

Delphix 6.0.8 and More

May 2021

Product Marketing



Delphix Data Platform 6.0.8: **What's New**

- **High Level Update**

- Ransomware Protection
- Full Name Algorithm

- **Data Control Tower Updates**

- Engine Inventory API

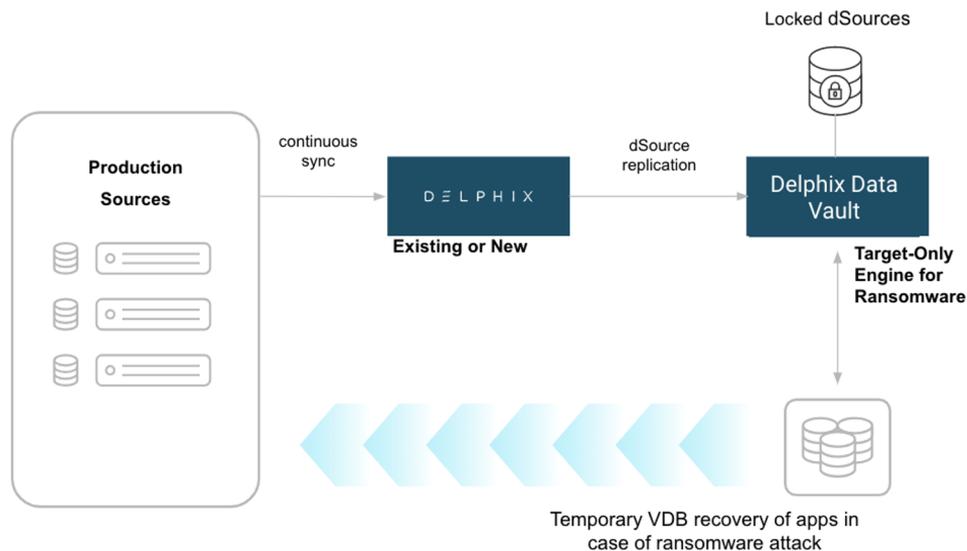
- **Technical Updates**

- Expansion of Retention Period on Replicated Objects
- dxi Encrypted Credentials and Executables
- Chaining Algorithm
- Db2 Staging Push Enhancements
- HANA Plugin Staged Architecture
- Escape Character Support for Delimited File Masking
- System Tunable Interface
- Oracle Customized Full Backups

- **End of Life**

Ransomware Protection

- We added the ability to replicate dSources stored on Delphix engines to a new target engine called Data Vault.
- Once securely stored on the Data Vault, the replicated data can be used to recover business applications after a ransomware attack with very low RTO and RPO.



Full Name Algorithm

- We introduced a new full name algorithm that masks full name strings so that they closely match first names and last names that are stored separately.

Full Name Algorithm Instance configuration examples (via API and via UI)

```
{  
  "algorithmName": "demo-FullName",  
  "algorithmType": "COMPONENT",  
  "description": "This is a new style FullName algorithm",  
  "frameworkId": 3,  
  "algorithmExtension":  
  {  
    "lastNameSeparators": [ ",", " " ],  
    "maxNumberFirstNames": 2,  
    "firstNameAlgorithmRef": { "name": "dlpx-core:FirstName" },  
    "lastNameAlgorithmRef": { "name": "dlpx-core:LastName" }  
  }  
}
```

Select Framework

Secure Lookup

Character Mapping

Payment Card

Date

Name

Full Name

Dependent Date Shift

Segment Mapping (legacy)

Mapping

Binary Lookup

Tokenization

Min Max

Data Clearing

Free Text Redaction

Full Name Algorithm [Learn More](#)

Algorithm Name
dlpx-core:FullName

Description
This FullName algorithm masks input full name, delegating the masking of the first name(s) and last name to corresponding chained Name algorithms.

First Name Algorithm
dlpx-core:FirstName

Last Name Algorithm
dlpx-core:LastName

Maximum First Names
2

Maximum Masked Full Name Length
0

Full Name Convention
First-Middle-Last

Consider Single Word Input as Last Name

Last Name Separators Separated by comma()
 Include comma

Cancel

Engine Inventory API

```
1 openapi: 3.0.0
2 info:
3     title: inventory
4     version: '2020-06-09'
5     description: Inventory of Managed
6 Assets
7 servers:
8     - url: 'http://localhost:3000'
9 paths:
10     /engines:
11     get:
12         summary: Engines
13         tags:
14         - engine
15         response:
16         '200'
```

Engine

GET /engines Engines

GET /engines/virtualization Virtualization Engines

GET /engines/masking Masking Engines

GET /engines/{id} Engines Details

Virtualization

GET /engines/virtualization Virtualization Engines

Masking

GET /engines/masking Masking Engines

Expansion of Retention on Replicated Objects

- In the past, when snapshots on the replication source engine are deleted, either due to retention policies or user action, the next replication job will delete those snapshots on the replication target engine.
- Now you can extend the retention period of replicated objects on target engines while keeping the original retention at the source.
- Once the object in the target engine reaches its retention period, it will be flagged and deleted by the policy agent based on a daily schedule.

The image displays two screenshots from a management interface. The left screenshot shows the 'Replica Retention Policy' configuration page. It features a sidebar with 'Policy', 'Datasets', and 'Summary' options. The main content area is titled 'Policy' and includes a description: 'Replica Retention policies determine how long snapshots are retained for replicated namespaces, groups, dSources, and VDBs after they have been deleted on the replication source.' Below this, there is a 'Policy Name' field containing 'Six Months' and a 'Keep Snapshots for' field with a value of '6' and a unit dropdown set to 'Month(s)'. A note at the bottom states: 'Policies that keep snapshots for 0 days will get deleted during the next replication process if they have been deleted on the replication source.' At the bottom of the page are 'Cancel', 'Back', 'Next', and 'Submit' buttons.

The right screenshot shows the 'Policies' overview page. The top navigation bar includes 'DELPHIX MANAGEMENT', 'Manage', 'Resources', 'System', and 'Help'. The 'Policies' section has tabs for 'SnapSync', 'VDB Snapshot', 'Retention', 'VDB Refresh', and 'Replica Retention'. A '+ Replica Retention' button is visible. The main content area shows a policy named 'Six Months' with a trash icon and an edit icon. Below the policy name, it states: 'Snapshots will be removed after 6 month(s)'. The policy is applied to 1 namespace(s): 'namespace-extended'. It is also applied to 1 group(s): 'namespace-extended | Untitled'. Finally, it is applied to 2 dataset(s): 'namespace-extended | Untitled' and 'Vdbd_VNS | dbdhcpt1'. An upward arrow is at the bottom of the list.

dxi Encrypted Credentials and Executables

- We distributed the dxi CLI as Windows, macOS and RHEL binaries.
- This will simplify your adoption and remove the requirement on Python.
- Critically, we've also added encryption for the login credentials.

dxi cluster support

- dxi now supports the 'dxi database ingest' command (creation of a dsource) for Oracle RAC and Windows AG Cluster as well as provisioning to Oracle multi-tenant PDBs.

Chaining Algorithm

- We will now support building algorithms from a combination of existing algorithms.
- The Chaining algorithm framework allows the you to create masking algorithms that apply different behavior (called actions) to each input based on whether it matches one or more regular expressions, as well as to apply different actions to *parts* of each input based on which sub-group of the regular expression it matches. Actions range from simple value removal, redaction and preservation, to applying another algorithm to the input.
- There are no predefined or "default" instances of the algorithm shipped with the masking engine. There is also currently no UI to assist in constructing new instances.

The algorithm uses JSON for this example:

```
{
  "maskPatterns" : [
    { "regex" : "[Nn]ew [Hh]ampshire|[nN][hH]", "actions" :
      [ { "type" : "REDACT", "redactString" : "NH" } ] },
    { "regex" : "[Mm]assachusetts|[mM][aA]|[mM]ass\.", "actions" :
      [ { "type" : "REDACT", "redactString" : "MA" } ] }
  ],
  "fallbackAction" : { "type" : "PRESERVE" },
  "requireMask" : false
}
```

This example shows how the Chaining algorithm can be used to perform Data Cleaning tasks. We have several variations for how the states "New Hampshire" and "Massachusetts" can be written, and they are each replaced with the standard two-letter postal code.

Note the final example - the pattern must match the *entire* input.

Masking Results

Input	Mask Result
Nh	NH
new Hampshire	NH
mass.	MA
mA	MA
something else	something else
nh nh nh	nh nh nh

Db2 Staging Push Enhancements

- As part of the porting of the Db2 plugin to the Python vSDK, the plugin will run additional validation checks on databases mounted through Staging Push.
- The checks will verify whether a virtual database can be provisioned from the source.

HANA Plugin Staged Architecture

- We introduced a staging architecture for HANA virtualization, making it consistent with other virtualizing data sources.
- These changes will allow us to support those of you with various SAP-certified, 3rd-party backup applications for HANA.
- We will continue to support the pull ingestion method with HANA native backups and logs.
- Lastly, like Db2, we have ported the HANA Plugin to the virtualization SDK.

Add dSource



- **Source**
- dSource Configuration
- Data Management
- Policies
- Hooks
- Summary

Data Source

Filter: none

- **DEMO**
rhel-81-hana2-sp53
AppData
- **DEMO1**
rhel-81-hana2-sp53
AppData

Data Type
AppData

NFS Mount Location *
Mount Location on Staging Host

SAP HANA Staging Tenant Database Name *
The name of the Tenant Database to be created

Staging Database Parameters
Database parameters are required

Staging SystemDB User Name *

Staging SystemDB Password *
Double quote, single quote or space characters are not valid.

Staging Tenant Database ADMIN User Name *
The tenant database admin user which will be recovered from the source database.

Staging Tenant Database ADMIN Password *
Double quote, single quote or space characters are not valid.

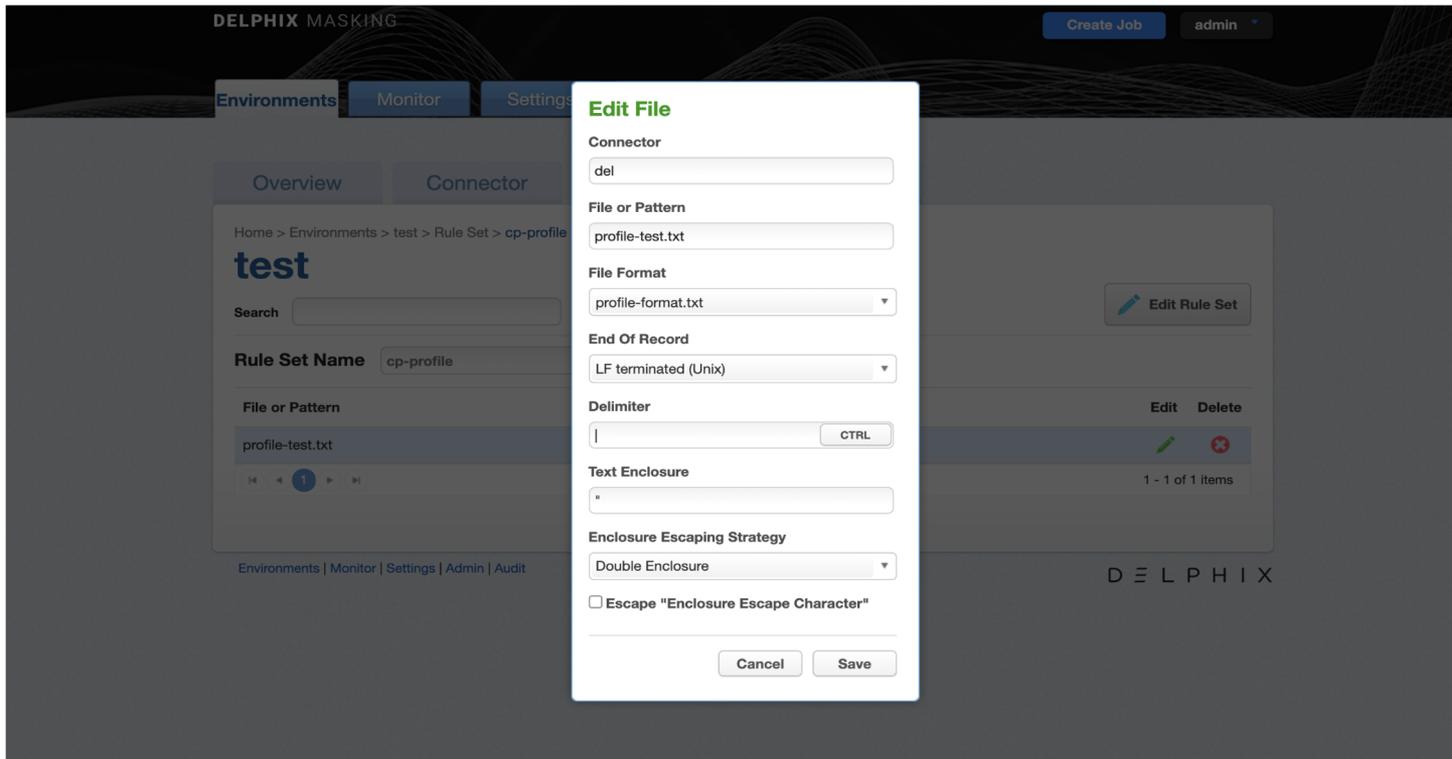
Customer Provided Data Backup Location *
Please provide the path to the data backup files on the staging host

Customer Provided Log Backup Location
Please provide the path to the backup files on the staging host. Please do not use spaces to be used

Cancel Back **Next** Submit

Escape Character Support for Delimited File Masking

- We will now support the configuration of escape characters for enclosure characters in file masking.



System Tunable Interface

- To better support you, we will be exposing a set of commonly used system settings.
- You will be able to use this with Support's direction.

```
[sysadmins-MacBook-Pro-4:appliance blewis$ ssh sysadmin@tunables.dcoll.delphix.com
Password:
[tunables.dcoll] system/advancedSettings
Changing advanced settings without consulting Delphix support is risky and not recommended.
[tunables.dcoll system advancedSettings] getTunable; set subsystem=virtualization; set name=aaa.lockout_failed_attempts; commit
3
[tunables.dcoll system advancedSettings] setTunable; set subsystem=virtualization; set name=aaa.lockout_failed_attempts; set value=2; commit
Successfully set "aaa.lockout_failed_attempts" to "2".
[tunables.dcoll system advancedSettings] ls
Properties
  type: AdvancedSettingsInfo
  modifiedTunables:
    0:
      type: Tunable
      name: aaa.lockout_failed_attempts
      subsystem: virtualization
      value: 2
```

Operations

```
getTunable
setTunable
```

```
[tunables.dcoll system advancedSettings] █
```

Oracle Customized Full Backups

- There is a rarely-seen bug in Oracle that results in some blocks not being written to the datafiles during an Oracle SnapSync operation. When this happens, the datafiles can become incomplete and provisioning/refreshing from that snapshot might fail.
- We are providing a SnapSync option via CLI that you can customize to accept all datafiles during an Oracle SnapSync operation to prevent this error.

The new option can be used from the CLI as follows:

```
ip-10-110-214-146 database> select CDOMLOSR1CEBPDB3  
ip-10-110-214-146 database 'CDOMLOSR1CEBPDB3'> sync  
ip-10-110-214-146 database 'CDOMLOSR1CEBPDB3' sync *> set  
filesForFullBackup=1,2  
ip-10-110-214-146 database 'CDOMLOSR1CEBPDB3' sync *> commit
```

End of Life

- ASE 15.0.3
- Windows 2008 R2
- SQL Server 2008/2008R2