



TRIFACTA

Admin Guide

Version: 9.2

Doc Build Date: 07/29/2022

Copyright © Trifacta Inc. 2022 - All Rights Reserved. CONFIDENTIAL

These materials (the “Documentation”) are the confidential and proprietary information of Trifacta Inc. and may not be reproduced, modified, or distributed without the prior written permission of Trifacta Inc.

EXCEPT AS OTHERWISE PROVIDED IN AN EXPRESS WRITTEN AGREEMENT, TRIFACTA INC. PROVIDES THIS DOCUMENTATION AS-IS AND WITHOUT WARRANTY AND TRIFACTA INC. DISCLAIMS ALL EXPRESS AND IMPLIED WARRANTIES TO THE EXTENT PERMITTED, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE AND UNDER NO CIRCUMSTANCES WILL TRIFACTA INC. BE LIABLE FOR ANY AMOUNT GREATER THAN ONE HUNDRED DOLLARS (\$100) BASED ON ANY USE OF THE DOCUMENTATION.

For third-party license information, please select **About Trifacta** from the Help menu.

1. Admin	4
1.1 Admin Tasks	5
1.1.1 Operations Tasks	6
1.1.1.1 Configure Running Environments	7
1.1.1.2 Verify Operations	9
1.1.2 Access Management Tasks	11
1.1.2.1 Enable Access to S3 and AWS Resources	12
1.1.2.2 Insert Trust Relationship in AWS IAM Role	16
1.1.3 User Admin Tasks	18
1.1.3.1 Create User Account	19
1.1.3.2 Configure Password Criteria	21
1.1.3.3 Create Admin Account	23
1.1.3.4 Manage Users under SSO	25
1.1.3.5 Invite Users	27
1.1.3.6 Create Role	30
1.1.4 Application Management Tasks	35
1.1.4.1 Manage Downloading	36
1.1.4.2 Manage Schedules	38
1.2 System Services and Logs	39
1.3 Storage Maintenance	48
1.4 Backup and Recovery	51
1.4.1 Platform Rollback	57
1.5 Admin Reference	64
1.5.1 Deployment Manager Page	65
1.5.2 Admin Console	68
1.5.2.1 Users Page	70
1.5.2.1.1 User Details Page	73
1.5.2.2 Roles Page	75
1.5.2.2.1 Create Role Dialog	77
1.5.2.2.2 Role Details Page	78
1.5.2.3 Workspace Settings Page	80
1.5.2.4 Admin Settings Page	94
1.5.2.5 AWS Settings Page	98
1.5.2.6 Environment Parameters Page	101
1.5.2.7 OAuth 2.0 Clients Page	103
1.5.3 Workspace Admin Permissions	104
1.5.4 Admin Download Logs Dialog	106
1.5.5 Required AWS Account Permissions	109
1.5.6 Privileges and Roles Reference	114

Admin

After initial deployment, most admin work for Trifacta® can be managed through the Trifacta application, including user provisioning, data access, and project or workspace configuration settings.

Before you begin:

Before you begin using the Trifacta application, please verify that the product has been properly deployed, including any required prerequisites.

NOTE: Access to the datastores and running environments within your enterprise infrastructure may require configuration outside of Trifacta.

If basic connectivity and integration has been completed, you can perform a few simple checks to verify that data can be loaded, transformed, and published. For more information, see *Verify Operations*.

Admin Tasks

This section covers administrative tasks for configuring storage and running environments, setting up the basic parameters and preferences, creating IAM roles, registering users, and assigning levels of privileges.

Operations Tasks

This section contains administrative tasks related to the operations of Trifacta®, including job execution within your enterprise ecosystem.

Configure Running Environments

Contents:

- *Trifacta Photon*
 - *EMR*
 - *Snowflake*
 - *Other Running Environments*
-

This section provides overview information on how to configure the running environments accessible from your deployment of the Trifacta application.

A **running environment** is the set of services that are used to execute a job.

- A job can include tasks to do the following:
 - Ingest data
 - Transform data
 - Profile data
 - Sample data
 - Generate results
- A running environment can be hosted on the Trifacta node or across a cluster that is connected to the product.

Trifacta Photon

Hosted on the Trifacta node, Trifacta Photon is an in-memory running environment designed for high performance on small- to medium-sized jobs.

Configuration:

Trifacta Photon may require enablement in your project or workspace:

- For more information, see *Workspace Settings Page*.

EMR

Amazon Elastic Map Reduce (EMR) is a managed-cluster data platform for processing large volumes of disparate sources of data. This scalable platform is used for running jobs from Trifacta and can handle data processing tasks of any size.

Configuration:

- The Trifacta application must be connected to an EMR cluster. For more information, see *Configure for EMR*.
- If you are accessing AWS resources using IAM roles, those roles must contain policies to run jobs on EMR. For more information, see *Required AWS Account Permissions*.

Snowflake

Snowflake provides cloud-based data storage and analytics as a service. If all of your source datasets and outputs are in Snowflake locations and other conditions are met, then the entire execution of the transformations can occur in Snowflake. For more information, see <https://www.snowflake.com>.

For datasets and outputs that are hosted in Snowflake, you can configure the Trifacta application to perform the transformation steps of your job in Snowflake. In this manner, no data needs to be transferred to and from the data warehouse, and performance should be significantly better.

Tip: For jobs that are executed in Snowflake, you can optionally enable the execution of the visual profile in Snowflake, too. Some limitations may apply. This option is enabled for individual flows. For more information, see *Flow Optimization Settings Dialog*.

Limitations:

NOTE: Snowflake is not a running environment that you explicitly select or specify as part of a job. If all of the requirements are met, then the job is executed in Snowflake. For more information on limitations, see *Overview of Job Execution*.

Configuration:

For more information, see *Snowflake Running Environment*.

Other Running Environments

Depending on your deployment of Trifacta, additional running environments may be available. For more information, see *Running Environment Options*.

Verify Operations

Contents:

- *Prepare Your Sample Dataset*
 - *Store Your Dataset*
 - *Verification Steps*
-

After you have applied a configuration change to the platform and restarted, you can use the following steps to verify that Trifacta® is working correctly.

If your configuration change was applied to `trifacta-conf.json`, you should restart the platform before continuing. See *Start and Stop the Platform*.

Prepare Your Sample Dataset

To complete this test, you should locate or create a simple dataset. Your dataset should be created in the format that you wish to test.

Tip: The simplest way to test is to create a two-column CSV file with at least 25 non-empty rows of data. This data can be uploaded through the application.

Characteristics:

- Two or more columns.
- If there are specific data types that you would like to test, please be sure to include them in the dataset.
- A minimum of 25 rows is required for best results of type inference.
- Ideally, your dataset is a single file or sheet.

Store Your Dataset

If you are testing an integration, you should store your dataset in the datastore with which the product is integrated.

Tip: Uploading datasets is always available as a means of importing datasets.

- You may need to create a connection between the platform and the datastore.
- Read and write permissions must be enabled for the connecting user to the datastore.
- For more information, see *Connections Page*.

Verification Steps

Steps:

1. Login to the application.
2. In the application menu bar, click **Library**.
3. Click **Import Data**. See *Import Data Page*.
 - a. Select the connection where the dataset is stored. For datasets stored on your local desktop, click **Upload**.
 - b. Select the dataset.
 - c. Click **Continue**.

4. The initial sample of the dataset is opened in the Transformer page, where you can edit your recipe to transform the dataset.
 - a. In the Transformer page, some steps are automatically added to the recipe for you. So, you can run the job immediately.
 - b. You can add additional steps if desired. See *Transformer Page*.
5. Click **Run**.
 - a. If options are presented, select the defaults.
 - b. To generate results in other formats or output locations, click **Add Publishing Destination**. Configure the output formats and locations.
 - c. To test dataset profiling, click the Profile Results checkbox. Note that profiling runs as a separate job and may take considerably longer.
 - d. See *Run Job Page*.
6. When the job completes, you should see a success message under the Jobs tab in the Flow View page.
 - a. **Troubleshooting:** Either the Transform job or the Profiling job may break. To localize the problem, try re-running a job by deselecting the broken job type or running the job on a different running environment (if available). You can also download the log files to try to identify the problem. See *Job Details Page*.
7. Click **View Results** from the context menu for the job listing. In the Job Details page, you can see a visual profile of the generated results. See *Job Details Page*.
8. In the Output Destinations tab, click a link to download the results to your local desktop.
9. Load these results into a local application to verify that the content looks ok.

Checkpoint: You have verified importing from the selected datastore and transforming a dataset. If your job was successfully executed, you have verified that the product is connected to the job running environment and can write results to the defined output location. Optionally, you may have tested profiling of job results. If all of the above tasks completed, the product is operational end-to-end.

Access Management Tasks

This section contains admin tasks to configure and maintain access to your enterprise datastores from Trifacta®.

Enable Access to S3 and AWS Resources

Contents:

- *AWS Overview*
- *Technical Setup*
 - *Create policy to grant access to S3 bucket*
 - *Update policy to accommodate SSE-KMS if necessary*
 - *Add policy for Redshift access*
 - *Whitelist the IP address range of the Trifacta Service, if necessary*

If you plan to use S3 as the default storage environment, the following sections outline the AWS configuration prerequisites and requirements.

Tip: This section should be shared with your S3 administrator, who can provide the required information.

AWS Overview

Below are the AWS objects that are required for S3 setup.

AWS object	Required?	Description
AWS account	Y	To create these objects are part of the setup process, you must have an AWS account. For more information, see https://aws.amazon.com/ .
Valid email address	Y	To validate your registration for a new workspace, you must have a valid email address to which the product can deliver the registration email.
Choice: cross-account role access or key-secret access	Y	<p>To integrate with your existing S3 resources, you must choose a method of authentication. Choices:</p> <ul style="list-style-type: none">• cross-account role: This method uses IAM roles to define the permissions used by the product for S3 access. <div>Tip: This method is recommended.</div> <ul style="list-style-type: none">• key-secret access: This method uses an IAM access keys to provide S3 access.
IAM policy	Y	<p>An IAM (Identity and Access Management) policy is an AWS resource used to define the low-level permissions for access to a specific resource. You can use an IAM policy for the product to use for either access method.</p> <p>For more information, see "Create policy to grant access to S3 bucket" below.</p>
cross-account role access: IAM role	Y	An IAM role contains one or more IAM policies that can be used to define the set of available AWS services and the level of access to them for a user. In this case, the user is the Trifacta application.
key-secret access: AWS key-secret	Y	An older AWS access method, the key-secret combination is essentially a username and password combination to one or more S3 buckets.
S3 bucket	Y	S3 (Simplified Storage Service) is a cloud-based file storage system hosted in AWS. An S3 bucket contains your data files and their organizing folders.

S3 bucket: encryption	N	<p>For better security, your S3 bucket may be encrypted, which means that the data is stored inside of S3 in a way that is not human-readable.</p> <p>NOTE: The product can optionally integrate with encrypted S3 buckets. The following S3 encryption methods are supported: sse-s3 and sse-kms.</p> <p>NOTE: If your bucket is encrypted with ss3-kms, additional configuration is required. See "Update policy to accommodate SSE-KMS if necessary" below.</p> <p>For more information on your bucket's encryption, please contact your S3 administrator.</p>
S3 bucket: storage location	N	<p>If needed, you can change the location where results are stored in S3.</p> <p>NOTE: The product must have write permission to this location. If you are changing the location from the default, please verify with your S3 administrator that the preferred location is enabled for writing through your access method.</p>
IAM role: Account ID	N	<p>The account ID identifies in the trust policy that Trifacta AWS account can use your IAM role.</p> <p>Tip: This identifier is provided to you during registration and setup.</p>
IAM role: External ID	N	<p>The external ID identifies in the trust policy that Trifacta can use your IAM role only on your behalf.</p> <p>Tip: This identifier is provided to you during registration and setup.</p>

Technical Setup

The following sections should be provided to your AWS administrator for setting up access to these resources, if required.

Create policy to grant access to S3 bucket

To use your own S3 bucket(s) with Trifacta, create a policy and assign it to either the user or IAM Role selected to grant access to AWS resources. In this section, you create the policy. Later, it will be applied.

- For more information on creating policies, see <https://console.aws.amazon.com/iam/home#/policies>.

Below is an example policy template. You should use this template to create the policy.

NOTE: You should not simply use one of the predefined AWS policies or an existing policy you have as it will likely give access to more resources than required.

Template Notes:

1. One of the statements grants access to the public demo asset buckets.
2. Replace `<my_default_s3_bucket>` with the name of your default S3 bucket.
3. To grant access to multiple buckets within your account, you can extend the resources list to accommodate the additional buckets.

Policy Template

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:DeleteObject",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::<my_default_S3_bucket>",
        "arn:aws:s3:::<my_default_S3_bucket>/*"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::aws-saas-samples-prod",
        "arn:aws:s3:::aws-saas-samples-prod/*",
        "arn:aws:s3:::aws-saas-datasets",
        "arn:aws:s3:::aws-saas-datasets/*",
        "arn:aws:s3:::3fac-data-public",
        "arn:aws:s3:::3fac-data-public/*",
        "arn:aws:s3:::trifacta-public-datasets",
        "arn:aws:s3:::trifacta-public-datasets/*"
      ]
    }
  ]
}
```

Update policy to accommodate SSE-KMS if necessary

If any accessible bucket is encrypted with SSE-KMS, another policy must be deployed. See <https://docs.aws.amazon.com/kms/latest/developerguide/iam-policies.html>.

Add policy for Redshift access

If you are connecting to Redshift databases through your workspace, you can enable access by creating a `GetClusterCredentials` policy. This policy is additive to the the S3 access policies. All of these policies can be captured in a single IAM role.

Example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GetClusterCredsStatement",
      "Effect": "Allow",
      "Action": [
        "redshift:GetClusterCredentials"
      ],
      "Resource": [
        "arn:aws:redshift:us-west-2:123456789012:dbuser:examplecluster/${redshift:DbUser}",
        "arn:aws:redshift:us-west-2:123456789012:dbname:examplecluster/testdb",
        "arn:aws:redshift:us-west-2:123456789012:dbgroup:examplecluster/common_group"
      ],
      "Condition": {
        "StringEquals": {
          "aws:userid": "AIDIODR4TAW7CSEXAMPLE:${redshift:DbUser}@yourdomain.com"
        }
      }
    }
  ]
}
```

For more information on these permissions, see *Required AWS Account Permissions*.

Whitelist the IP address range of the Trifacta Service, if necessary

If you are enabling any relational source, including Redshift, you must whitelist the IP address range of the Trifacta service in the relevant security groups.

NOTE: The database to which you are connecting must be available from the Trifacta service over the public Internet.

The IP address range of the Trifacta service is:

```
35.245.35.240/28
```

For Redshift:

For Redshift, there are two ways to whitelist the IP range depending on if you are using EC2-VPC or EC2-Classic (not common).

- **EC2-VPC (Security group):** Add the IP address range to the inbound rule for the security group associated with the cluster. For more information, see <https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-authorize-cluster-access.html#rs-gsg-how-to-authorize-access-vpc-security-group>
- **EC2-Classic:** Add the IP address range to the inbound rule for the security group associated with the EC2 instance. For more information, see <https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-authorize-cluster-access.html#rs-gsg-how-to-authorize-access-cluster-security-group>

For details on this process with RDS in general, see <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html>

For more information, please contact *Alteryx Support*.

Insert Trust Relationship in AWS IAM Role

If you are using per-user authentication through an AWS IAM role, you must insert a trust relationship into the role so that Trifacta® can leverage it.

Prerequisites:

NOTE: These steps should be performed by an AWS administrator.

Please acquire the following information:

- **IAM role:** The AWS IAM role that Trifacta should use.
- **EC2 instance role:** If the EC2 instance role is to be used to assume the AWS role, then please acquire the following:
 - AWS account ID
 - EC2 instance role
 - Details on the above are listed below.

For more information on the AWS Principal options described below, please review https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_principal.html.

Steps:

1. You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.
2. Locate the following parameter and retrieve its value (`true` or `false`):

```
"aws.ec2InstanceRoleForAssumeRole"
```

3. Login to the AWS console.
4. Open the IAM role for use with Trifacta.
5. If `aws.ec2InstanceRoleForAssumeRole=true`, then the EC2 instance role is used for assuming the provided AWS role. Paste the following into the IAM role for the trust relationship:

Property	Description
<awsAccountId>	AWS account identifier for which the EC2 instance role is assumed
<ec2InstanceRole>	EC2 instance role to use

6. If `aws.ec2InstanceRoleForAssumeRole=false`, then the AWS user associated with the provided AWS key and secret is assumed. Paste the following into the IAM role for the trust relationship:


```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::862753480162:user/sample-user"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

7. Save the IAM role definition.

User Admin Tasks

This section contains admin tasks on provisioning and maintaining users, their roles, and their permissions within Trifacta®.

Create User Account

Contents:

- *Creating your own user account*
 - *Creating users when self-registration is disabled*
 - *Troubleshooting*
 - *Account Not Configured login error*
-

By default, users can create their own accounts. As needed, self-registration can be disabled, so that all users must be created by an administrator. See *Configure User Self-Registration*.

Creating your own user account

Steps:

1. Users may self-register at the following address:
`http://<host_name>:<port_number>`
where:
`<host_name>` is the host of the Trifacta® application.
`<port_number>` is the port number to use. Default is 3005.
2. Click the Register link.
3. Enter your credentials in the spaces provided. A valid email address is required.
4. As soon as the account is created, you may login at the first address. See *Login*.

Creating users when self-registration is disabled

When self-registration is disabled, an administrator must manually create the accounts for users. Administrators can create accounts at the following address:

`http://<host_name>:<port_number>/register`

NOTE: If SSO or secure impersonation is enabled in your environment, administrators must apply a principal value to each newly created user. See *Users Page*.

When a new account is created, an email is sent to the address for the created user.

Troubleshooting

Account Not Configured login error

If you have created a user account, you may see the following error message when you try to login:



Account Not Configured

Your Trifacta user account has not been completely configured.

[Configure storage settings](#) or contact your Trifacta Administrator.

[Return to Sign In](#)

Figure: Account Not Configured

In this case, the account may require additional configuration. In SSO or Kerberos environments, an administrator may need to provision a SSO or Hadoop principal value to be applied to the user account. See *Admin Settings Page*

Configure Password Criteria

By default, the Trifacta® application enforces very few requirements on password length, capitalization, or special characters. Users who are setting or resetting their passwords are permitted to create a password of one character in length with no additional requirements.

NOTE: When passwords are set or reset, the platform does perform an assessment of the quality of the password and reports it to the user before saving. For more information, see *User Profile Page*.

Before you permit users to create accounts, you should review the password requirements for your enterprise and, where needed, apply them to the Trifacta application.

Enable

To enable enforcement of password criteria, please enable the following parameter.

Steps:

You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.

Locate the following parameter and set it to `true`:

```
"feature.enablePasswordCriteria": true,
```

When enabled, submitted changes to user passwords are evaluated based on the configuration settings defined below.

Configure

The following parameters govern the password criteria enforced by the Trifacta application when the feature is enabled.

Parameter	Description	Default
webapp.passwordCriteria.length.min	Minimum length of a password to Trifacta application	0
webapp.passwordCriteria.length.max	Maximum length of a password to Trifacta application	100
webapp.passwordCriteria.description	Text describing the criteria that a password must meet. Specify this value last.	
webapp.passwordCriteria.contains.uppercase	Defines whether the password must contain uppercase characters	undefined
webapp.passwordCriteria.contains.symbols	Defines whether the password must contain symbols	undefined
webapp.passwordCriteria.contains.spaces	Defines whether the password must contain space characters	undefined
webapp.passwordCriteria.contains.lowercase	Defines whether the password must contain lowercase characters	undefined
webapp.passwordCriteria.contains.letters	Defines whether the password must contain letters (a-z)	undefined
webapp.passwordCriteria.contains.digits	Defines whether the password must contain digits (0-9)	undefined

Criteria settings:

Some of the criteria settings support the following options:

Setting	Description
enforce	Each password must pass this requirement.
forbid	Passwords cannot have this requirement.
undefined	(default) This requirement is disabled. Users may choose to include or not include this requirement in their passwords.

Create Admin Account

Contents:

- *Default admin account*
- *Create admin accounts*
- *Create admin account outside the UI*
 - *Without SSO*
 - *With SSO*

You can create additional administrator accounts to the Trifacta® platform using one of the following methods.

Default admin account

When the Trifacta platform is installed, a default admin account is created for you. For licensing purposes, this account is counted as a valid user.

The password for the default admin account should be changed as soon as you have access to the application. See *Change Admin Password*.

NOTE: Do not delete the default admin account. To ensure that there is always one admin account that is accessible, this account is automatically recreated if you delete it.

NOTE: Since this account cannot be mapped to a valid email address within a customer domain, it cannot be used in an SSO environment.

Create admin accounts

Steps:

1. Login using another admin account.
2. Create the account normally. See *Create User Account*.
3. Select **User menu > Admin console > Users**.
4. For the newly created user, select **Edit** from the user's context menu.
5. Admin roles:
 - a. To enable administration of workspace users, roles and other settings, select **Workspace admin** from the Roles drop-down.
 - b. To enable administration of platform settings, click the **Platform admin** checkbox.
6. Save changes.
7. Login to the account and verify that the Admin console pages are available.

Create admin account outside the UI

If you do not have access to an admin account through the application, you can create admin accounts for users from the Trifacta node using the `webapp/bin/ensure-user` command.

Without SSO

If Single Sign-On (SSO) is not enabled, use the following command:

```
<install_dir>/webapp/bin/ensure-user --admin "<FirstName LastName>" <e-mail> <password>
```

With SSO

If the environment uses SSO, the following command can create the admin user based on an Active Directory login:

```
<install_dir>/webapp/bin/ensure-user --admin "<FirstName LastName>" <e-mail> <password> <AD_LOGIN>
```

where:

<AD_LOGIN> is the active directory login for the user.

Manage Users under SSO

Contents:

- *Enable SSO*
- *Configure Auto-Registration*
 - *User Management with Auto-Registration*
 - *Disable Auto-Registration*
 - *Provision new users under SSO without auto-registration*
 - *User access for reverse proxy method*

This section covers additional requirements for managing users of the Trifacta® platform in SSO environments.

Enable SSO

The Trifacta platform requires additional configuration to integrate with your SSO provider. Available methods:

Method	Description
SAML IDP	Integrate the platform with enterprise SAML identity provider. See <i>Configure SSO for SAML</i> .
Native LDAP-AD	Using native functionality in the platform, it can integrate with enterprise LDAP/AD. For more information, see <i>Configure SSO for AD-LDAP</i> .
LDAP-AD via reverse proxy	<div>A reverse proxy server outside of the platform can be set up for integration with enterprise LDAP/AD.</div> <div>NOTE: This method is likely to be deprecated in a future release.</div> <div>For more information, see <i>Configure SSO for AD-LDAP</i>.</div>

Configure Auto-Registration

Tip: By default, user auto-registration is enabled. It is recommended.

How users are managed depends on whether auto-registration is enabled:

- If auto-registration is enabled, after users provide their credentials, the account is automatically created for them.
- If auto-registration is disabled, a Trifacta administrator must still provision a user account before it is available. See below.

User Management with Auto-Registration

After SSO with auto-registration has been enabled, you can still manage users through the Trifacta application, with the following provisions:

- The Trifacta platform does not recheck for attribute values on each login. If attribute values change with your identity provider, they must be updated in the configuration.
 - For more information, see *Configure SSO for AD-LDAP*
 - For more information, see *Configure SSO for SAML*.
- If the user has been removed from AD, the user cannot sign in to the platform.

- If you need to remove a user from the platform, you should just disable the user through the Trifacta application.
 - If the user is deleted, then if the user returns to the platform in the future, a new account is created for the user.

For more information, See *Users Page*.

Disable Auto-Registration

To disable auto-provisioning in the platform, please verify the following property:

1. You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.
2. Set the following property:

```
"webapp.sso.enableAutoRegistration" : false,
```

3. Save your changes and restart the platform.
4. New users of the Trifacta platform must be provisioned by a Trifacta administrator. See below.

Provision new users under SSO without auto-registration

If SSO auto-registration is disabled, admin users can provision new users of the platform through the following URL:

```
https://<hostname>:<sso_port_number>/register
```

where:

- `<hostname>` is the host of the Trifacta platform
- `<sso_port_number>` is the port number.

The user's password is unnecessary in an SSO environment. You must provide the SSO principal value, which is typically the Active Directory login for the user.

- If you are connected to a Hadoop cluster, you must provision the Hadoop principal value.
- See *Create User Account*.

User access for reverse proxy method

Users access the application through the Trifacta node using the standard hostname and the port that you specified:

NOTE: All users must use this URL to access the Trifacta application. If they use the non-SSO URL, they may receive an Unprovisioned User error.

```
https://<hostname>:<sso_port_number>
```

Invite Users

Contents:

- *Invite User*
 - *Edit User*
 - *Edit roles*
 - *Change workspace admin*
 - *Assign roles*
 - *Disable User*
 - *Remove User*
-

Administrators can manage the users who are permitted to use Trifacta®.

All of these functions are available through the Admin console. For more information, see *Admin Console*.

Invite User

To permit a user to access the Trifacta application, an administrator must complete the following steps.

NOTE: When a user accepts your invitation, the additional user counts toward the maximum number of permitted users.

NOTE: If you are re-inviting a user who has been removed, you must wait 14 days to invite the user back to the same project or workspace and retain the user's data. If restoring the user's flows and recipes is not important, please contact *Alteryx Support* for immediate re-instatement.

Steps:

1. Login to the Trifacta application as an administrator.
2. From the left navigation bar, select **User menu > Admin console > Users**.
3. In the Users page, click **Invite users**.
4. In the Invite users dialog, enter a comma-separated list of email addresses to which to send invites.
 - a. These addresses become the user identifier for logging into the Trifacta application.
 - b. Avoid sending invites to email aliases.
 - c. Example:

```
joe.smith@example.com, mary.jones@example.com
```

5. To invite the list of users, click **Invite users**.
6. An email is sent to each valid user email address that you listed. The receiving user must click the link in the email to accept the invitation.

The user is invited via email and created in the Trifacta application. You can modify the user account as needed before the user chooses to log in. See below.

For more information, see *Users Page*.

Edit User

Edit roles

Steps:

1. In the Users page, locate the user to review.
2. On the right side of the row for the user, click the Actions menu.
3. Select **Edit user**.
4. In the dialog, you can add and remove roles for the user account.
5. When finished, click **Edit User**.

Change workspace admin

By default, a new user account is assigned a non-admin role. If needed, you can assign the user to be a workspace admin.

Use the following steps to change a user's workspace role between non-admin and admin.

Steps:

1. In the Users page, locate the user whom you are promoting to admin.
2. On the right side of the row for the user, click the Actions menu.
3. Select **Change admin role**.
4. In the Change admin role dialog, select the workspace role:
 - a. **Member** - standard user account, which is not permitted access to the Admin console and its functions.
 - b. **Admin** - administrator account, which can access all available features of the workspace.

You should avoid assigning the admin role to a large number of users.

5. Click **Save**.
6. The user's workspace role is immediately updated.

Assign roles

When the account is created, it is automatically assigned the `Default` role. You should review the permissions associated with this role and to determine if the user needs to be assigned a different one. For more information, see *Roles Page*.

Disable User

NOTE: Disabled users still count toward workspace limits on number of users.

If needed, a user's account can be disabled from accessing Trifacta. When a user account is disabled:

- The user can no longer log in to the Trifacta application or use any available API endpoints.
- The user's assets in Trifacta application are retained. They can be accessed by other users who have been granted permission.

To disable a user, please complete the following steps:

Steps:

1. In the Users page, locate the user to disable.

2. On the right side of the row for the user, click the Actions menu.
 - a. To disable **Disable**. Click **Disable** to confirm.
 - b. To reactivate a disable member, click **Enable**.
3. Effective immediately, the user cannot log in to the application.

Remove User

To remove a user completely, please complete the following steps.

When a user is removed from Trifacta, any assets that are owned by the user must be reassigned to other users, or they are lost and no longer accessible even by an administrator.

Steps:

1. In the Users page, locate the user to remove.
2. On the right side of the row for the user, click the Actions menu.
3. Select **Remove**.
4. If the user owns assets, you can choose to assign them to another user. If you do not assign them, these assets are lost.
5. Confirm that you wish to remove the user.

If you must recover a removed user or that user's assets, please contact *Alteryx Support* within 14 days of the deletion.

Create Role

Contents:

- *Create Role*
 - *Example - Read-only access role*
 - *Example - Flows-only access role*
 - *Example - Empty role*
 - *Assign Role*
 - *Modify Role*
 - *Example - Modify default role*
 - *Unassign Role*
 - *Delete Role*
-

Administrators can create and assign roles to users to govern access to user-created objects in Trifacta®.

- A **role** is a set of privileges that can be assigned to users.
- A **privilege** governs access level to a type of object.
- By default, all users are assigned the `default` role, which allows users to use the user-created object types.
- For more information, see *Privileges and Roles Reference*.

As needed, you can create user roles to define different access levels for different object types.

NOTE: You must be an administrator to create new roles.

NOTE: Roles are additive. If you assign multiple roles to a user account, the user receives the highest level of access for each privilege among the assigned roles.

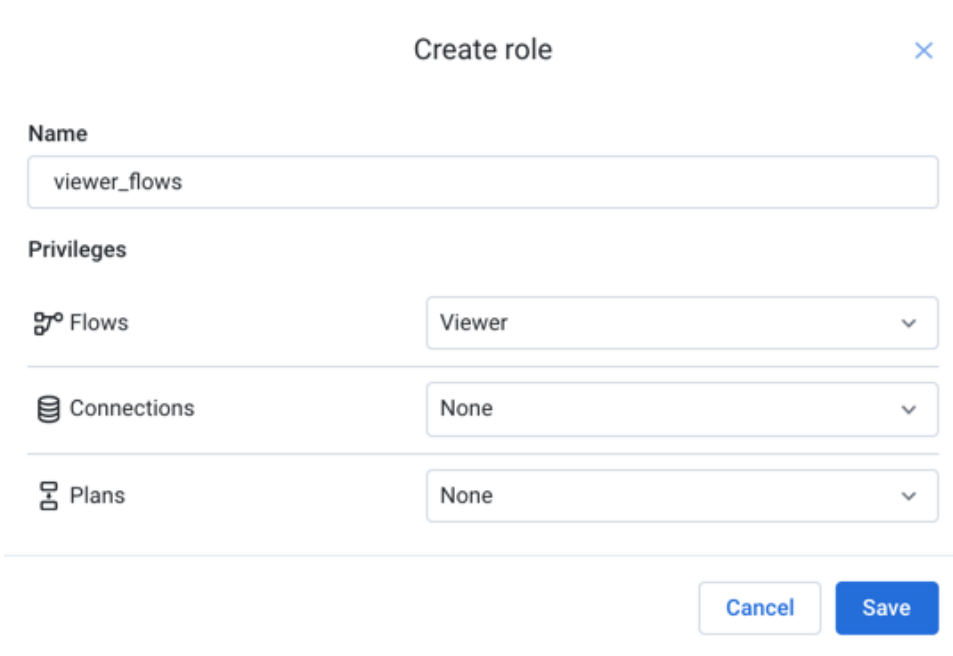
NOTE: When a role is assigned, unassigned, or modified, the changes to privileges are immediately applied to the associated user accounts. A new login is not required.

Create Role

To create a new role, please complete the following steps.

Steps:

1. In the left nav bar, select **User menu > Admin console > Roles**.
2. In the Roles page, review the list of available roles. For more information, see *Roles Page*.
3. To create a new role, click **Create Role**.





Create role


Name

viewer_flows

Privileges

 Flows Viewer

 Connections None

 Plans None

Cancel Save

Figure: Create Role dialog

4. In the Create Role dialog, specify the following:
 - a. **Name:** Enter a name for your role. This value must be unique among available roles.
 - b. **Privileges:**
 - i. For each of the available object types, specify the access level for the role.
 - ii. For more information on these privileges, see *Privileges and Roles Reference*.
 - c. To create the role, click **Save**.
5. The role is now available and can be assigned to users. See below.

For more information, see *Create Role Dialog*.

Example - Read-only access role

Suppose you wish to limit a set of users to read-only access to role-based objects.

Steps:

1. In the Roles page, click **Create role**.
2. In the Create Role dialog, enter the following:
 - a. Name: read-only
 - b. Privileges: For each available privilege, select viewer.

NOTE: Some privileges may not have a viewer access level. For these privileges, you should select none. However, users with such a role cannot access the pages where these objects are listed.

3. Click **Save**.
4. The role is now available and can be assigned to users. See below.

Example - Flows-only access role

Suppose you wish to limit a set of users to only be able to work with flows. These users should be able to view, share, edit, schedule, run jobs, and delete flows.

Steps:

1. In the Roles page, click **Create role**.
2. In the Create Role dialog, enter the following:
 - a. Name: `flows-only`
 - b. Privileges:
 - c. For the flows privilege, select `author`.
 - d. For every other privilege, select `none`.
3. Click **Save**.
4. The role is now available and can be assigned to users. See below.

Example - Empty role

In some circumstances, you may wish to assign an empty role to a user. For example, you may wish to limit some administrators to only be able to configure aspects of the platform without providing access to any user-created objects.

Steps:

1. In the Roles page, click **Create role**.
2. In the Create Role dialog, enter the following:
 - a. Name: `empty`
 - b. Privileges:
 - c. For every privilege, select `none`.
3. Click **Save**.
4. The role is now available and can be assigned to users. See below.

Assign Role

After a role has been created, you can assign it to users.

NOTE: Assigning a role adds the role to the user's account. It does not replace any role that is already present in the account.

NOTE: When you assign or unassign a role, the privileges are immediately applied to the assigned user's account. The user does not need to re-login to see the changes.

Steps:

1. In the Roles page, locate the role to assign.
2. On the right side of the screen, click the context menu for the role. Select **Assign role....**
3. In the Assign role dialog, enter a list of email addresses for users to whom you wish to assign the role.
4. Click **Assign**.
5. The role and its associated privileges are applied immediately to the user account(s).

Modify Role

After a role has been created, you can modify it as needed.

NOTE: When the privileges of a role are modified, the changes are applied immediately to all users who are currently assigned the role. Before making modifications, you should review the users who could be affected. See *Role Details Page*.

Steps:

1. In the Roles page, locate the role to modify.
2. In the context menu on the right side of the page, select **Edit**.
3. Review the privileges assigned to the role, and make any changes as necessary.
4. Click **Save**.
5. All users who currently have the role in their account immediately receive the changed privileges.

Example - Modify default role

Tip: If you are changing the privileges of a role, you might want to create a role that contains only the replaced privileges. For example, if you are changing access to flows in Role A from `author` to `editor`, you might create a new role first, which contains only the `author` privilege for flows. If a user needs to be able to create new flows, you can then assign the new role accordingly.

Suppose you wish to reduce privileges for the `default` role, which is assigned to all users. At the same time, some user should be assigned author-level access to the available objects. Here is the following general flow for managing this modification.

NOTE: You cannot modify the name of the `default` role.

Steps:

1. Before you begin, you might wish to inform users that you are making these changes. In some cases, users may lose access to objects that they have created.
2. Create new roles for author access to each object type. For more information, see *Roles Page*.
 - a. For example, you can create the `Flow Author` role, which has `author` privilege for flows and no other privilege. Optionally, for the other privileges, you could provide `viewer` access, which enables read-only access.
 - b. Repeat the above for each type of object for which there is a privilege.
 - c. At this point, the new roles have been created.
3. Assign these roles to users as needed. For example, for the `Flow Author` role, you can assign it to each user that must create flows.

Tip: Since roles are additive, you have not removed any privileges yet.

4. Now, you can modify the `default` role.
 - a. In this case, you should decide what is the baseline set of privileges that each new user should have. Set the privileges to the lowest level of common access.

Unassign Role

Use the following steps to remove a role from a user account.

NOTE: Removing a role from a user account may remove access to objects that the user has created. If the user is the owner of these objects, some access may be removed permanently, even if the object is shared. For more information, see *Overview of Sharing*.

Steps:

1. In the Roles page, locate the role that you wish to remove from one or more user accounts.
2. Select the role.
3. In the Role Details page, click the Users tab.
4. Locate the user to un-assign the role. In the context menu for the user, select **Unassign from role**.
5. The user no longer has the role in the account.

For more information, see *Role Details Page*.

Delete Role

You are permitted to delete roles that are still assigned to users. Deleting a role removes the role from all user accounts and cannot be undone. Before you delete a role, you should review the list of affected users and the objects to which they have access.

Steps:

1. In the Roles page, locate the role to delete. In the context menu, select **Delete**.
2. Confirm the deletion.
3. The role is deleted. All users who had the role can no longer access the privileges assigned in the role.

See *Roles Page*.

Application Management Tasks

This section describes how to manage specific aspects of the Trifacta® application.

Manage Downloading

Contents:

- *Job Results*
 - *Samples*
 - *Flows and Plans*
 - *Imported Datasets*
 - *Dataset previews*
-

For security reasons, you may need to apply controls to the ability of users to download data from the Trifacta® platform. You can choose the types of downloading you limit using the settings in the sections listed below.

Job Results

The Trifacta application allows users to download their job results up to a pre-defined limit. If you set to this limit 0, job results cannot be downloaded at all from the application. When downloading job results is disabled, all job results must be downloaded outside of the Trifacta application.

Tip: This parameter can also be used to increase the maximum size of permitted downloads from the Trifacta application.

NOTE: In general, you should avoid downloading files that are larger than 1 GB in size from the Trifacta application. Other limits, such as timeout settings, may be applied, which can cause download failures. This setting is also applied for other hard limits within the Trifacta application, so please modify with caution. For large files, please try to download through your storage layer.

Steps:

1. You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.
2. Locate the following parameter:

```
"webapp.maxQueryResultsSize": 1073741824,
```

3. Set the value to 0.
4. Save your changes and restart the platform.
5. Verify that job results cannot be downloaded. See *Job Details Page*.

Samples

In the Transformer page, users interact with samples of data from the entire dataset. New samples can be created as needed.

By default, samples can be downloaded to the local desktop. To disable sample downloading, please complete the following steps.

Steps:

1. You apply this change through the *Workspace Settings Page*. For more information, see *Platform Configuration Methods*.

2. Locate the **Sample downloads** parameter, and set it to `disabled`.
3. To verify:
 - a. Logout of the Trifacta application. Login again.
 - b. Check that the Download Sample as CSV context menu option is not available in the Recipe panel in the Transformer Page.

NOTE: Recipes can still be downloaded.

See *Recipe Panel*.

Flows and Plans

For backup and migration purposes, users can be permitted to download their flow and plan definitions.

NOTE: A flow definition does not contain any data from the referenced datasets.

If needed, you can disable the ability to export flows and plans from the Trifacta platform

Steps:

1. You apply this change through the *Workspace Settings Page*. For more information, see *Platform Configuration Methods*.
2. Locate the **Export** parameter, and set it to `disabled`.
3. To verify:
 - a. Logout of the Trifacta application. Login again.
 - b. Check that the context menu option for Export Flow is not available in Flow View. See *Flow View Page*.

Imported Datasets

The platform does not support downloading imported datasets.

Dataset previews

When you select an imported dataset, you can see a preview of the data. While no data is stored permanently on the local desktop, you may prefer to disable previews.

NOTE: When previews are disabled, selection of imported datasets requires that users know the contents based on filename, location, and size information.

Steps:

1. You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.
2. Locate the following parameter:

```
"webapp.client.previewLoadLimit": 128000,
```

3. Set the value to 0.
4. Save your changes and restart the platform.
5. Verify that when you select a file or table to import, no preview of it is displayed. For more information, see *Import Data Page*.

Manage Schedules

Contents:

- *Enable or Disable Schedule*
 - *Delete Schedule*
 - *Create Schedule*
-

Through the Schedules page, administrators of Trifacta® can manage all of the schedules in the deployment.

NOTE: The Schedules page is available to project owners and workspace administrators only.

Enable or Disable Schedule

Steps:

1. Login as an administrator.
2. In the left nav bar, click the Schedules icon.
3. In the Schedules page, locate the schedule to change.
4. For the schedule's entry, open the context menu on the right side of the page.
5. Select **Enable Schedule** or **Disable Schedule**.

Delete Schedule

Deleting a schedule cannot be undone.

Steps:

1. Login as an administrator.
2. In the left nav bar, click the Schedules icon.
3. In the Schedules page, locate the schedule to delete.
4. For the schedule's entry, open the context menu on the right side of the page.
5. Select **Enable Schedule** or **Disable Schedule**.

Create Schedule

A schedule is composed of:

- A schedule frequency
- A set of one or more scheduled outputs

These objects are created from within a flow in Flow View. See *Add Schedule Dialog*.

System Services and Logs

Contents:

- *Download Logs*
 - *Support logs*
 - *Log directory*
 - *Artifact Storage Service*
 - *Authorization Service*
 - *Batch Job Runner*
 - *Configuration Service*
 - *Connector Configuration Service*
 - *Conversion Service*
 - *Data Service*
 - *Java UDF Service*
 - *Java VFS Service*
 - *Job Metadata Service*
 - *Machine Learning Service*
 - *Migration*
 - *Nginx Service*
 - *Optimizer Service*
 - *Scheduling Service*
 - *Spark Job Service*
 - *Supervisord Server*
 - *Time-Based Trigger Service*
 - *VFS Service*
 - *Webapp Service*
 - *Additional logs*
 - *Job logs*
-

The Trifacta® platform provides the following major services. For each of the listed service, any relevant logs are listed.

The logging levels for many of these services can be modified through the Admin Settings page. See *Configure Logging for Services*.

Download Logs

Support logs

For support use, the most meaningful logs and configuration files can be downloaded from the application. Select **Help menu > Download logs**.

NOTE: If you are submitting an issue to *Alteryx Support*, please download these files through the application.

For more information, see *Download Logs Dialog*. The admin version of this dialog enables downloading logs by timeframe, job ID, or session ID. For more information, see *Admin Download Logs Dialog*.

Log directory

System logs are maintained in the following directory: `/opt/trifacta/logs`

Trifacta® administrators can access the logs through the Trifacta application. Use the following URL:

<hostname>:<port_number>/logs

Available logs

Filename	Description
jobgroups/	Directory of logs for transformation jobs by Id. <div>Tip: If you are troubleshooting a failed job, please acquire the job logs from the Job Details page when you contact <i>Alteryx Support</i>. See <i>Job Details Page</i>.</div>
jobs/	Directory of logs for other kinds of jobs, such as sampling or ingest, by Id.
nginx/	Temporary storage for nginx server. No log files are stored here.
artifact-storage-service.access.log	Access logs for the artifact storage service.
artifact-storage-service.log	Application logs for the artifact storage service.
authorization-service.access.log	Access logs for the authorization service.
authorization-service.log	Application logs for the authorization service.
batch-job-runner.access.log	Access logs for the batch job runner service. Batch job runner service manages transformation jobs and scheduling. More information is below.
batch-job-runner.job-status.log	Status information on batch job runner jobs.
batch-job-runner.log	Application logs for the batch job runner service.
configuration-service.access.log	Access logs for the configuration service service. Configuration service is used for managing configuration that can be changed at runtime for different workspaces and users. Some of these settings are available through the Trifacta application. See <i>Workspace Settings Page</i> .
configuration-service.log	Application logs for the configuration service.
connector-configuration-service.log	Application logs for the connector configuration service.
connector-configuration-service.access.log	Access logs for the connector configuration service.
conversion-service.access.log	Access logs for the conversion service. Conversion service is used for converting from various inputs formats and to various output formats.
conversion-service.log	Application logs for the conversion service.
data-service.access.log	Access logs for the data service. Data service is used for interacting with relational sources. More information is below.

data-service.log	Application logs for the data service.
java-vfs-service.access.log	Access logs for the Java VFS service.
java-vfs-service.log	Application logs for the Java VFS service.
job-metadata-service.access.log	Access logs for the job metadata service.
job-metadata-service.log	Application logs for the job metadata service.
migration.log	Application logs for database migrations performed for the webapp service.
ml_service.access.log	Access logs for the ml (machine-learning) service. Machine learning service is used for predictive interaction, suggestion ranking, pattern profiling, pattern suggestions, and collecting user action logs. More information is below.
ml_service.log	Application logs for the ml (machine-learning) service.
nginx_service.log	Application logs for the nginx service. More information is below.
optimizer-service.access.log	Access logs for the optimizer service.
optimizer-service.log	Application logs for the optimizer service.
protobuf-events.log	Client events around column values, user selections, and recipe editing.
proxy_access.log	Access logs for the nginx server.
proxy_error.log	Error logs for the nginx server.
scheduling-service.access.log	Access logs for the scheduling service. Scheduling service is used for scheduling jobs at a specific time. More information is below.
scheduling-service.log	Application logs for the scheduling service.
secure-token-service.access.log	Access logs for the secure token service. secure-token-service is used for securely storing tokens for some external services, such as Azure AD and Databricks and OAuth2-connected datastores.
secure-token-service.log	Application logs for the secure token service
spark-job-service.log	Application logs for the Spark job service. Spark job service is used for interfacing with cluster-based Spark service to plan and execute Spark jobs. More information is below.
supervisord.log	Logs for the supervisord system service. supervisord manages the starting, stopping, and restarting of services for the Trifacta platform. More information is below.
time-based-trigger-service.access.log	Access logs for the time-based trigger service. Time-based trigger service is used for managing the triggers for scheduled jobs. More information is below.
time-based-trigger-service.log	Application logs for the time-based trigger service.
vfs-service.access.log	Access logs for the VFS service.

log	VFS service is used for managing loading of files from various supported datastores. More information is below.
vfs-service.log	Application logs for the VFS service.
webapp.access.log	Access logs for the Webapp service. Webapp service serves the Trifacta application to users. More information is below.
webapp.log	Application logs for the Webapp service. More information is below.
webapp.sql-error.log	Error log for SQL issued from the Webapp service.
webworker.log	Event logs for webworkers running in browser clients. Webworkers run in the background of a browser client's Trifacta session and are used for predictive interaction, suggestion ranking, pattern profiling, pattern suggestions, and collecting user action logs.

Artifact Storage Service

Description: Manages storage of feature-specific usage data, such as value mappings.

Log File	Can Help With
artifact-storage-service-access.log	Access issues to the service.
artifact-storage-service.log	Transactions with the database for the features that use it.

Authorization Service

Description: Manages access permissions for workspace objects.

Log File	Can Help With
authorization-service-access.log	Access issues to the service.
authorization-storage-service.log	Transactions with the database for the features that use it.

Batch Job Runner

Description: This service manages the tracking of jobs submitted to the backend running environment.

Log File	Can Help With
batch-job-runner.log	<ul style="list-style-type: none"> Service errors and crashes Determine execution environment of the job. Search for: <ul style="list-style-type: none"> LocalJobRunner = local execution in Photon YARNRunner = execution in Spark Communication errors back from environment Status information on jobs Status information on counts of job retries

For more information on this service, see *Configure Batch Job Runner*.

Configuration Service

Description: Service for managing configurations at the user, workspace, and system levels, which can be changed at runtime.

--	--

Log File	Can Help With
configuration-service-access.log	Issues accessing the service.
configuration-service.log	Configuration problems.

Some configuration service options are surfaced in the Trifacta application. See *Workspace Settings Page*.

Connector Configuration Service

Description: Service for managing metadata for connector types.

Log File	Can Help With
connector-configuration-service-access.log	Issues accessing the service.
connector-configuration-service.log	Configuration problems.

Conversion Service

Description: Service converts some formats for input or output. For example, Microsoft Excel workbooks must be ingested through the conversion service and stored as separate CSVs for each worksheet.

Log File	Can Help With
conversion-service-access.log	Issues accessing the service.
conversion-service.log	Problems with ingest jobs for datasets whose sources must be converted.

Data Service

Description: Service prepares queries against JDBC interfaces, using internal REST API calls.

Log File	Can Help With
data-service.log	<ul style="list-style-type: none"> Initialization of communications through JDBC interface Query failures

For more information on this service, see *Configure Data Service*.

Java UDF Service

Description: Service enables the execution of Java-based user-defined functions within a transform recipe.

Log File	Can Help With
java-udf-service.log	<ul style="list-style-type: none"> Status of the service <div style="border: 1px solid green; padding: 10px; margin-top: 10px;"> <p>Tip: You can pass through messages on errors through Logger to this log, which can assist in diagnosing issues.</p> </div>

For more information, see *User-Defined Functions*.

Java VFS Service

Description: Service manages metadata associated with jobs during execution.

Log File	Can Help With
java-vfs-service.log	<ul style="list-style-type: none">• Help with issues accessing data on ADLS Gen2• Status of the service

For more information, see *Configure Java VFS Service*.

Job Metadata Service

Description: Service manages metadata associated with jobs during execution.

Log File	Can Help With
job-metadata-service.log	<ul style="list-style-type: none">• Help with job phases and status• Status of the service

Machine Learning Service

Description: ML service provides machine learning capabilities for the platform.

Log File	Can Help With
ml_service.log	<ul style="list-style-type: none">• This log is likely to contain information that is only useful if the ML service has crashed.

Migration

Description: Migration applies to database migrations executed as part of upgrading the platform.

Log File	Can Help With
migration.log	<ul style="list-style-type: none">• Connectivity errors and other issues that may have occurred during migration.

Nginx Service

Description: Nginx is a proxy server embedded in the platform that serves the web application and other resources.

Log File	Can Help With
nginx_service.log	<ul style="list-style-type: none">• This log may be useful in identifying any warnings that occurred when the nginx services starts.• The <code>nginx</code> server may contain log information for the server that provides HTTP access.

Optimizer Service

Description: Handles optimizations of the queries for data from relational sources.

Log File	Can Help With
<code>optimizer-service.log</code>	<ul style="list-style-type: none">• Service-related issues
<code>optimizer-service.access.log</code>	<ul style="list-style-type: none">• Gives information about accessed routes

Proxy

Description: The proxy (nginx) service manages requests from the user interface to the other components of the platform.

Log File	Can Help With
<code>proxy_access.log</code>	<ul style="list-style-type: none">• Shows any requests made through the nginx to the port used by the Trifacta platform.
<code>proxy_error.log</code>	<ul style="list-style-type: none">• Contains any errors thrown by the nginx service when a request is made to the port used by the Trifacta platform.

Scheduling Service

Description: Handles all metadata related to scheduling.

Log File	Can Help With
<code>scheduling-service.log</code>	<ul style="list-style-type: none">• Schedule-related issues
<code>scheduling-service.access.log</code>	<ul style="list-style-type: none">• Gives information about accessed routes

Spark Job Service

Description: Service that manages jobs processed on Spark.

Log File	Can Help With
<code>spark-job-service.log</code>	<ul style="list-style-type: none">• Status of the service• Debugging issues with Spark jobs• See <i>Configure for Spark</i>.

Supervisord Server

Description: Process that starts, stops, and restarts services in the platform.

--	--

Log File	Can Help With
<code>supervisord.log</code>	<ul style="list-style-type: none"> • Status information from platform services

Time-Based Trigger Service

Description: Handles all metadata related to the trigger service

Log File	Can Help With
<code>time-based-trigger-service.log</code>	<ul style="list-style-type: none"> • schedules not triggering correctly
<code>time-based-trigger-service.access.log</code>	<ul style="list-style-type: none"> • Gives information about accessed routes

VFS Service

Description: Loads data from the various filesystems supported by the platform, both in the front-end user interface and in batch mode when the Trifacta Photon running environment is enabled. For more information, see *Running Environment Options*.

Log File	Can Help With
<code>vfs-service.log</code>	<ul style="list-style-type: none"> • Client connection issues • Status issues with backend components • Information on batch jobs that cannot be started

Webapp Service

Description: Loads data from the various filesystems supported by the platform in the front-end user interface.

Log File	Can Help With
<code>webapp.log</code>	<ul style="list-style-type: none"> • Client connection issues • Status issues with backend components • Information on batch jobs that cannot be started
<code>webapp.access.log</code>	<ul style="list-style-type: none"> • Provides access information to the Webapp service and routes.

Additional logs

Job logs

The following sources of information may provide information related to job status and performance:

- job log
- spark log
- cdf script
- yarn application logs
(if log aggregation is enabled)
- platform configuration file
(`trifacta-conf.json`)
- batch job runner log

- spark service log
- hadoop conf directory

Storage Maintenance

Contents:

- *Trifacta Storage*
 - *Service logs*
 - *Job logs*
 - *Base Storage Layer*
 - *Temp files*
 - *Samples and profile statistics*
 - *Datasets*
 - *Storage for features*
-

This page provides some tips and guidelines for maintaining your backend storage.

NOTE: Except for temporary files that it creates as part of normal operations or storage used as part of feature execution, Trifacta® does not remove files from the backend storage for safety reasons. Unless resources have been provided to you by Alteryx, management of the backend datastore is the responsibility of the customer.

NOTE: Trifacta does not store data on the Trifacta node where the software is installed.

NOTE: Trifacta does not modify source data.

Trifacta Storage

Log files are stored by default in the following location on the Trifacta node:

```
/opt/trifacta/logs
```

Service logs

Service log files are automatically auto-rotated at 50 MB. For more information on configuring log rotation, see *Configure Logging for Services*.

Job logs

Logs related to job execution are not automatically rotated.

NOTE: Job log files can accumulate over time. As a good rule of thumb, you can set up a recurring job through an external scheduler to purge old job logs that are older than six months.

Job log files are stored in the following directories:

```
/opt/trifacta/logs/jobs  
/opt/trifacta/logs/jobgroups
```

They are organized by job identifier in sub-directories.

For more information on job logs, see *Diagnose Failed Jobs*.

Base Storage Layer

Temp files

Job temp files

Temporary files may be written to the temporary directory on the backend datastore, particularly during job execution.

```
/tmp
```

NOTE: These files may be purged during restarts of the platform.

Spark temp files

During execution of jobs, Spark may use the following directories on backend storage for storage of temporary files:

```
/user/<UserID>  
/trifacta/tempfiles
```

Samples and profile statistics

The Trifacta platform generates your samples and profiling statistics in one of the following directories for each user:

- The default directory:

```
/trifacta/queryResults/.trifacta
```

- The user-defined output directory

NOTE: These files should be removed on a periodic basis.

Datasets

While samples and job results may be retained on backend storage, the Trifacta platform does not store your source data.

NOTE: Datasets removed from the Library are removed as references to the product. The underlying data is not actually deleted.

Storage for features

The following features do store data on the base storage layer.

File conversion

Data sources that are stored in a binary format, such as PDF or Excel, or that require additional processing, such as JSON, must be converted to file format that can be natively ingested by the Trifacta platform. Typically these files are stored in the base storage layer in CSV format.

This feature is enabled by default.

JDBC ingestion

When JDBC ingestion is enabled, some objects used in sampling that are sourced from JDBC sources may be stored in the base storage layer for faster retrieval. After job execution, these objects are deleted, or if datasource caching is enabled, are moved to the appropriate datasource cache.

For more information, see *Configure JDBC Ingestion*.

Datasource caching

If datasource caching has been enabled, cached objects can be stored in either a global or user-specific cache. For more information, see *Configure Data Source Caching*.

Backup and Recovery

Contents:

- *Stop All Services*
- *Perform manual backups*
- *Restart*
- *Recovery*

This section provides overview information on the key data and metadata that should be managed by your enterprise backup and recovery policies.

NOTE: This section covers how to perform a basic cold backup of the product. Hot backups are not supported.

All backups should be performed in accordance with your enterprise's backup and recovery policies.

Stop All Services

Before you begin, the Trifacta platform and databases should be stopped. See *Start and Stop the Platform*.

Perform manual backups

Back up platform files

The following directories on the Trifacta node should be backed up on a regular basis:

Configuration files:

You can back up all key configuration files into the `/tmp` directory using the following commands:

```
cp -R /opt/trifacta/conf /tmp/conf
cp /etc/init.d/trifacta /tmp/trifacta.service
cp -R /opt/trifacta/pkg3p/tripache/conf/conf.d /tmp/conf.d
cp -R /opt/trifacta/services/data-service/build/conf/vendor /tmp/vendor
cp -R /opt/trifacta/hadoop-deps /tmp/trifacta-hadoop-deps
```

License file:

You should back up your license key:

```
cp /opt/trifacta/license/license.json /tmp/license.json
```

See *License Key*.

Log files:

Optionally, you can choose to back up your log files:

Tip: Trifacta platform upgrades may be faster if the log directory is empty. Before you upgrade, you may wish to back up this directory, empty it, and then restore your backup after the upgrade.

```
cp -R /opt/trifacta/logs /tmp/logs
```

Back up databases

The Trifacta platform utilizes the following databases as part of normal operations. These databases should be backed up on a regular basis:

Database Name	Databaseld	Description
Main DB	trifacta	Stores users and metadata for flows, including datasets, and recipes.
Jobs DB	trifacta-activiti	Stores and maintains job execution status and details.
Scheduling DB	trifactascheduling-service	Stores metadata for scheduled jobs.
Time-based Trigger DB	trifactatimebasedtrigger-service	Additional database required for scheduled jobs.
Configuration Service DB	trifactaconfiguration-service	Stores configuration settings for the workspace.
Artifact Storage Service DB	trifactaartifactstorage-service	Stores feature usage data such value mappings for the standardization feature.
Job Metadata Service DB	trifactajobmetadataservice	Stores metadata on job execution.
Authorization Service DB	trifactaauthorization-service	Storage of object permissions.
Orchestration Service DB	trifactaorchestration-service	Storage of plans, triggers, tasks, and snapshots.
Optimizer Service DB	trifactaoptimizer-service	Storage of SQL queries for optimization during job execution.
Secure Token Service DB	trifactasecuretokenservice	Storage of STS tokens for use in accessing third-party systems.
Connector Configuration Service DB	trifactaconnectorconfiguration-service	Storage of metadata information on connector types.

For more information on setting up these databases, see *Install Databases*.

Location of db tools - PostgreSQL

Depending on your operating system, you can find the backup tools in the following location.

CentOS/RHEL - PostgreSQL 12:

NOTE: These locations apply to PostgreSQL 12.

```
/usr/pgsql-12/bin/pg_dump  
/usr/pgsql-12/bin/psql
```

Ubuntu:

```
/usr/lib/postgresql/9.6/bin/pg_dump
/usr/lib/postgresql/9.6/bin/psql
```

Location of db tools - MySQL

Please locate the following programs in your MySQL distribution:

```
mysqldump
mysql
```

Manual backup commands

The following commands can be used to back up the databases.

PostgreSQL

For more information on command options, see <https://www.postgresql.org/docs/9.6/static/backup.html>.

NOTE: These commands must be executed as the `trifacta` user.

NOTE: The following commands are for PostgreSQL 12.3 for all supported operating systems. For specific commands for other versions, please see the database documentation.

Tip: You may see performance improvements by backing up and restoring using `.TAR` files. However, there is a risk that `.TAR` support could change in the future. For more information, please see the PostgreSQL documentation.

Trifacta DB:

NOTE: If you are providing a dump of the `trifacta` database to *Alteryx Support*, please include a dump of the `trifactaauthorizationservice` database, as well.

```
pg_dump trifacta > trif_triDB_bkp_<date>.sql
```

Jobs DB:

```
pg_dump trifacta-activiti > trif_actDB_bkp_<date>.sql
```

Scheduling DB:

```
pg_dump trifactaschedulingservice > trif_schDB_bkup_<date>.sql
```

Time-Based Trigger DB:

```
pg_dump trifactatimebasedtriggerservice > trif_tbtsDB_bkup_<date>.sql
```

Configuration Service DB:

```
pg_dump trifactaconfigurationsservice > trif_confsservDB_bkup_<date>.sql
```

Artifact Storage DB:

```
pg_dump trifactaartifactstorageservice > trif_artifactstorageservDB_bkup_<date>.sql
```

Job Metadata Service DB:

```
pg_dump trifactajobmetadataservice > trif_jobmetadataservDB_bkup_<date>.sql
```

Authorization Service DB:

```
pg_dump trifactaauthorizationsservice > trif_authorizationservDB_bkup_<date>.sql
```

Orchestration Service DB:

```
pg_dump trifactaorchestrationservice > trif_orchestrationservDB_bkup_<date>.sql
```

Optimizer Service DB:

```
pg_dump trifactaoptimizerservice > trif_optimizerservDB_bkup_<date>.sql
```

Secure Token Service DB:

```
pg_dump trifactasecuretokenservice > trif_securetokenservDB_bkup_<date>.sql
```

Connector Configuration Service DB:

```
pg_dump trifactaconnectorconfigurationsservice > trif_connectorconfigurationsservDB_bkup_<date>.sql
```

MySQL

For more information on command options, see <https://dev.mysql.com/doc/refman/5.7/en/mysqldump-sql-format.html>.

```
su - mysql
```

NOTE: The following commands are for MySQL 5.7 for all supported operating systems. For specific commands for other versions, please see the database documentation.

Trifacta DB:

NOTE: If you are providing a dump of the `trifacta` database to *Alteryx Support*, please include a dump of the `trifactaauthorizationsservice` database, as well.

```
mysqldump trifacta > trif_triDB_bkp_<date>.sql
```

Jobs DB:

```
mysqldump trifacta-activiti > trif_actDB_bkp_<date>.sql
```

Scheduling DB:

```
mysqldump trifactaschedulingservice > trif_schDB_bkup_<date>.sql
```

Time-Based Trigger DB:

```
mysqldump trifactatimebasedtriggerservice > trif_tbtsDB_bkup_<date>.sql
```

Configuration Service DB:

```
mysqldump trifactaconfigurationsservice > trif_confservDB_bkup_<date>.sql
```

Artifact Storage DB:

```
mysqldump trifactaartifactstorageservice > trif_artifactstorageservDB_bkup_<date>.sql
```

Job Metadata Service DB:

```
mysqldump trifactajobmetadataservice > trif_jobmetadataservDB_bkup_<date>.sql
```

Authorization Service DB:

```
mysqldump trifactaauthorizationsservice > trif_authorizationservDB_bkup_<date>.sql
```

Orchestration Service DB:

```
mysqldump trifactaorchestrationservice > trif_orchestrationservDB_bkup_<date>.sql
```

Optimizer Service DB:

```
mysqldump trifactaoptimizerservice > trif_optimizerservDB_bkup_<date>.sql
```

Secure Token Service DB:

```
mysqldump trifactasecuretokenservice > trif_securetokenservDB_bkup_<date>.sql
```

Connector Configuration Service DB:

```
mysqldump trifactaconnectorconfigurationsservice > trif_connectorconfigurationsservDB_bkup_<date>.sql
```

Scheduling

You can schedule nightly execution of these backups using a third-party scheduler such as cron.

Restart

You can restart the Trifacta platform now. See *Start and Stop the Platform*.

Recovery

See *Platform Rollback*.

Platform Rollback

Contents:

- *Prerequisites*
 - *Rollback Overview*
 - *Stop services*
 - *Create new versions of databases - PostgreSQL*
 - *Uninstall and reinstall Trifacta software*
 - *Restore databases and config files*
 - *Restart*
 - *Verify*
-

In the event that an upgrade or hotfix to your instance of the Trifacta® platform has run into issues that cannot be repaired in the upgraded instance, you can follow the steps in this section to rollback to your previous version.

NOTE: Before you perform a rollback, you should review the set of issues with Alteryx first. For more information, please contact *Alteryx Support*.

Prerequisites

In order to complete the rollback in a timely manner, please verify that you have access to the following:

Access

You must:

- Acquire root user access to the Trifacta node.
- Acquire database access to uninstall and reinstall the Trifacta databases.

Tip: You should communicate to any affected users the required maintenance and expected outage window.

Backups

If you do not have the following, you cannot perform a rollback. These items cannot be acquired from Alteryx.

- Backups of your pre-upgrade Trifacta configuration files
- Backups of your pre-upgrade Trifacta databases

The following can be acquired from Alteryx if you do not have them:

- RPM installers for the previous version. If any Hotfixes have been applied to the previous version, you should acquire and use the latest Hotfix RPM for your re-install.
- PDF documentation for the previous version.

Rollback Overview

To recover the Trifacta platform based on backups, please complete the following sections.

NOTE: When the databases are restored, internal identifiers such as job IDs, are reset in an order that may not correspond to the expected order. Consequently, references to specific identifiers may be corrupted. After restoring the databases, you should clear the job logs.

NOTE: If any of the hosts, pathnames, or credentials have changed since the backups were performed, these updates must be applied through `trifacta-conf.json` or through the Admin Settings page after the restoration is complete.

Stop services

You must stop the Trifacta platform.

Steps:

1. Login to the Trifacta node as root user.
2. Stop the Trifacta service:

```
service trifacta stop
```

Create new versions of databases - PostgreSQL

Before you restore, you must drop each database and create new versions of the databases that you wish to restore.

NOTE: The following assumes that the roles for each database are already created.

Login as a user that can run admin commands for PostgreSQL. This user may vary between deployments.

Trifacta database:

```
psql -c "DROP DATABASE trifacta;"
psql -c "CREATE DATABASE trifacta WITH OWNER trifacta;"
```

Jobs database:

NOTE: Please note the escaped quotes in the `CREATE DATABASE` command for this database.

```
psql -c "DROP DATABASE \"trifacta-activiti\";"
psql -c "CREATE DATABASE \"trifacta-activiti\" WITH OWNER trifactaactivit;"
```

Scheduling database:

```
psql -c "DROP DATABASE trifactascheduling;"
psql -c "CREATE DATABASE trifactascheduling WITH OWNER trifactascheduling;"
```

Time-based Trigger Service database:

```
psql -c "DROP DATABASE trifactatimebasedtriggerservice;"
psql -c "CREATE DATABASE trifactatimebasedtriggerservice WITH OWNER trifactatimebasedtriggerservice;"
```

Configuration Service database:

(Release 6.0 and later)

```
psql -c "DROP DATABASE trifactaconfigurationservice;"
psql -c "CREATE DATABASE trifactaconfigurationservice WITH OWNER trifactaconfigurationservice;"
```

Artifact Storage Service database:

(Release 6.0 and later)

```
psql -c "DROP DATABASE trifactaartifactstorageservice;"
psql -c "CREATE DATABASE trifactaartifactstorageservice WITH OWNER trifactaartifactstorageservice;"
```

Job Metadata Service database:

(Release 6.4 and later)

```
psql -c "DROP DATABASE trifactajobmetadataservice;"
psql -c "CREATE DATABASE trifactajobmetadataservice WITH OWNER trifactajobmetadataservice;"
```

Authorization Service database:

(Release 7.1 and later)

```
psql -c "DROP DATABASE trifactaauthorizationservice;"
psql -c "CREATE DATABASE trifactaauthorizationservice WITH OWNER trifactaauthorizationservice;"
```

Orchestration Service database:

(Release 7.1 and later)

```
psql -c "DROP DATABASE trifactaorchestrationservice;"
psql -c "CREATE DATABASE trifactaorchestrationservice WITH OWNER trifactaorchestrationservice;"
```

Optimizer Service database:

(Release 7.6 and later)

```
psql -c "DROP DATABASE trifactaoptimizerservice;"
psql -c "CREATE DATABASE trifactaoptimizerservice WITH OWNER trifactaoptimizerservice;"
```

Secure Token Service database:

(Release 8.1 and later)

```
psql -c "DROP DATABASE trifactasecuretokenservice;"
psql -c "CREATE DATABASE trifactasecuretokenservice WITH OWNER trifactasecuretokenservice;"
```

Connector Configuration Service database:

(Release 8.1 and later)

```
psql -c "DROP DATABASE trifactaconnectorconfigurationservice;"
psql -c "CREATE DATABASE trifactaconnectorconfigurationservice WITH OWNER
trifactaconnectorconfigurationservice;"
```

Uninstall and reinstall Trifacta software

Steps:

1. Uninstall the current version of the Trifacta software:

NOTE: All platform and cluster configuration files are preserved. User metadata is preserved in the Trifacta database.

CentOS/RHEL:

```
sudo rpm -e trifacta
```

Ubuntu:

```
sudo apt-get remove trifacta
```

2. Perform a clean install of the Trifacta software provided in your distribution. See *Install*.

Restore databases and config files

You can use the following commands in the sections for manual restores of:

- Databases
 - PostgreSQL
 - MySQL
- Configuration files

Restore databases - PostgreSQL

Trifacta database:

```
psql --dbname=trifacta < trif_triDB_bkp_<date>.sql
```

Jobs database:

NOTE: Please note the escaped quotes in the `CREATE DATABASE` command for this database.

```
psql --dbname="trifacta-activiti" < trif_actDB_bkp_<date>.sql
```

Scheduling database:

```
psql --dbname=trifactaschedulingservice < trif_schDB_bkup_<date>.sql
```

Time-based Trigger Service database:

```
psql --dbname=trifactatimebasedtriggerservice < trif_tbtSDB_bkup_<date>.sql
```

Configuration Service database:

(Release 6.0 and later)

```
psql --dbname=trifactaconfigurationservice < trif_confservDB_bkup_<date>.sql
```

Artifact Storage Service database:

(Release 6.0 and later)

```
psql --dbname=trifactaartifactstorageservice < trif_artifactstorageservDB_bkup_<date>.sql
```

Job Metadata Service database:

(Release 6.4 and later)

```
psql --dbname=trifactajobmetadataservice < trif_jobmetadataservDB_bkup_<date>.sql
```

Authorization Service database:

(Release 7.1 and later)

```
psql --dbname=trifactaauthorizationsservice < trif_authorizationservDB_bkup_<date>.sql
```

Orchestration Service database:

(Release 7.1 and later)

```
psql --dbname=trifactaorchestrationservice < trif_orchestrationservDB_bkup_<date>.sql
```

Optimizer Service database:

(Release 7.6 and later)

```
psql --dbname=trifactaoptimizerservice < trif_optimizerservDB_bkup_<date>.sql
```

Secure Token Service database:

(Release 8.1 and later)

```
psql --dbname=trifactasecuretokenservice < trif_securetokenservDB_bkup_<date>.sql
```

Connector Configuration Service database:

(Release 8.1 and later)

```
psql --dbname=trifactaconnectorconfigurationservice < trif_connectorconfigurationservDB_bkup_<date>.sql
```

Restore databases - MySQL

For details, see <https://dev.mysql.com/doc/refman/5.7/en/reloading-sql-format-dumps.html>.

Login:

```
su - mysql
```

Trifacta database:

```
mysql trifacta < trif_triDB_bkp_<date>.sql
```

Jobs database:

```
mysql trifacta-activiti < trif_actDB_bkp_<date>.sql
```

(Release 4.1 and later) Scheduling database:

```
mysql trifactaschedulingservice < trif_schDB_bkup_<date>.sql
```

(Release 4.1 and later) Time-based Trigger Service database:

```
mysql trifactatimebasedtriggerservice < trif_tbtsDB_bkup_<date>.sql
```

(Release 6.0 and later) Configuration Service database:

```
mysql trifactaconfigurationsservice < trif_confservDB_bkup_<date>.sql
```

(Release 6.0 and later) Artifact Storage Service database:

```
mysql trifactaartifactstorageservice < trif_artifactstorageservDB_bkup_<date>.sql
```

(Release 6.4 and later) Job Metadata Service database:

```
mysql trifactajobmetadataservice < trif_jobmetadataservDB_bkup_<date>.sql
```

(Release 7.1 and later) Authorization Service database:

```
mysql trifactaauthorizationsservice < trif_authorizationservDB_bkup_<date>.sql
```

(Release 7.1 and later) Orchestration Service database:

```
mysql trifactaorchestrationservice < trif_orchestrationservDB_bkup_<date>.sql
```

(Release 7.6 and later) Optimizer Service database:

```
mysql trifactaoptimizerservice < trif_optimizerservDB_bkup_<date>.sql
```

(Release 8.1 and later) Secure Token Service database:

```
mysql trifactasecuretokenservice < trif_securetokenservDB_bkup_<date>.sql
```

(Release 8.1 and later) Connector Service database:

```
mysql trifactaconnectorconfigurationservice < trif_connectorconfigurationservDB_bkup_<date>.sql
```

Restore config files

Restore your configuration files. The following commands assume that they were backed up to the `/tmp` directory on the node:

```
cp /tmp/trifacta-conf.json /opt/trifacta/conf/trifacta-conf.json
cp /tmp/env.sh /opt/trifacta/conf/env.sh
cp /tmp/trifacta.service /etc/init.d/trifacta
```

Restart

Apply any patches or maintenance updates that may have been provided to you.

Restart the platform. See *Start and Stop the Platform*.

Verify

Login and verify operations. See *Verify Operations*.

Admin Reference

This section describes the pages for admin users who want to modify settings and users at the project or workspace level for Trifacta® . Most of these pages are available through the Admin console in the Trifacta application. Additional admin reference materials are included.

Deployment Manager Page

Contents:

- *Access*
- *Deployment Hierarchy*
- *Deployments View*
- *Releases View*
- *Flows View*

Through the Deployment Manager page, you interact with flows that you have imported into your Production instance of the Trifacta® platform. Through this interface, you can activate Production versions of your flows or rollback to previous versions as needed.

NOTE: The Deployment Manager is available only in a Production environment, which is a special instance of the Trifacta platform designed to support production use of your flows. For more information, see *Overview of Deployment Manager*.

Access

A Production environment can be accessed in either of the following ways:

- You are given access to a separate instance of the Trifacta platform configured for Production use only.
- On any instance, the Deployment role is added to your user account by a Trifacta administrator.

For more information, see *Configure Deployment Manager*.

Deployment Hierarchy

In a Production environment, a **deployment** is version-managed flow and all of its dependencies, including other dependent flows. Through the Deployment Manager, an individual deployment is structured in the following hierarchy:

Hierarchy Level	Object	Description
1	deployment	When you open the Deployment Manager, you can review all of the deployments that have been created in the environment. A deployment is container for releases.
2	release	<p>When you select a deployment, you can explore its releases. A release is an individual instance of an imported flow and its dependencies (an import package). Each time that the import package is re-imported into the Production instance, a new release is created and made the active release for the deployment.</p> <p>You can activate previous releases as needed through the context menus in Deployment Manager.</p>
3	flow or flows	<p>Within a release, you can explore the flows that were included in the import package for the release:</p> <ul style="list-style-type: none">• The primary flow is the one that is executed when:<ul style="list-style-type: none">• Its release is the active one for the deployment• The job for the deployment is executed• Any secondary flows are the flows on which the primary flow depends for data.<ul style="list-style-type: none">• During export from the source instance, all objects in secondary flows are included in the package. There may be objects in a secondary flow that are unused in the Production instance.

Deployments View

When you open the Deployment Manager, you can explore all of the deployments in the Production instance.



Name	Releases	Last Updated By	Last Updated At
Deployment 02	0 Releases	SteveO	2019-03-15 10:36:57
Deployment 01	2 Releases	SteveO	2019-03-15 10:27:54

Figure: Deployment Manager

To create a new deployment, click **Create**.

1. Enter a name for the deployment, and click **Create**.
2. To create a new release, click the created deployment. See Releases View below.

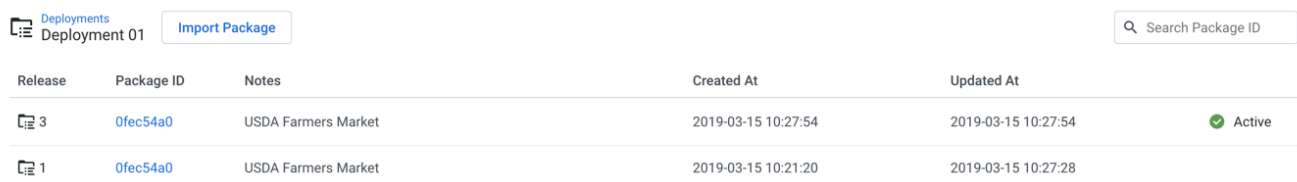
Actions:

- **Search:** Enter values in the search textbox to search deployment names. Matching occurs in real-time.
- **Edit name:** You can change the name of your deployment as needed.
- **Delete:** Select this option to remove the deployment, all of its releases, and all of the flows within each release.

You cannot undo deleting a deployment. Any results generated from jobs run for the deployment are not removed from the output location and are still accessible through the Jobs page of the Production instance.

Releases View

Through Releases view, you can import new packages to create new releases and activate them, roll back to previous releases, and remove releases that are no longer in use.



Release	Package ID	Notes	Created At	Updated At
3	0fec54a0	USDA Farmers Market	2019-03-15 10:27:54	2019-03-15 10:27:54 Active
1	0fec54a0	USDA Farmers Market	2019-03-15 10:21:20	2019-03-15 10:27:28

Figure: Releases View

To create a new release, click **Import Package**.

NOTE: Before you import a package, you must apply any import mapping rules to the deployment. These rules map values and objects in the package to corresponding values in the new instance. For more information, see *Define Import Mapping Rules*.

1. Navigate your local environment to select the ZIP file containing the flow and its dependencies from the source instance.
2. Click **Import**.
3. The release is added to Releases view.

For more information on importing, see *Import Flow*.

Actions:

Action	Description
Search	Enter values in the search textbox to search package identifiers. Matching occurs in real-time.
Activate	Make the selected release the active one for the deployment. When jobs are executed at the deployment level, the primary flow for the active release is executed.
Export	Export the release for use in another instance of the platform. See <i>Export Flow</i> .
Delete	<p>Delete the release from the Production instance.</p> <div> <p>NOTE: If the release was imported from a Dev instance on the same platform, the Dev instance of the release is not removed.</p> <p>NOTE: Deletion of a release does not remove any results generated from it. Those results are still accessible through the Jobs page.</p> </div>

Flows View

When you open a release, you can review the flows contained in it. You can explore the flow or flows that were included in the import package for the selected release.

Deployments > Deployment 01 0fec54a0 (Release 3)			Search name...
Name	Datasets	Updated At	
 USDA Farmers Market 2014 Flow	2 Datasets	Today at 10:27 AM	

Figure: Flows View

Actions:

- **Search:** Enter values in the search textbox to search flow names. Matching occurs in real-time.
- Click a flow name to explore the flow in Flow View.

NOTE: Avoid making changes to a flow in a Production instance. You can run ad-hoc jobs, but you should avoid making changes to the objects or their structure through Flow View in Production instances. Scheduling should be done through the command line.

For more information, see *Flow View Page*.

Admin Console

Contents:

- *Users*
 - *Roles*
 - *Workspace settings*
 - *Admin settings*
 - *AWS settings*
 - *Environment parameters*
 - *OAuth 2.0 clients*
-

Through the Admin console, admin users can modify settings and users at the system and workspace level, as well as run health checks and manage the license for Trifacta®. Select **User menu > Admin console**.

NOTE: You must be an administrator to access this feature.

Users

Invite, disable, and remove users to Trifacta. Change roles, as needed. For more information, see *Users Page*.

Roles

Create roles and assign permissions to them for access to objects created in Trifacta. For more information, see *Roles Page*.

Workspace settings

Review and edit settings applicable to the workspace. For more information, see *Workspace Settings Page*.

Admin settings

NOTE: The Admin Settings page is only available to administrators of the Trifacta platform.

Platform configuration settings:

- Review and manage configuration for the Trifacta® platform.
- Configure settings for external services.
- Manage user accounts.
- Diagnostics and license management
- Platform restart

For more information, see *Admin Settings Page*.

AWS settings

If per-user access to AWS has been enabled, individual users must apply personal access credentials to their account to gain access to resources on S3 through AWS. For more information, see *AWS Settings Page*.

Environment parameters

Define parameters that apply to the entire environment and are available for use by all users. For more information, see *Environment Parameters Page*.

OAuth 2.0 clients

Administrators can create and manage clients for accessing an OAuth 2.0 app in an external platform such as a relational datastore.

NOTE: Before you create an OAuth 2.0 client, you must have created an OAuth 2.0 app in the target system, to which your client can connect. For more information, see *Enable OAuth 2.0 Authentication*.

For more information, see *OAuth 2.0 Clients Page*.

Users Page

The Users page enables adding, disabling, or removing users from your project or workspace. You can also reset passwords and change roles.

Users				<div><div></div><div><</div><div>1</div><div>2</div><div>></div><div></div></div> <div>Test User</div> <div>X</div>
All	Enabled	Disabled		
Name	Email	Status	Last log in	
<div>TU</div> Test User 1786321687		Enabled	Today at 7:41 AM	
<div>TU</div> Test User 2292989979		Enabled	Today at 7:32 AM	
<div>TU</div> Test User19933547		Enabled	Today at 12:02 PM	
<div>TU</div> Test User21197568		Enabled	Today at 11:31 AM	
<div>TU</div> Test User22395996		Enabled	Today at 11:39 AM	
<div>TU</div> Test User29277976		Enabled	Today at 12:01 PM	
<div>TU</div> Test User33101513		Enabled	Today at 11:59 AM	
<div>TU</div> Test User46556047		Enabled	Today at 12:00 PM	
<div>TU</div> Test User57276476		Enabled	Today at 12:04 PM	
<div>TU</div> Test User57741240		Enabled	Today at 11:37 AM	

Figure: Users Page

Tabs:

- Click one of the tabs to display all users or a filtered list based on user status.

Fields:

- **Name:** Display name for the user. Click the name of the user to review details about the user account. See *User Details Page*.
- **Email:** Username (email address of users)
- **Status:** Current status of the user. See "Status" below.
- **Last login:** Timestamp for the last time that the user logged in to the Trifacta application

Actions:

- **Search:** Enter text to begin searching for specific usernames or email addresses.

Context menu actions:

For each user, you can perform the following actions in the context menu:

- **Configure storage:** If per-user access is enabled, you can configure the access credentials for individual users, either using key-secret combinations or IAM roles. For more information, see *Configure Your Access to S3*.
- **Edit:** Modify user properties, including platform roles. See "Edit Users" below.

- **Reset password:** Self-service password reset is enabled by default. If enabled, click this option to send an email to the user to reset his or her password.

NOTE: Only platform administrators can reset a user's password. Workspace admins cannot.

Disable: When a user is disabled, the user cannot access the Trifacta application.

- The disabled user still counts against the project or workspace limit.
- All of the user's flows and datasets are retained.

NOTE: Schedules owned by a disabled user continue to execute. An admin can disable the schedule. See *Schedules Page*.

- Resources such as connections and flows that are owned by the user become inaccessible to other users that have access.
- To permit access again, select **Enable**.

Status

Users can be set to one of the following statuses:

- **Enabled:** User can log in and use the Trifacta application normally.
- **Disabled:** User account has been disabled by an administrator. User cannot use the project or workspace.

NOTE: A disabled user's flows and datasets are still stored within the Trifacta application. However, the user cannot access them. Ownership of these objects has not been transferred. An administrator has ownership privileges on the user's objects.

Edit Users

To modify a user account, please complete the following steps.

NOTE: For security reasons, an administrator is not permitted to edit some settings in the administrator's own account.

Steps:

1. Locate the user in the list of users.
2. In the context menu on the right side of the user's listing, select **Edit**.
3. In the Edit User dialog, modify the following properties as needed:

Name: The display name of the user.

Email: The email address is used as the login identifier. This value cannot be modified.

Roles: Select or remove the roles to assign to the user. For more information, see *Roles Page*.

SSO Principal: If SSO is enabled, set this value to be the SSO principal value associated with this user.

NOTE: Required value for each user if SSO is enabled. See *Configure SSO for AD-LDAP*.

Hadoop Principal: If secure impersonation is enabled, set this value to be the Hadoop principal value associated with this user.

NOTE: The user principal value should not include the realm.

NOTE: Hadoop principal is a required value if secure impersonation is enabled. See *Configure for Secure Impersonation*.

NOTE: If Kerberos is enabled, verify that all user principals that use the platform are also members of the group of the keytab user.

Deployment management: When selected, this user is assigned the deployment role in the platform. In a Development environment, this role can be added to a user's account to enable access to the Deployment Manager.

NOTE: Deployment management user accounts are intended for managing production execution of flows. These users have a different and limited user interface in the Trifacta application. There should be a limited number of these accounts.

NOTE: Only platform administrators can assign the Deployment management role. Workspace admins cannot.

Tip: A deployment user should be assigned the flow author role. Lesser flow roles may prevent the deployment user from properly importing and managing flows. See *Roles Page*.

- In a Production environment where the Deployment Manager applies to the entire instance, this role does not apply.
- For more information, see *Configure Deployment Manager*.
- For more information on Deployment Manager, see *Overview of Deployment Manager*.

Platform admin: When selected, the user is granted admin privileges over the platform. These privileges include user administration, ability to modify platform settings, and permissions to use admin-only API endpoints.

NOTE: Avoid providing the Platform admin permission to a large number of users.

To save your changes, click **Edit user**.

User Details Page

Contents:

- *Group membership*
- *Roles*
- *Privileges*
- *User Details*

Review details about the selected user's account.

Users

N

nobody

nobody@trifacta.com

Edit

...

Roles

Default

Privileges

Flow author

Create, view, edit, delete, share, and execute flows

Connection author

Create, view, edit, delete, and share connections

Plan author

Create, view, edit, delete, and execute plans

User details

Status

Invited

Last log in

Today at 11:24 AM

Created on

Today at 11:24 AM

Figure: User Details Page

Context menu actions:

For each user, you can perform the following actions in the context menu:

- **Edit:** Modify user properties, including platform roles. See *Users Page*.
- **Configure storage:** If per-user access is enabled for the workspace, you can configure the access credentials for individual users, either using key-secret combinations or IAM roles. For more information, see *Configure Your Access to S3*.
- **Reset password:** Self-service password reset is enabled by default. If enabled, click this option to send an email to the user to reset his or her password.

NOTE: Only platform administrators can reset a user's password. Workspace admins cannot.

Disable: When a user is disabled, the user cannot access the Trifacta application.

- The disabled user still counts against the workspace limit.
- All of the user's flows and datasets are retained.
- Resources such as connections and flows that are owned by the user become inaccessible to other users that have access.

NOTE: Schedules owned by a disabled user continue to execute. An admin can disable the schedule. See *Schedules Page*.

- To permit access again, select **Enable**.

Group membership

Any group assignments are listed in this section.

NOTE: This feature is in Beta release.

For more information on groups, see *Configure Users and Groups*.

Roles

The roles assigned to the user are listed. For more information, see *Roles Page*.

Privileges

In this section, you can review the maximal set of privileges that are assigned to the user.

- Privileges are additive.
- For more information, see *Privileges and Roles Reference*.

User Details

Information on the current status and recent activity of the user. If the user has any platform roles, they are listed here. These roles can be enabled or disabled when you edit the user. For more information, see *Users Page*.

Roles Page

Through the Roles page, an admin can create roles and assign one or more of them to Trifacta users.

- A **role** is a set of privileges that can be assigned to one or more Trifacta users.
 - A **privilege** is a level of access to a type of user-generated object, such as flows.
 - For more information on these terms, see *Overview of Authorization*.
- For more information on managing roles, see *Create Role*.

You can also apply roles to groups that are synched from your enterprise LDAP provider. For more information, see *Configure Users and Groups*.




Roles Create role		
Name	Privileges	Last Updated
 default	Connection author, Plan author, Flow author	Last Sunday at 3:21 PM
 viewer_connections	Connection viewer	Today at 10:38 AM
 viewer_flows	Flow viewer	Today at 10:38 AM

Figure: Roles Page

The list of current roles is displayed in the Roles page. To create a new role, click **Create role**. See *Create Role Dialog*.

Columns:

- **Name:** The name of the role must be unique within the project or workspace.
- **Privileges:** The comma-separated list of privileges associated with the role. When a user is assigned the role, these privileges are available to the user.

Tip: Hover over the entry in the Privileges column to see additional detail on the privileges assigned to this role.

- **Last Updated:** Timestamp of when the role was most recently updated.

Context menu:

On the right side of the screen, you can select from a context menu for each available role.

- **Edit:** Modify the role. See *Create Role Dialog*.

NOTE: You cannot modify the admin role.

NOTE: All new and existing users are assigned the `default` role. Changes to this role may affect all existing users and any users that are invited in the future.

- **Assign role:** Assign the role to users.

NOTE: When you assign or un-assign a role, the privileges are immediately applied to the assigned user's account. The user does not need to re-login to see the changes.

- You can un-assign a role from users through the Role Details page. Select the role, and then click the Users tab. For more information, see *Role Details Page*.
- **Delete:** Delete the role.

You are permitted to delete roles that are currently assigned to users. Deleting a role may remove privileges from one or more users. This action cannot be undone. Before deleting, you should verify the list of users assigned to the role. For more information, see *Role Details Page*.

NOTE: You cannot delete the default or admin roles.

Create Role Dialog

To create a new role that you can assign to users, click **Create role** in the Roles page.

Create role

Name

viewer_flows

Privileges

Flows

Viewer

Connections

None

Plans

None

Cancel

Save

Figure: Create Role Dialog

In this dialog, you assign one or more privileges to the defined role.

Name: Enter a name for your role. This name must be unique within the roles in the current project or workspace.

Privileges:

Tip: You can create a role with no privileges, which may be useful for disabling access to objects without disabling the account itself. In this case, all other roles would need to be removed from the assigned user.

- **Flows:** These privileges govern the actions that users can perform on flows.
- **Connections:** These privileges govern the actions that users can perform on connections.
- **Plans:** These privileges govern the actions that users can perform on plans.
- For more information, see *Privileges and Roles Reference*.

To finish creating your role, click **Save**.

This role is now available for assigning to users. See *Roles Page*.

Role Details Page

Contents:

- *Overview tab*
 - *Users tab*
 - *Groups tab*
-

Through the Role Details page, you can review the privileges assigned to the role and assign the role to Trifacta users.

Actions:

The following actions are available.

- **Edit:** Modify the role.

NOTE: All new and existing users are assigned the `default` role. Changes to this role may affect all existing users and any users that are added in the future.

NOTE: You cannot edit the admin role.

See *Create Role Dialog*.

- **Assign role:** Assign the role to one or more users.

NOTE: When you assign or un-assign a role, the privileges are immediately applied to the assigned user's account. The user does not need to re-login to see the changes.

- **Delete:** Delete the role.

You are permitted to delete roles that are currently assigned to users. Deleting a role may remove privileges from one or more users. This action cannot be undone. Before deleting, you should verify the list of users assigned to the role in the Users tab.

NOTE: You cannot delete the `default` or admin roles.

Overview tab

In the Overview tab, you can review the privileges for the role and the current number of users that have been assigned the role.

Roles
Flow-Author

Edit
...

Overview
Users

Privileges

Flow author
View, execute, create, edit, delete, and share flows

Role Details

Users	1
Created On	07/02/2020
Last Updated	07/02/2020

Figure: Role Details Page - Overview tab

For more information on the listed privileges, see *Privileges and Roles Reference*.

Roles can be created through the Roles page. For more information, see *Roles Page*.

Users tab

In the Users tab, you can review the list of users who have been assigned the role.

Roles
Flow-Author

Edit
...

Overview
Users

Name	Email
Testuser 1234	testuser1234@trifacta.com

Figure: Role Details Page - Users tab

Columns:

- Name:** Display name of the user.

NOTE: You cannot modify the Name value for the default role.

- Email:** Email address for the user, which is used to login to the Trifacta application.

Context menu:

On the right side of the screen, you can select the following options from the context menu for each user:

- Unassign from role:** Select this option to remove the role from the user.

Groups tab

In the Groups tab, you can review the list of groups to which the role has been assigned.

NOTE: Optionally, users of the platform can be synched with your enterprise LDAP service provider. For more information, see *Configure Users and Groups*.

Workspace Settings Page

Contents:

- *General*
 - *Hide underlying file system to users*
 - *Locale*
 - *Self service password reset*
 - *Session Tags: Enable the use of session tags when assuming an IAM role*
 - *Session Tags: The name of the session tag that holds the username as its value*
 - *Session duration*
 - *Show file location*
 - *Storage directories*
 - *User messaging*
- *API*
 - *API Access Token*
 - *Allow users to generate access tokens*
 - *Maximum lifetime for user generated access tokens (days)*
- *Connectivity*
 - *Connectivity feature*
 - *Custom SQL query*
 - *Enable S3 connectivity*
 - *Enable conversion of standard JSON files via conversion service*
 - *Max endpoints per JDBC REST connection*
 - *Upgrade to connectivity*
 - *Upgrade to upload large files*
- *Flows, recipes and plans*
 - *Collaborative suggestions*
 - *Column from examples*
 - *Editor Scheduling*
 - *Export*
 - *Import*
 - *Maximum number of files to read in a directory for the initial sample*
 - *Plan feature*
 - *Sample downloads*
 - *Schematized output*
 - *Sharing*
 - *UI for range join*
 - *Webhooks*
- *Job execution*
 - *Combine Spark Transform and Profile jobs into one.*
 - *Custom Spark Options Feature*
 - *Databricks Cluster Policies*
 - *Databricks Job Management*
 - *Databricks Job Runs Submit Fallback*
 - *Logical and physical optimization of jobs*
 - *SQL Scripts*
 - *Schema validation*
 - *Schema validation: stop job if schema changes are found*
 - *Trifacta Photon execution*
 - *Spark Whitelist Properties*
- *Scheduling and parameterization*
 - *Include Hidden Files in Parameterization*
 - *Parameterization*
 - *Schedule list*
 - *Scheduling feature*
- *Publishing*
 - *Avro output format*

- CSV output format
- Hyper output format
- JSON output format
- Parquet output format
- Publish job results
- Publishing actions options
- Notifications
 - Email notifications
 - Email notifications: on job failure
 - Email notifications: on job success
 - Email notifications: on plan/flow share
- Experimental features
 - Cache data in the Transformer intelligently
 - Default language
 - Execution time threshold (in milliseconds) to control caching in the Transformer
 - Language localization
 - Show user language preference
 - Wrangle to Python Conversion

The following settings can be customized for the user experience in your workspace. When you modify a setting, the change is immediately applied to the workspace. To access the page, select **User menu > Admin console > Workspace settings**.

NOTE: Users may not experience the changed environment until each user refreshes the application page or logs out and in again.

Enablement Options:

NOTE: Any values specified in the Workspace Settings page applies exclusively to the specific workspace and override any system-level defaults.

Option	Description
Default	<p>The default value is applied. This value may be inherited from higher level configuration.</p> <div> <p>Tip: You can review the default value as part of the help text.</p> </div>
Enabled	<p>The setting is enabled.</p> <div> <p>NOTE: If the setting applies to a feature, the feature is enabled. Additional configuration may be required. See below.</p> </div>
Disabled	The setting is disabled.
Edit	Click Edit to enter a specific value for the setting.

General

Hide underlying file system to users

When enabled, workspace users cannot see locations in the default storage layer.

Locale

Set the locale to use for inferring or validating data in the application, such as numeric values or dates. The default is `United States`.

NOTE: After saving changes to your locale, refresh your page. Subsequent executions of the data inference service use the new locale settings.

For more information, see *Locale Settings*.

Self service password reset

When enabled, workspace users can reset their own passwords via link on the login page.

Session Tags: Enable the use of session tags when assuming an IAM role

If you are using IAM roles to request temporary credentials for access to AWS resources, you can enable the use of session tags to make those requests. When a **session tag** is submitted, the Trifacta user is provided access to AWS resources based on the user's corresponding permissions within AWS, instead of having to specify those permissions in Trifacta. This method leverages the existing permission infrastructure in your enterprise and simplifies the use of IAM roles in the Trifacta application.

NOTE: After enabling the use of session tags, you must spin up a new EMR cluster, which forces EMR to use the newly deployed credential provider JAR file.

NOTE: Additional configuration is required. For more information, see *Configure AWS Per-User Auth for Temporary Credentials*.

Session Tags: The name of the session tag that holds the username as its value

When Session Tags: Enable the use of session tags when assuming an IAM role is enabled, you must specify the name of the session tag to be submitted to AWS containing the username of the Trifacta user requesting resources. Default value is `trifacta-user`.

For more information, see *Configure AWS Per-User Auth for Temporary Credentials*.

Session duration

Specify the length of time in minutes before a session expires. Default is `10080` (one week).

Show file location

When enabled, workspace users can see the locations of source and output files within the application.

Storage directories

Allow members of the workspace to change paths to their upload and output results locations through their user profile.

For more information, see *Storage Config Page*.

User messaging

When enabled, workspace users can explore content through the Trifacta application.

API

API Access Token

When accessing the REST APIs, you can optionally use a token for simpler use and enhanced security.

NOTE: This feature may not be available in all environments.

NOTE: API access tokens must be enabled to use the API reference documentation available through the User menu.

For more information, see *Access Tokens Page*.

Allow users to generate access tokens

When enabled, individual workspace users can generate their own personal access tokens, which enable access to REST APIs. For more information, see *Manage API Access Tokens*.

Maximum lifetime for user generated access tokens (days)

Defines the maximum number of days that a user-generated access token is permitted for use in the product.

Tip: To permit generation of access tokens that never expire, set this value to -1.

For more information, see *Manage API Access Tokens*.

Connectivity

Connectivity feature

When enabled, workspace users can create connections to relational datasources.

NOTE: Disabling this feature hides existing relational connections.

See *Relational Access*.

Custom SQL query

When enabled, users can create custom SQL queries to import datasets from relational tables. For more information, see *Enable Custom SQL Query*.

Enable S3 connectivity

When enabled, base connectivity to S3 is enabled for workspace users.

NOTE: Additional platform configuration is required. See *S3 Access*.

Enable conversion of standard JSON files via conversion service

When enabled, the Trifacta application utilizes the conversion service to ingest JSON files and convert them to a tabular format that is easier to import into the application. For more information see *Working with JSON v2*.

NOTE: This feature is enabled by default but can be disabled as needed. The conversion process performs cleanup and re-organization of the ingested data for display in tabular format.

When disabled, the Trifacta application uses the old version of JSON import, which does not restructure the data and may require additional recipe steps to manually structure it into tabular format.

NOTE: The legacy version of JSON import is required if you are working with compressed JSON files or only Newline JSON files.

NOTE: Although imported datasets and recipes created under v1 of the JSON importer continue to work without interruption, the v1 version is likely to be deprecated in a future release. You should switch your old imported datasets and recipes to using the new version. Instructions to migrate are provided at the link below.

For more information, see *Working with JSON v1*.

Max endpoints per JDBC REST connection

For a REST API connection to a JDBC source, this parameter defines the maximum number of endpoints that can be defined to use the connection.

Avoid modifying this value unless you are experiencing timeouts or failures to connect.

For more information, see *REST API Connections*.

Upgrade to connectivity

When enabled, workspace users are presented with the option to upgrade to a plan that supports connection to external data sources, if the feature is current disabled.

Upgrade to upload large files

When enabled, workspace users are presented with the option to upgrade to a plan that supports uploading large files, if the feature is current disabled.

Flows, recipes and plans

Collaborative suggestions

If desired, you can enable the inclusion of suggestion cards that are generated from recent use of the Trifacta application. As the application gathers more information about how you or members of your workspace apply transformations to your data, the suggestions become more meaningful for the data that you are processing.

NOTE: No data is shared with Alteryx or any system outside of the Trifacta platform.

These collaborative suggestion cards can be generated from individual usage or from workspace level usage. These suggestions appear under the Recently Used heading in the side panel.

NOTE: This feature requires the machine learning service, which is enabled by default. For more information, see *Miscellaneous Configuration*.

When this feature is enabled, individual users can still choose to opt-out of sharing their usage data with this feature. See *User Profile Page*.

Option	Description
disabled	Collaborative suggestions are not surfaced in the application.
personal	Collaborative suggestions are based on the individual user's previous transformations.
workspace	Collaborative suggestions are based on the transformations from all users in the workspace.
Default	The default setting for the workspace is applied.

For more information, see *Overview of Predictive Transformation*.

Column from examples

When enabled, users can access a tool through the column menus that enables creation of new columns based on example mappings from the selected column. For more information, see *Overview of TBE*.

Editor Scheduling

When enabled, flow editors are also permitted to create and edit schedules. For more information, see *Flow View Page*.

NOTE: The Scheduling feature may need to be enabled in your environment. When enabled, flow owners can always create and edit schedules.

When this feature is enabled, plan collaborators are also permitted to create and edit schedules. For more information, see *Plan View Page*.

Export

When enabled, workspace users are permitted to export their flows and plans. Exported flows can be imported into other workspaces or product editions.

NOTE: If plans are been enabled in your workspace, enabling this flag applies to flows and plans.

- For more information, see *Export Flow*.
- For more information, see *Export Plan*.

Import

When enabled, workspace users are permitted to import exported flows and plans.

NOTE: If plans have been enabled in your workspace, enabling this flag applies to flows and plans.

- For more information, see *Import Flow*.
- For more information, see *Import Plan*.

Maximum number of files to read in a directory for the initial sample

When the Trifacta application is generating an initial sample of data for your dataset from a set of source files, you can define the maximum number of files in a directory from which the sample is generated. This limit is applied to reduce the overhead of reading in a new file, which improves performance in the Transformer page.

Tip: The initial sample type for files is generated by reading one file after another from the source. If the source is multiple files or a directory, this limit caps the maximum number of files that can be scanned for sampling purposes.

NOTE: If the files in the directory are small, the initial sample may contain the maximum number of files and less than the maximum size permitted for a sample. You may see fewer rows than expected.

If the generated sample is unsatisfactory, you can generate a new sample using a different method. In that case, this limit no longer applies. For more information, see *Overview of Sampling*.

Plan feature

When enabled, users can create plans to execute sequences of recipes across one or more flows. For more information, see *Plans Page*.

For more information on plans and orchestration, see *Overview of Operationalization*.

Sample downloads

When enabled, members can download the contents of the Transformer page at any time. For an individual step, a member can download the current sample, as modified by the current recipe up to the point of the current step. For more information, see *Recipe Panel*.

Schematized output

When enabled, all output columns for all types of outputs are typecast to their annotated types. This feature is enabled by default.

For non-schematized outputs, Trifacta enforces casting of all values to the annotated data type of the column by default. For example, if the output value is `-3.4` and the data type for the output column is Integer, the platform enforces Integer type casting and writes a null value instead.

- `true`: All output values must match the data type of the output columns, or a null value is written.
- `false`: All output values are written in their output form, regardless of the column's data type.

Sharing

When enabled, workspace users are permitted to share flows and plans with other users in the workspace.

NOTE: If plans have been enabled in your workspace, enabling this flag applies to flows and plans.

- For more information, see *Share a Flow*.
- For more information, see *Share a Plan*.

UI for range join

When enabled, workspace users can specify join key matching across a range of values. For more information, see *Configure Range Join*.

Webhooks

When enabled, webhook notification tasks can be configured on a per-flow basis in Flow View page. Webhook notifications allow you to deliver messages to third-party applications based on the success or failure of your job executions. For more information, see *Create Flow Webhook Task*.

NOTE: Additional configuration may be required. See *Configure Webhooks*.

Job execution

Combine Spark Transform and Profile jobs into one.

When enabled, the transform and profiling tasks of a job executed on the Spark running environment are combined. The profiling task is executed as a part of the transform task, which eliminates any time spent orchestrating the profiling task and accessing the profiler input file on storage.

NOTE: When these two tasks are combined, publishing actions are not undertaken if the profiling task fails.

For more information on these actions, see *Configure for Spark*.

In the Job Details page, combined jobs appear in a `Transform with profile` card. See *Job Details Page*.

Custom Spark Options Feature

When enabled, users can override Spark configuration options for output objects before running Spark jobs.

Tip: When enabled, a default set of Spark configuration options is available for users. Additional properties can be specified through the Spark Whitelist Properties setting.

See *Enable Spark Job Overrides*.

Databricks Cluster Policies

When enabled, this feature allows the Trifacta platform to leverage cluster policies that you have created for use when creating new Databricks clusters for job execution.

NOTE: You must create cluster policies before enabling this feature. Each user may select a cluster policy to use. Additional configuration and considerations may be required. If no policy is selected, jobs may still be executed.

For more information:

- *Configure for Azure Databricks*
- *Configure for AWS Databricks*

Databricks Job Management

Enables job execution on Databricks through a secondary method. When enabled, Databricks jobs are executed via the run/submit API endpoint, which avoids the job quota limitation imposed by Databricks clusters. This flag also enables deletion of Databricks jobs from the Databricks workspace.

For more information on these options:

- *Configure for Azure Databricks*
- *Configure for AWS Databricks*

Databricks Job Runs Submit Fallback

When this flag is enabled, users can execute Databricks jobs using the runs/submit API method as a fallback when the job quota limit is reached for a Databricks workspace.

For more information:

- *Configure for Azure Databricks*
- *Configure for AWS Databricks*

Logical and physical optimization of jobs

When enabled, the Trifacta application attempts to optimize job execution through logical optimizations of your recipe and physical optimizations of your recipes interactions with data.

NOTE: This feature requires the optimizer service, which is enabled by default, and the optimizer service database, which is installed by default. For more information on installing the database, see *Install Databases*.

This workspace setting can be overridden for individual flows. For more information, see *Flow Optimization Settings Dialog*.

SQL Scripts

When enabled, users may define SQL scripts to execute as part of a job's run. Scripts can be executed before data ingestion, after output publication, or both through any write-supported relational connection to which the user has access.

For more information, see *Create Output SQL Scripts*.

Schema validation

When enabled, by default the structure and ordering of columns in your import datasets are checked for changes before data is ingested for job execution.

Tip: This setting can be overridden for individual jobs, even if it is disabled. For more information, see *Run Job Page*.

Errors are immediately reported in the Job Details page. See *Job Details Page*.

For more information on schema validation, see *Overview of Schema Management*.

Schema validation: stop job if schema changes are found

When schema validation is enabled, this setting specifies the default behavior when schema changes are found.

- When enabled, jobs are failed when schema changes are found, and error messages are surfaced in the Trifacta application.
- When disabled, jobs are permitted to continue.
 - Jobs may ultimately fail due to schema changes.
 - Jobs may result in bad data being written in outputs.
 - Job failures may be more challenging to debug.

Tip: Setting this value to `Disabled` matches the behavior of the Trifacta application from before schema validation was possible.

Tip: This setting can be overridden for individual jobs, even if it is disabled. For more information, see *Run Job Page*.

Errors are immediately reported in the Job Details page. See *Job Details Page*.

For more information on schema validation, see *Overview of Schema Management*.

Trifacta Photon execution

When enabled, users can choose to execute their jobs on Trifacta Photon, a proprietary running environment built for execution of small- to medium-sized jobs in memory on the Trifacta node.

Tip: When enabled, you can select to run jobs on Photon through the Run Jobs page. The default running environment is the one that is best for the size of your job.

When Trifacta Photon is disabled:

- You cannot run jobs on the local running environment. All jobs must be executed on a clustered running environment.
- Trifacta Photon is used for Quick Scan sampling jobs. If Trifacta Photon is disabled, the Trifacta application attempts to run the Quick Scan job on another available running environment. If that job fails or no suitable running environment is available, the Quick Scan sampling job fails.

For more information, see *Run Job Page*.

Spark Whitelist Properties

Comma-separated list of additional Spark properties to be whitelisted for configuration of output objects while running Spark jobs.

NOTE: The Custom Spark Options feature must be enabled.

See *Enable Spark Job Overrides*.

Scheduling and parameterization

Include Hidden Files in Parameterization

When enabled, hidden files and hidden directories can be searched for matches for wildcard- or pattern-based parameters when importing datasets.

Tip: This can be useful for importing data from generated profiles, which are stored in the `.profiler` folder in a job output directory.

NOTE: Scanning hidden folders may impact performance. For existing imported datasets with parameters, you should enable the inclusion of hidden folders on individual datasets and run a test job to evaluate impact.

For more information, see *Parameterize Files for Import*.

Parameterization

By default, Trifacta supports the application of parameters to imported datasets. Datetime, wildcard, or variable parameters can be used to operationalize execution of jobs on different versions of the same dataset.

When enabled, users can create parameters, which can be applied to import, creating sample, and outputs. For more information, see *Overview of Parameterization*.

Schedule list

When enabled, administrators and workspace administrators can see a list of all schedules in the workspace.

Scheduling feature

When enabled, workspace users can schedule the execution of flows. See *Add Schedule Dialog*.

Publishing

Avro output format

When enabled, members can generate outputs in Avro format.

CSV output format

When enabled, members can generate outputs in CSV format.

Hyper output format

When enabled, members can generate outputs in Hyper format for publication and use on Tableau Server.

JSON output format

When enabled, members can generate outputs in JSON format.

Parquet output format

When enabled, members can generate outputs in Parquet format.

Publish job results

When enabled, workspace users are permitted to publish results through the Output Destinations tab in the Job Details page to external datastores.

NOTE: These external datastores must be enabled and configured. See *Connection Types*.

For more information, see *Job Details Page*.

Publishing actions options

When enabled, users are permitted to create custom publishing actions for their jobs.

When disabled, users must accept the default publishing actions.

For more information, see *Run Job Page*.

Notifications

Email notifications

When enabled, Trifacta can send email notifications to users based on the success or failure of jobs.

NOTE: This feature requires access to an SMTP server to send emails. For more information, see *Enable SMTP Email Server Integration*.

Email notifications: on job failure

When email notifications are enabled, you can configure the default setting for the types of failed jobs that generate an email to interested stakeholders. The value set here is the default value for each flow in the workspace.

Settings:

Setting	Description
Never	Email notifications are never sent for job failures.
Scheduled	Notifications are sent only when scheduled jobs fail.
Ad hoc	Notifications are sent only when ad-hoc (manually executed) jobs fail. <div>Tip: Jobs executed via API are Ad hoc jobs.</div>
All	Notifications are sent for all job failures.
Default	Notifications are sent based on the default settings for the product.

Default setting for job failures is *scheduled*.

Individual users can opt out of receiving notifications or configure a different email address. See *Email Notifications Page*.

Emailed stakeholders are configured by individual flow. For more information, see *Manage Flow Notifications Dialog*.

Email notifications: on job success

When email notifications are enabled, you can configure the default setting for the types of successful jobs that generate an email to interested stakeholders. The value set here is the default value for each flow in the workspace.

For more information on the settings, see the previous section. Default setting is *never*.

Individual users can opt out of receiving notifications or configure a different email address. See *Email Notifications Page*.

Emailed stakeholders are configured by individual flow. For more information, see *Manage Flow Notifications Dialog*.

Email notifications: on plan/flow share

When email notifications are enabled, users automatically receive notifications whenever an owner shares the plan or flow with the user.

Individual users can opt out of receiving notifications. For more information, see *Preferences Page*.

Experimental features

These experimental features are not supported.

Experimental features are in active development. Their functionality may change from release to release, and they may be removed from the product at any time. Do not use experimental features in a production environment.

These settings may or may not change application behavior.

Cache data in the Transformer intelligently

NOTE: This feature is in Beta release.

When enabled, this feature allows the Trifacta application to cache data from the Transformer page periodically based on Trifacta Photon execution time. This feature enables users to move faster between recipe steps.

Default language

Select the default language to use in the Trifacta application.

Execution time threshold (in milliseconds) to control caching in the Transformer

NOTE: This feature is in Beta release.

When intelligent caching in the Transformer is enabled, you can set the threshold time in milliseconds for when Trifacta Photon updates the cache. At each threshold of execution time in Trifacta Photon, the output of the intermediate recipe (CDF) steps are cached in-memory, which speeds up movements between recipe steps in the Trifacta application.

Language localization

When enabled, the Trifacta application is permitted to display text in the selected language.

Show user language preference

When enabled, individual users can select a preferred language in which to display text in the Trifacta application.

NOTE: This experimental feature requires installation of a language resource file on the Trifacta node. For this release, only U.S. English (default) and Korean are supported. For more information, please contact *Alteryx Support*.

Users can make personal language selections through their preferences. See *Account Settings Page*.

Wrangle to Python Conversion

Alpha feature: When enabled, you can use an API endpoint to generate Python Pandas code that completes the steps required to generate an output in Python.

This feature may be modified or removed in a future release without warning. It is intended for demonstration purposes only and should not be enabled in a production environment.

For more information, see *API Workflow - Wrangle Output to Python*.

Tip: You can download and install the Python SDK to integrate use of the Trifacta application in your Python environment. Use the visual tools of the Trifacta application to build your transformations, and then generate Python Pandas code for use in your Python data pipelines. For more information, see *Python SDK*.

Admin Settings Page

Contents:

- Platform Settings
- External Service Settings
 - AWS EMR Cluster ID
 - AWS Region
 - Resource Bucket
 - Resource Path
- Services
 - View Logs
- Tricheck
- SMTP settings
- Upload License
- Restart

Admin users of the Trifacta® platform can change settings through the Trifacta application. Login as an admin user. Select **User menu > Admin Console > Admin Settings**.

NOTE: You must be an administrator to access this feature.

Platform Settings

Do not modify settings through the Admin Settings page and through `trifacta-conf.json` at the same time. Saving changes in one interface wipes out any unsaved changes in the other interface. Each requires a platform restart to apply the changes.

Platform administrators can change any parameter value that is available through the web application. Enter some or all of parameter name into the search box to see a set of possible matches.

Do not modify parameters with which you are unfamiliar or have not been instructed to change. Some changes can be harmful to the system. In particular, changing the database connection parameters can break access to the application and the Admin Settings page.

Search:

Tip: You can copy setting names from the documentation to search the available set. Search retrieves matches from the setting name and the help text for the parameter. Do not paste in double quotes from documentation samples.

If your search for a parameter comes up empty and you know that the parameter exists, you must make changes on the Trifacta node in `trifacta-conf.json`. See *Required Platform Configuration*.

Search groupings:

If you search for the following strings, which may appear in property descriptions, you can review groups of settings pertaining to the configuration areas listed below.

NOTE: Do not perform configuration of these areas by simply reviewing and modifying the settings in these parameter groups. Additional configuration may be required. Some required settings may not be grouped, and some of these settings may not be documented. Please review the related documentation sections.

Search string	Setting group
[CORE]	Core platform settings.
[DISTRO]	Settings pertaining to specific distributions. See <i>Configure for Cloudera</i> .
[CLUSTER]	Settings that affect how the platform interacts with the integrated backend cluster. See <i>Prepare Hadoop for Integration with the Platform</i> .
[HIVE]	Settings pertaining to integration with the connected instance of Hive. See <i>Configure for Hive</i> .
[HA]	Settings pertaining to integration with cluster high availability for the Trifacta platform. See <i>Enable Integration with Cluster High Availability</i> .
[SECURITY]	General settings pertaining to security. See <i>Configure Security</i> .
[SSL]	Settings pertaining to SSL access to the platform. See the SSL section in <i>Configure Security</i> .
[ADVANCED]	Advanced settings.

When you modify a setting, your change is validated against the data type or set of accepted values. String-based entries cannot be validated.

Notes:

- Sensitive information is obfuscated in the display values in the Admin Settings page.
- To save changes, click **Save**.

NOTE: Saving changes forces an automatic type validation of the configuration and a restart of the platform, which terminates any active user sessions.

NOTE: Environmental validation is not performed as part of changes in this user interface. For example, you can change the port number for the Trifacta application to an invalid value and save the configuration change. However, when the platform is restarted, the application fails to start, and you cannot continue. In this case, you must fix the problem in `trifacta-conf.json`.

External Service Settings

AWS EMR Cluster ID

If you have deployed your instance of the Trifacta platform on to Amazon Web Services (AWS) and are connected to an Elastic Map Reduce (EMR) cluster, you can review or modify the cluster identifier in this location. For example, in the event of prolonged outage or failure of the original cluster, you can insert the cluster ID of a secondary cluster to effectively failover to the new cluster.

NOTE: When you first install and integrate with an EMR cluster, this identifier is stored in the Trifacta database for you. It should be modified only if you need to switch to a different EMR cluster. Only one EMR cluster can be active at any time.

NOTE: If this cluster ID is modified, you must modify any other EMR-related settings to match the corresponding values for the new cluster. Please search for `emr` among the admin settings.

When you have entered a new cluster ID, click **Save**.

NOTE: For this setting, you do not have to click the Save button at the bottom of the screen, which restarts the Trifacta platform.

AWS Region

Enter the AWS region code where the EMR cluster is hosted. For example:

```
us-east-1
```

For a list of available regions, see <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-available-regions>.

Resource Bucket

The name of the default S3 bucket where platform resources are stored

Resource Path

The path in the default S3 bucket to where resources are stored

After you have made any changes to the AWS properties, click **Save**.

Services

You can review overall status of the Trifacta platform.

View Logs

Click the View Logs link to review and download the logs maintained on the Trifacta node.

For more information on these logs, see *System Services and Logs*.

For more information, see *Configure Logging for Services*.

Tricheck

Tricheck performs a variety of tests of your environment to determine its suitability for use with the Trifacta platform.

Tip: Tricheck should be run immediately after the Trifacta software has been installed or upgraded or whenever there are significant changes to the node or its connected cluster.

Checks include but are not limited to:

- Sufficient hardware resources on the Trifacta node
- Supported versions of software installed on the Trifacta node
- Access to required ports and all nodes of the cluster
- Trifacta node system profiling

NOTE: Tricheck performs no data-dependent checking. It cannot assess suitability of the environment for specific data volumes, connections, or data types.

Click **Run Tricheck** to run checks and download the output log.

SMTP settings

Use this option to send a test email to the specified address through the SMTP email server that has been configured for the Trifacta application to use.

NOTE: The SMTP email server to use must be configured. For more information, see *Enable SMTP Email Server Integration*.

Steps:

1. To test the configured settings, click **Check email (SMTP) settings**.
2. Enter a valid email address. Then, click **Check settings**.
3. If the SMTP server is configured properly, a test email is sent to the specified email address.

Upload License

NOTE: For more information on acquiring an updated license file, please contact *Alteryx Support*.

You can update the license file stored on the Trifacta node. Click **Upload License** to browse for and select the license file.

NOTE: By default, the platform checks for a valid license once per hour. To apply your uploaded license immediately, please restart the platform.

For more information on your license, see *License Key*.

Restart

Click **Restart Trifacta** to immediately restart the platform.

Tip: The Restart Trifacta button is the preferred method for restarting the platform. A restart is automatically executed when you save changes to the platform settings.

NOTE: This button may not be available in high availability environments. In those deployments, please restart individual services or use the command line command. For more information, see *Start and Stop the Platform*.

AWS Settings Page

Contents:

- *Workspace Mode*
 - *IAM Role Settings*
 - *AWS Key and Secret Settings*
- *Per-User Mode*
- *Common Settings*
 - *S3 Buckets*
 - *Server-side Encryption*

In the AWS Settings page, workspace administrators can define the AWS credentials mode for the workspace and apply settings for the selected mode, including selecting the credential provider. From the left menu, select **User menu > Admin console > AWS settings**.

NOTE: Before you begin, some information must be gathered from AWS. See *Enable Access to S3 and AWS Resources*.

NOTE: This configuration section applies only if Trifacta® is integrated with Amazon Web Services.

AWS Mode:

Mode	Description
Workspace	<p>In Workspace mode, the workspace administrator applies a single set of AWS credentials for all users in the workspace. These credentials are used by each member of the workspace to authenticate with AWS and to gain access to S3 buckets.</p> <p>Tip: This mode requires more up-front setup but is easy to manage. However, all members of the workspace have the same set of access controls.</p>
Per User	<p>In Per User mode, individual members of the workspace must apply their AWS credentials to their accounts.</p> <p>Tip: This mode is easy to set up but turns responsibility for access controls over to the individual members. If members encounter problems gaining access to S3 assets, the workspace administrator may not be able to troubleshoot them.</p>

Credential Provider:

For workspace or per-user mode, the following provider methods can be used to manage authentication with AWS.

Credential Provider	Description
IAM Role	<p>Trifacta can use any IAM roles that have been defined for workspace members to access AWS data sources, such as S3 and Redshift.</p> <p>Tip: This credential provider method is recommended.</p>

AWS Key and Secret	You can apply key and secret combinations to gate access to AWS data sources. These combinations can be applied in workspace mode or in per-user mode by individual members.
--------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Workspace Mode

In workspace mode, you must select the credential provider and then specify the relevant settings.

IAM Role Settings

Prerequisites:

- The IAM roles must include a trust relationship for Trifacta. For more information, see *Insert Trust Relationship in AWS IAM Role*.
- If you want workspace members to be able to use the on-boarding walkthrough, they must have access to the Trifacta assets required for the walkthrough. For more information, see *Required AWS Account Permissions*.

Apply the following settings to define the IAM roles and related settings.

Setting	Description
Account ID	This value is pre-populated when the workspace is created. <div data-bbox="334 873 1458 957"> NOTE: Do not modify. </div>
External ID	This value is pre-populated when the workspace is created. <div data-bbox="334 1041 1458 1125"> NOTE: Do not modify. </div>
Available IAM Role ARNs	You can specify the set of IAM Role ARNs from which users can select for their access to AWS resources. <div data-bbox="334 1209 1458 1310"> NOTE: These roles cannot be modified if SAML passthrough authentication has been enabled. For more information, see <i>Configure for AWS SAML Passthrough Authentication</i>. </div>
Select Default IAM Role ARN	From the available IAM Role ARNs, you can specify the default one.

AWS Key and Secret Settings

For key-secret authentication to AWS, please specify the following settings.

NOTE: The AWS key and secret must provide read/write access to the default S3 bucket at least.

The account must have the ListAllMyBuckets ACL among its permissions.

Setting	Description
AWS Access Key	The AWS access key to use.
AWS Secret Key	The AWS secret associated with the access key.

Per-User Mode

For per-user mode:

- The workspace administrator must specify the encryption settings only. See below.
- Individual users configure all of the other AWS access settings through the Storage configuration page.

Common Settings

S3 Buckets

For key-secret authentication to AWS, please specify the following settings.

Setting	Description
Default S3 bucket	<p>Specify the name of the default S3 bucket.</p> <div>NOTE: Specify the top-level bucket name only. There should not be any backslashes in your entry.</div>
Extra S3 buckets	You can specify any additional S3 buckets in a comma-separated list of names.

Server-side Encryption

Trifacta supports the use of server-side encryption when writing results.

NOTE: When encryption is enabled, all buckets to which you are writing must share the same encryption policy.

Setting	Description
Encryption Type	<p>Supported encryption types:</p> <ul style="list-style-type: none">• None• SSE-S3• SSE-KMS
KMS Key ID	If SSE-KMS has been selected, you can paste the KMS Key ID value in this field.

Environment Parameters Page

Through the Environment Parameters page, you can create parameters that are defined for use throughout the workspace or project. Environment parameters can be exported as a set from one environment and imported for use in another.

NOTE: You must be a project owner or workspace administrator to access this page.

Environment parameters

Create

...

Search parameters /

Name	Default value	Last updated
<> env.bucket_key20210611185636	qa	Today at 9:57 AM
<> env.year120210611185636	2010	Today at 9:57 AM
<> env.bucket_key20210611185213	qa	Today at 9:53 AM
<> env.year120210611185213	2010	Today at 9:53 AM
<> env.year20210611183222	2010	Today at 9:33 AM
<> env.bucket_key20210611183222	qa	Today at 9:33 AM

Figure: Environment Parameters Page

Columns:

- **Name:** Internal name of the environment parameter.
- **Default value:** The default value for the parameter.
- **Last updated:** Timestamp for which the parameter was last modified.

Actions:

- **Create:** Click to create a new environment parameter. Specify:
 - **Name:** Name of the new environment parameter.

NOTE: To distinguish them from other parameters, the prefix `env.` is added to all environment parameter names.

- **Default value:** The default value for the new environment parameter.
- To save your new environment parameter, click **Save**.
- **Import:** Select **Import** to import a ZIP file containing a set of environment parameters that were exported from another environment. For more information, see *Manage Environment Parameters*.
- **Export:** Select **Export** to download a ZIP file containing the definitions and default values for the environment parameters in this environment.
- **Search:** Enter a search string to locate environment parameters by name.

Context menu:

The following options are available for parameters that have been created.

- **Edit value:** Enter a new default value for the environment parameter. Then, click **Save**.
- **Delete:** Delete the environment parameter.

When an environment parameter is deleted, all references to the environment parameter are rendered as an empty string, which may result in broken imported datasets and outputs and other unpredictable issues.

OAuth 2.0 Clients Page

Through the OAuth 2.0 Clients page, workspace administrators can create and manage OAuth 2.0 clients for authentication with third-party systems. In the Admin console, select **OAuth 2.0 Clients**.

OAuth 2.0 clients

Register OAuth 2.0 client

Connect to other applications using the OAuth 2.0 authorization framework. [Learn more about OAuth 2.0 client creation.](#)

Name	Type	Last updated
Test2	snowflake	Today at 2:35 PM
Test	snowflake	Today at 2:35 PM

Figure: OAuth 2.0 Clients page

Columns:

- **Name:** Display name of the client.
- **Type:** Type of client. For more information on supported clients, see *Create OAuth2 Client*.
- **Last update:** Timestamp for last modification to the listed client.

Actions:

- To create an OAuth 2.0 client, click **Register OAuth 2.0 Client**. For more information, see *Create OAuth2 Client*.
- To delete a client, select **Delete OAuth 2.0 Client** from the context menu.

Deleting an OAuth 2.0 client cannot be undone. When a client is deleted, any connections that utilized the client no longer work. Datasets and output locations may no longer be accessible through the application.

Workspace Admin Permissions

Contents:

- *Configuration*
 - *User Management*
 - *Object Access*
 - *Data Access*
-

The workspace admin user is super-user for the entire workspace.

NOTE: In Trifacta, any user who is granted the admin role is also granted the workspace admin role, which enables owner-level access to some object types in the workspace. Details are below.

Configuration

The workspace admin can enable and disable features and capabilities in the workspace. For more information, see *Workspace Settings Page*.

User Management

A workspace admin can administer all of the other users of the workspace, including disabling or deleting the user.

NOTE: In Trifacta, the workspace admin can also edit the platform roles assigned to individual users.

See *Users Page*.

Object Access

A workspace admin has owner-level access to objects in the workspace.

NOTE: A workspace admin can access these objects like their owners, even if the objects have not been shared.

This access applies, but is not limited, to the following types of objects:

- Flows
- Connections (see below)
- Output objects
- Job profiles and results
- Plans and tasks

A workspace admin has collaborator-level access to the following objects:

- Imported datasets
- Macros
- Schedules

A workspace admin does not have any changed permissions for access to the following object types:

- Deployments and releases

Data Access

The workspace admin can access the data of individual users under the following conditions.

NOTE: Workspace admin privileges do not affect access permissions on outside storage systems. Those systems can prevent data access by the workspace admin user.

Connections with credentials:

If the data is accessed through a connection that requires a specific set of credentials, then the workspace admin can access all data available through the connection when the credentials are shared.

If connection credential sharing is disabled after a connection has already been shared with credentials, then the connection remains accessible to the workspace admin and to all users who were previously shared the connection. Workspace admins created in the future also inherit this access. The sharing of a connection's credentials cannot be revoked, except by deleting the connection.

A workspace admin cannot:

1. Modify or remove the shared credentials.
2. Change the credential sharing on another user's connection.

If a connection with shared credentials remains after credential sharing has been disabled, you can do one of the following for the connection:

- Edit the connection to use credentials that are safe to share with all affected users.
- Create a duplicate connection with private credentials. Delete the old shared connection.

For more information on credential sharing, see *Configure Sharing*. **File-based backend storage:**

Source datasets and job results that are stored on file-based backend storage systems for individual users can be accessed by the workspace admin except in the following situations:

- If users have user-specific access controls to the storage, such as secure impersonation, the workspace admin can only access a user's data if the admin's own permissions enable it.
- Directory permissions on user directories may prevent the workspace admin from accessing a user's data. For example, the workspace admin user can see the link to a user's job results that were written on the backend storage. However, when the workspace admin attempts to download those results, a permissions error is displayed, since the workspace admin user does not have permissions on the directory.

Relational connections:

Data is accessible under the connections with credentials limitations described above.

Admin Download Logs Dialog

Contents:

- *Logs by Timeframe*
- *Logs by Job ID*
- *Logs by Session ID*
- *Download*

Administrators of Trifacta® can download log files based on a user's session identifier, a job identifier, or across the Trifacta platform for a specified time period. From the Help menu, select **Download logs**. The data downloaded from this dialog is encrypted by default.

NOTE: The files download through this dialog are always unencrypted.

NOTE: For more information on disabling this feature, see *Configure Support Bundling*.

Non-administrators can download logs for their current session. For more information, see *Download Logs Dialog*.

Download logs ×

Collect logs by

Time frame ▾

Collect all log files available for a specific time period

Time frame

Last ▾

2 Hours ▾

Log file size limit ⓘ

1000000 Bytes

Cancel Download logs

Figure: Download Logs Dialog for Admins

Collect logs by: Select the method by which log files for the Trifacta platform are collected.

Log file size limit: You can specify the size limit of individual log files in bytes. The default size is 1 MB.

Logs by Timeframe

You can download logs by specific time period. You can select the last few hours or even customize the date range to download the log files.

Tip: Try to narrow the time frame if possible. Larger time frames are more likely to run up against the size limit for individual log files.

Steps:

1. To download the log files by time frame, select the Time frame option from the **Collect logs by** drop-down.
2. Select the required option from the **Time frame** drop-down. The following are the available options:
 - a. **Last:** Preceding number of minutes, hours, or days.
 - b. **Between:** Use the date and time tools to specify the starting (top) and ending (bottom) dates for the range.
3. To control the size of the log files, you can specify the size limit of individual log files in bytes. The default size is 1 MB.
4. To download logs, click **Download logs**. Logs are downloaded as a ZIP file.

Logs by Job ID

You can download log files for a specific job ID.

Tip: Administrators can review the IDs for all accessible jobs in the Jobs page. See *Jobs Page*.

Steps:

1. To download the log files by Job ID, select the Job ID option from the **Collect logs by** drop-down.
2. In the Job ID field, enter the required job ID.

NOTE: The Job ID is a string of numbers that can be found in the Jobs page or Job emails.

3. To control the size of the log files, you can specify the size limit of individual log files in bytes. The default size is 1 MB.
4. To download logs, click **Download logs**. Logs are downloaded as a ZIP file.

Logs by Session ID

You can download the logs files for a specified user session.

Tip: Non-admin users can retrieve their session ID from the application. For more information, see *Download Logs Dialog*.

Steps:

1. To download the log files by Job ID, select the Job ID option from the **Collect logs by** drop-down.
2. From the Session ID drop-down, select any one of the available options:
 - a. **I've got a session ID:** If you know the session ID, enter the same in the **Enter Session ID** field.
 - b. **Use my current session ID:** If you want to use your current session ID, select this option and the current session ID is automatically populated.

3. To control the size of the log files, you can specify the size limit of individual log files in bytes. The default size is 1 MB.
4. To download logs, click **Download logs**. Logs are downloaded as a ZIP file.

Download

To download the specified set of logs, click **Download logs**.

For more information on the contents of this download, see *Support Bundle Contents*.

For more information on configuring the contents of the support bundle, see *Configure Support Bundling*.

Required AWS Account Permissions

Contents:

- S3
 - *Read-only access policies*
 - *Write access policies*
 - *Other AWS policies for S3*
 - Redshift
 - Snowflake
 - EMR
-

To access the following AWS resources, you must configure your AWS account or accounts with the listed permissions. These permissions can be applied through AWS access key/secret combinations or through IAM roles applied to the account.

S3

All access to S3 sources occurs through a single AWS account (system mode) or through an individual user's account (user mode). For either mode, the AWS access key and secret combination must provide access to the default bucket associated with the account.

Read-only access policies

NOTE: To enable viewing and browsing of all folders within a bucket, the following permissions are required:

- The system account or individual user accounts must have the `ListAllMyBuckets` access permission for the bucket.
- All objects to be browsed within the bucket must have Get access enabled.

The policy statement to enable read-only access to your default S3 bucket should look similar to the following. Replace `3c-my-s3-bucket` with the name of your bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::3c-my-s3-bucket",
        "arn:aws:s3:::3c-my-s3-bucket/*"
      ]
    }
  ]
}
```

Write access policies

Write access is enabled by adding the `PutObject` and `DeleteObject` actions to the above. Replace `3c-my-s3-bucket` with the name of your bucket:

Other AWS policies for S3

Policy for access to Trifacta public buckets

To access S3 assets that are created by Alteryx, you must apply the following policy definition to any IAM role that is used to access the Trifacta. These buckets contain demo assets:

NOTE: Product walkthroughs must be enabled. For more information, see *Workspace Settings Page*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::trifacta-public-datasets/*",
        "arn:aws:s3:::trifacta-public-datasets"
      ]
    }
  ]
}
```

For more information on creating policies, see <https://console.aws.amazon.com/iam/home#/policies>.

KMS policy

If any accessible bucket is encrypted with KMS-SSE, another policy must be deployed. For more information, see <https://docs.aws.amazon.com/kms/latest/developerguide/iam-policies.html>.

Attribute-based access to S3


If you are using IAM roles to provide per-user access to S3, you can apply AWS session tags to any request for S3 resources, which allows you to leverage your enterprise permissioning to S3 based on the user identifier. IAM policies must be modified, and this feature must be enabled. For more information, see *Configure AWS Per-User Auth for Temporary Credentials*.

Redshift

Since Redshift requires S3 to be used, to enable read/write access to Redshift using an IAM role, the sole additional requirement to the above is to add the `GetClusterCredentials` permission to the IAM role used for S3. A policy statement similar to the following example needs to be included as part of any IAM role used by the Trifacta platform users to access AWS resources.

The following example policy adds the `GetClusterCredentials` permission for the specified AWS user (`aws:userid`). This user is permitted to get cluster credentials for three different resources:

- a personal Redshift cluster
- The `testdb` cluster
- The `common_group` cluster



For more information on `getClusterCredentials`, see https://docs.aws.amazon.com/redshift/latest/APIReference/API_GetClusterCredentials.html.

Snowflake

If you are creating a connection to your AWS-based Snowflake deployment, you must specify the following policies in the operative IAM role(s) for each S3 bucket:

Stage bucket

If you are creating your own Snowflake stage, it must point to the default S3 bucket in use by Trifacta. The policy that you created for read-write access to S3 should be applied to the Snowflake user.

NOTE: If users in your deployment are using IAM roles in user mode for AWS access, then the Snowflake stage must have permissions to write to the user's S3 bucket.

Snowflake bucket

You must create a separate policy to permit access to the S3 bucket that backs your AWS-based Snowflake deployment. The following example permission provides the minimum set of permissions.

Notes:

- The `s3:GetBucketLocation` is required for access to the S3 bucket that Snowflake requires for itself.
- The additional `s3:PutObject` and `s3:DeleteObject` permissions are required only if you plan to unload files to the bucket or automatically purge the files after loading them into a table.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws:s3:::<snowflake_bucket_name>/<prefix>/*"
    },
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::<snowflake_bucket_name>",
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "<prefix>/*"
          ]
        }
      }
    }
  ]
}
```

Where:

- `<snowflake_bucket_name>` = the name of the S3 bucket that is used by Snowflake
- `<prefix>` = the folder path prefix within the bucket. This value can be omitted if it is not required.
 - The above `StringLike` definition grants access to all prefixes on the bucket.

NOTE: If your bucket or prefixed path contains more than 1000 files, you may encounter the following error: Access Denied (Status Code: 403; Error Code: AccessDenied) .

- To address the above error, specify the `StringLike` condition with the following change. This change allows access to all files while eliminating the condition that causes the above error:

```
"Condition": {
  "StringLike": {
    "s3:prefix": [
      "*"
    ]
  }
}
```



```
}  
  }  
]
```

For more information, see <https://docs.snowflake.com/en/user-guide/data-load-s3-config-aws-iam-user.html>.

EMR

Additional permissions to access EMR depend on how the Trifacta deployment is configured to interact with EMR. For more information, see *Configure for EMR*.

Privileges and Roles Reference

Contents:

- *Privileges*
 - *Flows*
 - *Connections*
 - *Plans*
 - *Standard Roles*
 - *default*
 - *Workspace admin*
-

In the Trifacta® application, you can create and assign roles, each of which consists of one or more privileges. A **privilege** is a level of access to a type of object, such as flows.

Below, you can review the available privileges, including the supported levels for each.

For more information on privileges and roles, see *Overview of Authorization*.

Privileges

Flows

The flows privilege governs access to flow objects.

Access Level	Name	Description
0	none	Assigned role cannot see or use flows, including the pages where flows are available.
1	viewer	Assigned user can access Flows page and Flow View page for flows that the user owns or has been shared. User can also run jobs on the user's own flows. User cannot make changes to any flows.
2	editor	All of the above, plus: Assigned user can edit, share, and run jobs on flows to which the user has access. <div>NOTE: By default, editors can also schedule flows. This option can be disabled by an administrator.</div>
3	author	All of the above, plus: Assigned user can create new flows, schedule flows, and delete flows.

Tip: If you have enabled deployment management, a deployment user should be assigned author-level access. Lesser flow roles may prevent the deployment user from properly importing and managing flows. See *Roles Page*.

Connections

The connections privilege governs access to connection objects.

Access Level	Name	Description
--------------	------	-------------

0	none	Assigned role cannot see or use connections, including the pages where connections are available.
1	viewer	Assigned user can access Connections page for connections that the user owns or has been shared. User can share connections. User cannot make changes to any connections.
2	editor	All of the above, plus: Assigned user can edit and share connections to which the user has access.
3	author	All of the above, plus: Assigned user can create new connections and delete connections.

Plans

The plans privilege manages access to plan objects.

Access Level	Name	Description
0	none	Assigned role cannot see or use plans, including the pages where plans are available.
1	viewer	Assigned user can access Plans page and Plan View page for plans that the user owns or has been shared. User can also run jobs on the user's own plans. User cannot make changes to any plans.
2	editor	All of the above, plus: Assigned user can edit, share, and run jobs on plans to which the user has access. NOTE: By default, editors can also schedule plans. This option can be disabled by an administrator.
3	author	All of the above, plus: Assigned user can create new plans, schedule plans, and delete plans.

Standard Roles

The following roles are provided with the product.

NOTE: The following roles cannot be removed.

default

The default role is assigned to each user when the user is initially created. This role contains the following permissions:

Privilege	Access Level/Name
Flows	3 - author
Connections	3 - author
Plans	3 - author
User defined functions	3 - author

Tip: You can modify the default role if you want to set a lower level of base access for each new user of the product. For more information, see *Overview of Authorization*.

Workspace admin

This role provides super-user privileges to the assigned user.

NOTE: This role enables for the user owner-level access to all objects in the project or workspace and access to all admin-level settings and configuration pages in the admin console. This role should not be assigned to many users. At least one user should always have this role.

NOTE: You cannot modify or delete this role.



Copyright © 2022 - Trifacta, Inc.
All rights reserved.