Delphix Platform Extensibility Using the Virtualization SDK
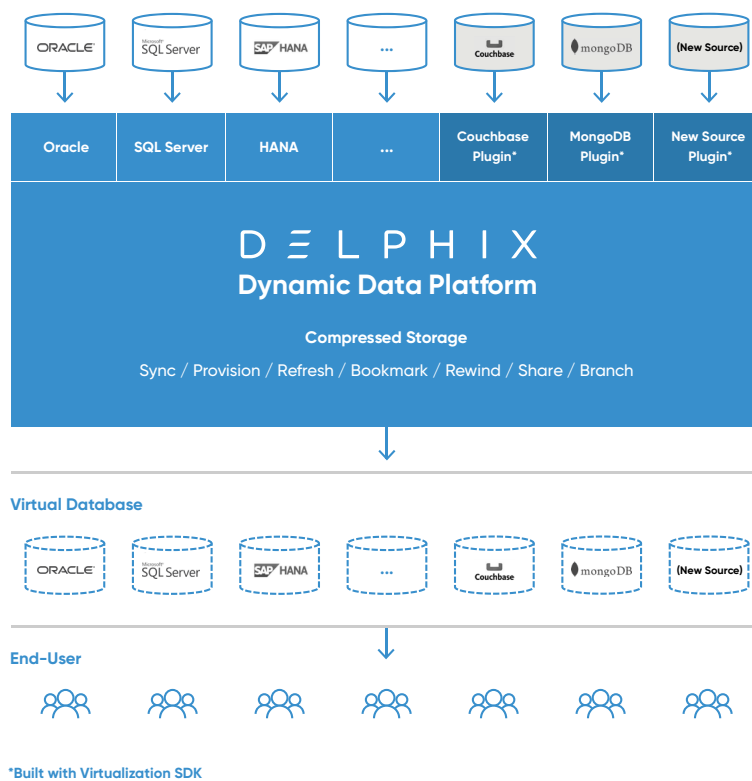

FIGURE 1 The Delphix Dynamic Data Platform with Delphix Virtualization SDK

Applications are the new competitive edge — and the best way to build and test modern applications is with copies of production data. However, access to production-quality data is hampered by the variety of data sources and complexity of tools and processes required for access to those sources.

The Delphix Dynamic Data Platform addresses the complexity of heterogeneous data sources by providing a consistent set of abstractions for data ingestion, management, and delivery. Delphix users have access to virtual databases that are lightweight; designed for developers, testers, and other users with needs for fresh data; and easy to rewind, branch, and share.

The Delphix platform has built-in support for databases such as Oracle, SQL Server, and ASE. The platform expands its coverage with available virtualization plugins for a wide variety of SQL and NoSQL databases — such as PostgreSQL or MongoDB — whether managed directly on-premises or available as PaaS cloud services.

The Delphix Virtualization SDK expands support for data sources even further by enabling organizations to build their own plugins. The SDK can be used to develop powerful data management solutions for a variety of databases.

DELPHIX

How Delphix Virtualization SDK Plugins Work

Plugins extend the Delphix platform's capabilities by telling it how to run essential virtual data operations on these datasets, including starting and stopping a dataset, where to store the data, and making virtual copies of datasets.

BUILDING A PLUGIN

The Delphix platform provides pre-defined wrappers for a plugin. An organization provides the python code needed to access the datasource, and details such as a port number or admin credentials.

The Delphix platform assembles these elements into a plugin. Using the code and information in the plugin, the Delphix platform can ingest the dataset from the original data source.

INGESTION

The Delphix platform uses the plugin to ingest data from a source environment, either by pulling data directly in using rsync, or using a staging environment, such as backup files.

To Ingest a datasource, the Virtualization SDK needs plugins to supply information and operations for discovery and snapshotting of the data source.

The Delphix platform then creates a compressed copy of the ingested data. The compressed copy can be updated on a schedule set by the administrator.

PROVISIONING

The Delphix platform can provision the ingested data to users as virtual datasets — just as it would with built-in data sources.

First the dataset is cloned onto the Delphix engine, then this newly-cloned data is mounted as a virtual dataset onto a target environment. This new virtual dataset can be updated by its end users while preserving the state of the original snapshot on the Delphix platform. The snapshot on the Delphix platform is regularly updated from the source dataset, and can be used to provision and update virtual databases to match the current state of the source dataset.

For provisioning, the Delphix platform needs plugins to supply information and operations to configure, reconfigure, and unconfigure the virtual source.

SELF-SERVICE

Virtual database end users are able to use the Delphix platform to refresh, rewind, replicate, branch, and share a virtual database with other users.

To perform these operations, the platform requires the organization to provide information and operations to start, mount, return status, snapshot, and stop the virtual database.

ABOUT DELPHIX

Delphix's mission is to free companies from data friction and accelerate innovation. Fortune 100 companies use the Delphix Dynamic Data Platform to connect, virtualize, secure and manage data in the cloud and in on-premise environments. For more information visit www.delphix.com.

DELPHIX