What's New

New Features for July 2022 & 6.0.15

Product Offerings

Connectors

Add-ons

Base products

Application

Database

Replication
(Active / Passive)

Continuous Vault

Continuous Data
(Virtualization)

Data Control Tower

Application

Database

Hyperscale Compliance

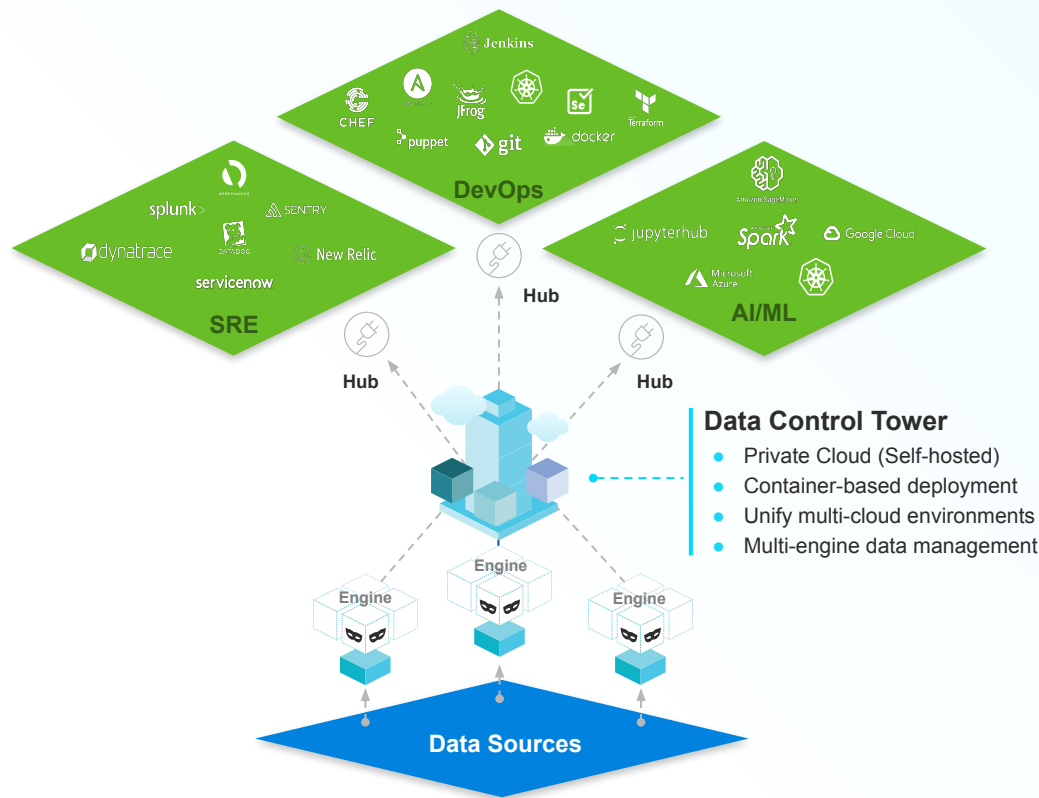Continuous Compliance
(Masking)

DevOps
Data Platform
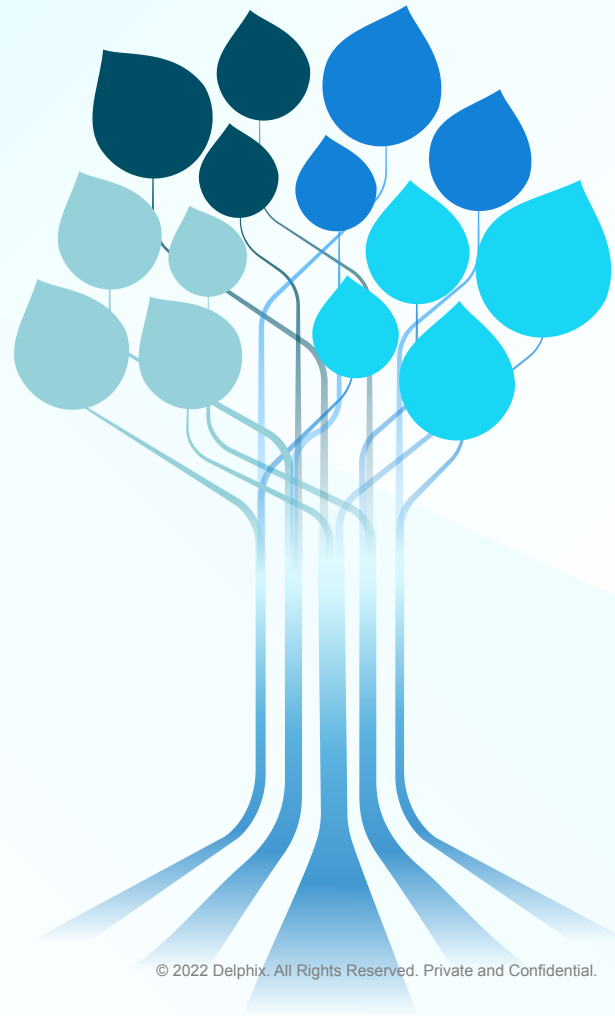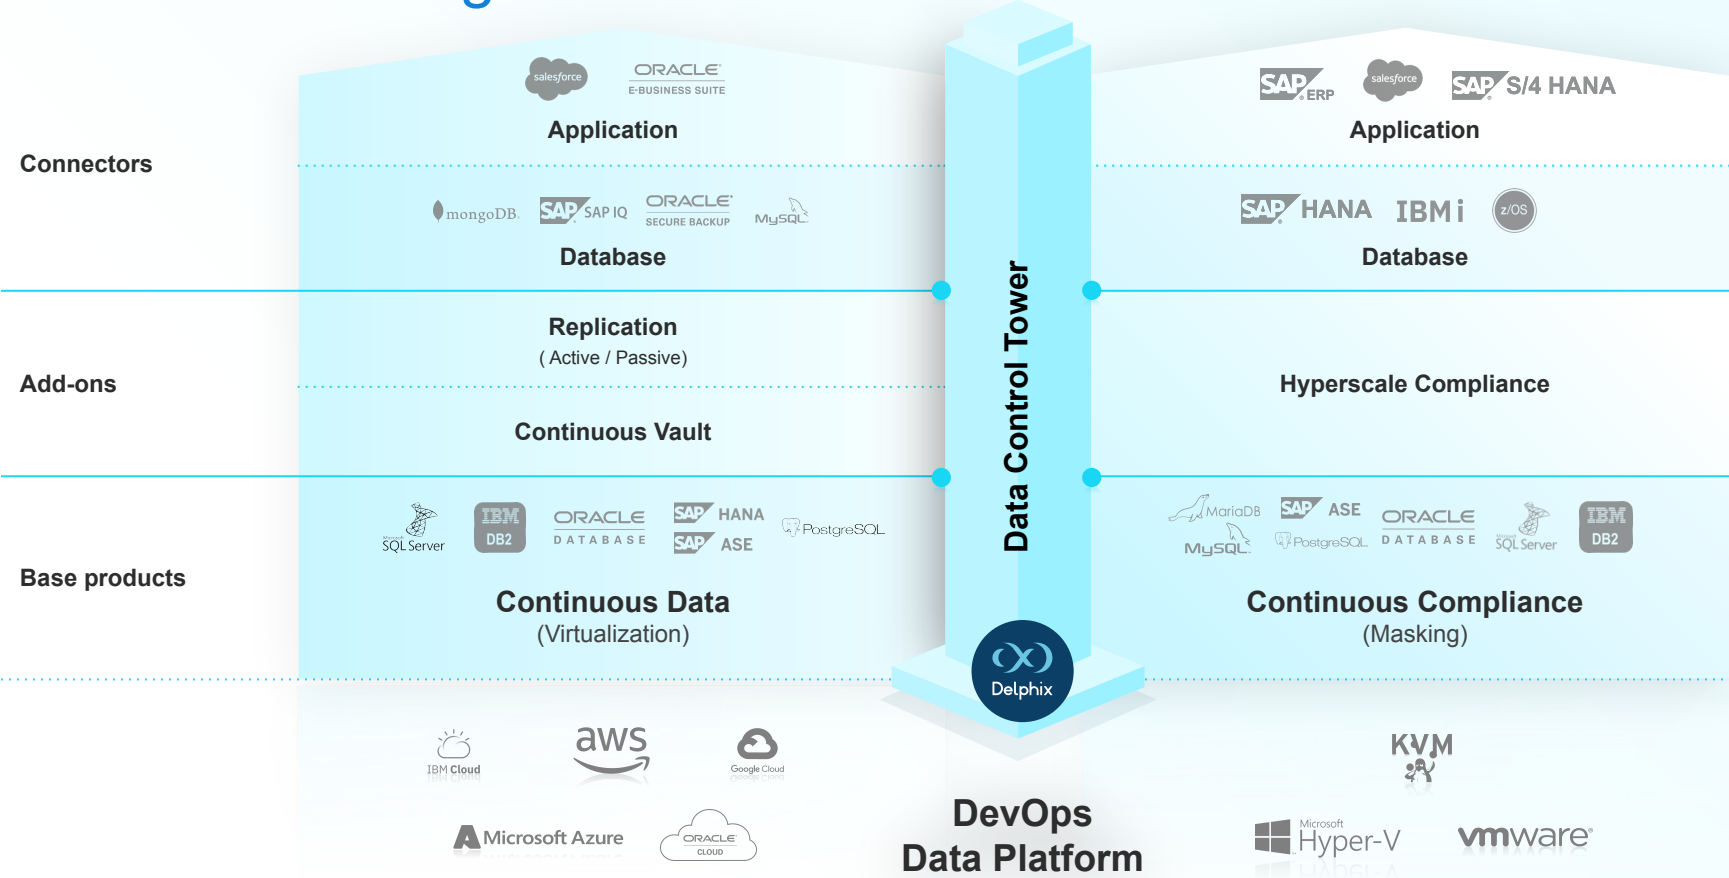
Data Control Tower

New & Updated for July 2022

- DCT 2.0 with over 30 new APIs

- Terraform Provider using DCT

- CI/CD APIs for Ephemeral Test Data Management (TDM) collaboration workflow

- Batch refresh VDB using tags for an application workflow

Data Control Tower 2.0

Data Control Tower
- Private Cloud (Self-hosted)
- Container-based deployment
- Unify multi-cloud environments
- Multi-engine data management

Generally Available

- **API-Driven Data Automation to Accelerate DevOps** - Single point of integration for automation toolchains across many Delphix engines

- **Centrally Govern User Data Access** - Control access to sensitive data for compliance and automate data entitlement authorization

- **Unified Engine Administration with Enterprise-wide Visibility** - Leverage a central point of management and reporting for multiple Delphix engines

Data Control Tower APIs

Legend:
- New in 2.0
- Existing in 1.0.6

Tags
- Create tag for VDB
- Remove tag from VDB
- List all tags
- List all keys
- List tags by key
- List tags by value
- List tags by key-value
- List values for a tag key
- Create tag for engine
- Delete tag for engine
- Create tag for environment
- Delete tag for environment
- Create tag for dSource
- Delete tag for dSource

Environment
- Enable environment
- Disable environment
- Refresh environment
- Create environment
- Delete environment by ID
- List all environments
- Show environment by ID

VDB Groups
- List bookmarks by group
- List all VDB groups
- List VDB group by ID
- Create VDB group
- Refresh VDB groups
- Delete VDB group
- Rollback VDB group

Reporting
- Storage capacity
- VDB inventory
- dSource inventory

VDB
- Refresh by snapshot
- Refresh by timestamp
- Provision by Bookmark
- Provision by timestamp
- Provision by snapshot
- List all VDBs
- Show VDB by ID
- Start VDB
- Stop VDB
- Enable VDB
- Disable VDB
- Delete VDB
- Rollback by snapshot
- Rollback by timestamp

Snapshots
- Delete snapshot by ID
- Create snapshot
- List snapshots
- Show snapshot by ID

Bookmarks
- VDB groups by bookmark
- Create bookmark
- Show bookmark by ID
- List bookmarks
- Rewind to bookmark
- Delete bookmark

Management
- List engines
- Show engine by ID
- Update engine
- Register engine
- Unregister engine
- Show Hashi vault by ID
- Configure Hashi vault
- List Hashicorp vaults
- Remove Hashi vault by ID

API Clients
- List API clients
- Show client by ID
- Create new API client
- Update client by ID
- Remove client by ID

Jobs
- Job by ID
- Abandon job by ID

dSources
- List all sources
- Show source by ID

Sources
- List all sources
- Show source by ID

Generally Available

- **30+ New APIs** for fast integration with CI/CD pipelines and Test Data Management (TDM) workflows

- **70+ Total APIs** to manage data, environments, engines, dSources, and VDBs

- **New tagging and VDB group APIs** to overlay business intelligence data

- **Containerized API gateway** for multi-cloud environments

- **Aggregated management point** to control all enterprise data assets and workflows

Terraform Provider using DCT

TERRAFORM ACTION	TRANSLATES TO	STATE OF INFRASTRUCTURE AFTER OPERATION
APPLY	Creates Infrastructure	Data Provisioned and Running
REFRESH	Update local state to match infrastructure	Update local with remote
DESTROY	Teardown infrastructure	Data destroyed
SHOW	Reads the status of infrastructure	No change
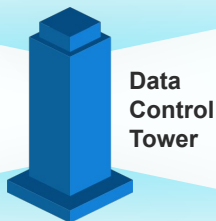
Data Control Tower

CREATE
READ
UPDATE
DELETE

ENVIRONMENT

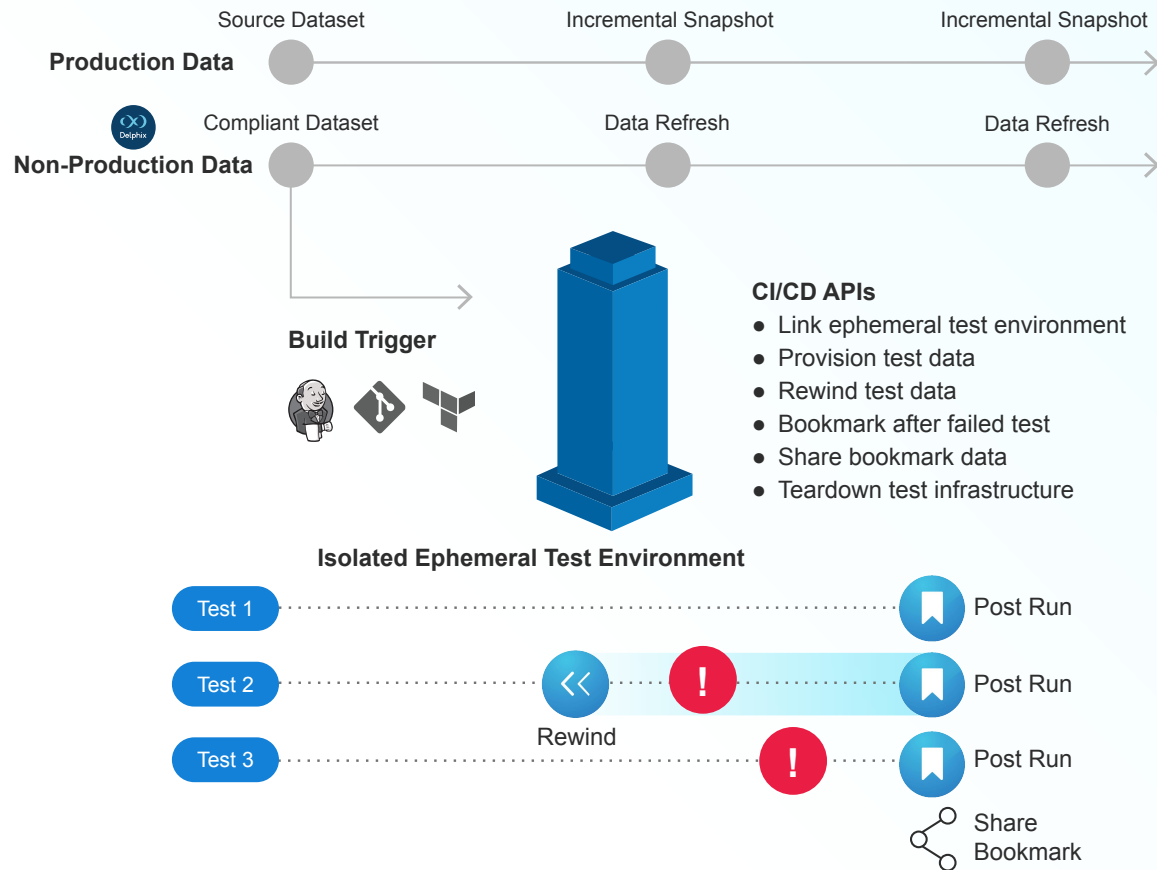**DATA**

Generally Available

- Ephemerally deploy data-as-code alongside infrastructure-as-code

- Instantly deploy data & infrastructure using native Terraform commands to accelerate DevOps workflows

- Ensure speed, quality, and efficiency of code releases

- Collaborate across multi-cloud Dev/Test/QA environments

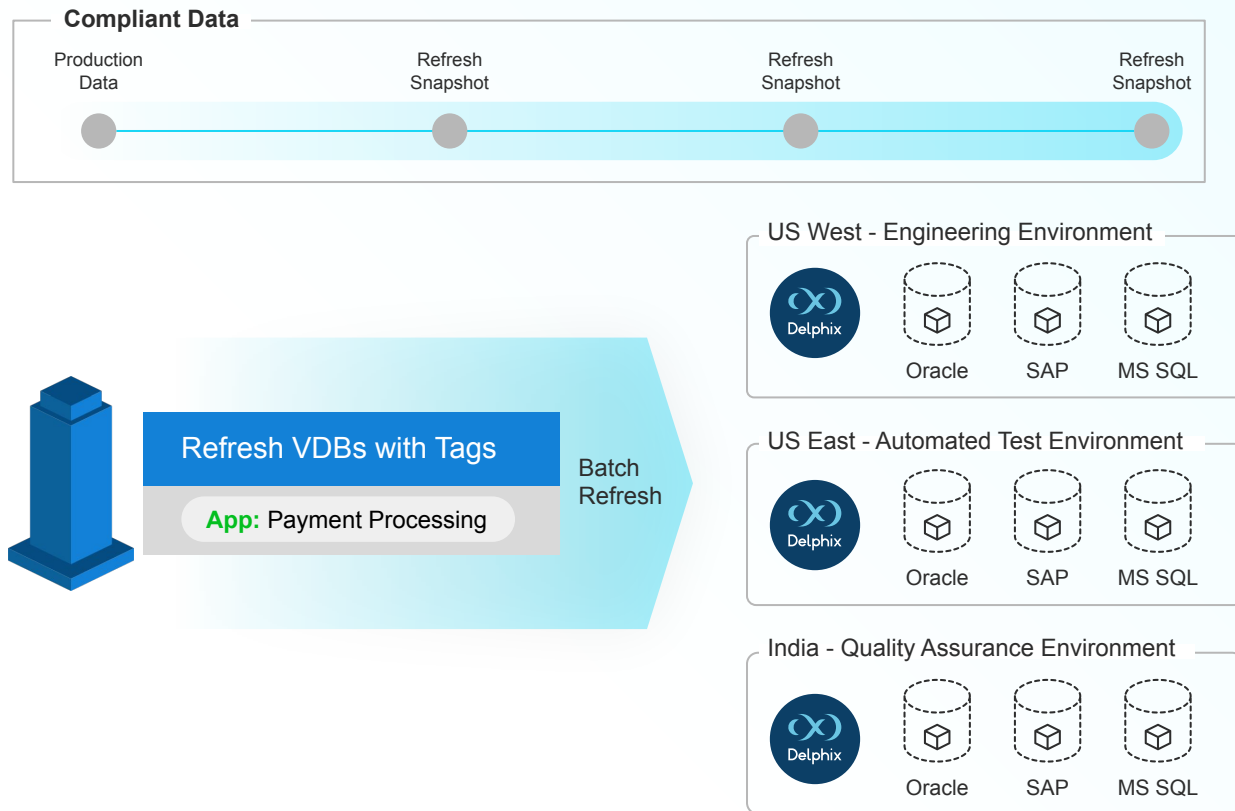- Download from Terraform Registry and link to DCT

DCT: TDM CI/CD Ephemeral Data-as-Code

DevOps TDM Workflow

Production Data

Source Dataset

Incremental Snapshot

Incremental Snapshot

Non-Production Data

Compliant Dataset

Data Refresh

Data Refresh

Build Trigger

CI/CD APIs
- Link ephemeral test environment
- Provision test data
- Rewind test data
- Bookmark after failed test
- Share bookmark data
- Teardown test infrastructure

Isolated Ephemeral Test Environment

Test 1 — Post Run

Test 2 — Post Run

Rewind

Test 3 — Post Run

Share Bookmark

- Quickly deploy data as code in CI/CD Test Data environments using DCT

- Collaborate preserved post-run data between Dev/Test/QA teams

- Quickly rewind data to pre-run state in ephemeral isolated environments

- Synchronize and refresh test data to non-production environments as production data changes

DCT: DevOps VDB Batch Refresh by Tag

Compliant Data
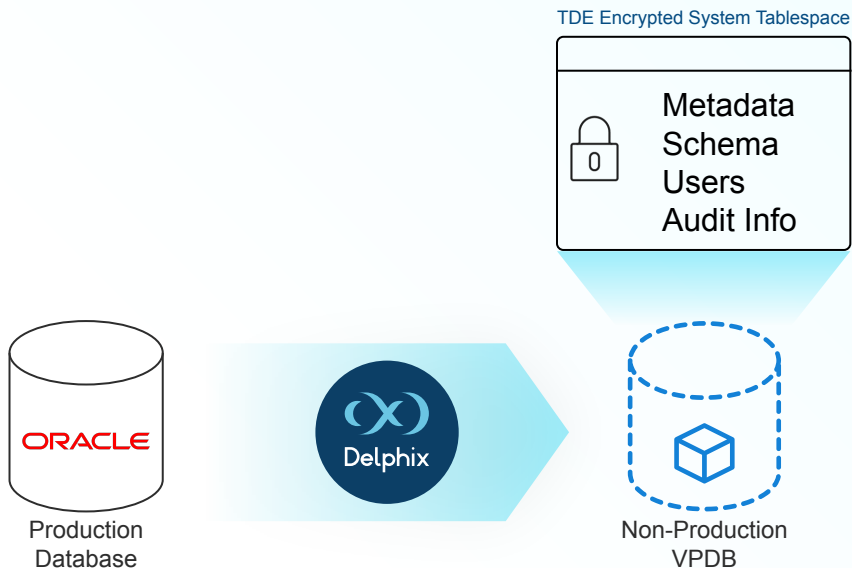
Production Data — Refresh Snapshot — Refresh Snapshot — Refresh Snapshot

Refresh VDBs with Tags

App: Payment Processing

Batch Refresh

US West - Engineering Environment
Delphix | Oracle | SAP | MS SQL

US East - Automated Test Environment
Delphix | Oracle | SAP | MS SQL

India - Quality Assurance Environment
Delphix | Oracle | SAP | MS SQL

DevOps Workflow

- Collect a list of deployed Virtual Databases based on tag selection via HTTPS GET

- Concurrently refresh Virtual Databases serving an app in different environments from the same point in time via HTTPS POST

- Concurrently refresh multiple databases of different database types from the same point in time using APIs

- Overlay Business Intelligence (BI) information using tags and groups for data/infra entities

Continuous Data

New & Updated for July 2022

- Oracle MT: TDM Encrypted System Tablespaces

- Staging Push: Postgres

- Staging Push: IBM Db2 HADR

- Cloud Engines: Cache Analysis UI

- Storage Device Removal
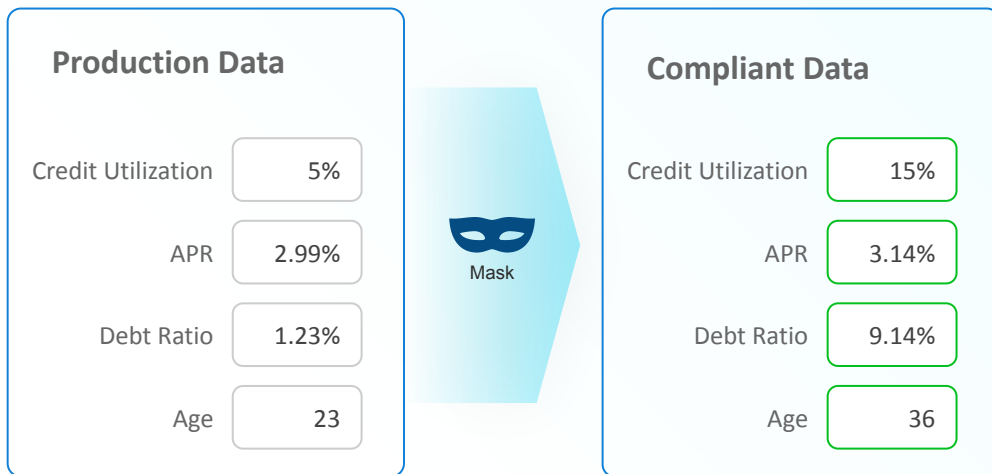
Oracle MT Encrypted System Tablespaces

New for 6.0.15

TDE Encrypted System Tablespace

Metadata
Schema
Users
Audit Info

Production
Database

Non-Production
VPDB

- Provision vPDBs with encrypted system tablespaces

- Prevent database metadata, schema, and user access tampering in non-prod environments

- Run tighter security controls in downstream Dev/Test/QA environments

- Sync & deploy fully encrypted Exadata Cloud at Customer (ExaCC) data in non-prod

Staging Push for Postgres

Production Environment

Downstream Non-Production Environments

Production DB
Postgres

Backup/Replica

Staging DB
Postgres

Staging Push

dSource

Deploy/Refresh

VDB
Postgres
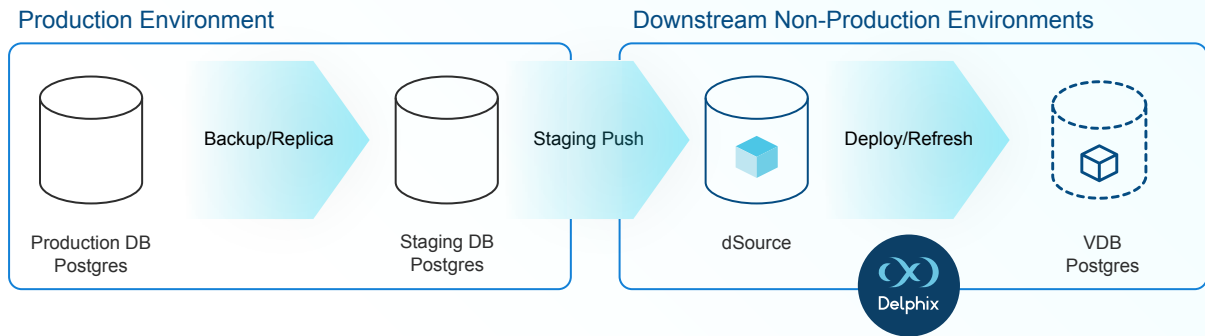
Delphix

- Push new data to the Delphix DevOps Data Platform instead of requesting data on a polling interval

- Isolate ingest to Delphix from production instances

- Ingest Postgres for use as VDBs in Dev/Test/QA
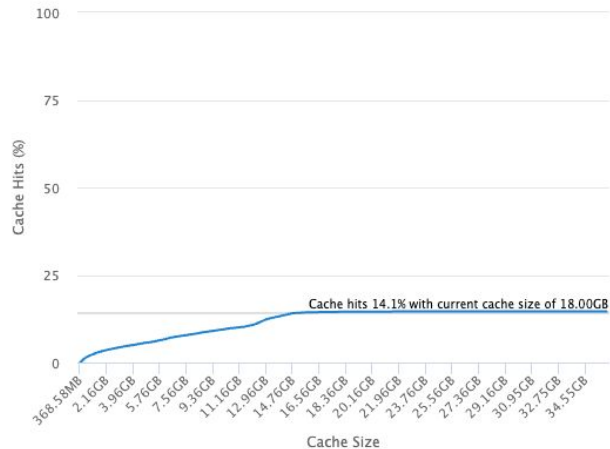
Staging Push for IBM DB2 HADR

New for 6.0.15

Production Environment

Production DB
IBM DB2 HADR

Backup/Replica

Staging DB
IBM DB2 HADR

Staging Push

Downstream Non-Production Environments

dSource

Deploy/Refresh

VDB
IBM DB2 HADR
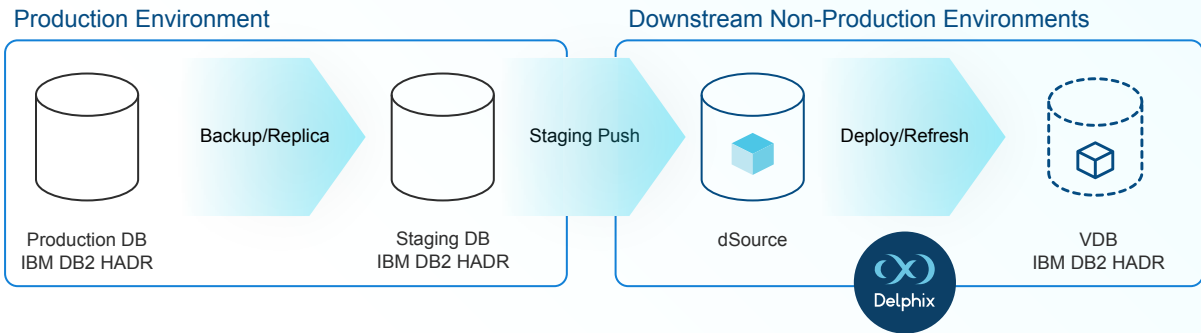
Delphix

- Push new data into the Delphix DevOps Data Platform instead of requesting data during a polling interval

- Isolate ingest to Delphix from production instances

- Confirmed compatibility with IBM DB2 HADR

Cloud Engines Cache Analysis UI

Storage ⓘ   Modify

Object Storage for Data
Enabled   12.00TB

Block Storage for Cache   **24.00GB**
Enabled   24.00GB
Disabled   0.00B

Name ⌃	Size	Enabled
xvdb	8.00GB	Yes
xvdc	8.00GB	Yes
xvdd	8.00GB	Yes

Block Storage Cache Access

Cache Hits (%) vs Cache Size

Cache hits 14.1% with current cache size of 18.00GB

Cache Size: 368.58MB, 2.16GB, 3.96GB, 5.76GB, 7.56GB, 9.36GB, 11.16GB, 12.96GB, 14.76GB, 16.56GB, 18.36GB, 20.16GB, 21.96GB, 23.76GB, 25.56GB, 27.36GB, 29.16GB, 30.95GB, 32.75GB, 34.55GB

Cache data collection start time
Jun 8, 2022 11:38:51 PM

Cache data collection end time
Jun 9, 2022 10:59:11 AM

Cache lookups
3046753

- Quickly right-size Delphix Cloud Engine cache size based on cache-hit percentages

- Manage relative cost-performance of ephemeral engines

- UI provides more in-depth visual inspection of relative performance information at specific cache sizes

Storage Device Removal UI

- Quickly remove multiple storage devices from within the UI

- Easily reduce the total block-based storage attached

- Previously only available through the CLI

Continuous Compliance

New & Updated for July 2022

- Select Connector for SAP

- Hyperscale Oracle Connector

- Segmented Mapping Algorithm V2

- JSON File Masking

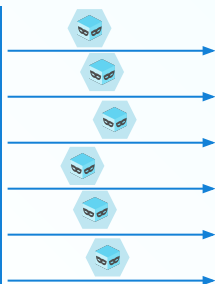- Numeric Expression Algorithm

# Hyperscale Compliance for Oracle

New for July

Continuous Compliance

High-Performance Masking Orchestration

Connector Plug-in

ORACLE
DATABASE

Masking Job

Masking
Orchestrator

- Quickly mask Oracle databases used for applications that require massive datasets for TDM

- Easily bring databases with billions of rows containing PII, PHI, or sensitive data up to compliance

- Automatically distribute masking load across multiple engines attached to an Orchestrator

- Greatly increase data compliance speed for massive databases in non-prod (Dev/Test/QA)

Segmented Mapping Algorithm V2

Production Data

VIN WP0CA29814L001158

Mask

Compliant Data

VIN 3KPA24AD6NE432307
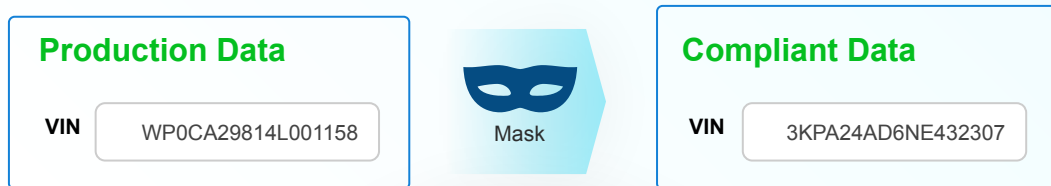
Segment Mapping Rules

1) 3 Alpha-Numeric -> Map to WP0 Replace with 3KP

2) 8 Alpha-Numeric -> Different Range & Values: A-Z, 0-9

3) 6 Numeric -> Map to Value in Range 000001-999999

- Enhanced algorithm performance, extensibility, and portability

- Divide values into smaller segments that are handled separate masking rules

- Preserve or remove specific characters in a value

- Quickly handle non-conforming characters and map specific character replacements

JSON File Masking

JSON Input

```
{
  "data": [{
    "Name": "Layla",
    "State": "NM",
    "City": Albuquerque",
    "Phone": "505-867-5309"
  }],
  "Context":[{
    "Created": "1995-11-22",
    "Updated": "2022-03-17",
  }]
}
```
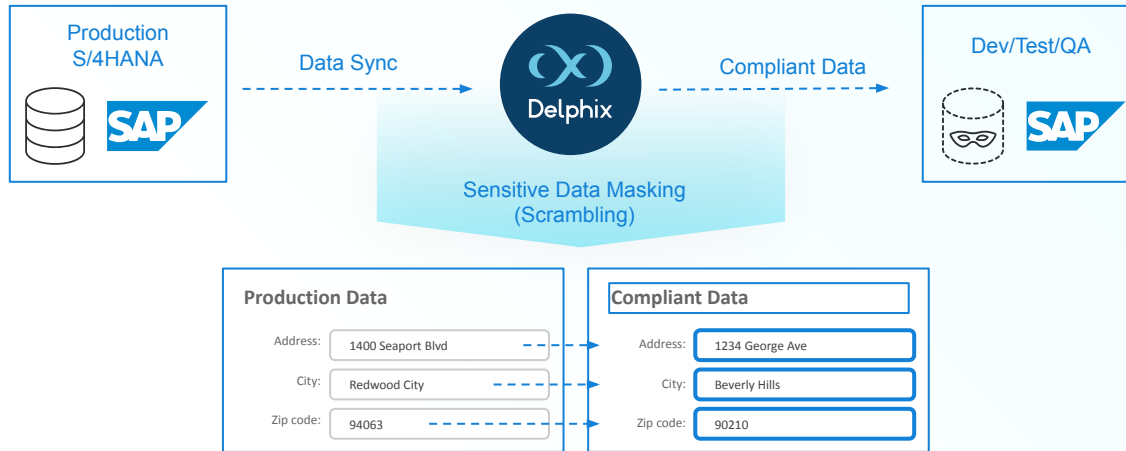
Masking
Job

JSON Output

```
{
  "data": [{
    "Name": "Katherine",
    "State": "NM",
    "City": Albuquerque",
    "Phone": "404-555-2368"
  }],
  "Context":[{
    "Created": "1995-11-22",
    "Updated": "2022-03-17"
  }]
}
```

- Mask JSON data fields using the Delphix algorithm library and customizable frameworks

- Quickly make complex JSON files data compliant for developers and system integrators

Multi-Cloud Applications

New & Updated for July 2022

- Compliance Accelerator for SAP

- Salesforce: Package Management

- Salesforce: Oauth 2.0 Connections

Salesforce Package Management

Salesforce Packages
Validation Rules
Workflows
Process Builders
Triggers
Field History Tracking
Feed Tracking
Duplicate & Matching Rules
Required & Unique Constraints
Picklicks
Filters

Disable
Packages

Masking
Operations

Enable
Packages

salesforce

Compliant Data
Sandbox

- Fast data compliance for Salesforce Sandboxes

- Automate tedious manual processes for managing packages that interfere with masking

- Added packages: Duplicate & Matching Rules, Required & Unique Constraints, Picklists, and Filters

- Quickly deliver compliant data sandboxes with production-realistic data to developers
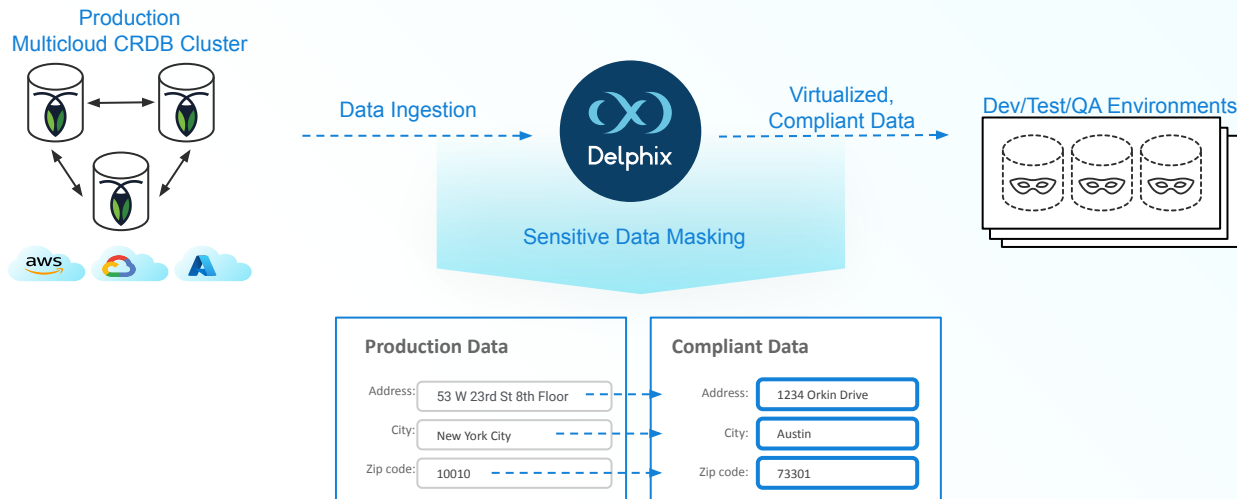
OAUTH 2.0 Connections for Salesforce

- Improve connection security using OAUTH 2.0 between Salesforce and Delphix Engines

- Provides a secure alternative to username/password and security tokens

Select Connector for CockroachDB

Bring the power of Delphix Continuous Data and
Continuous Compliance to modern, cloud-native applications

Production
Multicloud CRDB Cluster

Data Ingestion

Virtualized,
Compliant Data

Dev/Test/QA Environments

Sensitive Data Masking

Production Data		Compliant Data	
Address:	53 W 23rd St 8th Floor	Address:	1234 Orkin Drive
City:	New York City	City:	Austin
Zip code:	10010	Zip code:	73301

CockroachDB

Recent Updates

6.0.13 Released March, 2022	**6.0.14** Released May 20, 2022	**6.0.15** Released July 7, 2022

Continuous Data (Virtualization)
- Staging Push for Microsoft SQL Server
- SAP (Sybase) IQ for RHEL 7.0/7.4/7.9
- Oracle E-Business Suite 12.2.11 in OCI
- FluentD Event & Metric Output
- MS SQL Server on Microsoft Windows Server 2022 (Standalone & Cluster)
- Select Connector for SAPIQ

Continuous Compliance (Masking)
- Tokenization Algorithm V2
- Zip+4 Algorithm
- Enhanced Profiling Opt-In (Early Adopter)
- Masking monitoring improvements
- Masking IBM DB2 12 for z/OS
- Masking PostgreSQL 13/14

Multi-Cloud Applications
- Salesforce Data Protection for production

Continuous Data (Virtualization)
- Delphix Cloud Engines on AWS using S3
- Single Engine Continuous Vault
- SQL Server V2P improvements
- Postgres single-db ingest from multi-db
- Certification: Oracle & Plugins RHEL 8.5

Continuous Compliance (Masking)
- Data Cleansing Algorithm V2
- "Optional" Multi-Column Algorithm
- MinMax Algorithm V2
- Core Segment Mapping UI
- Sensitive Data Discovery: +40 profile sets
- Certification for VMware ESXi 7.0 U3c, and SAP HANA 2.0 SP 05

Data Control Tower (DCT)
- 2.0 Beta:Tagging, Reporting, VDB Groups

Cloud Applications
- Salesforce Data Protection - Performance

Continuous Data (Virtualization)
- Staging Push: Postgres & IBM DB2 HADR
- Cloud Engines Cache Analyzer
- Oracle MT Encrypted System Tablespaces
- Certifications for Postgres 13 on SUSE 15, Postgres 14 on RHEL 8.5, and Oracle EBS with TDE for Oracle 19C.

Continuous Compliance (Masking)
- Hyperscale 3.0 for Oracle
- Segment Mapping Algorithm V2
- JSON File Masking
- Numeric Expression Algorithm

Data Control Tower (DCT)
- Expanded Tagging, Reporting, VDB Groups

Multi-Cloud Applications
- Compliance Accelerator for SAP
- Salesforce Package Management
- Salesforce OAUTH 2.0 Connections
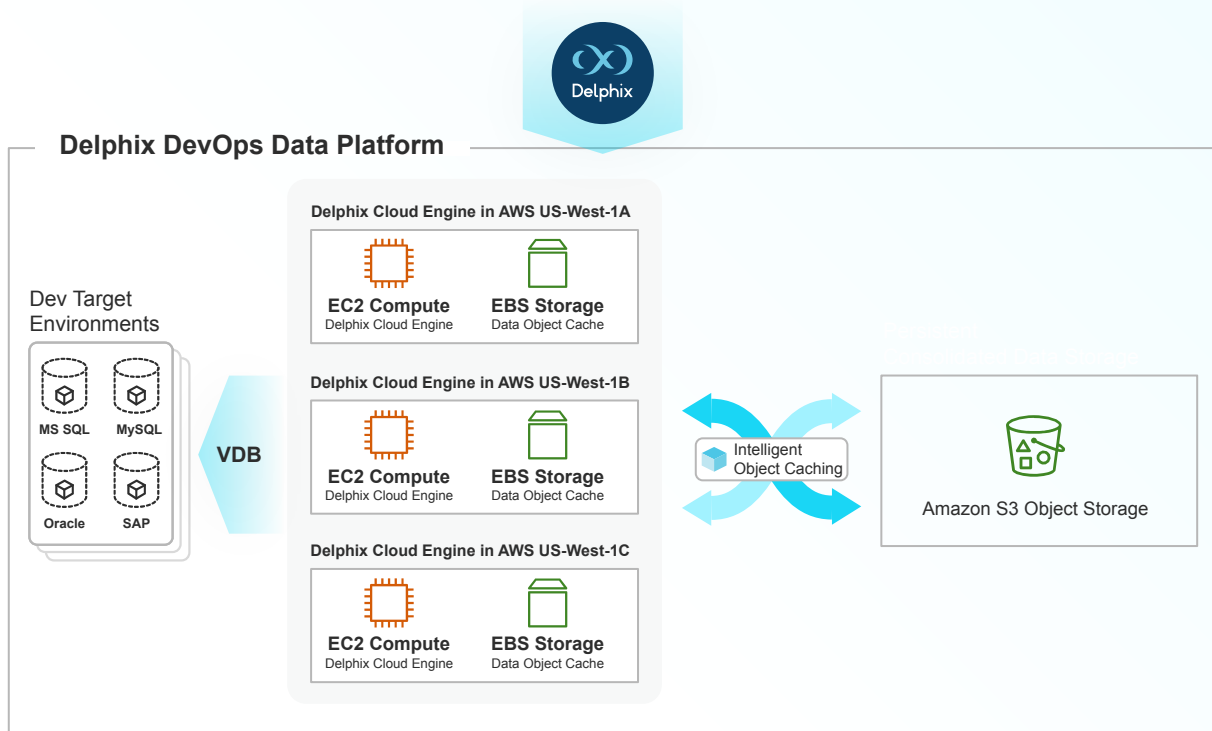- Select Connector for CockroachDB

Appendix

Previously Released Items

- Delphix Cloud Engines (DCE)