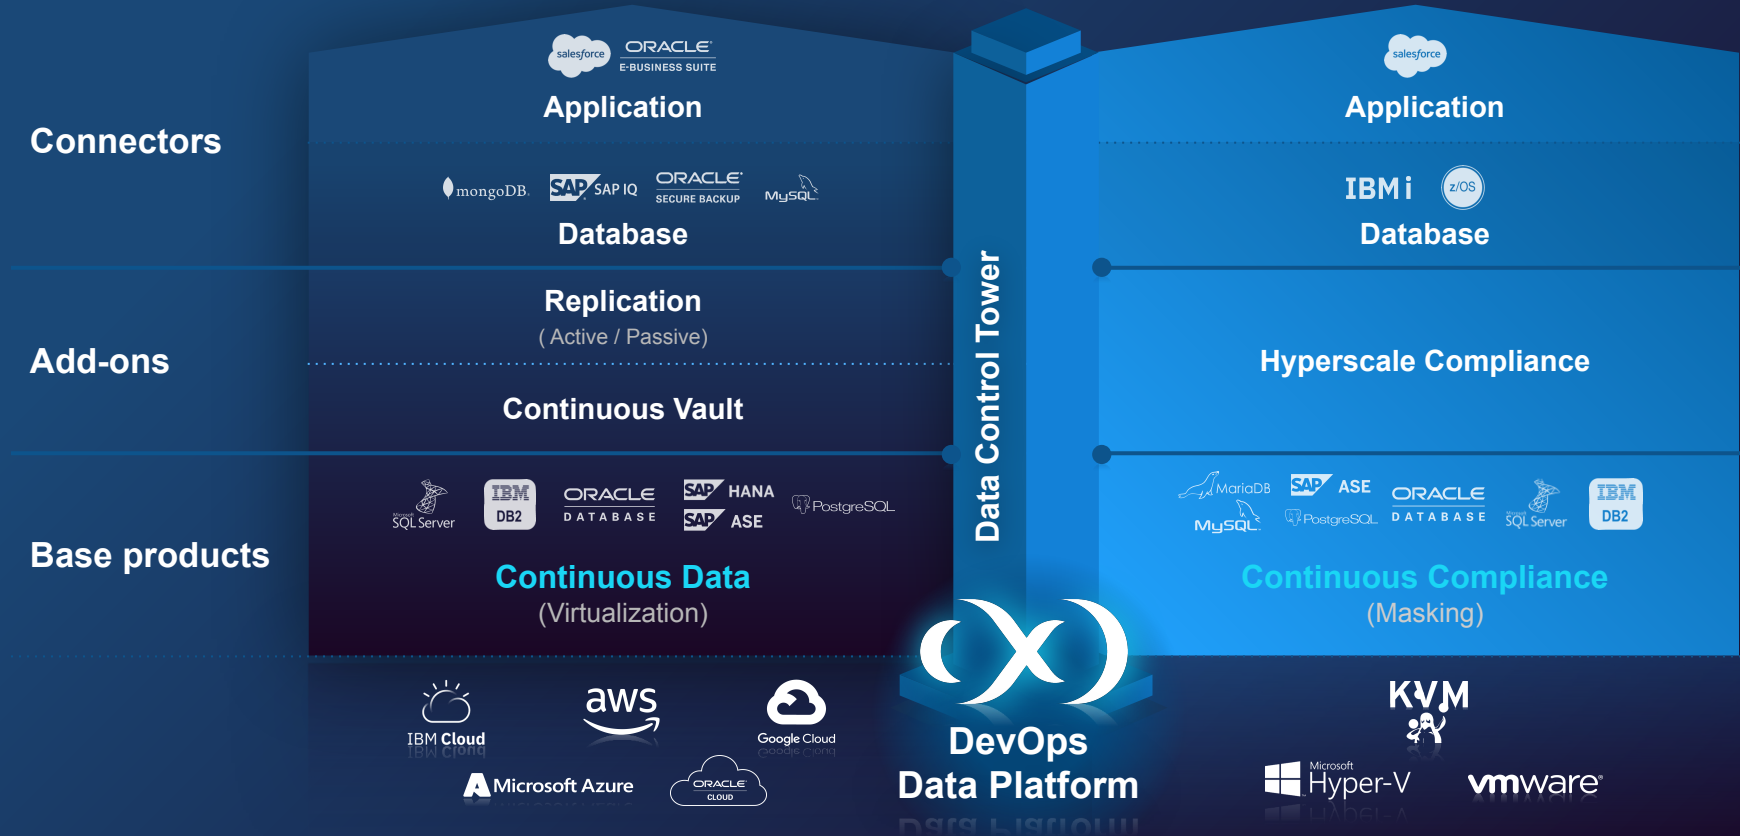


What's New

New Features in May 2022 & 6.0.14



Product Name Changes



Continuous Data

New & Updated in May 2022 & 6.0.14

Delphix Cloud Engines (DCE) on Amazon using S3 Object Storage

Single Engine Continuous Vault

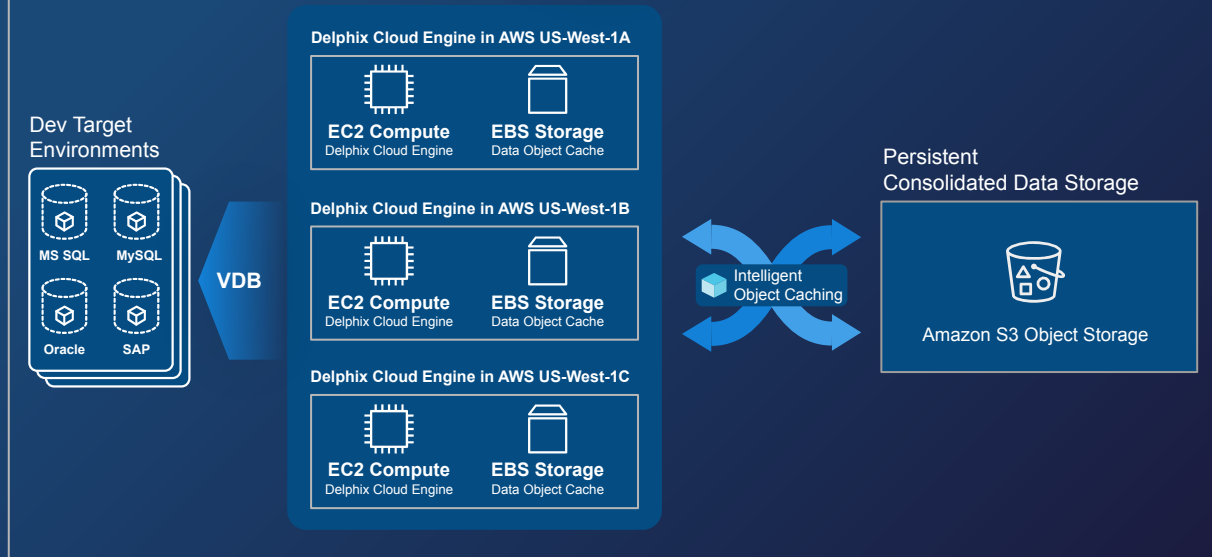
V2P for MS SQL Improvements

Postgres Logical Backup Ingest from PaaS (RDS)

Delphix Cloud Engines on Amazon S3 Object Storage



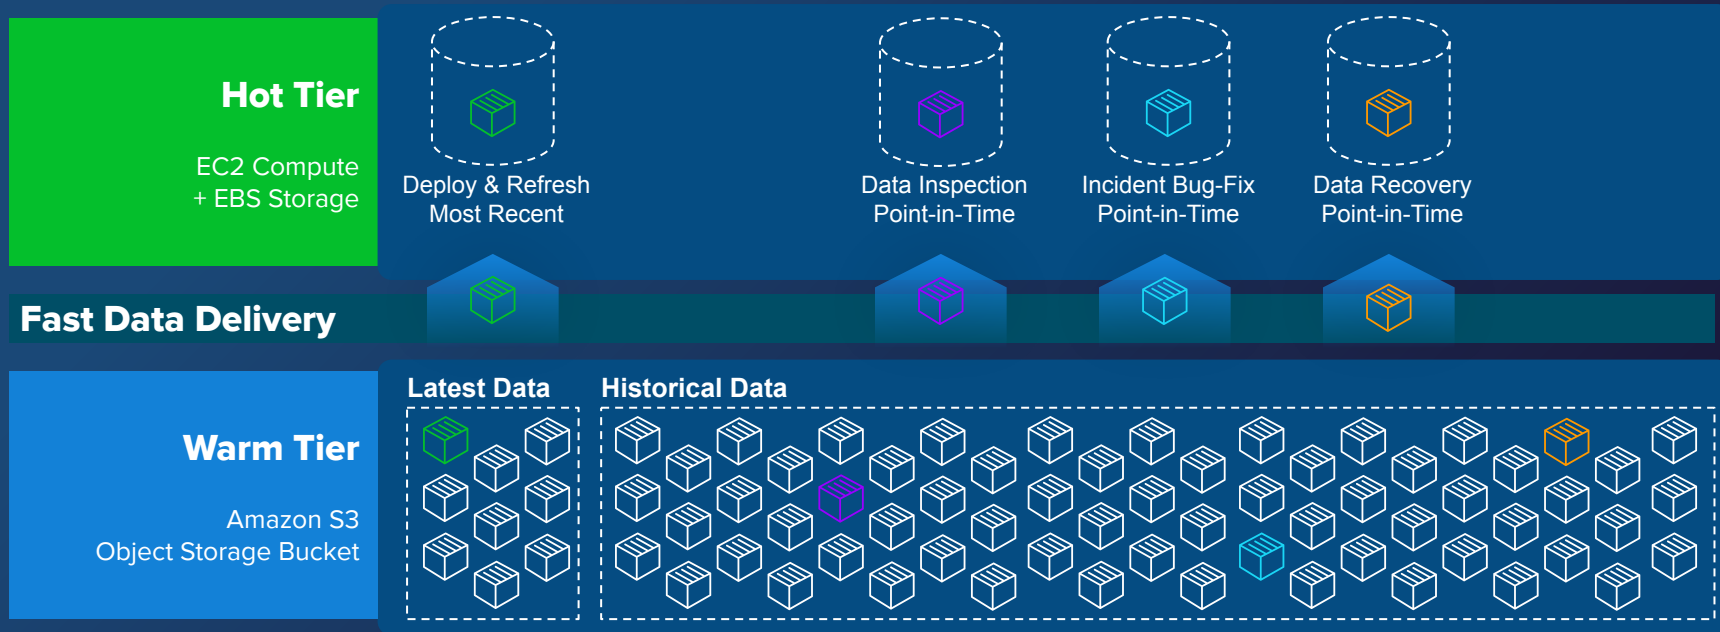
Delphix DevOps Data Platform



New in 6.0.14

- Use Amazon S3 object storage in place of block storage to reduce data estate costs
- Cache only necessary data to engines to maintain consistent VDB performance
- Elastically expand storage footprint as-you-go in the cloud
- Quickly deploy data to engines across availability zones
- Potential cost savings of over 50%

DCE: Exceptional Price-Performance for Ephemeral Non-Prod Data



DCE leverages cloud native architectures, services, and storage tier switching to provide exceptional price-performance for non-production ephemeral data.

DCE Cost Savings Estimates

Production engine example

Modeled on customer engine:

- Total allocated storage to engine: 146TB
- Storage used: 92TB
- Unused storage: 54TB

Additional details:

- Volume type: GP2
- Volume size: 5TB (approx. 30 volumes)
- IOPS: 15360 per volume
- Object count: 2 dSources and 25 VDBs

\$210,800/
year

EBS over
allocated
storage

54TB

EBS storage
used

92TB

R5n.4XLarge

\$59,197/
year
(71%
Savings)

5% EBS Cache

S3 storage
used

92TB

R5n.4XLarge

\$69,430/
year (69%
Savings)

15% EBS Cache

S3 storage
used

92TB

R5n.4XLarge

\$105,200/
year
(50%
Savings)

50% EBS Cache

S3 storage
used

92TB

R5n.4XLarge

\$156,400/year
(25% Savings)

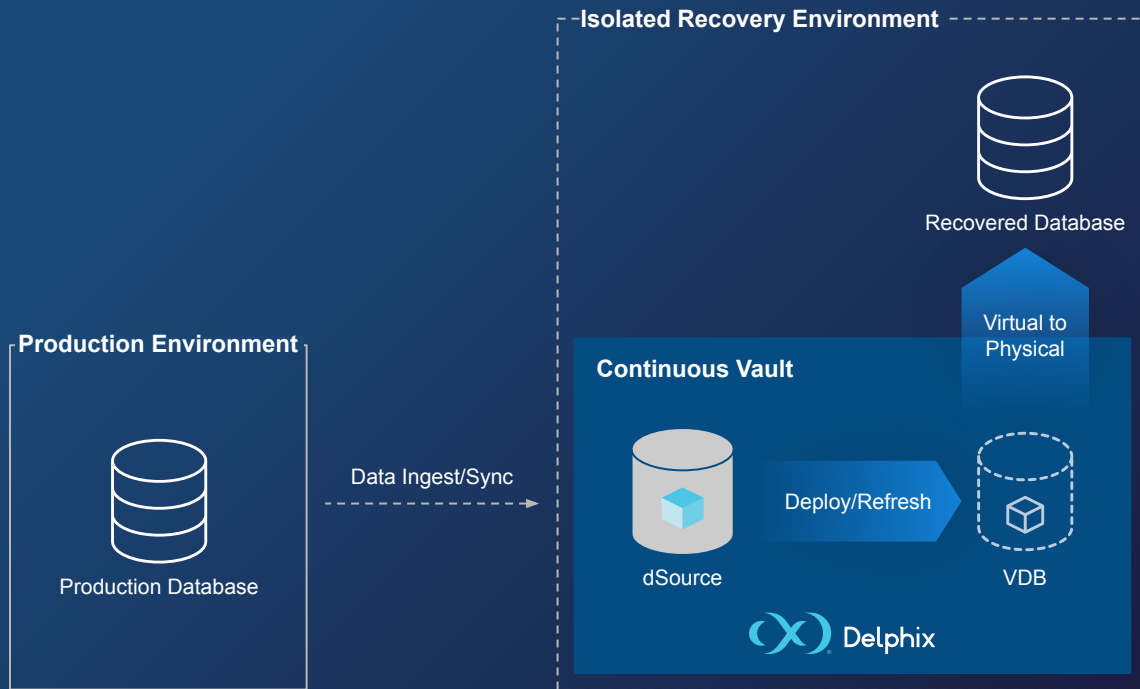
100% EBS Cache

S3 storage
used

92TB

R5n.4XLarge

Single Engine Continuous Vault



New in 6.0.14

- Quickly restore data from any point-in-time to isolated recovery environments
- Protect data and data changes with locked policy-based retention controls
- Instantly deploy ephemeral VDBs from any point in time
- Directly ingest production data to a Continuous Vault
- Provides a streamlined data protection architecture for non DevOps TDM configurations

V2P for MS SQL Server Improvements



Production DB Instance



VDB

Redeployed Infrastructure & Data as Code



Multi-Drive Storage



Production DB Instance


Update in 6.0.14

- Split MS SQL Server archive logs to multiple drives for Virtual to Physical (V2P)
- Enhanced V2P performance
- Improved data deployment to new full size test instances for performance benchmarking
- Improved lift-and-shift cloud migration performance
- Faster recovery from ransomware attacks or business continuity events

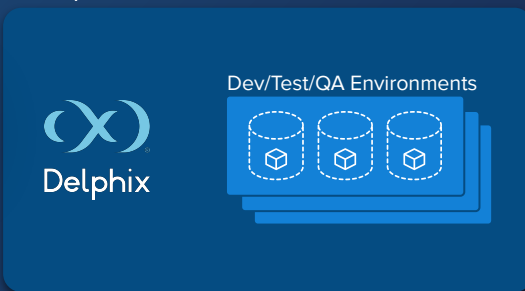
PostgreSQL Logical Backup Ingest & Restore

New in 6.0.14

Selective Ingest from Multi-DB Instance



DevOps Data Platform



PaaS Source




Data
Ingest



Data
Deploy

Data Recovery to IaaS



- Quickly deploy data downstream in non-prod environments from PaaS
- Ingest and restore Postgres logical backups using native `pg_dump/pg_restore`
- Deploy data from PaaS instances into clean IaaS instances for inspection
- Selectively ingest single databases from multi-db Postgres instances
- Enable more cost-efficient ephemeral deployment of Postgres data for non-prod

Data Control Tower

New & Updated in May 2022

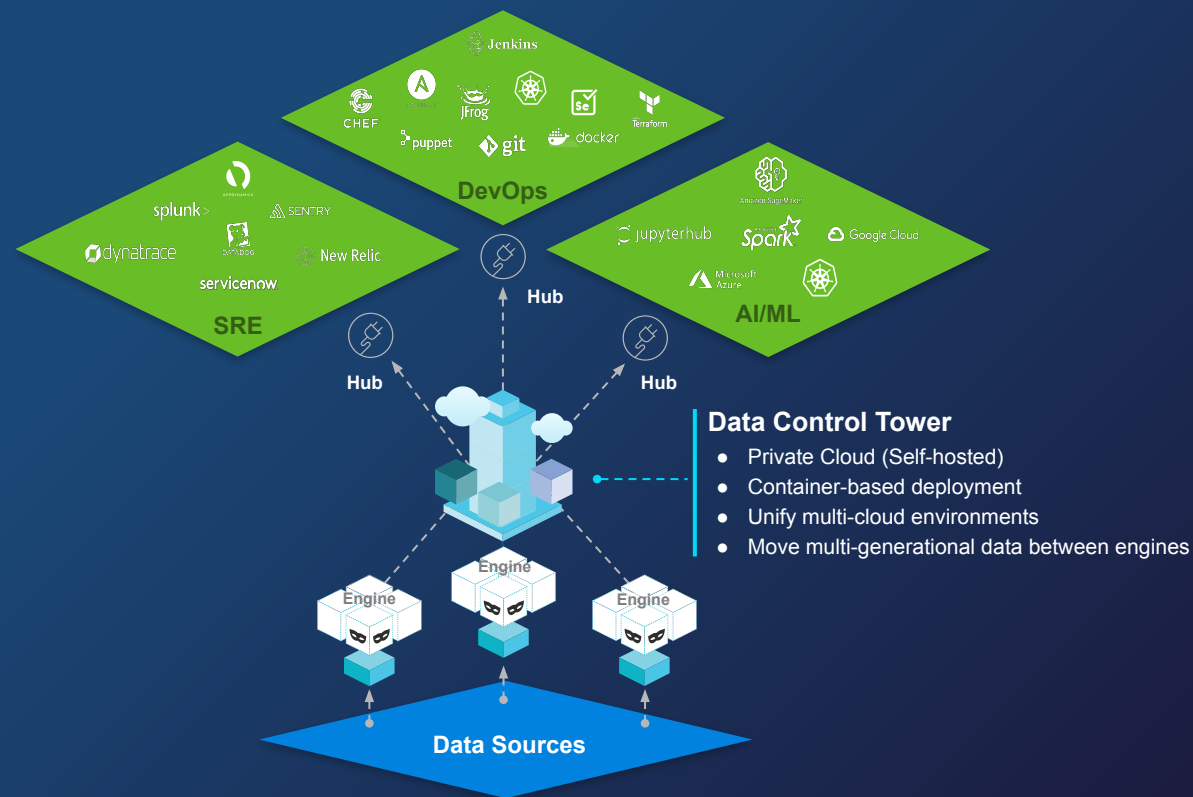
DCT 2.0 Beta with 19 all new APIs

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 Beta



Early Adopter Beta

- **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

— New in 2.0 Beta
— Updated in 2.0 Beta
— Existing in 1.0.6

Environment

PUT Enable environment Disable environment Refresh environment	DELETE Delete environment by ID POST Create environment	GET List all environments Show environment by ID
--	--	---

VDB Groups

GET List bookmarks by group List all VDB groups List VDB group by ID	POST Create VDB group Refresh VDB groups Rollback VDB group	DELETE Delete VDB group
--	---	-----------------------------------

Bookmarks

GET List VDB groups List bookmarks Show bookmark by ID	POST Rewind to bookmark Create bookmark	DELETE Delete bookmark
--	--	----------------------------------

Snapshots

GET List snapshots Show snapshot by ID	DELETE Delete snapshot by ID
---	--

Reporting

GET Storage capacity VDB inventory dSource inventory
--

Virtual Databases (VDB)

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

Tags

POST Create tag on VDB	DELETE Remove tag from VDB
----------------------------------	--------------------------------------

Sources

GET List all sources Show source by ID

Jobs

GET Job by ID	POST Abandon Job by ID
-------------------------	----------------------------------

Management

GET List engines Show engine by ID List Hashicorp vaults Show Hashi vault by ID	POST Register engine Configure Hashi vault Put Update engine	DELETE Unregister engine Remove Hashi vault by ID
--	--	--

API Clients

GET List API clients Show client by ID	POST Create new API client PUT Update client by ID	DELETE Remove client by ID
---	---	--------------------------------------

dSources

GET List all sources Show source by ID

Early Adopter Beta

- 19 New CI/CD APIs available now in DCT 2.0 Beta
- 61 APIs total in DCT to manage environments, engines, dSources, & VDBs
- Containerized API gateway for multi-cloud environments
- Aggregated management point to control all enterprise data assets and workflows
- Fast integration with CI/CD pipelines and Test Data Management (TDM) workflows

Terraform Provider using DCT Beta



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



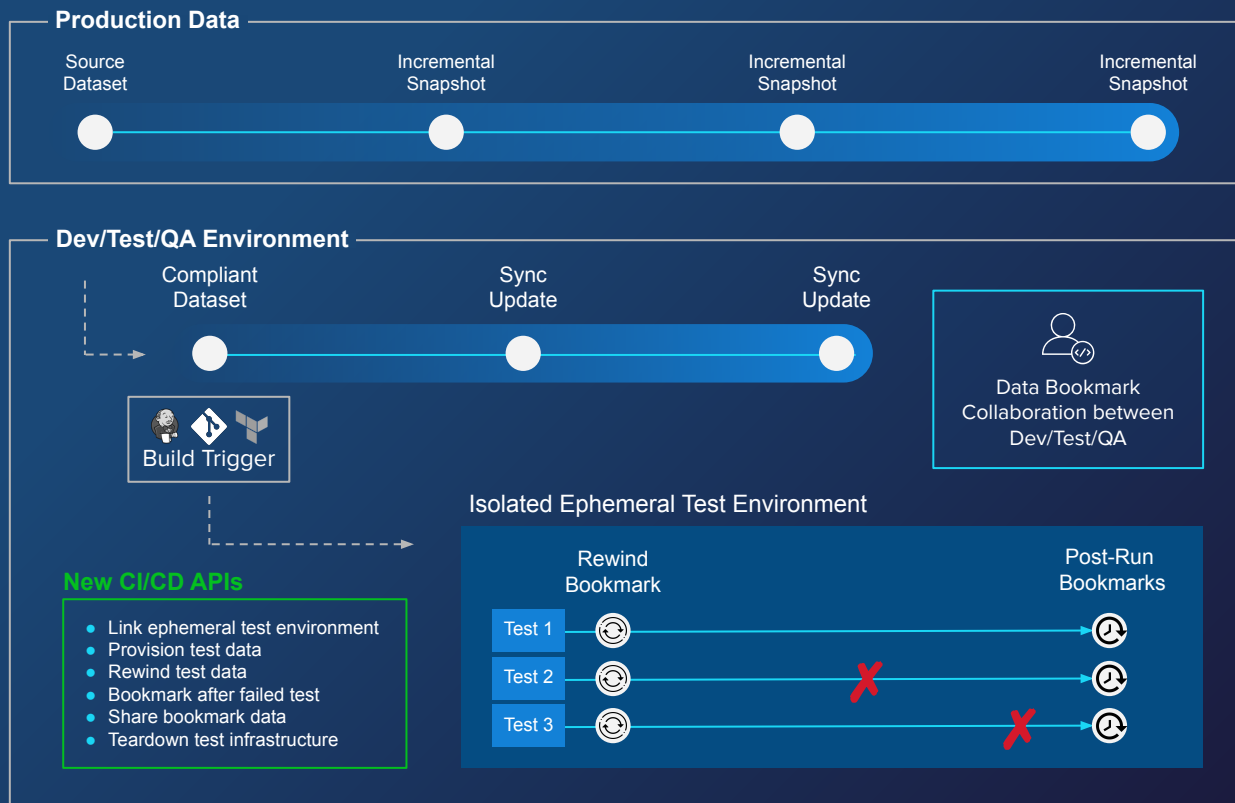
ENVIRONMENT	DATA	SOURCE	BOOKMARK	SNAPSHOT
CREATE READ UPDATE DELETE	CREATE READ UPDATE DELETE	CREATE READ UPDATE DELETE	CREATE READ UPDATE DELETE	CREATE READ UPDATE DELETE

Early Adopter Beta

- Ephemeral deploy data-as-code alongside infrastructure-as-code
- Instantly deploy data & infra using 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
- Contact your CSM to opt-in to this early-adopter feature

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

Early Adopter Beta



- 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



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

Early Adopter Beta

- 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 Compliance

New & Updated in May 2022 & 6.0.14

Segment Mapping UI Updates

Automated Sensitive Data Discovery Updates

Optional Column in Multi-Column Masking Operations

MinMax V2 Algorithm

Data Cleansing Algorithm

Masking UI Improvements

Segment Mapping UI

New in 6.0.14

Select a segment to edit its masking characteristics.

1

2

3

+

<

>

Length

Segment Treatment

4

Mask

[Delete Segment](#)

?

Character Type

Alpha-Numeric

Alpha Ranges and Values

Numeric Ranges and Values

If original values are

A - E, X, W - Z

or numbers between

001 - 999

Map to

Different Ranges and Values

Alpha Ranges and Values

Numeric Ranges and Values

Then replace with

A - E, X, W - Z

or numbers between

001 - 999

Note: If replacement range spans are different lengths than the originals, they will not be reversable.

Ignore

Automatically Ignore Special Characters

Nonconforming Data Handling

Mark as succeeded

Mark as succeeded

Nonconforming Data Handling

- Vastly improved user interface for configuring segment mapping
- Quickly configure alphanumeric character mapping

Automated Sensitive Data Discovery

Input Table

Email_1 [Numeric]	Email_2 [String(8)]	Email_3 [DATETIME]	Email_4 [STRING(255)]
87127892376	kMFOhH3R	2007-03-31	CleonDayBelisarius@Empire.com
58712897237	qqT5y5xQ	2022-03-04	HariSeldon@Foundation.com
83618405831	1d%teGwc	2021-04-05	GaalDornick@Synnax.com
44008675309	IL4eO7t%	2020-03-05	SalvorHardin@Terminus.com

Email Matching Criteria

Type: STRING
Minimum Length: 20

Matching Sensitive Column

Email_4
[STRING(255)]

CleonDayBelisarius@Empire.com

HariSeldon@Foundation.com

GaalDornick@Synnax.com

SalvorHardin@Terminus.com

New in 6.0.14

- Improved sensitive data discovery-matching accuracy and speed
- **40 new profile expressions** utilizing type constraints
- Reduces false positives associated with data type mismatches during inspection

“Optional” Column in Multi-Column Masking

Multi-Column Input

Address #	Street Name	City	State	Zip Code	Zip+4 (optional)
650	Gough St	San Francisco	CA	94102	1234
30	Water St	New York	NY	10004	-
3708	Las Vegas Blvd	Las Vegas	NV	89109	-
955	Fulton Market	Chicago	IL	60607	-

- Multi-Column masking now allows for some fields to be optional
- Optional field(s) enable real-time adaptation to available columns

Optional Column(s)

Column Name
Zip+4

Mask with “optional” column

Masked Full Address
2779 Mission St, San Francisco, CA 94110-4321
113 St Marks Pl, New York, NY 10009
953 E Sahara Ave Ste A5, Las Vegas, NV 89104
809 W Randolph St, Chicago, IL 60607

MinMax V2

Input Data

Monthly Revenue
\$1
\$20,000
\$10,000
\$15,000
\$100,000

MinMax

Normalized Output

Monthly Revenue
\$17,890
\$20,000
\$16,321
\$15,000
\$18,675

Algorithm Configuration

Range: \$15,000-\$20,000

Update in 6.0.14

- Normalize data that contains outliers to within a specified range
- Removes sensitive numbers or dates that reveal outlier significance
- Improved algorithm performance

Data Cleansing Algorithm V2

Multi-Column Input

City
Hanalolo
Hon
Hon-HI
Honolulu_____

Lookup Table

Dirty-Value	Clean-Value
Hanalolo	Honolulu
Hon	Honolulu
Hon-HI	Honolulu
Honolulu_____	Honolulu

Data Output

Cleansed Values
Honolulu
Honolulu
Honolulu
Honolulu

- Standardize terms across spellings, misspellings, and abbreviations
- Remove trailing whitespace behind values
- Map dirty to clean values using lookup table pairs

Masking UI Improvements

Update in 6.0.14

The screenshot displays the 'DELPHIX CONTINUOUS COMPLIANCE' interface. At the top, there are tabs for 'Environments', 'Monitor' (selected), 'Settings', 'Admin', and 'Audit'. Below the tabs, a 'Home > Monitor' breadcrumb is shown. A summary section indicates '1/7 Jobs Running' and '1024/6045 Memory Usage (MB)'. A search bar is present with filters for 'Start Date', 'End Date', 'Any Status', and 'All Jobs'. The main table lists jobs with columns for Environment, Job Name, Submit Time, Start Time, Type, Progress, and Status. A detailed view for the 'demo_M_IP' job is shown on the right, including a 'Job Type' dropdown, 'Mask' status, and a list of steps like 'Init Execution', 'Collecting Job Configurations', 'Preparing Execution', 'Start Execution', 'Pre SQL Script', 'Post SQL Script', 'Execute Post Execution Custom Driver Task', 'Collecting Job Information', and 'Execution Finished'. The 'Status' column shows 'RUNNING' for the selected job.

Environment	Job Name	Submit Time	Start Time	Type	Progress	Status
src_env	demo_M_IP	2022-04-06 06:19	2022-04-06 06:19	MASK	2 of 10 - Collecting...	RUNNING
src_env	My_big_name_job...	2022-04-05 14:08	2022-04-05 14:08	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	My_big_name_job...	2022-04-05 09:14	2022-04-05 09:14	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	My_big_name_job...	2022-04-05 09:12	2022-04-05 09:12	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	M_delim_IP	2022-04-05 08:52	2022-04-05 08:52	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	M_delim_IP	2022-04-05 08:51	2022-04-05 08:51	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	M_vsam_IP	2022-04-05 08:50	2022-04-05 08:50	MASK	10 of 10 - Execution...	FAILED
src_env	demo_M_IP	2022-04-05 07:24	2022-04-05 07:24	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	demo_M_IP	2022-04-05 07:11	2022-04-05 07:11	MASK	10 of 10 - Execution...	SUCCEEDED
src_env	M_vsam_IP	2022-04-05 07:02	2022-04-05 07:02	MASK	10 of 10 - Execution...	FAILED
src_env	M_vsam_IP	2022-04-05 07:01	2022-04-05 07:01	MASK	10 of 10 - Execution...	FAILED
src_env	demo_M_IP	2022-04-05 07:00	2022-04-05 07:00	MASK	10 of 10 - Execution...	SUCCEEDED

- Monitoring UI now shows job step is underway
- Redesigned execution monitor screen
- Improve masking job management for intermediate steps

The screenshot displays the 'Waiting' tab in the 'Execution Monitor' section. It shows a list of jobs with columns for ID, Name, Progress, and Rows Remaining. A detailed view for the 'test1' job is shown on the right, including a 'Job Type' dropdown, 'Mask' status, and a list of steps like 'Init Execution', 'Collecting Job Configurations', 'Preparing Execution', 'Start Execution', 'Pre SQL Script', 'Post SQL Script', 'Execute Post Execution Custom Driver Task', 'Collecting Job Information', and 'Execution Finished'. The 'Status' column shows 'WAITING' for the selected job.

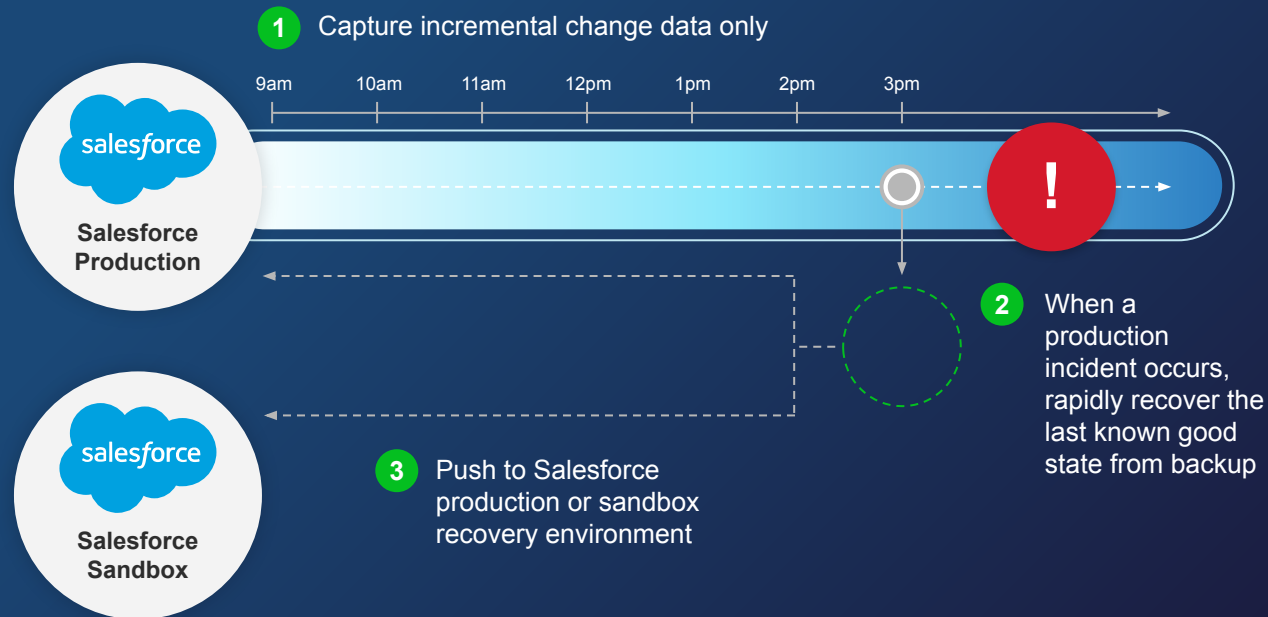
ID	Name	Progress	Rows Remaining
21031	test1	0%	3
21039	test10	0%	3
21032	test11	0%	3

Multi-Cloud & Apps

New & Updated in May 2022 & 6.0.14

Update: Salesforce Data Protection performance improvements

Enhancements to Salesforce Data Protection



Update in May 2022

- Multi-threaded uploads and 10x increase in parallelization expedite recovery
- Efficient API calls accelerate backup processes and improve resource utilization
- Automated processes eliminate manual effort if a record fails to restore

Recent Updates

6.0.12

Released Jan, 2021

Continuous Data (Virtualization)

- Multiple Storage Removal
- OAuth 2.0 Web UI & token logins
- SAP HANA Staging Push
- Not susceptible to Log4j vulnerability
- Log4j libraries upgraded to 2.17.1

Continuous Compliance (Masking)

- Microsoft SQL Masking
- Secure Lookup V2
- Free Text Redaction V2

Developer

- Delphix SDK 4.0 with Python 3.8

6.0.13

Released March 10, 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
- MS SQL Server on Microsoft Windows Server 2022 (Standalone & Cluster)

Continuous Compliance (Masking)

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

Integration & Monitoring

- FluentD Event & Metric Output
- Masking monitoring improvements

Cloud Applications

- Salesforce Data Protection for production

6.0.14

Released May 20, 2022

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
- Oracle & Plugins for RHEL 8.5 Certification

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 CockroachDB, VMware ESXi 7.0 U3c, and SAP HANA 2.0 SP 05

Data Control Tower (DCT)

- Tagging, reporting, and Batch Refresh APIs

Cloud Applications

- Salesforce Data Protection - Performance