

# Oracle® Big Data Discovery

Getting Started Guide

Version 1.2.2 • June 2016

## Copyright and disclaimer

Copyright © 2015, 2016, Oracle and/or its affiliates. All rights reserved.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. UNIX is a registered trademark of The Open Group.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software or hardware and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

# Table of Contents

<b>Copyright and disclaimer</b> .....	<b>2</b>
<b>Preface</b> .....	<b>5</b>
About this guide .....	5
Audience .....	5
Conventions .....	5
Contacting Oracle Customer Support .....	6
<b>Chapter 1: Welcome to Big Data Discovery</b> .....	<b>7</b>
Ways of using Big Data Discovery .....	7
Value of using Big Data Discovery .....	8
Addressing your goals and needs .....	9
Results of working with Big Data Discovery .....	9
About Studio .....	10
About Data Processing .....	10
About the Dgraph .....	11
<b>Chapter 2: Taking a Tour of Studio</b> .....	<b>13</b>
Workflow in Studio .....	13
Finding data sets .....	14
Profiling and Enriching Data .....	16
Using Explore .....	16
Using Transform .....	17
Using Discover .....	19
Using filtering, Guided Navigation, and Search .....	20
<b>Chapter 3: Data Sets in Big Data Discovery</b> .....	<b>22</b>
About data sets .....	22
Data sets in Catalog vs data sets in projects .....	23
Data Set Manager .....	23
Sampled and full data sets .....	25
Data set lifecycle in Studio .....	25
Projects and applications in Big Data Discovery .....	27
Data set access, access to projects, and user roles .....	29
<b>Chapter 4: Data Loading and Updates</b> .....	<b>30</b>
Data loading options .....	30
Data loading and sample size .....	31
Studio-loaded files: data update diagram .....	32
DP CLI-loaded files: data update diagram .....	34
Data update options .....	35

---

<b>Chapter 5: Where You Go from Here</b> .....	<b>37</b>
Quick reference to areas of interest .....	37

## Preface

Oracle Big Data Discovery is a set of end-to-end visual analytic capabilities that leverage the power of Apache Spark to turn raw data into business insight in minutes, without the need to learn specialist big data tools or rely only on highly skilled resources. The visual user interface empowers business analysts to find, explore, transform, blend and analyze big data, and then easily share results.

## About this guide

This guide introduces Oracle Big Data Discovery, and orients you how to use it for your needs, from loading data to exploring, transforming, and updating it.

Along the way, the guide introduces key terms, components and user interfaces in Big Data Discovery. This guide is complementary to the Getting Started video series.

## Audience

This guide is for business analysts, data scientists, and data engineers who work with big data.

## Conventions

The following conventions are used in this document.

### Typographic conventions

The following table describes the typographic conventions used in this document.

Typeface	Meaning
<b>User Interface Elements</b>	This formatting is used for graphical user interface elements such as pages, dialog boxes, buttons, and fields.
Code Sample	This formatting is used for sample code segments within a paragraph.
<i>Variable</i>	This formatting is used for variable values. For variables within a code sample, the formatting is <i>Variable</i> .
File Path	This formatting is used for file names and paths.

### Symbol conventions

The following table describes symbol conventions used in this document.

Symbol	Description	Example	Meaning
>	The right angle bracket, or greater-than sign, indicates menu item selections in a graphic user interface.	File > New > Project	From the File menu, choose New, then from the New submenu, choose Project.

## Path variable conventions

This table describes the path variable conventions used in this document.

Path variable	Meaning
\$ORACLE_HOME	Indicates the absolute path to your Oracle Middleware home directory, where BDD and WebLogic Server are installed.
\$BDD_HOME	Indicates the absolute path to your Oracle Big Data Discovery home directory, \$ORACLE_HOME/BDD- <i>&lt;version&gt;</i> .
\$DOMAIN_HOME	Indicates the absolute path to your WebLogic domain home directory. For example, if your domain is named <i>bdd-<i>&lt;version&gt;</i>_domain</i> , then \$DOMAIN_HOME is \$ORACLE_HOME/user_projects/domains/ <i>bdd-<i>&lt;version&gt;</i>_domain</i> .
\$DGRAPH_HOME	Indicates the absolute path to your Dgraph home directory, \$BDD_HOME/dgraph.

## Contacting Oracle Customer Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. This includes important information regarding Oracle software, implementation questions, product and solution help, as well as overall news and updates from Oracle.

You can contact Oracle Customer Support through Oracle's Support portal, My Oracle Support at <https://support.oracle.com>.



# Welcome to Big Data Discovery

This section introduces Big Data Discovery (BDD), and includes overviews of its components.

[Ways of using Big Data Discovery](#)

[Value of using Big Data Discovery](#)

[Addressing your goals and needs](#)

[Results of working with Big Data Discovery](#)

[About Studio](#)

[About Data Processing](#)

[About the Dgraph](#)

## Ways of using Big Data Discovery

In your organization, use BDD as the center of your data lab, as a unified environment for navigating and exploring all of your data sources in Hadoop, and to create projects and BDD applications.

- [Using BDD in the data lab on page 7](#)
- [Using BDD to navigate and explore data sets in Hadoop on page 8](#)
- [Using BDD to create BDD applications on page 8](#)

### Using BDD in the data lab

The data lab is a center of innovation that makes analytics creativity possible for everyone who works with data. It supports a large portfolio of data projects, and is integrated with the production environment for commercialization and feedback.

The data lab is a complete set of descriptive, diagnostic, predictive, and prescriptive analytics solutions. It is a place for collaboration between a select group of analysts who work together as a team and have easy access to multiple data sets. The data lab serves as an easy sandbox for ad-hoc data experiments.

When used in the data lab, Big Data Discovery lets you:

- Define projects in BDD and share them with others in your research data lab.
- Use BDD to quickly develop ideas, build prototypes and models, and invent ways of deriving value from data. For example, you can try out new approaches, quickly discard the ones that are not working, and move on to try new ways of working with your data.
- Shape data in your projects in multiple ways, from making easy single-row changes, such as trimming, editing, splitting, or null-filling, to advanced data shaping techniques, such as aggregations, joins, and custom transformations.

- Let others in your organization consume models created in the data lab, and publish insights to decision-making groups inside your organization.

## Using BDD to navigate and explore data sets in Hadoop

When used as the navigator on top of your data sources in Hadoop, BDD visually represents all data available to you in your environment.

You see all the data in Studio's Catalog and can find interesting data sets quickly. You can filter, edit metadata on data sets, and create new data sets for others in your team.

## Using BDD to create BDD applications

BDD lets you create **BDD applications** for well-established and known problems, to suit your business needs. For example, you can:

- Create BDD projects from scratch, improve them, and share them with wider groups of users in the organization.
- Serve as the author of all discovery solution elements: data sets, information models, discovery applications, and transformation scripts.
- Run discovery repeatedly and transparently to others in your group, and on different large-scale data sets arriving periodically from various sources.

## Value of using Big Data Discovery

In BDD, a wider number of people can work with big data, compared with traditional analytics tools. You spend less time on data loading and updates, and can focus on actual data analysis of big data.

In traditional analytics projects you spend most of your effort on data preparation. You must first predict all possible business questions, model the data to match, locate and get the data sources, and manipulate feeds to fit into the model. You build pipelines for extracting, transforming and loading data (ETL). Only after these tasks you can engage with the data. As a result, only a minimal effort is actually focused on data analysis.

The complexity of big data compounds the up-front costs of data loading. Big data increases the amount and kinds of data. You need more time to understand and manipulate new source data, especially unstructured sources, before it is ready for analysis.

Big data also changes fast. You must update your existing analytics projects with newer data on a regular basis. Such data loading and update tasks need qualified engineers with expert skills. As a result, you spend more time and money on up-front data loading and data updates.

Big Data Discovery addresses these problems:

- More users can engage with the data and contribute to big data analytics projects.
- Data preparation becomes a small part of your effort. You don't have to prepare data up-front, and can focus your time and effort on analyzing data, and improving your business.
- You can update BDD projects with new data, refresh already loaded data, or add newer data to existing projects.
- You can work with data as the data scientist, and rely on a range of ways for data shaping and visual analysis in BDD.



## Addressing your goals and needs

Big Data Discovery makes your job faster and easier than traditional analytics tools. This topic discusses your goals and needs in data analysis and shows how Big Data Discovery can address them.

As a data scientist or analyst, you:

- **Solve complex questions using a set of disjointed tools.** You start with many dynamic imprecise questions. You often have unpredictable needs for data, visualization, and discovery capabilities. To address them, you rely on open source and custom discovery tools. You use tools in combination with other tools, and often need to reopen the same tools several times.

In Big Data Discovery, this fragmented workflow between tools is replaced with a single workflow that is natively part of your ecosystem.

- **Need to collaborate.** Together with your team, you work with big data from many external and internal sources. Other team members may consume the results of your findings. You also improve and publish insights and prototypes of new offerings.

In Big Data Discovery, you can create personal projects, or create and share a project with your team.

- **Want to make sense of data.** You often need to create and use insights. You do this by collecting, cleaning and analyzing your data.

Big Data Discovery lets you make sense of data. You use it to collect, join, shape, and analyze your data.

- **Generate ideas and insights.** You want to create insights that lead to changes in business, or you want to enhance existing products, services, and operations. You also need to define and create prototypes (or models) for new data-driven products and services.

In Big Data Discovery, you arrive at insights by using many data visualization techniques. They include charts, statistical plots, maps, pivot tables, summarization bars, tag clouds, timelines, and others. You can save and share results of your discovery using snapshots and bookmarks.

- **Validate, trace back, tweak, and share your hypotheses.** You often need to provide new perspectives to address old problems. The path to the solution is usually exploratory and involves:
  - Hypotheses revision. You must explore several hypotheses in parallel. They are often based on many data sets and interim data products.
  - Hypotheses validation. You need to frame and test hypotheses. This requires validation and learning of experimental methods and results done by others.
  - Data recollection and transparency. You would like all stages of the analytical effort to be transparent so that your team can repeat them. You want to be able to recreate analytical workflows created before. You also would like to share your work. This requires linear history of all steps and activities.

Big Data Discovery helps you by saving your BDD projects, data sets, and transformation scripts. This lets you improve your projects, share them, and apply saved transformation scripts to other data sets.

## Results of working with Big Data Discovery

In BDD you can create an insight, an exploratory project, or a tool for long-term analysis of your data.

For example, you can create:

- Another data set that illustrates the business hypothesis, or answers a specific business question.

- A product, such as a set of transformations useful for similar data sets.
- One or more ad-hoc exploratory projects.
- An insight, such as a conclusion that you illustrate with several saved visualizations and share with your team.
- A long-term BDD application for further analysis of future data sets.
- An infrastructure (a set of tools that includes Big Data Discovery as a component).

As an outcome, you:

- reach insights for action
- gain understanding for building more sophisticated mathematical models
- create data products, such as new features and services for your organization.

## About Studio

**Studio** is a component in Big Data Discovery. Studio provides a business-user friendly user interface for a range of operations on data.

It combines search, guided navigation, and many data visualizations in a single environment.

Some aspects of what you see in Studio always appear. For example, Studio always includes search, **Explore**, **Transform**, and **Discover** areas. You can add other parts of your interface as needed. These include many types of data visualization components. For example, you can add charts, maps, pivot tables, summarization bars, timelines, and other components. You can also create custom visualization components.

Studio provides tools for loading, exploration, updating, and transforming data sets. It lets you create projects with one or more data sets that are linked. You can load more data into existing projects. This increases the corpus of analyzed data from a sample to a fully loaded set. You can also update data sets. You can transform data with simple and more complex transforms, and write transformation scripts.

Studio's administrators can control data set access and access to projects. They can set up user roles, and configure other Studio settings.

For more information, see [Taking a Tour of Studio on page 12](#).

## About Data Processing

The **Data Processing** component of BDD runs a set of processes and jobs. This set is called **data processing workflows**. Many of these processes run natively in Hadoop.

The Data Processing component controls several workflows in BDD. For example, you can create workflows for data loading, data updates, and others. The Data Processing component discovers new Hive tables and loads data into BDD. It runs data refresh operations and incremental updates. It also keeps BDD data sets in synch with Hive tables that BDD creates.

For example, during a data loading workflow, the Data Processing component performs these tasks:

- Discovery of data in Hive tables
- Creation of data sets in BDD
- Running a select set of enrichments on discovered data sets

- Profiling of the data sets
- Indexing of the data sets, by streaming data to the Dgraph.

To launch a data processing workflow when Big Data Discovery starts, you use the **Data Processing Command Line Interface (DP CLI)**.

## Data Processing CLI

The DP CLI is a shell Linux utility that launches data processing workflows in Hadoop. You can control their steps and behavior. You can run the DP CLI manually or from a `cron` job. You can run the data processing workflows on an individual Hive table, all tables within a Hive database, or all tables within Hive. This depends on DP CLI settings, such as a blacklist and a whitelist.

Here are some of the jobs you can run with DP CLI:

- Load data from Hive tables after installing Big Data Discovery (BDD). When you first install BDD, your existing Hive tables are not processed. You must use the DP CLI to launch a data processing operation on your tables.
- Run data updates. They include:
  - An operation to refresh data. It reloads an existing data set in a Studio project, replacing the contents of the data set with the latest data from Hive in its entirety.
  - An incremental update. It adds newer data to existing data sets in a Studio's project.
- Launch the BDD Hive Table Detector (this is a utility in Data Processing). The BDD Hive Table Detector detects if a new table is added to Hive. It then checks the whitelist and blacklist. If the table passes them, it creates a data set in BDD. It also deletes any BDD data set that does not have a corresponding source Hive table. This keeps BDD data sets in synch with data sets in Hive. For detailed information on how data sets are managed in BDD, see [Data set lifecycle in Studio on page 25](#).

For information on Data Processing and the DP CLI, see the *Data Processing Guide*.

## About the Dgraph

The Dgraph uses data structures and algorithms to issue real-time responses to queries.

**The Dgraph** stores the ingested data in the Dgraph databases it creates for the indexed data sets. The Dgraph receives client requests from Studio, queries its databases, and returns the results. The Dgraph is stateless. This design requires that a complete query is sent to it for each request. The stateless design facilitates the addition of Dgraph instances for load balancing and redundancy. Any instance of a Dgraph can reply to queries independently of other instances.

A Big Data Discovery cluster can include two or more Dgraph instances. They handle end-user query requests and store Dgraph databases on shared storage. For each data set, the leader Dgraph (once it is automatically appointed in BDD) handles all write operations, such as updates and configuration changes to that data set. The remaining Dgraphs serve as read-only followers for that data set. The Dgraph Gateway performs the routing of requests to the Dgraph nodes. It handles query caching, business logic, and cluster services for the Dgraph nodes.

The deployment of Dgraphs occurs at the installation time. For information, see the *Installation Guide*.

## The Dgraph Tracing Utility

The Dgraph Tracing Utility is a Dgraph diagnostic program used by Oracle Support. It stores the Dgraph trace data, which are useful in troubleshooting the Dgraph. It starts when the Dgraph starts and keeps track of all Dgraph operations. It stops when the Dgraph shuts down. You can save and download trace data to share it with Oracle Support.

## The `bdd-admin` script

Use the `bdd-admin` script to run administrative operations on the BDD cluster deployment, such as starting components, backing up, restoring, and updating the configuration. For information on the `bdd-admin` script, see the *Administrator's Guide*.

## The Dgraph HDFS Agent

The Dgraph HDFS Agent loads data into the Dgraph. It reads the Avro files from the data processing workflow and formats them for data loading. It then sends them to the Dgraph. For information on the Dgraph HDFS Agent, see the *Data Processing Guide*.



## Chapter 2

# Taking a Tour of Studio

This section walks you through Studio. It describes how you move through areas within Studio when working with big data.

[Workflow in Studio](#)

[Finding data sets](#)

[Profiling and Enriching Data](#)

[Using Explore](#)

[Using Transform](#)

[Using Discover](#)


[Using filtering, Guided Navigation, and Search](#)

## Workflow in Studio

Studio is the visual face of Big Data Discovery. It enables exploratory data analysis and is structured around **Explore**, **Transform** and **Discover** areas.

- After you log in, you see the Catalog. It lets you browse and discover the contents of the data lake. (A **data lake** is a storage repository that holds raw data in its native format until it is needed.)

In Catalog you can find projects you created before, or projects shared by other users. You can also browse data sets discovered in a Hive data base, or data sets that other users uploaded from a file or imported from an external JDBC data source:



- Next, if there are no **Discover** pages created yet by other users, the **Explore** page for the first data set opens. If **Discover** pages already have been created by other users, the first **Discover** page opens.
- If you move to **Explore**, you can pick data sets that are of interest to you for further exploration, and use the results of data profiling and enrichments to grasp basic data characteristics. For example, you can look for outliers. Outliers are values that are distant from other values for the attribute. You can also explore scatter plots to see correlations between values for two attributes, or link data sets.

**Explore** ▾ **Transform** ▾ **Discover** ▾

- If some pages by other users have already been created and you have access to them, then instead of **Explore**, you move to **Discover** for the project that is shared with you.

- If you started with **Explore**, then to move to **Transform**, or **Discover**, you must first add the data set to a project. Alternatively, you can find and select an existing project:

**Add to project**

Note that once your data set is in a project, you can continue using **Explore**.

- You can then move to **Transform**, where you transform the data and remove inconsistencies. For example, you can change data types, or create custom transformation scripts.

**Explore** ▾ **Transform** ▾ **Discover** ▾


- In **Discover** you arrive at insights about the data and save your projects for sharing with others.

**Explore** ▾ **Transform** ▾ **Discover** ▾

## Finding data sets

The number of data sets in any data repository can be huge. In Studio's Catalog, you can use search and Guided Navigation to find data sets and existing projects relevant to your analytical purpose.

After you log in, you see the **Catalog**:



Here, you can:

- Find data sets
- Find existing projects
- Add a data set to BDD

You can quickly look at a data set by clicking the tile to expand the Data Set Details panel. This displays the data set's enriched metadata to provide a summary to help you decide if you'd like to explore this data set further:

The screenshot shows a data set tile for 'Wine' with the following details:

- Wine**  
The wine data set we all know and love.
- Data Source:** default.wine  
(57.1k records) ⌚
- Preview** button and **New** button.

Below the tile, the expanded details panel for 'Wine (57,076 records)' is shown:

- Wine (57,076 records)**
- Buttons: **Data Set Info.** (selected), **Used in Projects (4)**
- Description: *The wine data set we all know and love.*
- Tags**  
*Click to add tags...*
- Attributes (15)**  
⌚
- Metadata table:
 

<b>Data source:</b>	default.wine
<b>Origin system type:</b>	Hive
<b>Origin system:</b>	default.wine
<b>Created on:</b>	1/31/2016 2:35:59 AM (UTC)
<b>Access:</b>	Public (Default Access)

Through the Catalog, you can:

- Navigate data sets using a familiar keyword search and Guided Navigation experience similar to that found on major commerce sites.
- Browse data sets based on name, metadata, attribute names, topics, dates, locations, or key terms.
- Find related data sets that are potentially relevant to your interests based on overlapping data or frequent use by other users.
- Filter displayed data sets based on data set metadata (such as the creation date, tags, notes, number of records, or contents of the data) and also by attribute metadata (such as semantic types).

## Adding tags and notes to data sets

When you expand a data set to view the data set details, you can optionally add tags or notes that describe the data. All users can then filter data sets based on those tags, or run keyword searches that include text from the data set notes.

To add notes, click the description field below the data set name.

To add searchable tags, click the text field below **Tags**. To help you choose consistent tag names, autocomplete results may display as you type.

## Profiling and Enriching Data

This topic summarizes how Big Data Discovery helps to profile and enrich new data.

You can add new data as files in HDFS using data integration technologies. You can also use Studio to upload files, such as Excel and CSV, or pull data from a database using your credentials, and import it into a personal sandbox in HDFS. In either case, when it loads data, Big Data Discovery performs these activities for you:

- Profiles the data, inferring data types, classifying content, and understanding value distributions.
- Lists most interesting data sets first and indicates what they contain.
- Decorates the data with metadata, by adding profile information.
- Enriches the data, by extracting terms, locations, sentiment, and topics, and storing them as new data in HDFS.
- Takes a random sample of the data of the specified size. (You can increase the sample size, or load full data.)
- Indexes the data and prepares it for fast search and analysis in BDD.

As an outcome, the data you need resides in HDFS, is indexed by the Dgraph and enriched by BDD, and is ready for your inspection and analysis.

## Using Explore

**Explore** provides you with an out-of-the-box guided analytics experience. It configures the canvas with the right visualizations to explain the data based on your goals. This allows you to get a summary of the data and explore data quality.

**Explore** ▾ **Transform** ▾ **Discover** ▾

**Explore** provides an attribute-focused visual summary of the data, summarizing value distributions, data quality gaps, and relationships.

**Explore** presents visualizations that give you the most insight into the data set, selecting the types most suitable to each attribute's data type and value distribution.

Visualizations are automatically composed, to save you time and effort at this early stage in the process. When you have a better understanding of the data set, you can compose your own visualizations on data that you have identified as worthy of further analysis.

You can think of the **Explore** area as a guided tour of new data sets, freeing you from the need to manually query the data or configure your own visualizations. It helps you immediately extract meaning from the results to better understand what's inside the data set.



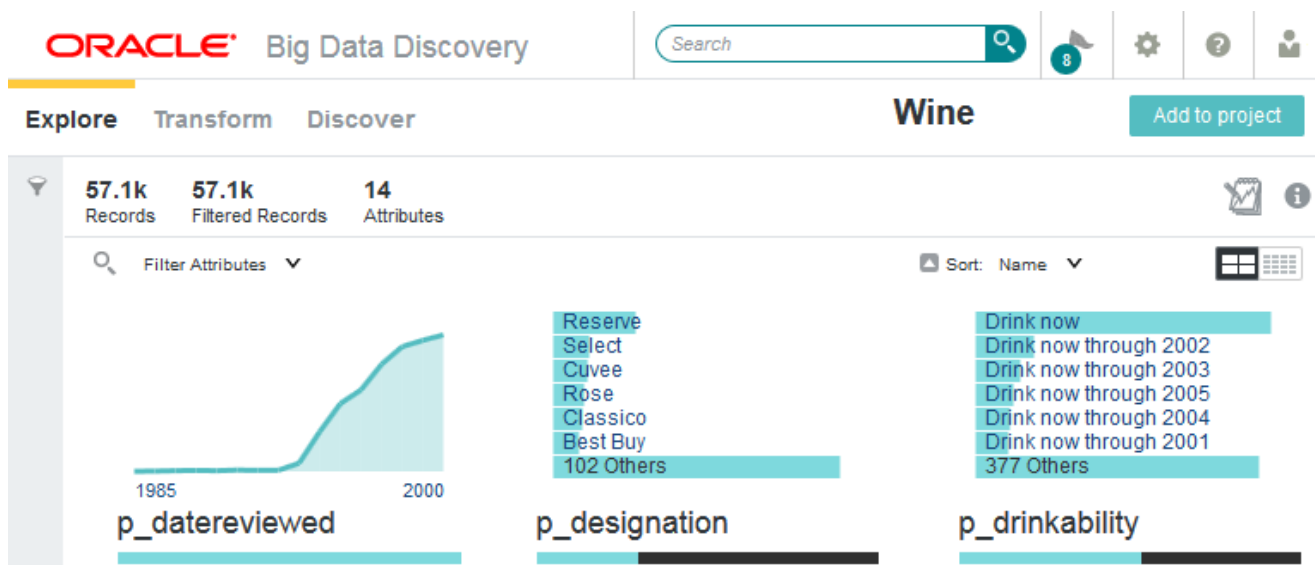
Here are some of the questions that Big Data Discovery provides answers to, when you use **Explore**:

- What fields are in my data? Do I understand what they are? **Explore** shows these fields as a series of visualizations.
- How are values distributed? **Explore** uses descriptive statistics (mean, median, mode, quintiles) and visuals (histograms, box plots).
- How dirty/messy is the data? Are there missing values or outliers? **Explore** uses the bar underneath each attribute, showing you at a glance whether any inconsistencies exist.
- What relationships exist between fields? **Explore** uses numeric correlations and visuals.

In addition, **Explore** lets you add your own notes and tags to data sets. You can later sort or search on those tags, or review your notes for a quick summary without having to walk through the data a second time.

As an outcome, you understand the data sets you've been exploring.

Here are some examples of **Explore** visualizations:



This image shows **Explore** for a data set with 57.1K records, with attributes sorted by name.



**Note:** Even if you could explore a large data set at full scale, you always want to start by exploring a representative sample, and later confirm your hypotheses or expand your analysis at full data scale.

## Using Transform

Through exploration, data quality issues often come to the surface. You must triage them and, in many cases, correct, clean, or change data in some way. The interactive **Transform** area in Studio helps you turn your data into data that is fit for advanced analysis.

**Transform** helps you isolate data quality problems and lets you quickly apply a number of data transformations to clean or restructure the data.

You can find **Transform** in Studio here:

**Explore** ▾ **Transform** ▾ **Discover** ▾

**Transform** presents a spreadsheet-style view of data, organized either by attribute's value list view, or row view. Here are examples of how you can transform data:

- You can focus on just a few basic transformations needed to support your visualization or analysis. You do this by running common data transformations, such as filling missing values, converting data types, collapsing values, reshaping, pivoting, grouping, merging, and extracting.
- You can join one or more data sets together within a project and also aggregate attribute values to create new derived attributes that sum, average, find minimum/maximum values, and run other operations.
- You can use live **Preview** to see the set of transformations in action before applying them to the entire data set in Studio.

Because data transformation and data exploration are two integral aspects of any discovery project, you can seamlessly transition between **Transform** and **Explore** within a unified interface, which promotes a smooth workflow.

- You can also share your transformation scripts with other BDD users in your team.

Here is an example of what **Transform** looks like in a project:

**Explore** ▾ **Transform** ▾ **Discover** ▾ **Test Project: Sales**

60.9K Records   60.9k Filtered Records   79 Attributes

Basic   Convert   Advanced   Shaping   Editor

|x| Absolute value

Filter Attributes ▾   Sort: Name ▾

sales_activitydates	# sales_calendarquarter	# sales_calendarseme...
2007-04-08 00:36:51 UTC, 20...	2	
2007-10-01 15:30:42 UTC, 20...	4	
2007-04-08 20:00:42 UTC, 20...	2	
2007-10-01 14:00:15 UTC, 20...	4	2

TRANSFORM SCRIPT

- \* sales\_activitydates - Extract ...
- \* sales\_monthname - Uppercase

Commit to Project

Context menu for # sales\_calendarquarter:

- Add to Favorites
- Hide
- Edit Details
- Sort

In this image, you can see the **Transform** area of Studio. Notice the drop-down attribute menu in the second column. It includes options to hide an attribute, mark it as favorite, edit it, and sort it. To the right, you can also see the **Transform Script** panel, with three transforms created by the project's user.

Also notice the transform menu just under the header. The Transform menu is made up of the tabs: **Basic**, **Convert**, **Advanced**, **Shaping**, and **Editor**.

The transforms that display under each tab vary depending on the data type of the attribute you select. For example, if you select a date time attribute, the **Basic** tab displays the **Truncate date** and **Extract date part** transforms. These are both transforms that are contextually appropriate for date time attributes. And if you select a numeric attribute, the **Basic** tab displays the **Absolute value** transform because the **Absolute value** transform is contextually appropriate for numeric attributes.

After you run the transform script by clicking **Commit to Project**, the data is transformed and ready for further analysis.

## Using Discover

Using **Discover** in Studio, you analyze the data and discover insights hidden in data. **Discover** lets you identify and pick meaningful visualizations and capture the results of the data exploration phase.

You can find the **Discover** area of Studio here:

**Explore** ▾ **Transform** ▾ **Discover** ▾

Here is an example of how **Discover** might look like for your project:

The screenshot displays the Oracle Big Data Discovery interface. At the top, the header reads "ORACLE Big Data Discovery" with a search bar and user icons. Below the header, navigation tabs include "Explore", "Transform", and "Discover". The main workspace is titled "alanTest: Page 1" and features a "MAP" component. The map shows a world view with several data points marked by numbered callouts (1-19) across continents like North America, Europe, and Asia. To the right of the map is a list of 7 records, each with a "Record detail" link. On the far right, an "ADD COMPONENT" panel lists various visualization options such as Box Plot, Chart, Column Container, Histogram Plot, iFrame, Map, Parallel Coordinates Plot, Pivot Table, Results List, Results Table, and Summarization Bar.

In this diagram, you can see the **Discover** area, with the Map visualization added. To the right, notice some of the other visualization components you can add to your project.

Some of the examples of the insights you can create with **Discover** are:

- Use search and Guided Navigation to identify subsets of data for deeper analysis.
- Use interactive visualizations to let the data speak for itself, in raw or aggregated forms, in text or in numbers. Statistical techniques (regression, correlation) and machine learning tools (classification, clustering) are used in Big Data Discovery to drive these visualizations.
- Join (or combine) multiple data sets for new insights.
- Create projects that serve as models for different outcomes. These models automatically understand the weight of different factors in contributing to given outcomes without requiring experts to hand-code rules or heuristics.
- Add new data to existing projects or run your projects on existing data that has been altered to examine new scenarios. (This is known as "what if" analysis.)

- Export your data sets from Big Data Discovery back into Hadoop, so that you can run complex models built in other tools (such as R, or Python) on your custom-created data sub-sets, such as on different customer segments.

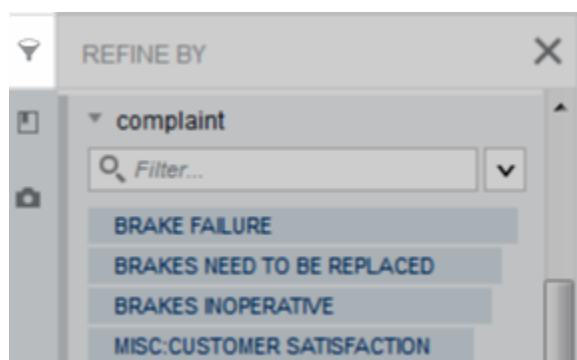
## Using filtering, Guided Navigation, and Search

The following examples explain how search filtering and Guided Navigation work in Big Data Discovery.

The Catalog, **Explore**, **Transform**, and **Discover** areas in Studio make use of powerful search and data-driven Guided Navigation capabilities of Big Data Discovery.


In the Catalog, refinements are always available on the left sidebar.

For data sets in projects, you can expand or collapse the sidebar using the funnel icon. If you expand it, it looks like this:



Using refinements, you can filter on attributes or on available metadata such as creation date. In the Catalog, this also includes any tags that users add to projects or data sets.

Selected refinements represent applied filters. They show up as breadcrumbs at the top of your screen. Here is an example of a selected refinement from the attribute `complaint = BRAKE FAILURE`. Notice that applying the refinement reduces the number of matching records to only 1.2k of the 37.8k total records in the data set.



Attribute filters apply across all parts of BDD. If you highlight or exclude attributes, your project reflects these selections when you move from the Catalog to **Explore** to **Transform** and back. For example, if you select all numeric attributes in a data set in **Explore**, and then switch to **Transform**, **Transform** is filtered by all numeric attributes.

You can also search data sets at any stage in your workflow using the search box:



In addition to searching across project and data set content, keyword search results from the Catalog also include any notes that users add to describe projects or data sets.

Like refinement filters, keyword search filters apply across both **Explore** and **Transform** if you switch between pages.



## Chapter 3

# Data Sets in Big Data Discovery

---

This section describes what is a data set in BDD, discusses sampled and full data sets, and access to data sets. It also discusses a data set's lifecycle in Studio, and BDD projects and applications.

[About data sets](#)

[Data sets in Catalog vs data sets in projects](#)

[Data Set Manager](#)

[Sampled and full data sets](#)

[Data set lifecycle in Studio](#)

[Projects and applications in Big Data Discovery](#)

[Data set access, access to projects, and user roles](#)

## About data sets

Data sets in BDD are called **BDD data sets**. This helps you to distinguish them from **source data sets** in Hive tables.

A BDD data set is the central concept in the product architecture. Data sets in BDD originate from these sources:

- Many data sets are loaded as a result of the data processing workflow for loading data. It runs when you launch DP CLI, after installing Big Data Discovery. This process adds data sets to Studio's Catalog.
- Other data sets appear in Studio because you load them from a personal file or a JDBC data source.
- Also, you can create a new BDD data set by transforming an existing data set, or by exporting a data set to HDFS as a new Hive table. Such data sets are called **derived data sets**.

When the Data Processing component runs data loading, or when you add a data set by uploading a file or data from a JDBC source, it appears in Studio's Catalog. You can use **Preview** to see the details for all data sets in Catalog.

You can also use Studio's **Data Set Manager** to see information about all data sets in a project. For each data set in a project, you can see its record count, when it was added, whether it is private to you, and other details.

## Data sets in Catalog vs data sets in projects

It is important to distinguish between data sets in the Catalog, and data sets you add to a project.

Here is why:

- The Catalog only contains data sets that are discovered in a data lake, or that you add to BDD from a personal file or import from a JDBC source. A data set in a project represents your own modified version of the data set. Once you move a data set to a project, it is similar to creating a personal branch version of the data set.
- You can edit data set metadata and some attribute metadata from the Catalog or the **Explore** view, but data set altering operations, such as transformations, enrichments, and other attribute metadata changes, require editing the data set in a project.
- You can perform some updates of data sets from the Catalog, such as reloading a newer version of a data set. In order to run scheduled updates using the DP CLI, a data set must be part of a project. This is beneficial if you want to run updates periodically, and automate them, by adding them to your scripts that use DP CLI.

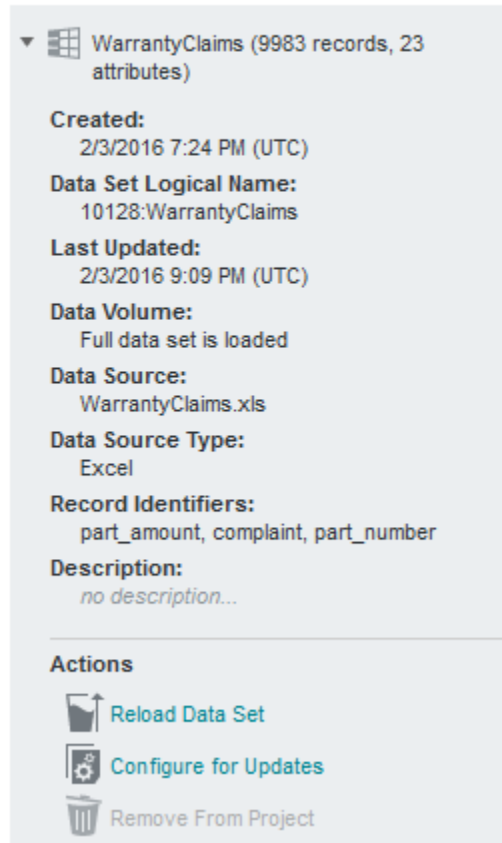
## Data Set Manager


In Studio, the **Data Set Manager** includes information about all project data sets. You access Data Set Manager once you are in a project, from **Project Settings**.

Here is an example of data set's details in the Data Set Manager:

## Data Set Manager

### Data Sets in This Project



▼  WarrantyClaims (9983 records, 23 attributes)

**Created:**  
2/3/2016 7:24 PM (UTC)

**Data Set Logical Name:**  
10128:WarrantyClaims

**Last Updated:**  
2/3/2016 9:09 PM (UTC)

**Data Volume:**  
Full data set is loaded

**Data Source:**  
WarrantyClaims.xls




**Data Source Type:**  
Excel

**Record Identifiers:**  
part\_amount, complaint, part\_number

**Description:**  
*no description...*

---

**Actions**

-  Reload Data Set
-  Configure for Updates
-  Remove From Project

You can see that the project includes one data set. You can also see the number of records and attributes in it.

The **Data Source Type** shows the type of the source file, such as Excel, or CSV, or it shows "Hive", if the data originated from a Hive table. In this example, the Data Source Type is Excel. This means the data set originated from the spreadsheet that was loaded in Studio and then was reloaded. You can tell this by observing the dates for creation and update.

The **Data Source** field shows the name of the source file. In this example, it is `WarrantyClaims.xls`.

Notice the **Data Set Logical Name**. This is the name that each data set has, whether it exists in Catalog or belongs to a project. When a data set is in Catalog, it has one data set logical name. When you move it into a project, the data set's logical name changes. To run scripted updates with DP CLI, you need to note the correct data set logical name, so that you know which data set you are going to update. For information on scripted updates, see the *Data Processing Guide*.



## Sampled and full data sets

Data sets in BDD can be sampled, or they can represent a full data set.

### Sampled data sets

A *sample data set* in BDD represents a random sample of a source data set in Hive. If the data set originates in Hive, you use the Data Processing CLI to load it. The DP CLI uses the default sample size of 1 million records. You can specify a different sample size at data loading time.

- If you specify a sample size that is less than the size of the source data set in Hive, a sample data set is loaded into BDD.
- If you specify a sample size that is greater than or equal to the size of the source data set, a full data set is loaded.

### Full data sets

A *full data set* in BDD represents a data set that contains all records, if you compare it to the source it was loaded from. For example, if a data set originates in Hive, and the sample size in DP CLI is greater than the record count in the source Hive table, this data set is loaded in full.

For a summary of how to get from a sample to a full data set, see [Data loading and sample size on page 31](#)

For more information on sampling and data set loading during data processing, see the *Data Processing Guide*.

For information on adding and managing data sets in Studio, including loading a full data set, see the *Studio User's Guide*.


## Data set lifecycle in Studio

As a data set flows through Big Data Discovery, it is useful to know what happens to it along the way.

Before we describe the data set lifecycle, here's how BDD interacts with source data sets it finds in Hive:

- BDD does not update or delete source Hive tables. When BDD runs, it only creates new Hive tables to represent BDD data sets. This way, your source Hive tables remain intact if you want to work with them outside of Big Data Discovery.
- Most of the actions in the BDD data set lifecycle take place because you select them. You control which actions you want to run. Indexing in BDD is a step that runs automatically.

This diagram shows stages in the data set's lifecycle as it flows through BDD:



In this diagram, the data set goes through these stages:

1. **Create** the data set. You create a data set in BDD in one of two ways:

- Uploading source data using Studio. You can upload source data in delimited files and upload from a JDBC data source. When you upload source data, BDD creates a corresponding Hive source table based on the source data file.
- Running the Data Processing CLI to discover Hive tables and create data sets in Studio based on source Hive tables. Each source Hive table has a corresponding data set in Studio.

The Catalog of data sets in Studio displays two types of data sets. Some data sets originate from personally-loaded files or a JDBC source. Other data sets are loaded by Data Processing from source Hive tables.

2. Optionally, you can choose to **enrich** the data set. The data enrichment step in a data processing workflow samples the data set and runs the data enrichment modules against it. For example, it can run the following enrichment modules: Language Detection, Term Frequency/Inverse Document Frequency (TF/IDF), Geocoding Address, Geocoding IP, and Reverse Geotagger. The results of the data enrichment process are stored in the data set in Studio and not in the Hive tables.



**Note:** The Data Processing (DP) component of BDD optionally performs this step as part of creating the data set.

3. **Create an index** of the data set. The Dgraph process creates binary files, called the Dgraph database, that represent the data set (and other configuration). The Dgraph accesses its database for each data set, to respond to Studio queries. You can now explore the data set.

4. **Manage** access to a data set. If you uploaded a data set, you have private access to it. You can change it to give access to other users. Data sets that originated from Hive are public. Studio's administrators can change these settings.

5. **Transform** the data set. To do this, you use various transformation options in **Transform**. In addition, in **Transform**, you can use two operations that affect the lifecycle of a data set: You can create a new data set (this creates a new Hive table), and you can commit a transform script to modify an existing data set.

If you commit a transform script's changes, Studio writes the changes to the Dgraph and stores the changes in the Dgraph's database for the data set. Studio does not create a new Hive table for the data set. You are modifying the data set in the Dgraph, but not the source Hive table itself.

6. **Update** the data set. To update the data set, you have several options. For example, if you loaded the data set from a personal data file or imported from a JDBC source, you can reload a newer version of this data set in Catalog. If the data set was loaded from Hive, you can use DP CLI to refresh data in the data set.

Also, you can **Load Full Data Set**. This option is useful for data sets that represent a sample. If a data set is in a project, you can also configure the data set for incremental updates with DP CLI.

7. **Export** the data set. When your data set is in a project, you can export it. For example, you can export a data set to HDFS, after you have applied transformations to it. This way, you can continue working with this data set using other tools. Also, you can export a data set and create a new data set in Catalog. Even though on this diagram **Export** is shown as the last step, you can export the data set at any stage in its lifecycle, after you add the data set to a project.

8. **Share** the data set. At any stage in the data set's lifecycle, you can share the data set with others.

9. **Remove** the data set. When you delete a data set from Studio, the data set is removed from the Catalog and is no longer accessible in Studio. However, deleting the data set does not remove the corresponding source Hive table that BDD created, when it loaded this data set.

It's important to note that BDD does not update or delete original source Hive tables. BDD only creates new Hive tables to represent BDD data sets. You may need to ask your Hive data base administrator to remove old tables as necessary to keep the Hive data base clean. If the Hive data base administrator deletes a Hive table from the data base, the Hive Table Detector detects that the table was deleted and removes the corresponding data set from the Catalog in Studio. The Hive Table Detector is a utility in the Data Processing component of BDD.

## Projects and applications in Big Data Discovery

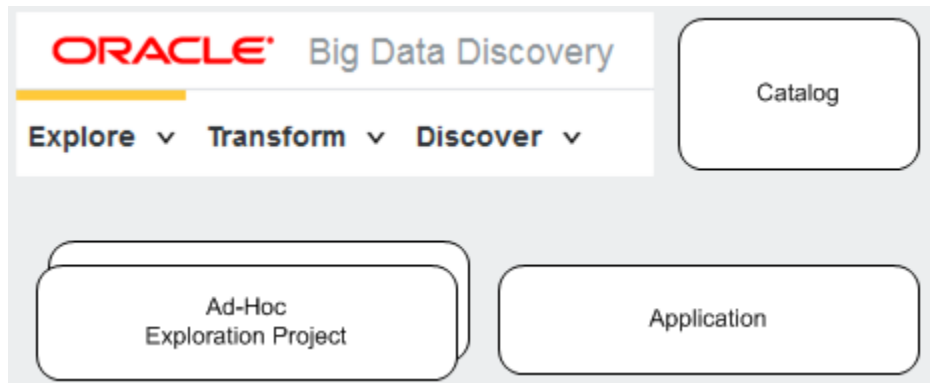
When you work with data sets in BDD, you put them in projects in Studio. These are known as BDD projects. Some BDD projects can become BDD applications.

BDD projects let you perform ad-hoc exploration and discovery using standard and advanced analytic techniques.

BDD applications expose an interactive analytic dashboard to a broad set of users.

While each BDD application is also a BDD project in Studio, it is referred to as a BDD application because it has a number of characteristics that make it special. In other words, you always start with projects in Studio; you can turn some projects into BDD applications.

The following diagram illustrates that BDD can contain one or more projects and an application:



In the diagram, the Catalog is also shown. It contains data sets that you select to add to your projects. Some of the projects you can later turn into BDD applications.

### BDD projects

**BDD projects** are created by each user and serve as personal sandboxes. Each BDD deployment supports many BDD projects at once. This lets everyone in the BDD analyst community explore their own data, try different sample data sets, and identify interesting data sets for future in-depth analysis.

BDD projects often, but not always, run on sample data and allow you to load newer versions of sample data into them. Each BDD deployment can support dozens of ad-hoc, exploratory BDD projects. You can turn the most interesting or popular BDD projects into BDD applications.

Here are characteristics of a BDD project. It:

- is often a short-lived project, used to try out an idea
- is owned by a single user who can maintain, edit, and delete the project once no longer needed
- typically, but not always, runs on a sample of data
- answers a simple analytics question, or helps with initial exploration of the data set in Catalog
- lets you identify and select data set candidates for in-depth analysis.

As opposed to BDD projects that any user in BDD can create, BDD administrators own and certify BDD analytic applications, which they can share with their teams.

## BDD applications

A **BDD application** includes a list of data sets that have been tweaked, linked and transformed, based on the goals of business analysis. Data analysts with power user permissions create and configure it. They can share it with other users in the BDD business analyst community who can use it for analysis.

Such an application answers a set of predefined questions and lets a group of users analyze the findings in it. A BDD application is often built on at least one full data set, with other data sets linking to it.

It is configured for periodic scripted data updates. You can either load newer versions of data onto it, or keep the existing data while periodically adding new data.

Each BDD deployment can support several BDD applications. The number of applications a BDD deployment can support depends on the capacity and sizing of your BDD deployment.

You can think of a BDD application as a "productized" or "certified" project that gained greater significance and is now maintained by power users (or by your IT organization) for use by a group of analysts within your team.

Here are the characteristics of a BDD application. It:

- is a longer-term project as opposed to other BDD projects
- often includes joins from multiple data sets
- often contains full data from at least one of the comprising data sets
- includes a predefined set of visualizations aimed at answering a set of specific questions
- allows you to run periodic data updates implemented with DP CLI
- has user access for a BDD application that is split between owners, curators and the team. Owners create and configure it, curators maintain it; the team can view and use visualizations. The team members may have different access to data sets in the application.
- has results and findings that are intended for sharing and viewing by a team (as opposed to exploratory BDD projects used by individual users who create them).

## Data set access, access to projects, and user roles

Studio Administrators and project authors manage access to data sets and projects in Studio.

### User roles

In Studio, two sets of user roles exist: a set of global permissions, and project-specific permission sets. For information on Studio's project access, data set access and user roles, see the *Administrator's Guide*.

### Data set access

If you have Read access to a data set, you can see it in search results, or by browsing the Catalog. You can also explore it or add it to a project and then modify the project-specific version of the data set.

Write access allows you to set data set metadata or change data set permissions.

For information on changing access to a data set, see the *Studio User's Guide*.

For information on changing the default Studio settings for data set access, see the *Administrator's Guide*.

### Project access

Projects are private by default. Only the project author and the Studio Administrators group can access them. If you are an Administrator or the project author, you can modify access to add user groups or individual users to your project. For information on changing access to a project, see the *Studio User's Guide*.



## Chapter 4

# Data Loading and Updates

This section discusses options for initial data loading and data updates. It illustrates how you can load files in Studio, or using Data Processing CLI.

*Data loading options*

*Data loading and sample size*

*Studio-loaded files: data update diagram*

*DP CLI-loaded files: data update diagram*

*Data update options*

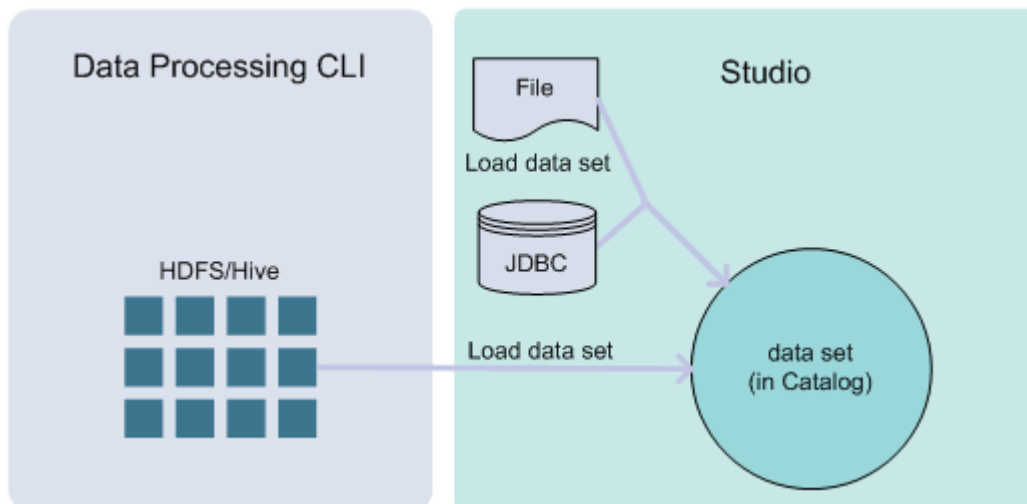
## Data loading options

BDD offers several options for data loading. You can load data by running the data loading workflow with DP CLI. Also, in Studio, you can upload a personal file or import data from a JDBC source.

For initial loading of data into BDD, three methods exist:

1. **Load of Hive tables.** When you run DP CLI, it runs the data processing workflow and loads data from Hive tables into BDD.
2. **Personal file upload.** In Studio you can upload a data set from a personal file.
3. **Import from a JDBC source.** In Studio you can import a data set from a JDBC source.

This diagram illustrates data loading options:



Also, you can load a sample or full data. Or, you can also start with a sample and then load full data. See [Data loading and sample size on page 31](#).

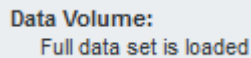
## Data loading and sample size

You can load either a sample or a full data set. If you load a sample, you can go to a full data set later. This topic summarizes how to get from a sample to a full data set.

These options are high-level summaries only. For detailed steps, see the referenced documentation for each option.

- **Controlling sample size when loading from Data Processing CLI.** DP CLI has a parameter for data size sample. The default sample size is 1 million records. When you use DP CLI for data loading, you can customize this parameter:

- If it is less than the record count of source records in Hive, a full data set is loaded. In this case you already loaded a full data set. This is indicated in the Data Set Manager in Studio, in the **Data Volume** field:



**Data Volume:**  
Full data set is loaded

- If it is greater than the number of records in Hive, a sampled data set is loaded, based on the sample size you specify. In this case, you can use DP CLI with `--Incremental` update flag, or you can use **Load Full Data Set** in Studio, to load the entire source data set from Hive. You will then have a full data set in BDD.

For detailed information on specifying the sample size with DP CLI, see the *Data Processing Guide*.

- **Controlling data set size when loading from a file or a JDBC source.**

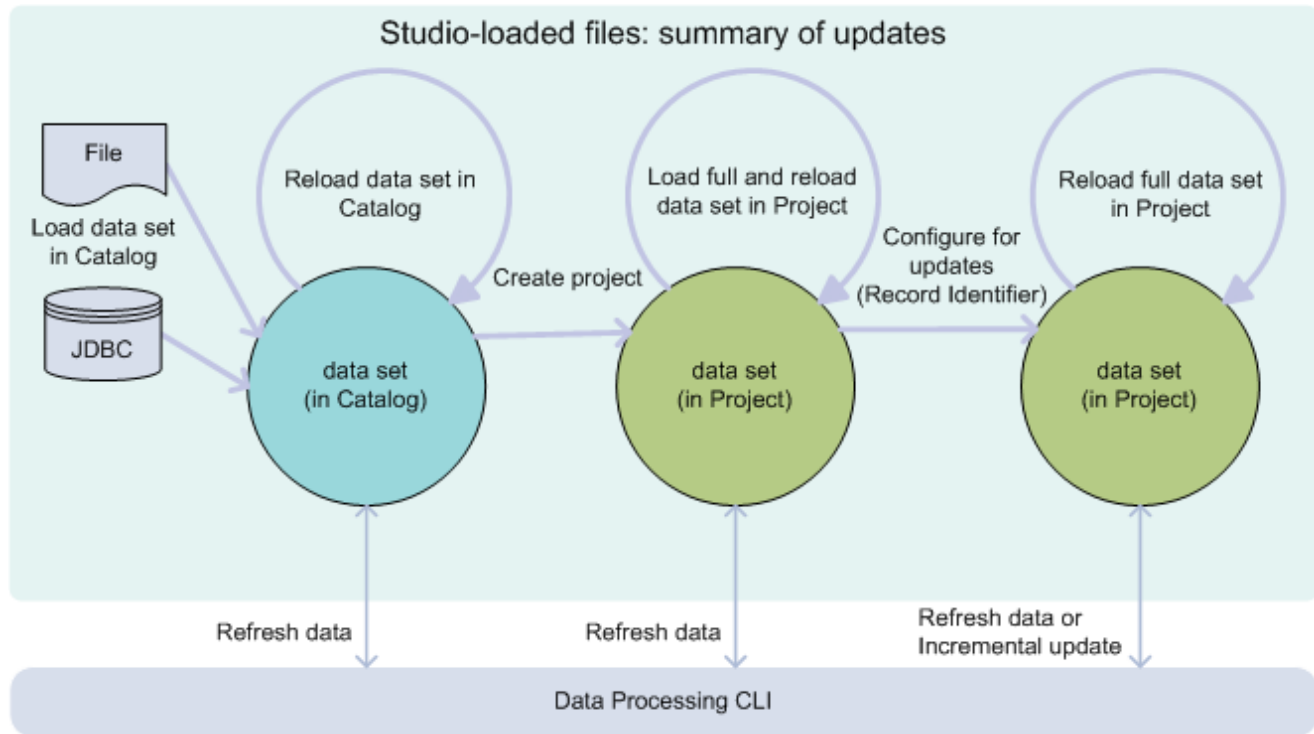
If you load a data set from a personal file or import it from a JDBC source, then all data is loaded. However, it may still be a sample if you compare it to the source data you may also have elsewhere on your system.

If you later want to add full data from the source, you can locate the Hive data set that BDD created when you loaded a file. Next, use the `drop` command to place that data set in Hue, and replace it with a production Hive table. You can then run **Load Full Data Set** on this table in Studio. This will load a full data set.

This process is known as creating a BDD application. For detailed steps on this procedure, see the topic on creating a BDD application in the *Studio User's Guide*.

## Studio-loaded files: data update diagram

The diagram in this topic shows data sets loaded in Studio by uploading a personal file or importing data from a JDBC source. It illustrates how you can reload this data set in Studio. Also, you can update the data set with DP CLI, and increase its size from sample to full.



In this diagram, the following actions take place, from left to right:

- You load data with the Studio's **Add data set** action. This lets you load a personal file, or data from a JDBC data source. Once loaded, the data sets appear in Catalog. BDD creates a Hive table for each data set. This Hive table is associated with the data set in Catalog.
- If you later have a newer version of these files, you can reload them with **Reload data set** action, found in data set's details, in Catalog.
- Next, if you'd like to transform, or to use **Discover** area to create visualizations, you should create a project. This is shown in the diagram in the middle circle. Notice that in this circle, the data set is now in your project and is no longer in Catalog. This data set can still be shown in Catalog, for others to use, if they have permissions to see it. However, now you have your own "version" of this data set, inside your project.
- Next, you can choose to load full data into this data set with **Load full data**, found in Data Set Manager. This replaces the data set sample with a fully loaded data set within your project. Or, if the data set was already fully loaded, you can reload the full data set in your project, by using this option again. This is shown as a circular arrow above the middle circle.
- Alternatively, you can use the options from Data Processing CLI to run scripted updates. There are two types of scripted updates: `Refresh Data` and `Incremental update`.



- To run `Refresh Data` with DP CLI, identify the data set's logical name in the properties for this data set in Studio. You must use that data set's logical name as a parameter for the data refresh command in DP CLI.
- To run `Incremental update` with DP CLI, you must specify a primary key (also known as record identifier in Studio). To add a primary key, use **Control Panel>Data Set Manager>Configure for Updates**. You also need the data set's logical name for the incremental update.
- Be careful to use the correct data set logical name. Notice that a data set in Catalog differs from the data set in a project.
- Notice that you can run **Reload data set** in Catalog only for data sets in Catalog. If you want to update a data set that is already added to a project, you need to use one of the scripted updates from DP CLI.

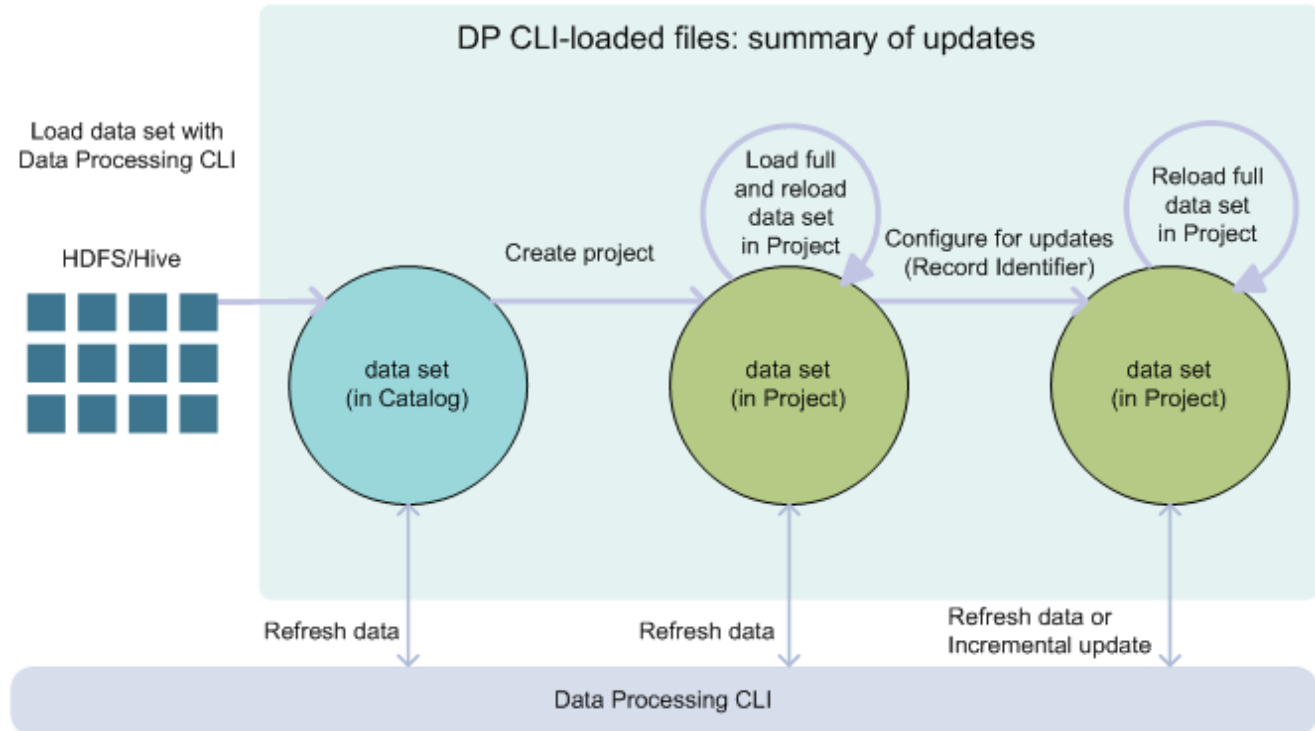
If you reload a data set in Catalog that has a private copy in a project, Studio alerts you to this when you open the project, and gives you a chance to accept these updates. If any transformations were made to the data set in the project, you can apply them to the data set in Catalog.

Note the following about this diagram:

- You can `Refresh Data` with DP CLI for data sets in Catalog and in projects. Typically, you use DP CLI `Refresh Data` for data sets in a project.
- You can run an `Incremental Update` with DP CLI after you specify a record identifier. For this, you must move the data set into a project in Studio.
- Once you load full data, this does not change the data set that appears in Catalog. Moving a data set to a project is similar to creating your personal version of the data set. Next, you can load full data and write scripted updates to this data set using DP CLI commands `Refresh data` and `Incremental update`. You can run these updates periodically, as `cron` jobs. The updates will run on your personal version of this data set in this project. This way, your version of this data set is independent of the data set's version that appears in Catalog.

## DP CLI-loaded files: data update diagram

The diagram in this topic shows data sets loaded by Data Processing component of BDD, from Hive. The diagram illustrates how you can update this data set using DP CLI, and increase its size from sample to full.



In this diagram, from left to right, the following actions take place:

- You load a data set from Hive using the Data Processing workflow (DP CLI). The data set appears in Catalog in Studio.
- You can now create a BDD project from this data set. Notice that a data set in your project does not eliminate this data set in Catalog. That is, other users may still see this data set in Catalog, if they have permissions. However, you have now taken this data set with you into your project. You can think of it as a different version of the data set, or your project's private version of the data set.
- In a project, you can load full data into it with the **Load full data set** action in Studio.
- Using Data Processing CLI, you can also run scripted updates to this data set. Two types of scripted updates exist: `Refresh Data` and `Incremental Update`.

To run scripted updates, you need a data set logical name, found in the data set's properties in Studio. It is important to provide the correct data set logical name to the DP CLI. If the data set is in Catalog, it is not the same data set that you have in your project. Note the correct data set logical name.

Note the following about this diagram:

- You cannot run an `Incremental Update` with DP CLI for a data set in Catalog; you can only run it when you add the data set to a project.
- You can run an `Incremental Update` with DP CLI after you specify a record identifier. For this, you must move the data set into a project in Studio.

- Because this data set arrived from Hive, you cannot update it within Studio. Instead, you can use DP CLI for data set updates.
- Once you load full data, this does not change the data set that appears in Catalog. Moving a data set to a project and loading full data is similar to creating a personal version of the data set. Next, you can write scripted updates to this data set, using DP CLI commands `Refresh data` and `Incremental update`. You can run these updates periodically, as `cron` jobs. The updates will run on your personal version of this data set in this project. This way, your version of this data set is independent of the data set's version that appears in Catalog.

With this workflow, you create a project of your own, based on this data set, where you can run scripted updates with DP CLI. This approach works well for BDD projects that you want to keep around and populate with newer data.

This way, you can continue using the configuration and visualizations you built in Studio before, and analyze newer data as it arrives.

## Data update options

Here is a summary of how you can update data loaded into BDD, and when each type of update is useful to use.

### Options for data updates

To update already loaded data, you have these options:

- [Reload data set in Studio on page 35](#)
- [Refresh data with DP CLI on page 35](#)
- [Run an incremental update with DP CLI on page ???](#)
- [When to use each type of update? on page ???](#)

Updates that you run with DP CLI are also called **scripted updates**.

### Reload data set in Studio

The **Reload data set** option in Studio is useful when you want to reload a newer version of the data than you loaded before. It applies to personally uploaded files and to data imported from a JDBC source. Note that this option works only for data sets in Studio's Catalog.

For a diagram of updating data sets that were loaded in Studio, see [Studio-loaded files: data update diagram on page 32](#) in this guide.

For detailed procedures for loading and reloading data in Studio, see the sections in the *Studio User's Guide*.

### Refresh data with DP CLI

The `Refresh data` operation from DP CLI reloads an existing data set in a Studio project, replacing the contents of the data set with the latest data from Hive in its entirety. If the schema in the source Hive table changes, so does the newly-referenced data set. In this type of update, old data is removed and is replaced with new data. New attributes may be added, or attributes may be deleted. Also, the data type for an attribute may change.

For a diagram of updating data sets that were loaded with DP CLI, see [DP CLI-loaded files: data update diagram on page 34](#) in this guide.

For detailed information on how to run scripted updates with DP CLI, see the *Data Processing Guide*.

## Run an incremental update with DP CLI

The `Incremental update` operation from DP CLI lets you add newer data to an existing BDD application, without removing already loaded data. In this type of update, the records' schema cannot change. An incremental update is most useful when you keep already loaded data, but would like to continue adding new data. For example, you can add more recent twitter feeds to the ones that are already loaded.

For a diagram of updating data sets that were loaded with DP CLI, see [DP CLI-loaded files: data update diagram on page 34](#) in this guide.

For detailed information on how to run scripted updates with DP CLI, see the *Data Processing Guide*.

## When to use each type of update

This table summarizes when it is useful to use each type of update.

Type of data update	Useful when...
<b>Reload data set in Catalog</b> (in Studio)	This update is useful when you want to replace the loaded file with an updated version. Similarly, if data in a JDBC source was updated, you can reload it this way.
Scripted updates with DP CLI (Refresh data and Incremental update)	<p>You can run scripted updates on files that originated from a Studio's upload, and on files that BDD discovers in Hive, when you run its data processing workflow for loading data using DP CLI.</p> <p>You can run either type of the scripted updates periodically, by writing update scripts and <code>cron</code> jobs on the Hadoop machines that utilize options for these updates from Data Processing CLI.</p> <p>Depending on characteristics of your source data, you may need to periodically run both types of scripted updates, or only one of them.</p> <p>For example, you may want to create a <code>cron</code> job that runs an incremental update nightly. This adds data from that day to the existing data set in a project in Studio.</p> <p>In addition to a periodic incremental update, you can run a <code>Refresh data</code> update weekly, to replace the data in the project wholesale with the new data that was collected in Hive during the week.</p> <p>A <code>Refresh data</code> update is also useful to run weekly because it lets you handle deletes from the source data set.</p>



## Chapter 5

# Where You Go from Here

*Quick reference to areas of interest*

## Quick reference to areas of interest

Use this table to locate information on a specific subject of BDD.

<b>BDD feature</b>	<b>Where to find information on it?</b>
List of addressed and known issues and most recent patches	Release notes for Big Data Discovery, available at the Oracle download site.
Installing BDD	<i>Installation Guide</i>
Upgrading BDD from a previous version	<i>Upgrade Guide</i>
Loading data into BDD with DP CLI	<i>Data Processing Guide</i>
Loading data in Studio from a personal file or a JDBC source	<i>Studio User's Guide</i>
Updating data sets with DP CLI: Refresh data and Incremental update	<i>Data Processing Guide</i>
Updating data sets in Studio	<i>Studio User's Guide</i>
Transforming data	<i>Studio User's Guide</i>
Hadoop-related information in Big Data Discovery, including default attribute types at loading time, sample size, DP CLI options, Dgraph HDFS Agent, and logs from the Data Processing component.	<i>Data Processing Guide</i>
Writing your own Custom Visualization Component	<i>Extensions Guide</i>
Administering BDD, including using <code>bdd-admin</code> script for starting, stopping and other lifecycle operations	<i>Administrator's Guide</i>
Understanding security features in BDD	<i>Security Guide</i>

# Glossary

# Index

## A

applications, in BDD 27

## B

BDD

- as the data lake navigator 8
- using in the data lab 7
- using to create BDD applications 8

Big Data Discovery

- data scientist's tasks 9
- value of using 8
- workflow 13

breadcrumbs 20

## D

data loading options 30

Data Processing 10

Data Set Manager 23

data sets

- about 22
- access 29
- Data Set Manager 22
- getting from a sample to a full data set 31
- in Catalog vs in project 23
- lifecycle in Studio 25
- loading 31
- logical name 24
- preview in Catalog 22
- private or public 29
- sampled and full 25

data source type 24

Dgraph 11

Discover, area of Studio 19

## E

enrichments 16, 26

Explore, area of Studio 16

## F

filtering 20

Find, portion of Studio 14

## G

Guided Navigation 20

## H

Hive Table Detector 27

## J

JDBC data source 30

## K

key

See data sets: logical name

## P

profiling 16

projects, in BDD 27

## R

refine by 20

results of working with BDD 9

## S

search 20

security of data sets 29

Studio

- about 10
- user interface, Discover 19
- user interface, Explore 16
- user interface, Find 14
- user interface, Transform 17

## T

Transform, area of Studio 17

transformations 18

## U

updates

- diagram for files loaded from Hive 32, 34
- options 35

## V

value of using BDD 8

visualizations 19

## W

where you go from here 37

workflow in Big Data Discovery 13