels": {
13               "heuristic": "available-mem"
14             },
15             "pendingAt": "2022-03-15T15:54:38.574990164Z",
16             "start": "2022-03-15T15:54:38.574990164Z",
17             "resolvedAt": null,
18             "leadup": 30000000000,
19             "value": 168.48828125
20           }
21         ]
22       }
23     }
24   }
25 }
```

Stop a process on an endpoint

After you establish a connection to an endpoint through Direct Connect, you can stop running processes on the endpoint. The following mutation stops a process named `notepad.exe` on an endpoint. You need the `connectionID` that is returned by the mutation to open the connection.

```
1 mutation {
2   killProcess(
3     input: {connectionID: "7212763a-20aa-4cdd-a8b2-6b20b3968f2a", name: "notepad.exe", pid:
4       7056, signal: SIGKILL}
5   ) {
6     result
7   }
}
```

Example response:

```
1 {
2   "data": {
3     "killProcess": {
4       "result": true
5     }
6   }
7 }
```

Close connection to an endpoint

The following mutation closes a Direct Connect connection to an endpoint. You need the `connectionID` that is returned by the mutation to open the connection.



Direct Connect connections close after two minutes of inactivity.

```
1 mutation ($connectionID: ID!) {
2   closeDirectConnection(input: {connectionID: $connectionID}) {
3     result
4   }
}
```

```
5 | }
```

Include the connection ID in the **QUERY VARIABLES** panel:

```
1 | {  
2 |   "connectionID": "5fc564d6-5767-47fc-abb6-25cba65409d8"  
3 | }
```

Example response:

```
1 | {  
2 |   "data": {  
3 |     "closeDirectConnection": {  
4 |       "result": true  
5 |     }  
6 |   }  
7 | }
```