cloudera®

Altus Shared Data Experience (SDX)

Important Notice

The information in this document is subject to change without notice. Cloudera shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

Cloudera, Inc.
395 Page Mill Road
Palo Alto, CA 94306
info@cloudera.com
US: 1-888-789-1488
Intl: 1-650-362-0488
www.cloudera.com

Release Information

Version: Cloudera Altus
Date: May 8, 2019

Table of Contents

Overview of Altus SDX

Altus Shared Data Experience (SDX) is a service that provides a consistent view of data for clusters and workloads running in the cloud. Altus SDX externalizes cluster metadata into a shared, long-running service available to multiple clusters and workloads running in the cloud.

Altus SDX enables you to externalize cluster metadata with the Altus SDX namespace. The Altus SDX namespace points to a database that stores metadata for the data accessed by clusters in the cloud, shareable among multiple clusters that access the same data. You can use an Altus SDX namespace for clusters that access data in Amazon S3 or in Azure Data Lake Store (ADLS).

Ordinarily, metadata does not persist beyond the lifecycle of a cluster. When you run workloads across multiple clusters, you lose the metadata when you terminate a cluster and start another. You must define the data structures each time you start a new cluster and access data. When you configure multiple clusters in the cloud to share the Altus SDX namespace, you enable the clusters to access the data without the need to recreate the metadata.

Configured SDX Namespaces

An Altus SDX namespace backed by Hive metastore and Sentry databases that you set up and manage is called a configured SDX namespace. You can create configured SDX namespaces for use with workloads that run on clusters that you create in Altus and clusters that you create with Altus Director or Cloudera Manager.

When you create an Altus cluster, you can specify a configured SDX namespace to use with the cluster. The configured SDX namespace can also be used by other clusters in the cloud that access the same dataset. Any metadata generated by your cluster is stored in the Hive metastore and Sentry database associated with the configured SDX namespace and can be used by other clusters that share the same Hive metastore and Sentry database or that use the same configured SDX namespace.

An Altus cluster can read metadata from or write metadata to only one SDX namespace. Clusters that use the same SDX namespace share only the metadata of the dataset that they access. Each cluster uses its own computing power to access the data and execute jobs in the cluster.

SDX Administrator

To create or delete an SDX namespace, you must be an Altus administrator or have the *SdxAdmin* role.

If you are an Altus administrator, you can assign Altus users the *SdxAdmin* role so that they can create SDX namespaces to use for their clusters or to be shared with clusters created by other users. You can assign the *SdxAdmin* role to an Altus user, machine user, or an Altus group.

For more information about assigning the *SdxAdmin* role to a user or group, see Assigning a Role.

SDX Sentry Administrator Group

Altus uses Apache Sentry as the authorization service for user access to data and metadata stored in the Hive metastore database. When you create a configured SDX namespace, Altus creates an SDX Sentry administrator group for the SDX namespace. The SDX Sentry administrator group is an Altus group which Altus adds to the Sentry server as an administrator group. You, as creator of the configured SDX namespace, are automatically a member and the administrator of the Altus group.

Altus also assigns your user account the *IamGroupAdmin* resource role for the group, which makes you the group membership administrator for the SDX Sentry administrator group. You can add users to or remove users from the group.

Although you can use the SDX Sentry administrator group the same way as other Altus groups, Cloudera recommends that you treat the SDX Sentry administrator group as a special-use group and manage it differently than other Altus groups. For guidelines on using the SDX Sentry administrator group, see Guidelines for Using the SDX Sentry Administrator Group on page 6.

SDX Sentry Administrator Group Privileges

When you create a configured SDX namespace, you can select what privileges Altus grants the associated SDX Sentry administrator group.

By default, the SDX Sentry administrator group has administrative privileges in Sentry, which includes the privilege to create roles and grant privileges to groups in Sentry. You can select to grant the SDX Sentry administrator group ALL privileges in addition to the default administrative privileges.

On the Altus console, when you create a configured SDX namespace, Altus provides the following options for setting the SDX Sentry Administrator Group privileges:

- **Yes, automatically grant ALL Sentry privileges to the admin group**

When you select this option, Altus creates a role in Sentry with the same name as the SDX Sentry administrator group with ALL privileges on the Sentry server. Altus then assigns the role to the SDX Sentry administrator group.

With this option, a member of the SDX Sentry administrator group can perform the following tasks:

- Create roles and grant privileges to groups in Sentry.
- Create and manage database schemas.

You might want to select this option if you need to create databases immediately after you create the cluster you associate with the configured SDX namespace, such as for testing or demonstration purposes.

- **No, I will set up roles and privileges in Sentry**

  When you select this option, Altus does not create a role for the SDX Sentry administrator group to grant ALL privileges on the Sentry server. A member of the SDX Sentry administrator group has Sentry administrative privileges but cannot create or manage database schemas.

  With this option, you set up and manage groups and roles in Sentry to grant database privileges.

  Select this option if you already have Sentry authorization policies in place and you do not want to provide additional privileges to new groups. For example, the Sentry database that you use for the configured SDX namespace is one that you created previously and already has groups and roles set up with authorization to specific databases and tables. You might not want to assign a role to the SDX Sentry administrator group that provides members of the group unrestricted access to all databases and tables.

> **Note:** Because members of the SDX Sentry Administrator Group have administrative privileges in Sentry, they can also grant ALL privileges to the SDX Sentry administrator group and any other group in the Sentry server.

Guidelines for Using the SDX Sentry Administrator Group

Use the following guidelines for using and managing the SDX Sentry administrator group:

Any user who is a member of the SDX Sentry administrator group is a Sentry administrator for the configured SDX namespace.

When you add a user account to the SDX Sentry administrator group, the user becomes a Sentry administrator and can grant privileges to users who access data and metadata stored in the configured SDX namespace. Likewise, if you remove a user from the SDX Sentry administrator group, you revoke the user's Sentry administrative privileges.

Make sure that the users that you add to the SDX Sentry administrator group are only those users that require the SDX Sentry administrator group privileges to do their jobs.

Altus uses a naming convention for the SDX Sentry Administrator Group name.

Altus creates the SDX Sentry administrator group with the following naming convention:

admin*PartOfSDXNamespaceName_UniqueID*

Altus includes the first nine alphanumeric characters, excluding special characters, of the SDX namespace name in the SDX Sentry administrator group name.

Altus lists SDX Sentry administrator groups among other Altus groups in the in the **Groups** page on the Altus console. You can determine whether the group is an SDX Sentry administrator group and which SDX namespace it is associated with by looking at the group name.

When you delete the configured SDX namespace, Altus deletes the SDX Sentry Administrator Group.

Do not delete the SDX Sentry administrator group. The configured SDX namespace requires its associated SDX Sentry Administrator Group to work with Sentry.

If you no longer require the configured namespace, delete the configured SDX namespace. Altus deletes the namespace and the associated SDX Sentry Administrator Group.

Creating a Configured SDX Namespace

You must be an Altus administrator or have the *SdxAdmin* role to create an SDX namespace.

To create an SDX namespace on the console:

1. Sign in to the Cloudera Altus console:

   https://console.altus.cloudera.com/

2. On the side navigation panel, click **SDX Namespaces**.

   The **SDX Namespaces** page displays the list of SDX namespaces available in the Altus account.

3. Click **Create Namespace** and select **Configured Namespace**.
4. On the **Create Configured Namespace** page, set the name of the configured namespace

   The name of the SDX namespace must be unique within the Altus account. The name is case-sensitive and can have a maximum of 128 characters. It can contain only alphanumeric characters, hyphens (-), and underscores (_).

5. In the **Hive Metastore Settings** section, set the following parameters:

Property	Description
JDBC URI	Connection URL that the Altus SDX service uses to connect to the Hive metastore database that you want to use for the Altus cluster.
User Name	User account to use to log in to the Hive metastore database.
Password	Password for the user account used to log in to the Hive metastore database.
Confirm Password	Type the password again to confirm.

6. In the **Sentry Settings** section, set the following parameters:

Property	Description
JDBC URI	Connection URL that the Altus SDX service uses to connect to the Sentry database that stores the database security policies you want to use for the Altus cluster.
User Name	User account to use to log in to the Sentry database.
Password	Password for the user account used to log in to the Sentry database.
Confirm Password	Type the password again to confirm.

   When you create a configured SDX namespace, Altus creates an associated Altus group and adds it to Sentry as an administrator group. By default, you, as creator of the configured SDX namespace, are a member and administrator of the group and have the administrative privileges to create roles and grant privileges in Sentry.

7. In the **Sentry Admin Group** section, select what privileges Altus grants to the SDX Sentry administrator group.

   Select one of the following options:

Option	Description
Yes, automatically grant ALL Sentry privileges to the admin group	Allows Altus to grant ALL privileges on the Sentry server to the SDX Sentry Administrator Group in addition to administrative privileges.

Option	Description
	When you select this option, the SDX Sentry Administrator group can create roles and grant privileges in Sentry. The SDX Sentry Administrator group can also create and manage database schemas.
No, I will set up roles and privileges in Sentry	Does not allow Altus to grant ALL privileges on the Sentry server to the SDX Sentry Administrator Group. When you select this option, the SDX Sentry Administrator group can create roles and grant privileges in Sentry but cannot create or manage database schemas.

For more information about the SDX Sentry Administrator Group privileges, see SDX Sentry Administrator Group Privileges on page 5.

8. Click **Create** to create the SDX namespace.

   The SDX namespace you create displays in the list of SDX namespaces with type **Configured**.

To delete an Altus SDX namespace, click the name of an SDX namespace. On the SDX namespace page, click **Actions** > **Delete namespace**. Click **OK** to confirm that you want to delete the SDX namespace. Altus deletes the SDX namespace and the SDX Sentry administrator group associated with the SDX namespace.

Configured SDX Namespace Use Cases

The configured SDX namespace enables you to share your Hive metastore and Sentry policies among clusters that you launch in the cloud. For example, if you set up a Hive metastore and Sentry database for clusters that you create using Altus Director, you can use the same Hive metastore and Sentry databases for clusters that you create with Altus services.

A common use case for configured SDX namespaces is when you need additional capacity for a workload. Typically, you maintain a large cluster with capacity to handle any spike in workload volume. For example, you have a workload that runs every day. Most days of the month you might use only 50% of capacity to run the workload. But at the end of the month, the volume of data you need to process is higher and you need much more capacity to be able to meet workload SLAs.

An Altus SDX namespace enables you to use Altus clusters to meet the capacity requirements at the end of the month without having to maintain the maximum capacity throughout the month. When your cluster requirements go up at the end of the month, you can use Altus clusters to increase your capacity to run your workload. You can create Altus clusters and set them up to use a configured SDX namespace that point to the Hive metastore and Sentry databases used by the other clusters in the workload. The Altus cluster can immediately access the dataset using the metadata in the SDX namespace without the need to first define the metadata for the dataset accessed by the workload.

After the monthly spike, when you no longer need the extra capacity, you can delete the Altus cluster and run the workload with the cluster capacity that you regularly need.

Guidelines for Using SDX Namespaces

Use the following guidelines when you use configured SDX namespaces with Altus clusters:

When you create databases and tables, use the *LOCATION* attribute to write the data to cloud storage.

When you create a database or table that you want to make accessible to other clusters that share the SDX namespace, you must create the database or table in cloud storage. Use the *LOCATION* attribute to indicate the location in cloud storage where you want to create the database or table.

If you do not provide the location, the database or table is created in a default location in HDFS in the cluster. When the cluster is terminated, the HDFS databases and tables are lost.

The following examples show the CREATE DATABASE statement with the LOCATION attribute pointing to S3 and ADLS locations:

```
CREATE DATABASE databasename LOCATION s3a://path-to-aws-s3/dir
CREATE DATABASE databasename LOCATION adl://path-to-azure-adlgen1/dir
CREATE DATABASE databasename LOCATION abfs(s)://path-to-azure-adlgen2/dir
```

For more information, see CREATE DATABASE Statement.

The following examples show the CREATE TABLE statement with the LOCATION attribute pointing to S3 and ADLS locations:

```
CREATE EXTERNAL TABLE tablename LOCATION s3a://path-to-aws-s3/dir/table_data
CREATE EXTERNAL TABLE tablename LOCATION adl://path-to-azure-adlgen1/dir/table_data
CREATE EXTERNAL TABLE tablename LOCATION abfs(s)://path-to-azure-adlgen2/dir/table_data
```

For more information, see CREATE TABLE Statement.

To view the location attribute of a database, use the DESCRIBE DATABASE statement:

```
DESCRIBE DATABASE databasename
```

When you create interim tables in HDFS, use unique names.

Problems can arise when different clusters using the same SDX namespace create interim Hive tables with the same name in HDFS. For example, a cluster creates a table named *Temp* to store temporary data in HDFS. When a job in another cluster also creates a table named *Temp* to store temporary data in HDFS, the job can fail or it can overwrite data in the table.

To avoid naming conflicts, use names unique to the cluster. An easy way to make a table name unique is to include the cluster ID in the table name.

Avoid concurrent updates by multiple clusters to the same schema, table, or partitions in a table.

The Altus SDX service does not manage the metadata updates made by different clusters. It does not have a mechanism to lock the metadata to prevent simultaneous updates by multiple clusters. Data conflicts and errors can arise if multiple clusters sharing an SDX namespace access a dataset at the same time and perform conflicting updates.

For example, problems can arise if multiple clusters concurrently update the same table or partitions within a table or add or change the same schema or database.

Run your workloads in a way that ensures that multiple clusters do not make overlapping data or metadata changes.

Refresh Altus Data Warehouse clusters after other clusters make changes to the metadata.

If you configure an Altus Data Warehouse cluster to use an SDX namespace that is also used by other clusters, you must take into consideration that the Impala service in the Altus Data Warehouse cluster keeps a local HDFS metadata cache.

If another cluster modifies the dataset for the Altus Data Warehouse cluster, Altus SDX updates the SDX namespace with the change in the metadata. However, because the change is made in another cluster, the Altus Data Warehouse cluster is not updated. You must run refresh or invalidate metadata operations to incorporate the latest updates.

For example, you create an Altus Data Warehouse cluster and configure it to use an SDX namespace. Then you start another cluster configured to use the same SDX namespace and add a column to a table in the dataset. Altus SDX updates the metadata in the SDX namespace with the new column. However, because the change to the table is done outside the Altus Data Warehouse cluster, the metadata cache in the Altus Data Warehouse cluster is not updated with the new column.

To avoid errors with obsolete metadata, refresh or invalidate the metadata in the Altus Data Warehouse cluster to get the latest metadata from the SDX namespace.

Ensure that interim local tables in HDFS are deleted before you terminate a cluster.

When you write interim data to a table stored in HDFS in a cluster, the metadata for the interim files is stored in the SDX namespace. If you terminate the cluster, Altus SDX does not delete the metadata for these tables from the SDX namespace.

The metadata for the HDFS tables remain in the SDX namespace and can cause data conflicts and errors for other clusters. For example, a job in another cluster that uses the same SDX namespace might try to read data in the tables and encounter errors because the HDFS locations do not exist or are not be valid.

To avoid errors with orphaned metadata in an SDX namespace, delete all tables created in HDFS before you terminate a cluster.

Do not delete an SDX namespace that is being used by an Altus cluster.

When you delete a configured SDX namespace, Altus deletes the SDX namespace in Altus and the associated SDX Sentry administrator group. Altus does not delete the Hive metastore or the Sentry database that the configured SDX namespace points to. If you delete a configured SDX namespace, Altus clusters lose access to the Hive metastore and Sentry database that the SDX namespace points to. Before you delete an SDX namespace, verify that the namespace is not used by any Altus cluster.

Other CDH clusters that share the same Hive metastore and Sentry database, such as clusters created using Director, are not affected when the configured SDX namespace is deleted and can still access the metadata.

When you use S3Guard for clusters that store data in AWS S3, use the same S3Guard instance for all clusters that share an SDX namespace.

Using the same S3Guard instance for clusters that share the same SDX namespace ensures that the clusters see the metadata stored in AWS S3 in a consistent manner.

You can set clusters to use S3Guard by enabling S3Guard consistency in the Altus environment that you use to create them. When you create clusters using the same Altus environment with S3Guard consistency enabled, the clusters you create use the same S3Guard instance.

If you create clusters using different Altus environments with S3Guard consistency enabled, set the environments to use the same S3Guard instance if you want the clusters to share the same SDX namespace.

Setting up a Cluster with a Configured SDX Namespace

You can use a configured SDX namespace with a secure Altus Data Warehouse or Data Engineering cluster. A secure Altus cluster uses an environment with the *Secure Clusters* option enabled.

To set up a cluster with an SDX namespace, complete the following steps:

1. **Create or identify the Hive metastore and Sentry databases that you want to use for the configured SDX namespace.**

   Altus SDX supports using a MySQL or PostgreSQL database for the Hive metastore and Sentry.

   You can manually set up the databases to use for the configured SDX namespace or you can use Altus Director to set up the databases:

   - **Manually setting up the databases**

     If you have the rights to set up databases on your cloud service provider, you can set up the databases for the Hive metastore and Sentry that you want to use for the configured SDX namespace. You must ensure that any cluster you create that uses the configured SDX namespace will have access to the Hive metastore and Sentry databases. When you create a cluster that uses the configured SDX namespace, Altus initializes the database schemas.

     For more information about setting up a MySQL database, see Install and Configure MySQL for Cloudera Software. For more information about creating a PostGres database, see Install and Configure PostgreSQL for Cloudera Software.

   - **Using Altus Director to set up the databases**

     If you use Altus Director to set up the external databases, you can also use Altus Director to create a cluster and generate the schema in the Hive metastore and Sentry databases.

     For more information about using Altus Director to create an external database, see Defining External Database Servers.

   Note the connection URI for the Hive metastore and Sentry databases. You need to provide the connection URI and the database user login credentials when you create the configured SDX namespace in Altus.

2. **Create a configured SDX namespace in Altus.**

   For more information about creating a configured SDX namespace, see Creating a Configured SDX Namespace on page 7.

3. **Set up a secure Altus Data Engineering or Altus Data Warehouse cluster that uses the configured SDX namespace.**

   When you create the cluster, set the following parameters :

   - ***SDX Namespace***. Specify the configured SDX namespace you created in step 2 on page 12.
   - ***Environment***. Specify an environment with the *Secure Clusters* option enabled. The Altus environment must also be set up to allow access from the clusters to your data in object storage and to your Hive metastore and Sentry databases.

   For more information about enabling the *Secure Clusters* option for an environment, see Enable Secure Clusters.

   For more information about creating an Altus Data Engineering cluster, see Creating and Working with Clusters on the Console or Creating and Working with Clusters Using the CLI.

   For more information about creating an Altus Data Warehouse cluster, see Altus Data Warehouse Clusters on the Console or Altus Data Warehouse Clusters in the CLI.

4. **Grant administrator privileges to the SDX Sentry administrator group.**

When you create a configured SDX namespace, Altus creates an Altus group and adds it to Sentry as an administrator group. You, as creator of the configured SDX namespace, are a member of the group and have administrative privileges in Sentry.

Altus provides the option to grant the Sentry administrator group ALL privileges on the Sentry server in addition to administrative privileges. Based on the option you select, Altus performs the following actions:

- **If you select the option to grant the Sentry administrator group ALL privileges:**

  Altus creates a role with ALL privileges on the Sentry server and assigns it to the Sentry administrator group when you create a cluster to use with the SDX namespace. You can immediately create the databases you need for your use. Cloudera recommends that you select this option when you create a configured SDX namespace to for testing or demonstration purposes.

- **If you do not select the option to grant the Sentry administrator group ALL privileges:**

  You must create the groups and roles with the privileges that you require to secure access to the databases and assign the appropriate role to the SDX Sentry administrator group.

  You can create a role with ALL privileges and assign the role to the SDX Sentry administrator group so that you can create and manage the databases that you need to work with. To grant all privileges to the SDX Sentry administrator group, run the following commands:

```
create role SentryAdminRoleForAltus;
grant all on server server1 to role SentryAdminRoleForAltus;
grant role SentryAdminRoleForAltus to group SDXSentryAdminGroup;
```

  For Altus Data Engineering clusters, submit a Hive job to run the commands.

  For Altus Data Warehouse clusters, use the Query Editor to run the commands.

For more information about the SDX Sentry administrator group privileges, see SDX Sentry Administrator Group Privileges on page 5.

Appendix: Apache License, Version 2.0

SPDX short identifier: Apache-2.0

Apache License
Version 2.0, January 2004
http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution.

You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

1. You must give any other recipients of the Work or Derivative Works a copy of this License; and
2. You must cause any modified files to carry prominent notices stating that You changed the files; and
3. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
4. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions.

Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks.

This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty.

Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability.

In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability.

While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

```
Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

  http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```