

# Deploying Catalogic® ECX® on Pure Storage FlashArray for Oracle® databases

**Contents**

---

- Executive Summary ..... 3**
- Goals and Objectives ..... 3**
- Audience ..... 3**
- Copy Data Management..... 4**
- Pure Storage Overview ..... 5**
- Catalogic Software – Solution Overview..... 6**
  - Catalog ..... 7
  - Automate ..... 7
  - Transform ..... 7
- System Requirements ..... 8**
- Pre-requisites ..... 9**
- Catalogic ECX – Software Overview ..... 11**
  - Installing and Configuring ECX ..... 12
- Copy Data Management Use Cases ..... 16**
  - Storage Workflows (Templates) ..... 16
  - Pure Storage Templates ..... 16
  - Copy Data Policies ..... 17
  - Use Data Policies ..... 22
    - Creation of Dev-Test and Business Analytic Environments ..... 23
    - Other Use Cases ..... 27
      - For DevOps Automation ..... 28
      - For Near-real-time Data Access for Business Analytics: ..... 28
  - Copy Data Analytics ..... 29
- Summary ..... 30**
- About the Author ..... 32**

## Executive Summary

---

Oracle database sprawl comes with significant cost and complexity. The complexity is the result of proliferation of database copies created for a range of uses including disaster recovery, provisioning of development and testing infrastructure, quality assurance (QA), DevOps in private or hybrid clouds to name a few. Catalogic Software's® ECX™ Copy Data Management platform deployed in conjunction with the Pure Storage® FlashArray™, allows businesses to manage, orchestrate and analyze Copy Data. The solution provides full lifecycle management of your Copy Data through automated workflows that allow you to streamline the creation, management and use of Oracle database copies.

## Goals and Objectives

---

This document is intended to discuss the specific ways to manage, orchestrate and analyze Copy Data in a Oracle database environment deployed on the Pure Storage FlashArray using Catalogic Software's intelligent Copy Data Management platform, ECX.

This document discusses critical operations in IT organizations and explains how ECX software provides the ability to get control over Oracle data copies across the enterprise, delivering the right data copy for the right business function, at the right time and in the right location--all within a single simplified and automated platform.

The document includes details on installation and setup as well as the creation of Copy Data workflows and Use Data workflows. These workflows are used to spin up instances of application consistent snapshots of an Oracle database running on the Pure Storage FlashArray, to be leveraged for different types of use cases.

## Audience

---

This paper is written for IT decision makers, storage administrators and Oracle database administrators interested in ECX technology integration with the Pure Storage FlashArray snapshot and replication technology to deliver superior copy management functionalities for Oracle databases. Familiarity with Oracle, Pure Storage and Catalogic ECX is recommended.

# Copy Data Management

---

Database administrators working extensively with Oracle are challenged when faced with mission-critical use cases such as Backup, Recovery, DevOps, and Business Analytics. This is especially true given that most Oracle databases expand in size and number over time need to be up and running 24x7x365.

Oracle DBAs struggle with the following:

- Backups are slow, complex and need constant management
- Backup processes impact production servers
- Recoveries are slow and complex
- Repurposing App consistent backups (clones) for DevOps and Business Analytics is slow, complex and storage inefficient
- Lack of automation exists for producing quick and secure clones required to accelerate DevOps
- Copy sprawl problems occur due to no central catalog of copies
- Unable to meet stringent RPO and RTO requirements

Catalogic ECX simplifies Oracle database copy management by enabling administrators to orchestrate App consistent copy creation, recover and clone databases in minutes instead of hours or days. ECX copy management leverages advanced storage features such as snapshot and replication to rapidly create, replicate, restore or clone copies of Oracle databases which are space and time efficient. ECX created copies and clones are space and time efficient ECX enables you to focus on the copy and recovery requirements of your business rather than the technical details of the underlying storage platforms.

## Pure Storage Overview

---

With Pure Storage FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient. Database copies can be created and managed at the storage array level by leveraging zero-footprint snapshot. Snapshots are an intrinsic part of the Purity Operating Environment. That is, read/write addressable volumes that are created "virtually" and can be mounted and used by servers. The goal is to provide access to data quickly while remaining extremely space efficient

The Pure Storage snapshot management interface is designed to provide flexibility, scale, and ease of use. The interface allows users to select one or multiple volumes simultaneously in order to create a consistent point-in-time snapshot of all the selected volumes. Snapshots can also be created instantaneously for all the volumes owned by a host or host group thereby providing a consistent view of data for all the volumes assigned to selected hosts at a given point in time. Purity snapshot also protect volumes from accidental deletion by creating an internal snapshot before a volume delete is triggered. Leveraging FlashArray's flexible protection policy management, a user can automate the creation and retention of snapshots for local data protection and recovery.

The Pure Storage REST API provides simple interfaces for many Purity commands. For example, to create snapshots of two volumes, say v5 and v6, with a custom suffix you would execute this request:

```
POST https://pure01.example.com/api/1.0/volume
{
  "snap": true,
  "source": [
    "v5",
    "v6"
  ],
  "suffix": "crackle"
}
```

And the Purity Operating Environment returns the following response:

```
[
```

```
{
  "created": "2014-03-25T20:44:37Z",
  "name": "v5.crackle",
  "serial": "31B3F6E02E4B691C00010009",
  "size": 5368709120, "source": "v5"
},
{
  "created": "2014-03-25T20:44:37Z",
  "name": "v6.crackle",
  "serial": "31B3F6E02E4B691C0001000A",
  "size": 1073741824,
  "source": "v6"
}
]
```

Snapshots deliver superior space efficiency, high scalability, and simplicity but are not centrally managed. If you have got three different arrays on your floor, you will have three different sets of snapshot information. Not the most efficient way to manage things, and it is very hard to get a high level view of what is going on when everything is in a separate array.

ECX automated workflows allow clients to streamline Copy Data management operations. It leverages the Pure Storage REST API snapshot features to create and manage snapshots as described in the following sections.

## Catalogic Software – Solution Overview

---

Catalogic Software delivers the first “In Place” Copy Data Management platform ECX, that works on your existing storage environment, driving operational efficiencies, cost savings and better leverage of your storage assets. In the modern IT environment, copies of production data can be more vital to the business than the production data itself. But with exponential Copy Data growth, and a mix of of existing tools and scripts to manage these copies, IT is often unable to meet the commitments to the business that depend on this data.

ECX immediately delivers on three key value areas by managing the full lifecycle of Copy Data. ECX allows IT to manage, orchestrate and analyze Copy Data across the enterprise and cloud.

## Catalog

The Catalog function discovers the assets in your environment to build a rich, meta-data catalog which you can search for numerous IT objects.

## Automate

The Automate function brings automation and ease of operations to common IT tasks, using a policy-based model.

## Transform

And finally, these features let you transform your IT environment by easily expanding into areas such as hybrid cloud and DevOps.

With this core functionality, ECX drives many use cases as shown in Figure 1. ECX allows IT to leverage application consistent data copies to drive use cases like enhanced protection and disaster recovery, automated Dev-Test, DevOps integration and near real-time data access for Business Analytics. At the heart of the platform is an actionable catalog and includes a robust policy engine to manage and orchestrate the Copy Data environment and associated workflows. ECX has detailed and customizable reports along with an advanced query engine providing deep analysis as well as real-time and historical service-level reporting.

# IT Modernization through “In Place” Copy Data Management

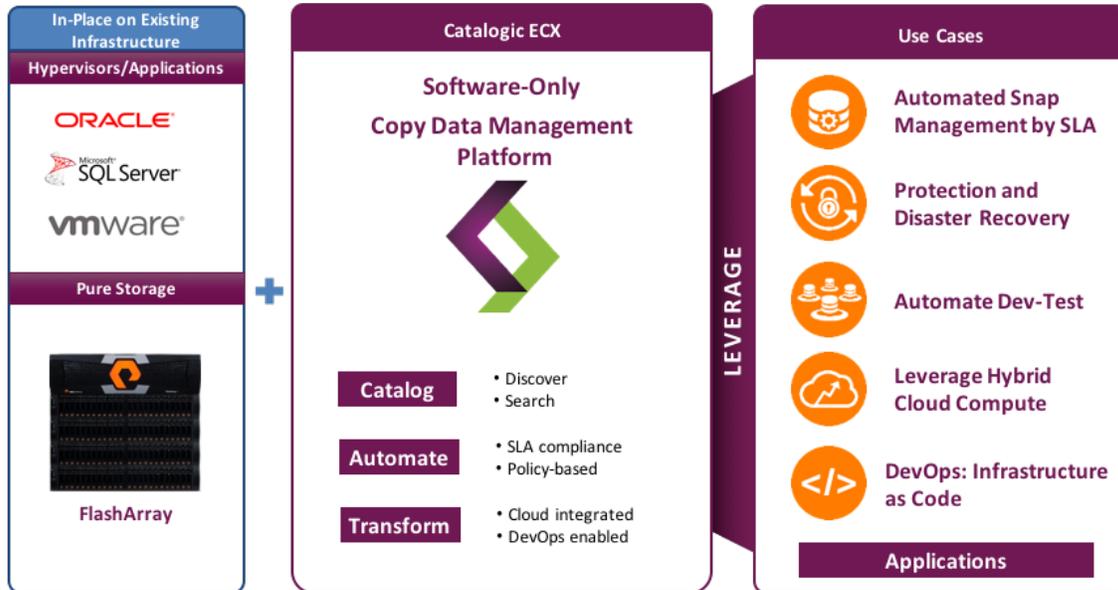


Figure 1. Catalogic Software Overview

## System Requirements

Ensure that you have the required system configuration to deploy and run ECX with Pure Storage.

- Pure FlashArray//m series
- Purity v4.2.3 or later
- Oracle 11gR2 or Oracle 12c
  - Single instance – a single instance running on a single server accessing a database
  - RAC (Real Application Clusters) leveraging ASM (Automatic Storage Management) – more than one instance running on multiple servers are accessing a database simultaneously

- Linux Physical servers
  - RHEL/CentOS 6.5+, RHEL/CentOS 7.0 and above
  - OEL 6.0+, 7.0+
  - SLES 11.0+, 12.0+
- Licensing:
  - Pure Storage arrays require no additional licensing for Catalogic Software ECX functionality
  - ECX is licensed per-FlashArray and costs vary by array model
- Protocol Support: Fibre Channel or iSCSI
- Catalogic Software ECX version 2.5 and above

## Pre-requisites

---

Ensure the following pre-requisites before walking through the configuration steps and use cases presented in the document

1. **Catalogic ECX Virtual Appliance** – ECX is a self contained virtual appliance. ECX deployment will automatically create a new virtual machine and install the necessary components. It is a software defined control plane that manages all the activity between applications, virtual machines and storage systems through API calls.
2. **Oracle Environment** - Oracle should be installed and running on physical server with data hosted on Pure storage.
3. **Pure Storage Volume** – One or more volumes needs to be connected to the Oracle servers.
4. **Security** – ECX requires the following privileges:
  - **Sudo privileges for the operating system** – This is required for tasks such as discovering storage information, rescanning iSCSI sessions, and mounting and unmounting disks.

- **Permissions to read the Oracle inventory** –This is required to discover and collect information about Oracle homes and databases.
- **SYSDBA privileges for database instances** –This is required to perform several database tasks such as querying instance details, starting and ending hot backup mode and performing RMAN cataloging during backups, as well as starting and stopping instances during restore.
- **SYSASM privileges for ASM instances** -This is required to perform several storage tasks such as querying ASM disk information as well as renaming, mounting, and unmounting diskgroups during restore.

# Catalogic ECX – Software Overview

ECX delivers is a robust in-place copy data management platform, giving IT a single enterprise-wide system that replaces the complicated set of products, tools and scripts that are collectively used today. ECX is a software-only solution that installs as a virtual machine, requires no agents and deploys within 15 minutes. ECX automated workflows allow clients to streamline Copy Data management operations.

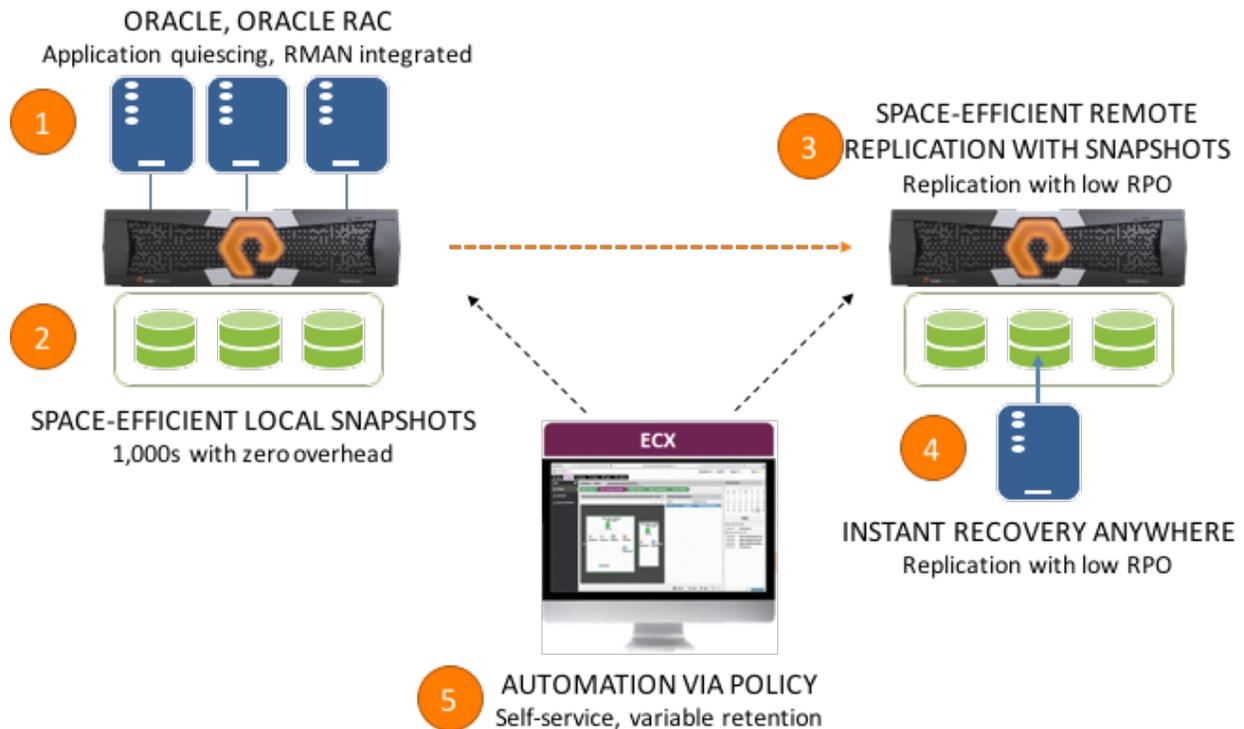


Figure 2. ECX Oracle copy management overview

Catalogic ECX catalogs all of the primary Copy Data in your Oracle Databases hosted on Pure Storage FlashArray. This allows you to orchestrate, analyze, search, and report on all your data, allowing you to take full advantage of your data assets. By cataloging and managing all tiers of array-based snapshots and Oracle objects, with an intuitive point and click interface, clients can automate and orchestrate oracle aware snapshots and instantly use them for test-dev, disaster recovery and analytics operations. Additionally, the power of Pure Storage snapshots ensures that the tasks are completed in seconds without losing any drop in performance.

The orchestration and automation in ECX copy data management solution

involves two operations, Copy Data Policies and Use Data Policies as shown in Figure 2.

The Copy Data policies enable application consistent in-place snapshots and array based replications. The first step in Copy Data policy is injection of a light weight agent into the servers running the database instances. Next step is to identify the mappings between the Oracle databases and the underlying storage volumes. Then Oracle database is then placed in a hot backup mode temporarily and a FlashRecover snapshot is created on the Pure Storage array. The application is then taken out of hot backup mode and the lightweight agent is removed. Additional options like log backup, RMAN catalog, data masking can be added to the workflow. Templates can be customized with storage workflows feature. Storage workflows define the operations performed on the storage array driven by API calls from ECX. Storage arrays include in-place FlashRecover snapshots, replications, storage quotas etc. Replication relationships and target storage provisioning are automatic and are detailed in Storage Workflows section.

Use data policies leverage these copies of production data and make them available instantly for multitude of use-cases including disaster recovery, dev-test automation and improved devops. Use data policies can be created for individual databases or multiple databases and assigned to end users based on Roles Based Access Control (RBAC).

## Installing and Configuring ECX

ECX can be deployed using a few clicks within a VMware infrastructure. In vSphere, simply specify the location of the .ova file and pick the Host and Network to run the appliance as shown in Figure 3. ECX comes pre-packaged with all of the required software, and once powered on, the console screen points to the link to the portal as seen in Figure 4.

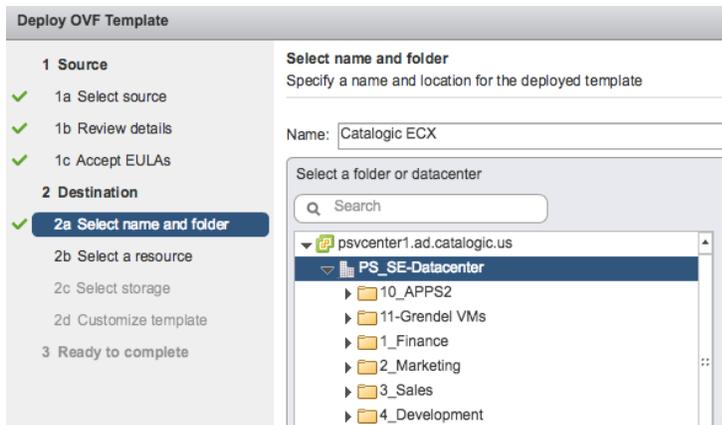


Figure 3. Deploy ECX Virtual Appliance

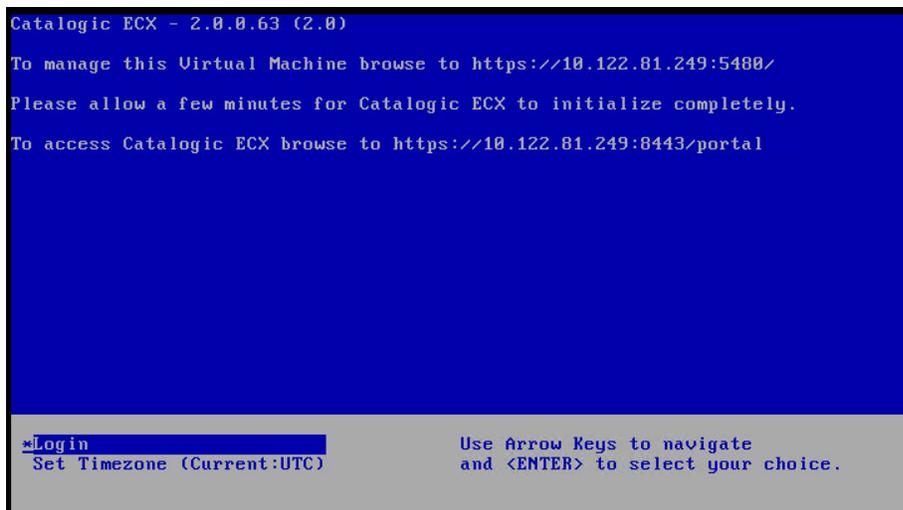


Figure 4. ECX appliance with link to the Web Portal



Figure 5. ECX Web Portal Login

The portal link can be accessed from any web browser. Figure 5 shows the login dialog to access the application.

The first step is to register the storage assets and Oracle resources and assign to their respective sites. This is a simple and intuitive one-time agentless registration process. To add a new instance, simply right-click and select Register, as shown in Figure 6. Select Register and fill in the details, as shown in Figure 7. The registration dialog box takes in credentials and a site selection. ECX has a site-based concept to identify resource location and also the capability of selecting a data copy to be used from any of the sites. For this whitepaper the setup has two sites: 'New York, NY' and 'San Francisco, CA'.

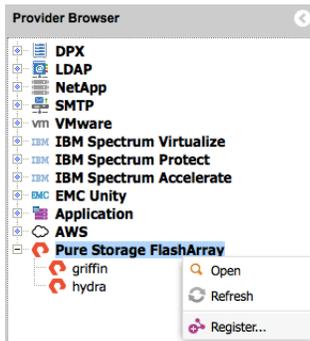


Figure 6. Register Resource

**Register Pure Storage Server** ✕

Site:

Name:

Host Address:

Port:

Use SSL

**Select** **New**

Name:

Username:

Password:

Comment:

[Need help?](#)

**OK** **Cancel**

**Register Application Server** ✕

Application Type:  Oracle  SQL

Site:

Name:

Host Address:

Catalog provider resources after registration

Authentication:  Key:  Password

**Select** **New**

Name ↑	Username	Type
<input type="radio"/> GJDFS	administrator@GJ.DFS	Domain
<input type="radio"/> PSDEMO-SQL01 Domain Ad...	mphan@devad.catalogic.us	Domain
<input checked="" type="radio"/> PureOra37Oracle	oracle	Domain
<input type="radio"/> SSSQLPURESTD Domain Login	administrator@gj.dfs	Domain

<< < Page 1 of 1 > >> | [Refresh](#)
Displaying 1 - 4 of 4

**OK** **Cancel**

Figure 7. Register Pure Storage and Oracle Resource

Once the resources are registered, ECX will automatically create a Catalog policy. This Catalog policy will discover high-level objects in Storage arrays like storage volumes, capacity etc. and for Oracle discover all the databases contained in that particular server. Figure 8 shows the Plan page, where all the policies can be created. We will explain them in detail in the following sections.

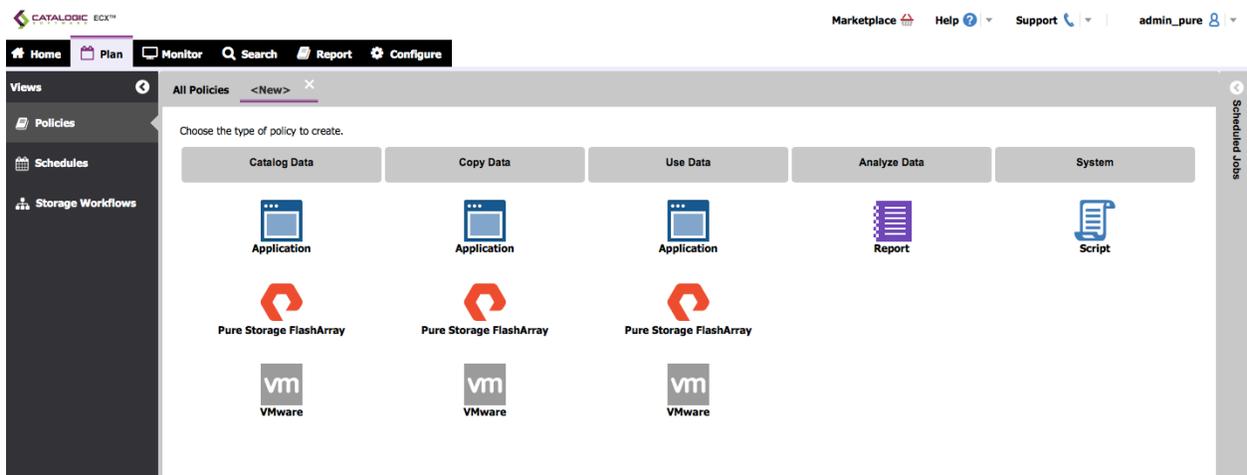


Figure 8. ECX Copy Management Plan screen

# Copy Data Management Use Cases

## Storage Workflows (Templates)

Storage workflows allow the storage administrator to create multiple templates, which define the sequence of operations, and to assign service levels to each of the templates (for example, a Gold/Silver/Bronze service level model). Having visibility into your existing data sets, applications, snaps and mirrors is important to help determine how your data is currently being used today, and to make decisions on how it should be leveraged as you move forward. However, that is only half the battle. The other half starts with the simplification of the existing data copy creation processes through automation, such that the data copies align with the SLAs of the enterprise, while ensuring compliance and auditability of the environment via exception-based reporting. As you introduce the creation of data copies in multiple destinations, this management and orchestration becomes impossible without the right tools.

In this section we will go through the steps to create Storage Workflows for Pure FlashArray systems that would allow the storage administrator to publish templates that can be used in the Copy data policies for applications like Oracle and SQL or VM owners.

## Pure Storage Templates

Storage admins can create multiple templates based on SLA requirements. In the **Plan** tab, select **Storage Workflows**, then click **New**. The **New Storage Workflow** pane opens as shown in Figure 9. For Type, select the Pure Storage FlashArray option, as shown in Figure 10. Then select the green circle icon with label Source. This gives you three options Add Snapshot, Add Replication and VM Copy.

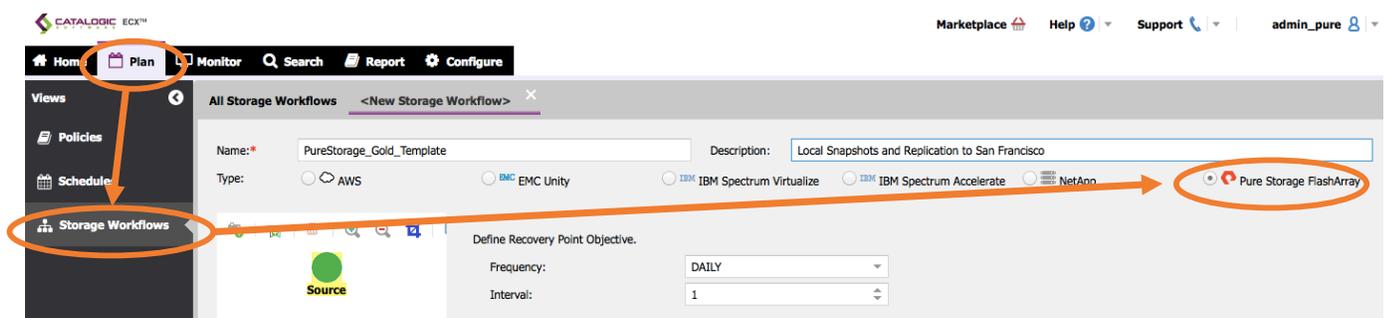


Figure 9. Create Storage Workflows

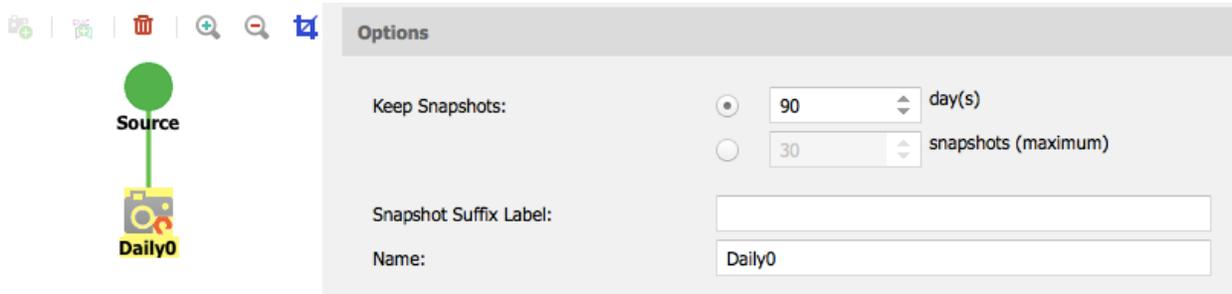


Figure 10. Pure Storage Workflow

Pure Storage FlashRecover snapshots deliver superior space efficiency and high scalability. Replication enables data-mirroring between two Pure Storage FlashArrays. In our example, we use the name “Pure Gold Template” and the description “Local Snapshot and Replication to San Francisco” as seen in Figure 9. For an ECX Copy Data policy that uses a replication relationship, ECX will create snapshots of the destination volume and then catalog these copies for subsequent use in ECX Use Data policies.

Figure 11 shows the storage target for Replication and the options to set retention and replication frequency. In a similar manner. Storage administrator can customize the workflows depending on the SLA’s and publish multiple templates such as Gold, Silver, Bronze, etc. These templates are then assigned to application owners based on Role Based Access Control (RBAC).

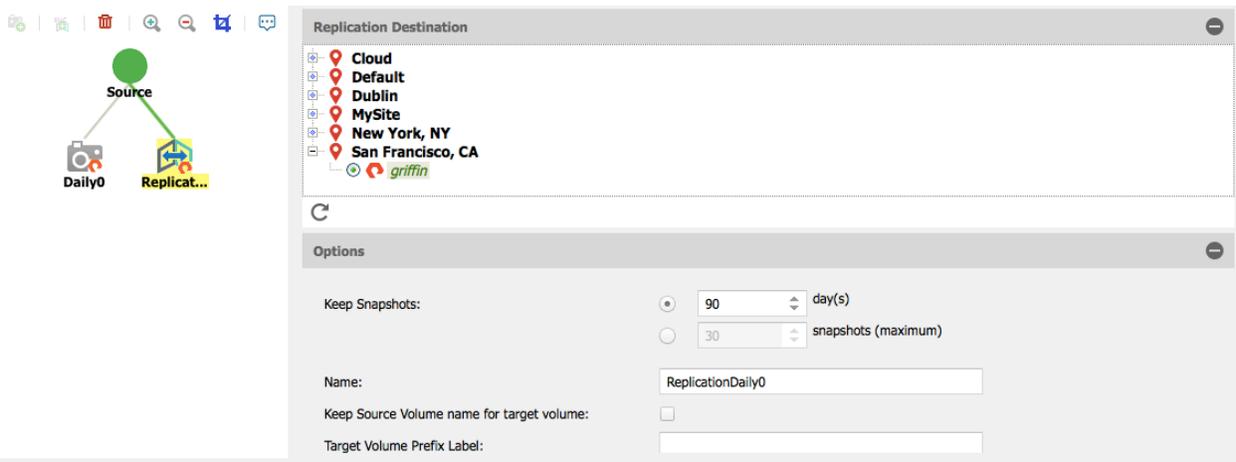


Figure 11. Pure Storage Replication Destination

## Copy Data Policies

Within ECX we will create application consistent Copy Data policies that are

aligned with existing best practices for different data sets. To do this we will use Oracle databases hosted on Pure Storage FlashArray. As shown in Figure 12, in the **Plan** section select **Policies** and click **New**. Select **Applications** in the **Copy Data** tab and click on **Oracle** as shown in Figure 13. The first step in creating Oracle copy policy is to identify the databases that need to be copied as show in Figure 14.

Note: Ensure that the Application Catalog Policy has run successfully before creating a new Copy Policy

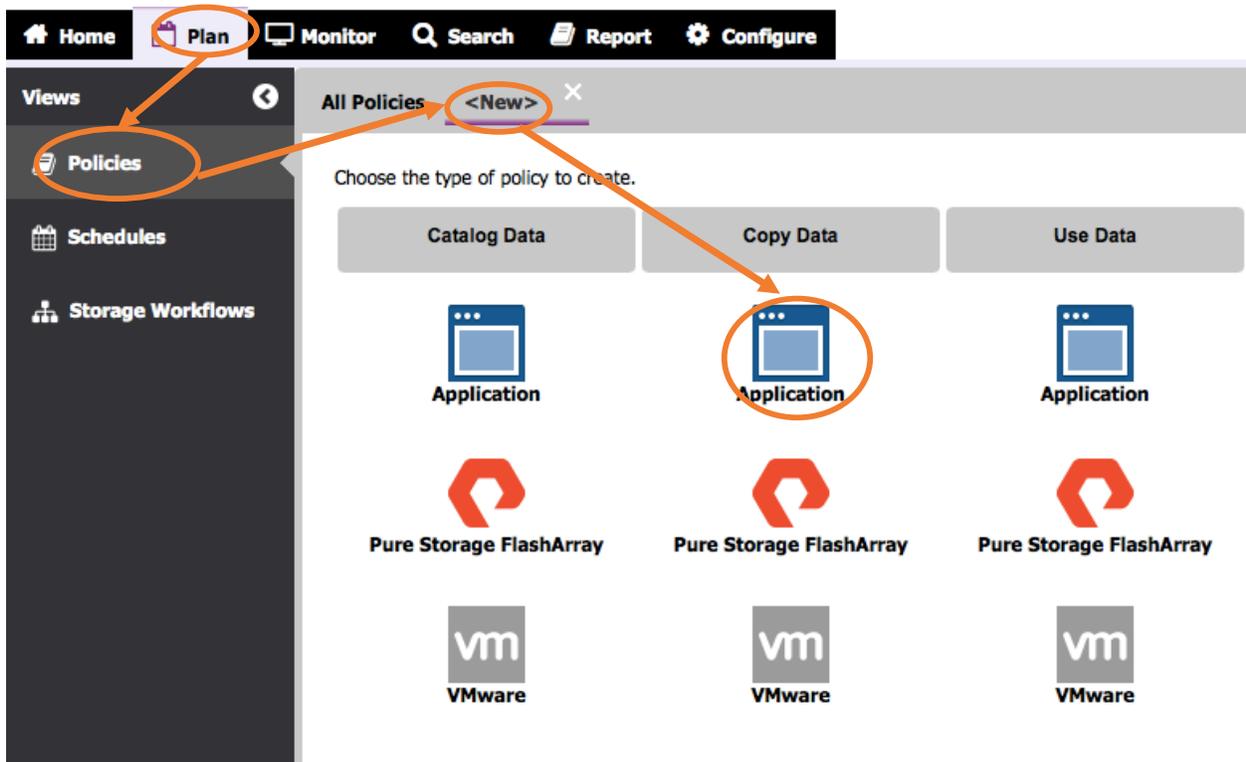


Figure 12. Creation of New Application Copy Policy

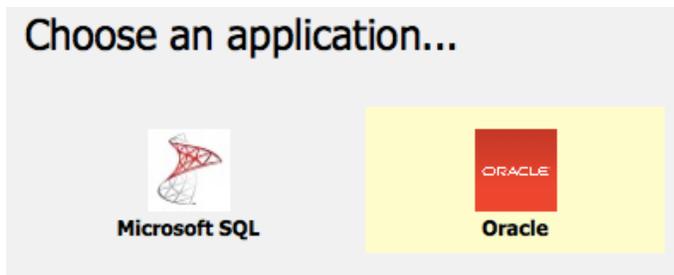


Figure 13. Selection of Oracle Application

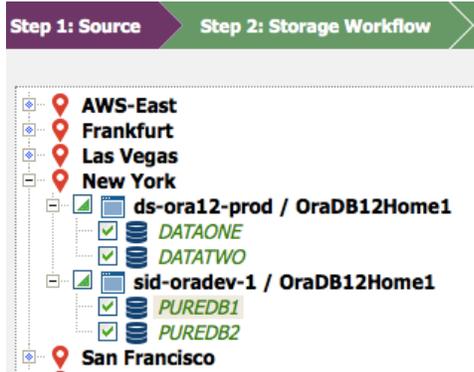


Figure 14. Oracle Database Selection

Once the databases are selected, we select “Step 2. Storage Workflow” button to see all the templates that were created in the previous section. In this example, we select the Pure Gold Template as shown in Figure 15. This template will orchestrate an Oracle consistent snapshot coordinated with a Pure Storage FlashRecover Snapshot, which is by far the most efficient means of gaining an application consistent copy of an Oracle database, followed by Pure Storage replication. The data will be replicated to the Pure Storage FlashArray in San Francisco.

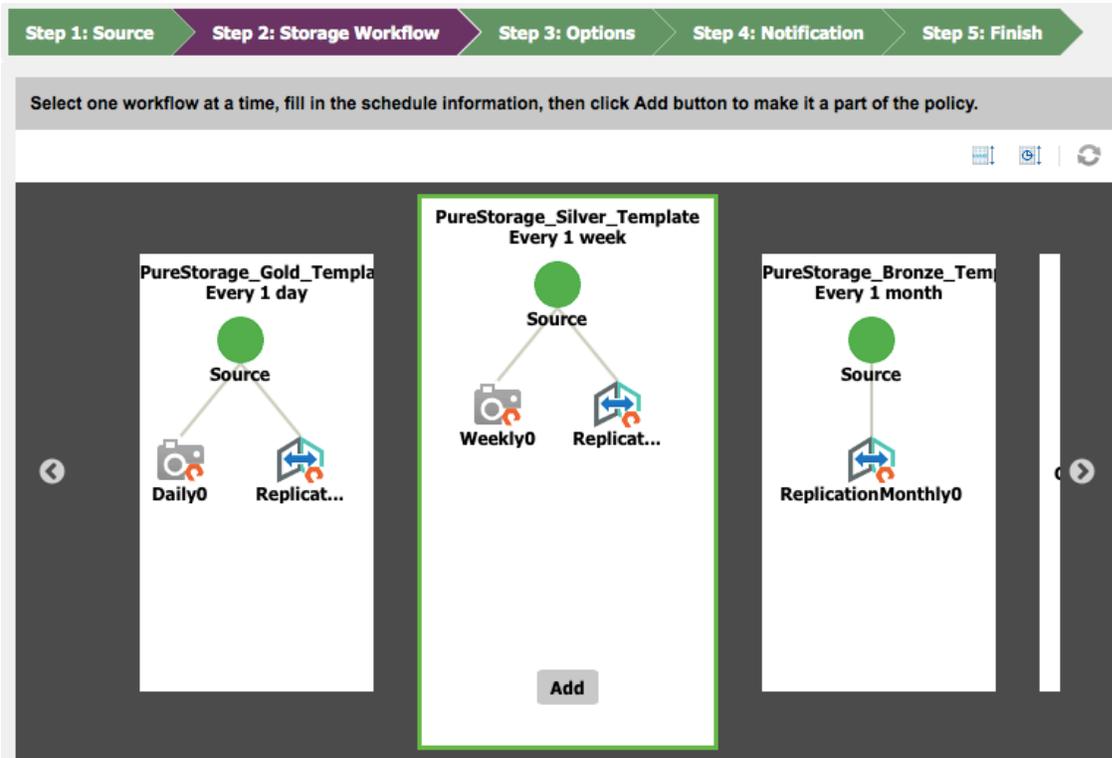


Figure 15. Selection of Storage Template

The next step is to apply additional options to the policy. ECX automates cataloging of Oracle database copies in the Recovery Manager (RMAN) recovery catalog. In “Step 3: Options” we can choose to catalog in the local RMAN repository or record copies in a pre-configured remote RMAN catalog as shown in

Figure 16.

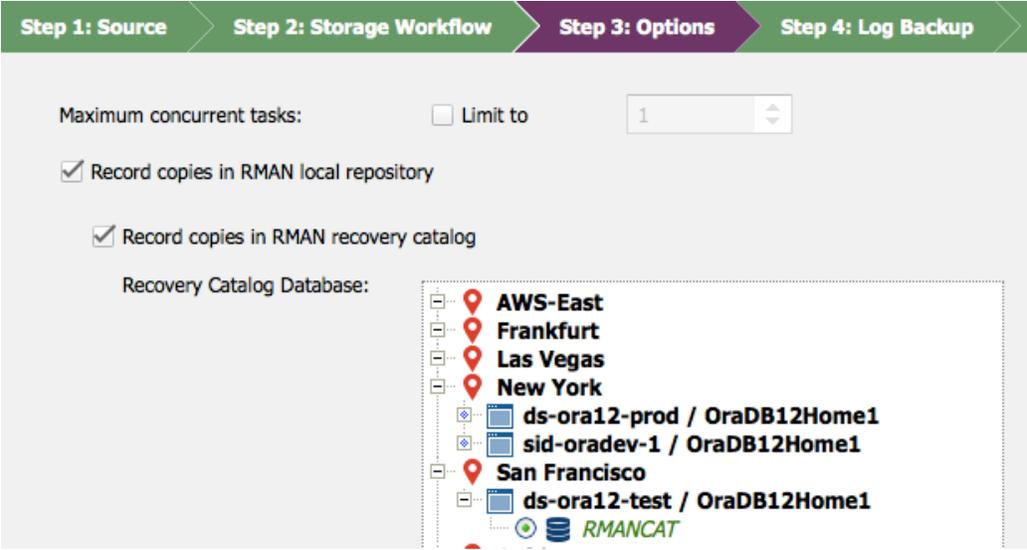


Figure 16. Oracle Recovery Manager (RMAN) catalog

Oracle databases creates transaction logs as part of its operation. In “Step 4: Log Backup” we select the option for Archive log files processing as shown in Figure 17. If we choose to protect archive logs, ECX enables copy of Archive logs to a specified destination providing lowest RPO (transaction level recoveries). Archive log retention is managed by settings in RMAN. ECX leverages archived logs to enable PIT recoveries of databases to facilitate fine grain RPOs.

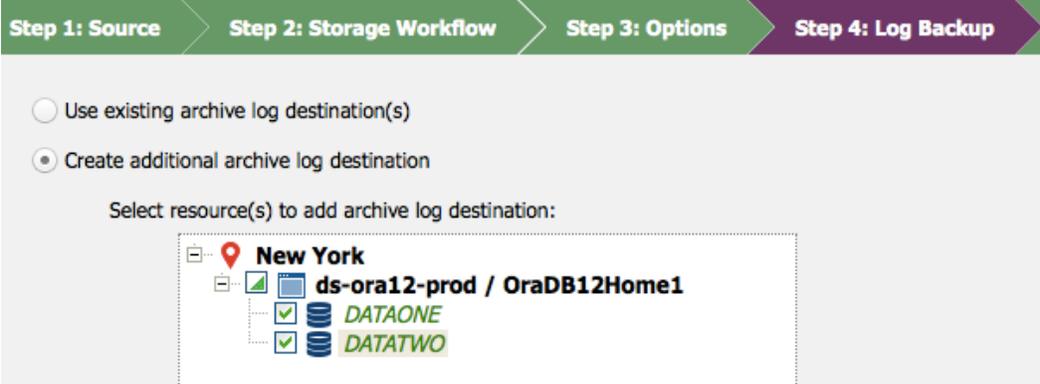


Figure 17. Archive log processing

ECX includes data masking workflows with the ability to leverage third party masking tools. In “Step 5: Data Masking” we can choose between enabling data masking or storing snapshots as is. If Enable Data Masking option is chosen, the user can select a local copy or a Replicated copy to be masked as shown in Figure 18. Masking routines vary from simple operations to more complex algorithms and are typically intensive and time consuming operations. In ECX we can offload the masking operations to an alternate server. The snapshot will be

mounted to the alternate server where the data masking scripts are located. Once masking operation is completed, they will be automatically marked in the ECX catalog.

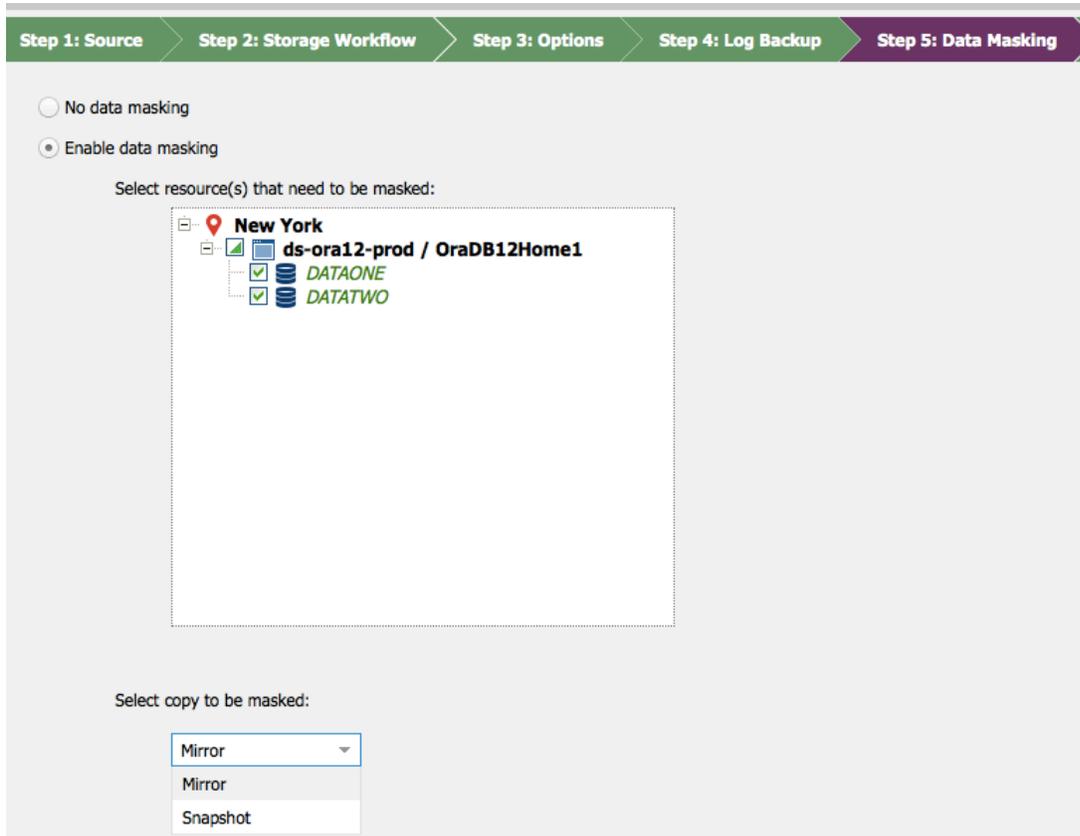


Figure 18. Selection of Data masking option

The Policies will be triggered based on the Schedule applied to them in the Storage workflow templates or can be manually initiated from the Monitor tab as well.

## Use Data Policies

The next step in the process is to create a Use Data workflow. These workflows allow business lines to access the data copies for multiple use cases. ECX Use Data workflows can use a local copy or a replicated copy to spin up databases in an Oracle server. In this section we will go through an example of setting up Oracle instances for Dev-Test environments. For this we will use a copy of couple of databases from the production Oracle server in **New York** and mount it to a test server in **San Francisco** using the replicated copy.

## Creation of Dev-Test and Business Analytic Environments

In the **Plan** section select **Policies** and click **New**. Select **Applications** as shown in Figure 19 in the **Use Data** tab and click on **Oracle** as shown in Figure 20. In this use data policy, we will select 2 Oracle databases to be spun up in the San Francisco site on a test Oracle server.

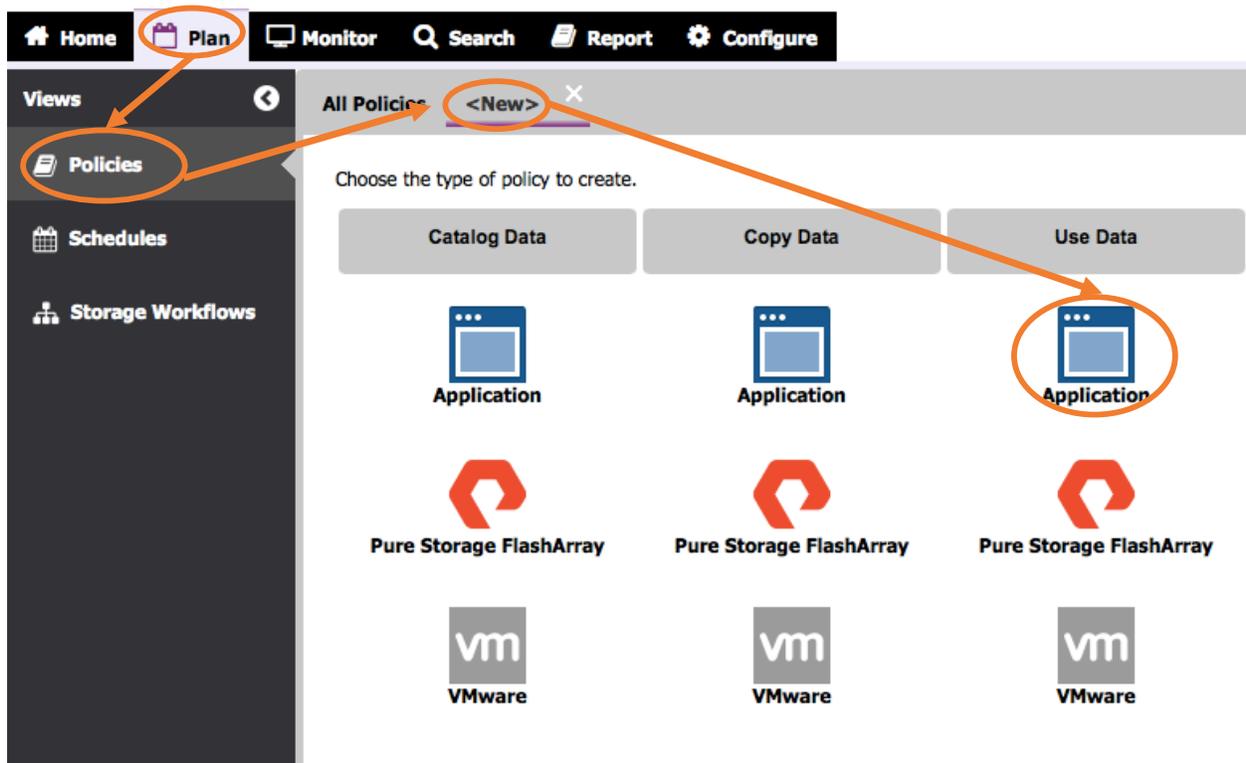


Figure 19. Creation of New Application Use Policy

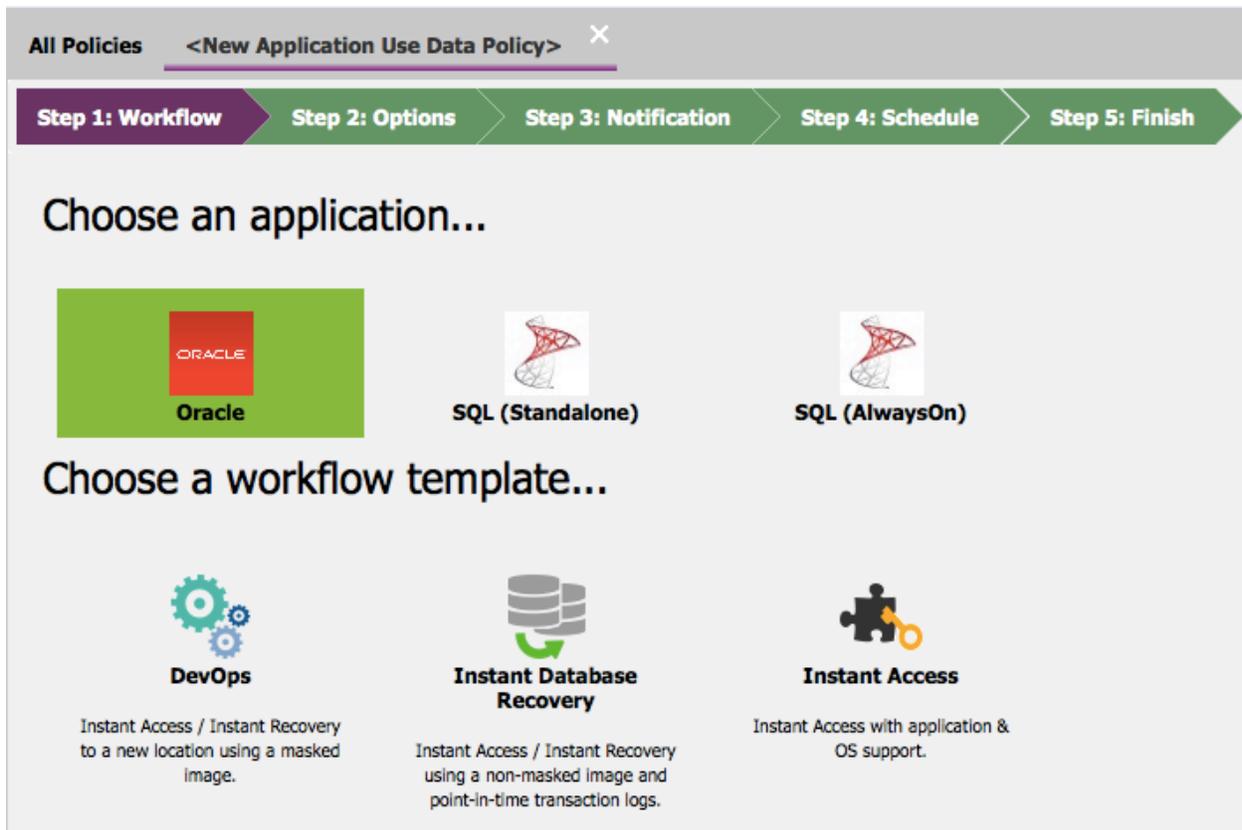


Figure 20. Selection of Oracle Application and Use workflow

There are three recovery templates based on user experience and access level. All three of them use the powerful Pure Storage Clone feature and can instantly create a clone of the database any Oracle server. The three templates are:

**DevOps** – This provides users access to data masked version of the database copies. Note: This feature is only available if Data masking feature was enabled during Copy Policies.

**Instant Database Recovery** – This feature allows databases to be instantly cloned either from a local snapshot or a replicated copy of the database.

**Instant Access** – This template is for advanced Oracle administrators, who would like to leverage RMAN for recoveries. Instant Access will provision a LUN from the snapshot and mount it to an Oracle server. DBAs can then run RMAN commands to recover the database.

DevOps and Instant Access are beyond the scope of this document and we will focus on Instant Database recovery template to use a database copy for Dev-test.

Commonly, dependencies between applications necessitate that they be recovered in a particular order. ECX allows the user to drag and drop the order in which customers would like to have the databases come up in. This ensures that all of the applications will work properly. Figure 21 shows how we can select the databases and put them in proper order. In this particular case we will select databases 'DATAONE' and 'DATATWO'

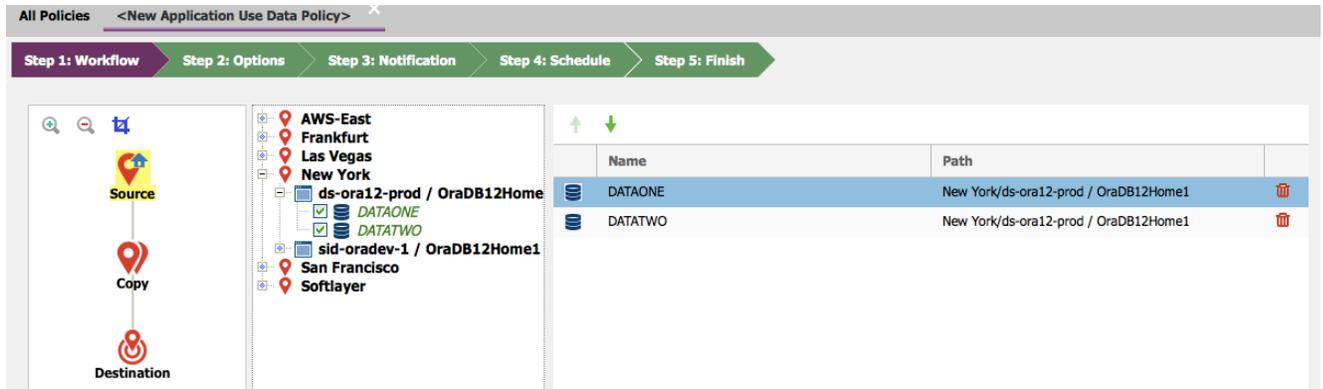


Figure 21. Use Policy Database list and sequence

Figure 22 shows the next step to choose the location where those resources exist. We will select the San Francisco site. We also have the ability to select a particular snapshots instance to recover or a checkbox which allows Point in Time recovery. We will select Use Latest snapshot. Once the location and the version is selected, we next select the Oracle Test server where the applications and data would be available, as shown in Figure 23. In a test scenario, databases can be mounted in a different name. ECX provides this functionality and we will name the two databases ONETEST and TWOTEST respectively

Copies of the selected data have been found at the following site(s). Select the site you would like to use.



**New York, NY**  
100% resources found



**San Francisco, CA**  
100% resources found

By default, the latest copy will be used. If you would like to choose a specific version, select a site then click the "Select Version" Button.

**Use Latest**

Allow Point-in-Time selection when job runs

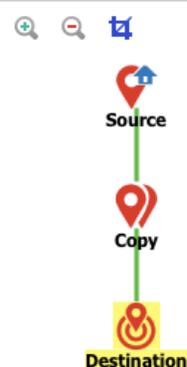
Select a specific version for the resources you want to recover.

Name	Type	Version
DATAONE	database	Use Latest
DATATWO	database	Loading versions...
		Use Latest
		July 26, 2016 at 3:40:37 AM PDT (RMAN Tag: ECX_20160726_104037)
		July 22, 2016 at 7:59:57 AM PDT (RMAN Tag: ECX_20160722_145957)
		July 20, 2016 at 2:24:43 AM PDT (RMAN Tag: ECX_20160720_092443)
		July 13, 2016 at 1:50:56 PM PDT (RMAN Tag: ECX_20160713_205056)
		July 13, 2016 at 7:17:07 AM PDT (RMAN Tag: ECX_20160713_141707)

Figure 22. Selection of Copy location and version

All Policies <New Application Use Data Policy>

Step 1: Workflow Step 2: Options Step 3: Notification Step 4: Schedule Step 5: Finish



- AWS-East
- Frankfurt
- Las Vegas
- New York
- San Francisco
- ds-ora12-test / OraDB
- Softlayer

Database	
Source	Destination
DATAONE	ONETEST
DATATWO	TWOTEST

Figure 23. Destination server selection and Alternate naming options

The policy to test the database environment can be scheduled to run as often as needed, even every day, as shown in Figure 24. Now that we have setup the policy it will run per schedule and the IT team can easily validate that the test environment is successfully running. In Figure 25, we can see the Monitor screen showing a dev-test policy and run in the San Francisco site, that spins up databases and refreshes the copies with the updated data.

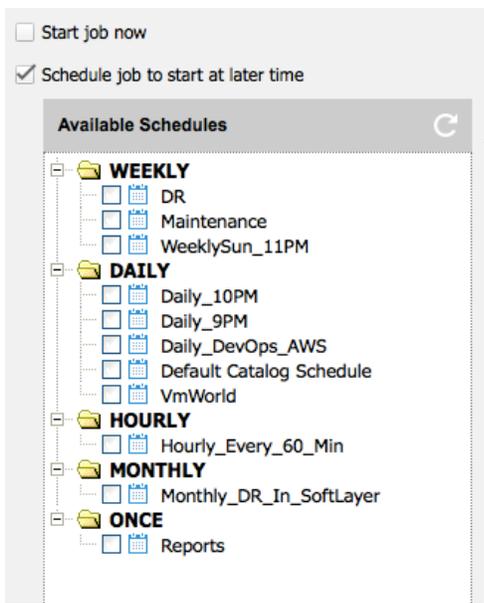


Figure 24. Daily Testing schedule

All Jobs <span>OraONE_TEST</span>					
	Start Time ↓	End Time	Duration	Description	Status
General	<a href="#">August 26, 2016 at 1:00:55 PM PDT</a>	August 26, 2016 at 1:16:53 PM PDT	0h 15m 57s		✓ COMPLETED
	<a href="#">August 24, 2016 at 7:48:13 AM PDT</a>	August 26, 2016 at 12:59:31 PM PDT	53h 11m 17s		✓ COMPLETED
	<a href="#">August 23, 2016 at 7:55:47 PM PDT</a>	August 24, 2016 at 7:47:48 AM PDT	11h 52m 1s		✓ COMPLETED
Activity	<a href="#">August 23, 2016 at 7:44:24 PM PDT</a>	August 23, 2016 at 7:55:02 PM PDT	0h 10m 37s		✓ COMPLETED
	<a href="#">August 23, 2016 at 7:35:43 PM PDT</a>	August 23, 2016 at 7:43:03 PM PDT	0h 7m 19s		✓ COMPLETED
	<a href="#">August 23, 2016 at 1:39:07 PM PDT</a>	August 23, 2016 at 1:57:51 PM PDT	0h 18m 44s		✓ COMPLETED
Policy	<a href="#">July 21, 2016 at 1:51:49 AM PDT</a>	July 21, 2016 at 2:04:15 AM PDT	0h 12m 25s		✓ COMPLETED
	<a href="#">July 20, 2016 at 9:58:01 AM PDT</a>	July 20, 2016 at 10:06:48 AM PDT	0h 8m 47s		✓ COMPLETED
History					

Figure 25. Monitor screen showing duration and completion status of DR policy

## Other Use Cases

The previous section focused on the applicability of leveraging ECX for Dev-Test environments, but the same analysis could apply to enhanced DevOps automation, as well as near real-time data access for business analytic workloads. In other words, while ECX may be initially deployed for one use case, it can deliver additional use cases by leveraging the same infrastructure. In fact,

once a user has access to application consistent replication, each use case has the same set of requirements for cataloging, reporting, managing and orchestrating the spinning up and spinning down of compute resources for effective data access.

### For DevOps Automation

As software development gets more and more agile, enterprise IT needs to provide increased agility and automation. The ECX workflows can provide the required orchestration and automation required to leverage enterprise storage in support of these faster development windows.

Because all ECX operations can be controlled via the ECX REST API, data workflows are easily integrated with DevOps tools such as Puppet, Chef and so on. Test infrastructure can be deployed via a simple API call from within the DevOps tool. This fits cleanly into the DevOps model of rapid application development, as data sets and infrastructure can be continuously updated to keep pace with development efforts. In this way, the combination of ECX and Pure Storage solves one of the biggest hurdles in DevOps, which is allowing DevOps to be run on data center resources.

### For Near-real-time Data Access for Business Analytics:

The same capabilities described in the previous sections could be utilized providing near real-time data access for analytics. We can follow the same steps to create test machines in a fenced off environment and use it to run Business Analytics, Application queries, Application-based reporting, ETL processes, etc., and not have to provision large amounts of storage assets. Once the analytics has been completed, we can let ECX tear down the environment, avoiding unnecessary data sprawl and cost.

## Copy Data Analytics

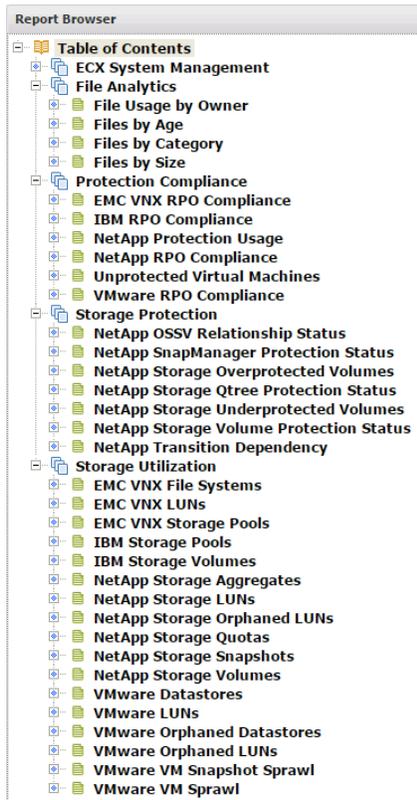


Figure 26. Reports

We can also run Analytics on the data collected within the ECX catalog. The Report tab provides access to several key Reporting options, as shown in **Error! Reference source not found.**, which could be leveraged for day to day activities.

We will go through a couple of reports in particular. Figure 27, shows Application RPO compliance, which shows the Oracle objects and the relation to underlying storage systems. It also displays the recovery points in New York and in San Francisco.

## Application RPO Compliance

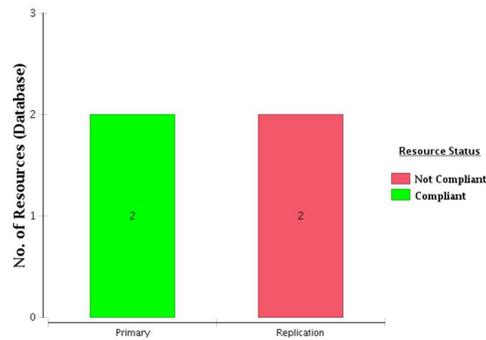
### Report Properties

Creation Date: September 26, 2016 8:07:22 PM UTC  
Report Generated By: admin

#### Report Filters:

Application Server: (All)  
Protection Type: (All)  
Display Resources That Are: Compliant, Not Compliant  
RPO Older Than: 1 Day

### Quick View



### Compliant for Primary Protection

Search:

Database	Application Server (Type)	Oracle Home / SQL Instance	Policy Name	Last Successful Protection Time	Compliance Time Remaining
PUREDB2	OracleServer2 (Oracle)	sid-oradev-1 / OraDB12Home1	ApplicationSnapshotOraclePure	Sep 26, 2016, 2:42 PM	19 Hour(s)
PUREDB1	OracleServer2 (Oracle)	sid-oradev-1 / OraDB12Home1	ApplicationSnapshotOraclePure	Sep 26, 2016, 2:42 PM	19 Hour(s)

Total Databases: 2

Figure 27. Application RPO Compliance Report

## Summary

In this whitepaper we have demonstrated how to manage, orchestrate and analyze Copy Data in an IT environment deployed on Oracle and Pure Storage FlashArray infrastructure using Catalogic Software's intelligent Copy Data management platform, ECX. The document referred to other whitepapers and User Guides that document detailed set up and testing of important pieces.

The document includes details on installation and setup as well as the creation of Copy Data workflows and Use Data workflows to spin up instances of application consistent snaps onto various sites to be leveraged for multiple use cases.

IT now has a more powerful way to harness the value Pure Storage infrastructure, without adding complexity or requiring added expertise.

## Trademark Information

Catalogic is a registered trademark of Catalogic Software Inc. ECX is a trademark of Catalogic Software Inc. All other company and product names used herein may be the trademarks of their respective companies.

© 2016 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Oracle is the registered trademark of Oracle Corporation in the U.S. and other countries. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

VMware is a registered trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

## About the Author

---



Twitter: @p\_jagannathan

Prashant Jagannathan is a Technical Manager at Catalogic with 10 years of experience in the IT industry. He is part of the Business Development team at Catalogic and is responsible for working with the technical teams of strategic business partners and identifying joint solutions and potential integration in products. Prior to Catalogic, Prashant was a APAC Team lead in Syncsort Inc, where he managed the technical sales operations for data protection group in the Asia-Pacific region.



Twitter:@alexinfanzon59

Alex Infanzon is an Oracle Solutions Manager at Pure Storage with over 30 years of professional experience in consulting, pre-sales and management roles. He has worked with fortune 500 companies in different vertical industries, both in a selling and consulting capacity. Alex has extensive experience in architecting and implementing Business Intelligence, Data Management, Enterprise Information Integration, Fuzzy Matching and High Performance Computing solutions