



CATALOGIC ECX INTEGRATION WITH PURE STORAGE

for VMware Virtualization

April 2017



TABLE OF CONTENTS

INTRODUCTION	2
AUDIENCE	2
BENEFITS	2
PURE STORAGE SNAPSHOT OVERVIEW	3
CATALOGIC SOFTWARE OVERVIEW	3
DATA MANAGEMENT CHALLENGES	4
SYSTEM REQUIREMENTS	5
PRE-REQUISITES	6
SOLUTION OVERVIEW	6
INSTALLATION AND CONFIGURATION	7
OUTLINE OF STEPS	13
TEST	14
PRODUCTION	14
CLONE	15
USE CASE I – ARRAY-BASED SNAPSHOT OF VIRTUAL MACHINES	15
USECASE II – INSTANT RECOVERY OF VIRTUAL MACHINES TO ORIGINAL PRODUCTION ENVIRONMENT	27
USECASE III – INSTANT RECOVERY OF VIRTUAL MACHINES TO DEV-TEST ENVIRONMENT	34
SUMMARY	43
REFERENCES	43
ABOUT THE AUTHOR	44

INTRODUCTION

Organizations of all sizes need to modernize their IT processes in order to enable critical new use cases, such as operational automation, DevOps, and integration of system-of-record data with scalable compute. They are equally challenged with improving management efficiencies for long-established IT processes, such as data protection, disaster recovery, reporting, and business analytics.

The common factor across multiple IT use cases is gaining access to copies of data. How best to get this done can be considered the data management challenge. Typically, access to data (such as on an enterprise storage array) or systems that connect to data (such as a virtual machine) require a lengthy request and provisioning process passed through gatekeepers in the IT organization. Data consumers can wait days or even weeks to get the data they request, which can limit their effectiveness.

Catalogic Software's® ECX™ Copy Data Management platform, deployed in conjunction with the Pure Storage FlashArray, allows clients to manage, orchestrate, and analyze data copies, providing full lifecycle management of your Copy Data. ECX provides automated workflows that allow you to streamline the creation and use of data copies for multiple business solutions, such as application data protection, automated DR, Dev-Test infrastructure provisioning, and DevOps in the hybrid cloud.

This document is intended to discuss the specific ways to manage, orchestrate, and analyze Copy Data in VMware® infrastructure 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 Copy Data 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.

AUDIENCE

The target audience for this document includes Pure Storage System Engineers (SEs), Technical Support Engineers (TSEs), Certified Partners system engineers, VMware administrators, and others who deploy and support Pure Storage FlashArray as storage for VMware virtual machines.

BENEFITS

There are many advantages to deploying Catalogic ECX with Pure Storage for VMware. Some of them are listed below:

- Create VMware-aware snapshots in the storage array with minimal impact on production resources
- Provision development environments which are periodically refreshed based on a routine schedule or on demand
- Populate test environments with near real-time production data
- Quickly perform root cause analysis using a copy of production data
- Test bug-fixes during development cycles

PURE STORAGE SNAPSHOT OVERVIEW

With Pure Storage FlashArray, 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 snapshots. These snapshots are an intrinsic part of the Purity Operating Environment: they are 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 to create a consistent point-in-time snapshot of all 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 snapshots 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 as well as remote data protection and recovery.

CATALOGIC SOFTWARE OVERVIEW

With ECX, Catalogic Software delivers the first “in place” Copy Data Management platform that works with your existing storage environment, driving operational efficiencies and cost savings and providing 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 existing tools and scripts to manage these copies, IT is often unable to meet commitments to business functions 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 metadata 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 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

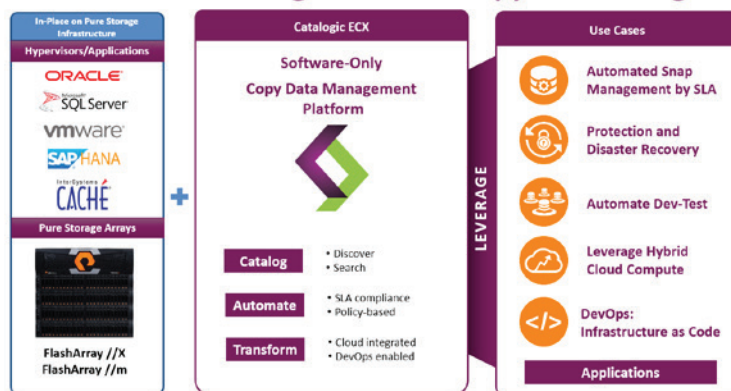


Figure 1. Catalogic Software overview

DATA MANAGEMENT CHALLENGES

Pure Storage FlashRecover Snapshots deliver superior space efficiency, high scalability, and simplicity. Catalogic ECX compliments Pure Storage by providing centralized management of these snapshots.

The process of making copies and then utilizing data copies is a very challenging and time-consuming process. For example, a typical data protection (DP) and disaster recovery (DR) solution involves two sites. These sites are normally in geographically different locations to assure data resiliency. Additionally, in the remote location, there is a mirror image (or a close resemblance thereof) of the physical infrastructure in the primary location. This is done to ensure that the business can meet all of its recovery service-level agreements (SLAs) for each of the applications in the event of a disaster.

A more granular breakdown of the steps involved in a data protection, or DR, scenario would involve the following operations:

- Use traditional backup utilities to move a large amount of data to secondary location, often a secondary disk storage device, Virtual Tape library or physical tapes
- Pre-synchronize storage resources between protected and recovery sites
- When it's time for restores or DR testing, shut down the virtual machines at the protected site and prepare them for migration to the recovery site
- Synchronize storage resources between protected and recovery sites
- Suspend non-critical virtual machines at recovery site to make room for migrated virtual machines from the protected site
- Change storage access at recovery site to “writable”
- Power on the virtual machines at the recovery sites (these virtual machines can be powered on in a pre-determined priority order if necessary)
- Just to test this DR setup, the process typically takes a full weekend and the process is driven by multiple scripts that need to be continually updated

Similarly, Development and Test environments pose their own challenges. Pure Storage FlashArray, along with Catalogic ECX, simplifies Oracle® database copy management by enabling administrators to orchestrate application-consistent copy creation, and recover and clone databases in minutes instead of hours or days. ECX copy management leverages Pure Storage Purity//Protect snapshots to rapidly create, replicate, restore, or clone copies of Oracle databases which are space- and time-efficient. ECX enables IT organizations to focus on the copy and recovery requirements of the business rather than the technical details of the underlying storage platforms.

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

SYSTEM REQUIREMENTS

The following are the system requirements to deploy and run ECX with Pure Storage and VMware.

STORAGE

- Pure FlashArray series
- Purity v4.8.2 or later
- REST API v1.5 or later

VMWARE VERSION DEPLOYED

- vSphere 5.1, 5.5, or 6.0

LICENSING

- Pure Storage FlashArray requires no licensing fee for any of its software functionality
- ECX is licensed per-FlashArray and costs vary by array model

PROTOCOL SUPPORT

- Fibre Channel
- iSCSI

CATALOGIC SOFTWARE ECX VERSION

- 2.5 or above

PRE-REQUISITES

Ensure the following pre-requisites before walking through the configuration steps and use cases presented in this 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 activity between applications, virtual machines, and storage systems through API calls.
- 2. VMware Environment** – vSphere needs to be installed with vCenter server and ESXi hosts. This paper uses an ESXi 6.0 on UCS B series servers.
 - This paper uses virtual machines with different operating systems such as Microsoft Windows 2012 R2, Centos 7.2, and SuSE Linux.
- 3. Pure Storage Volume** – One or more volumes needs to be connected to vCenter servers. This paper uses two Pure Storage FlashArrays – //M50 R1 and //M50 R2. All vCenter instances are connected via Fiber Channel to the FlashArrays.
- 4. Security** – ECX requires the following privileges:
 - **Administrator Access to the vCenter Servers** This is required to access the virtual machines which are associated with the vCenter Server.

SOLUTION OVERVIEW

ECX delivers 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.

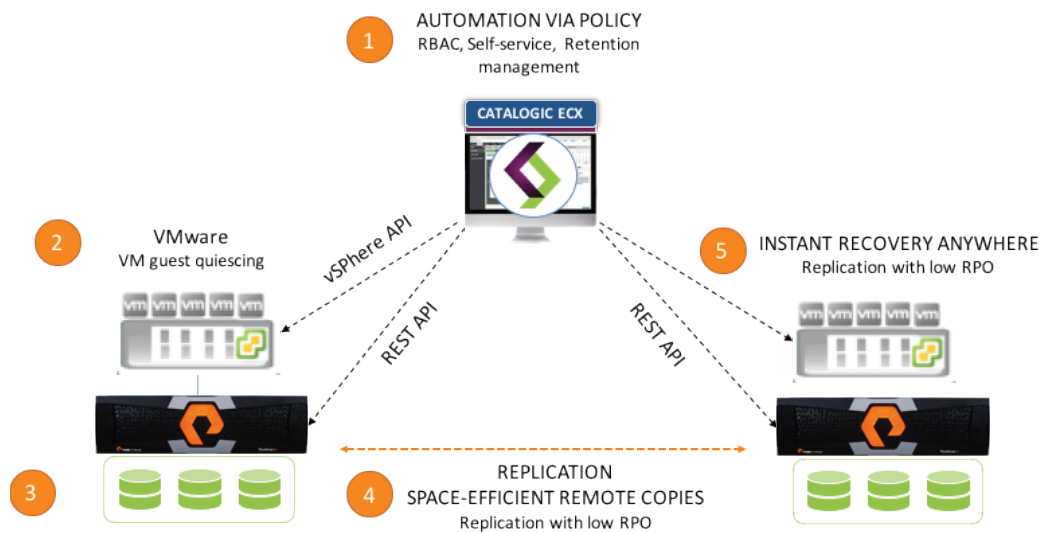
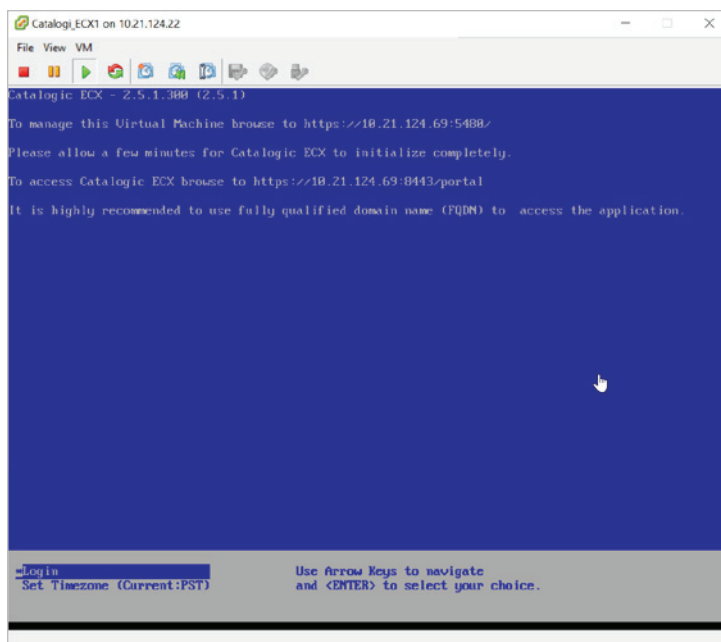
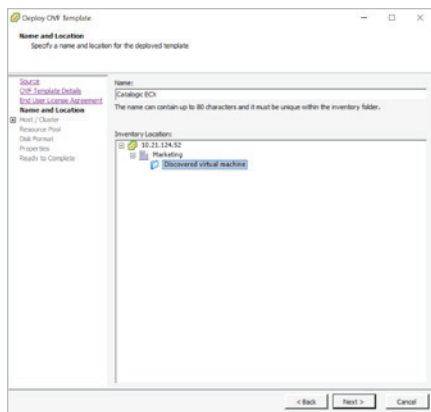


Figure 2. Topology overview

Catalogic ECX catalogs all 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 application data, enabling 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 Purity//Protect snapshots ensures that the tasks are completed in seconds without impacting performance.

INSTALLATION AND CONFIGURATION

ECX is packaged as a self-contained VMware OVA and installed 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 deploy the appliance, as shown in Figure 3. ECX comes pre-packaged with all the required software, and, once powered on, the console screen points to the link to the portal, as seen in Figure 4.



Once you've deployed the OVA, you will be requested to enter a license key. The key (ECX.lic) is applied to the appliance by logging into the management console.

Follow these instructions to apply the key:

1. From a web browser, enter the following URL to access the management console of the virtual machine where ECX is deployed:

https://(HOSTNAME):5480/ where (HOSTNAME) is the HOSTNAME or the IP address of the virtual machine where ECX is deployed.

2. In the login window, enter the user name and password to access the management console of the virtual machine:
User Name: administrator
Password: ecxadLG235
3. Click the **License** tab as shown in Figure 5.
4. In the **License** field, browse for **ECX.lic** on your computer, then click **Upload**.
You may have to authenticate again with username and password.

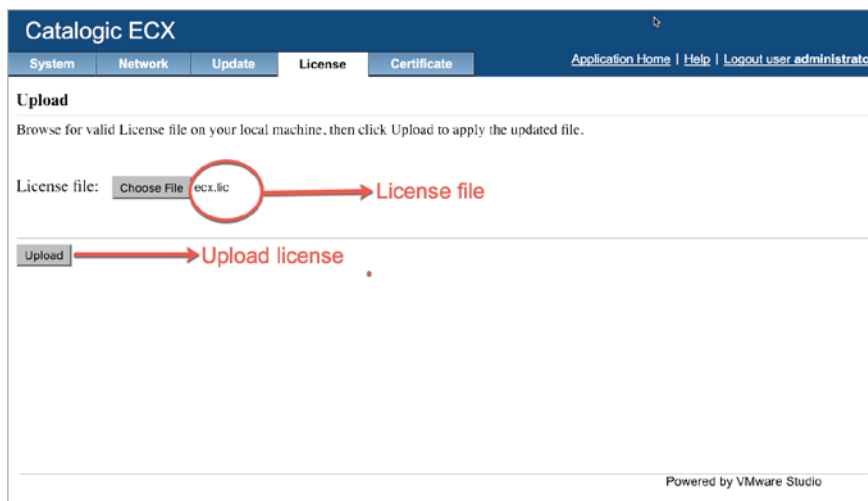


Figure 5. License key upload

Once the license key is uploaded Select the **Application Home** link to access the Catalogic ECX portal. Alternatively, the portal link can also be accessed by entering the URL **https://(HOSTNAME):5480/** where (HOSTNAME) is the HOSTNAME or the IP address of the virtual machine where ECX is deployed. Figure 6 shows the login dialog to access the application. Login using credentials

User Name: admin

Password: password

You will be prompted to change the password on first login.




Figure 6. ECX web portal login

The first step in configuration is to register the storage assets and vCenter resources and assign them to their respective sites. This is a simple and intuitive one-time, agentless

registration process. ECX has a site-based concept to identify resource location and the capability of selecting a data copy to be used from any of the sites. To create a site, click on the **Configure** tab, select **Sites**, click **New**, and enter the Site name as shown in Figure 7. For this paper, the setup has two sites: **Primary** and **Secondary**.

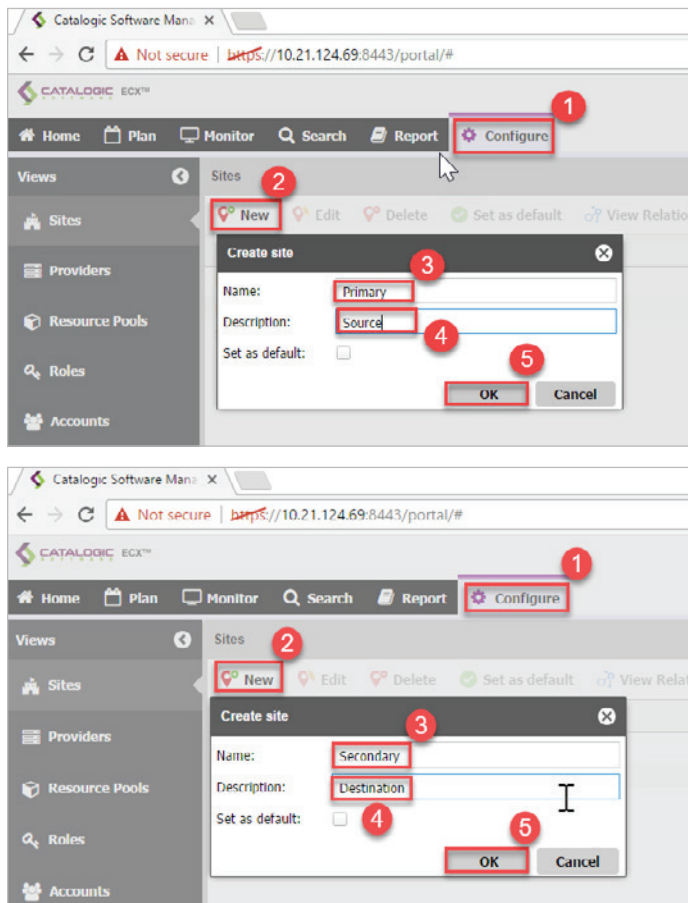


Figure 7. Create Primary and Secondary Sites

To add VMware and Pure Storage FlashArray resources, click the **Providers** sub-tab in the **Configure** tab, then right-click and select **Register**, as shown in Figure 8. Select the two Pure Storage FlashArray//M50 systems in the **Primary** and **Secondary** sites and fill in the details, as shown in Figure 8. The registration dialog box takes in credentials and a site selection.

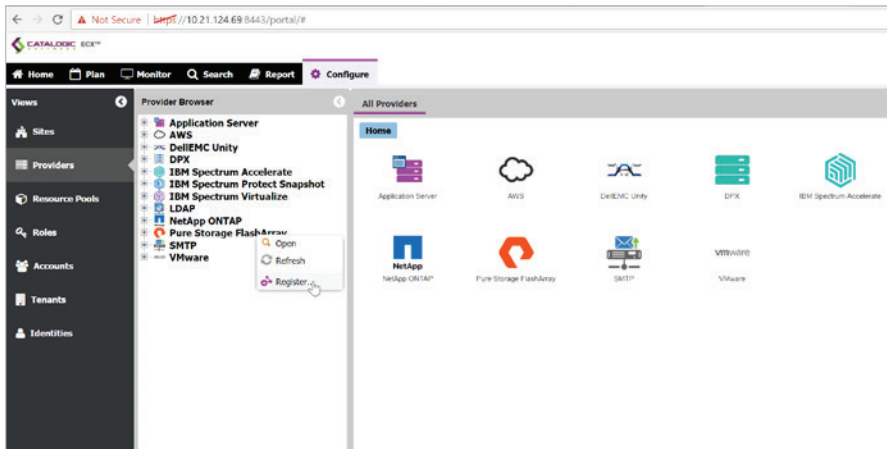


Figure 8. Configure Pure Storage FlashArray

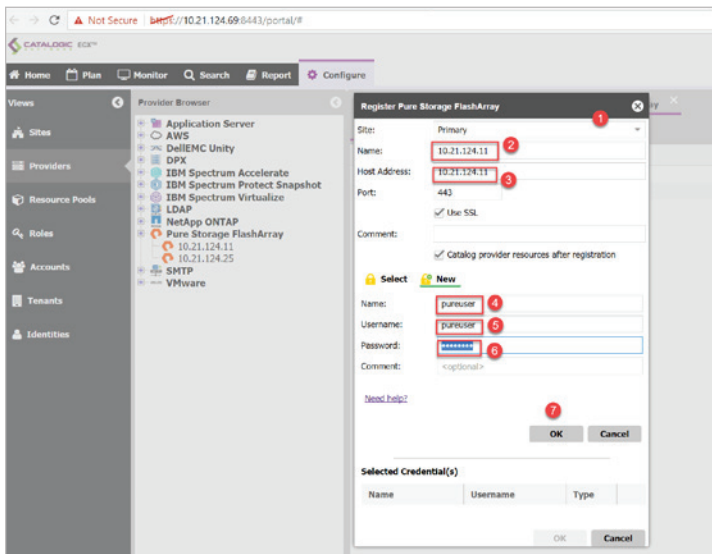


Figure 9. Register Pure Storage FlashArray

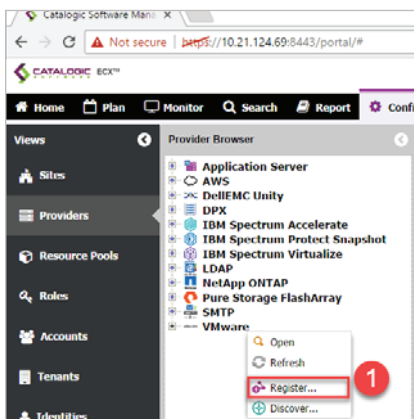


Figure 10. Register VMware (vCenter)

Similarly, right click on VMware, as shown in Figure 10, and register the two vCenter servers in the **Primary** and **Secondary** sites, as shown in Figure 11.

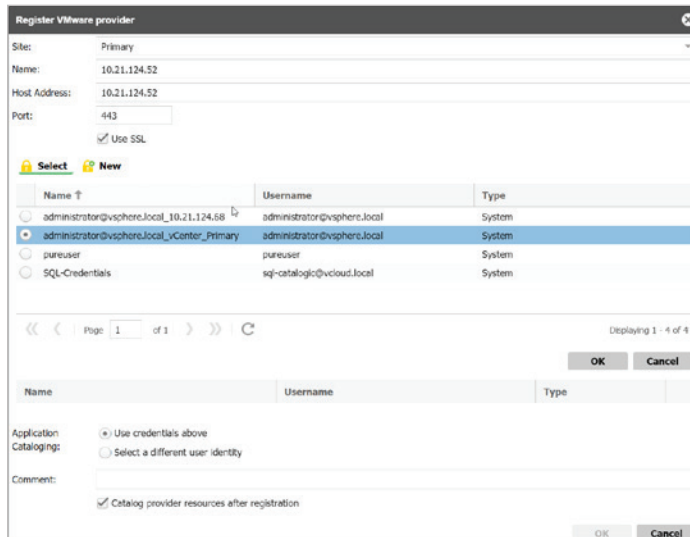


Figure 11. Register VMware Provider

Once the resources are registered, ECX will automatically discover and catalog objects like the storage volumes, virtual machines, and datastores, and identify the mapping between these resources.

Figure 12 shows the **Plan** page, where all the core management functionalities of ECX are performed. The **Plan** page is where the Copy and Use Data policies are created. We will explain them in detail in the upcoming sections.

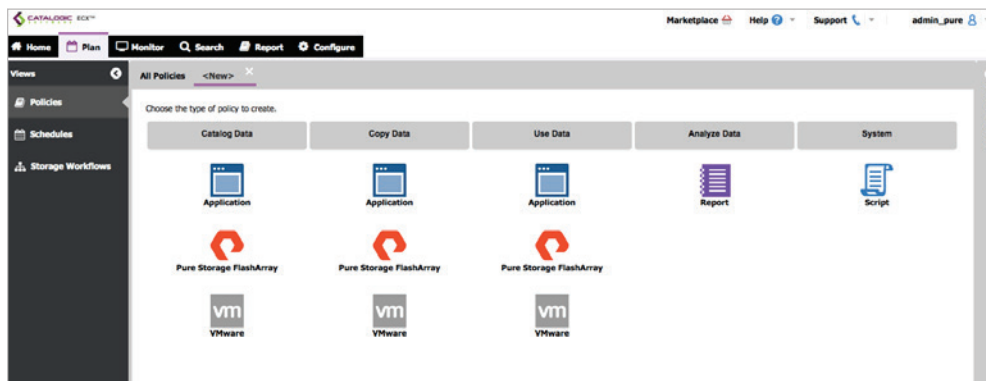


Figure 12. ECX copy management plan screen

OUTLINE OF STEPS

The orchestration and automation in ECX involves two core operations, Copy Data Policies and Use Data Policies.

Copy Data policies enable application-consistent, in-place snapshots and array-based replication. The first step in a Copy Data policy is to use VMware APIs to identify the virtual machines that need protection. As a next step, ECX automatically identifies the mappings between the virtual machines, the datastores that they are hosted on, and the underlying Pure Storage volumes. ECX then invokes VMware tools to take a virtual machine snapshot temporarily, and a Purity//Protect snapshot is created on the Pure Storage array. The virtual machine snapshot is then removed. Additional options like pre/post scripts, concurrency, etc., can be added to the policy. Templates can be customized with storage workflow features. Storage workflows define the operations performed on a Pure Storage FlashArray driven by REST API calls from ECX.

Storage workflows are templates created by storage administrators which define storage operations. Storage workflows include Snapshots, Replication, and VMCopy. Figure 13 shows the Storage Workflow within VMware Copy Data Policy.

Snapshots – This option drives in-place Pure Storage Purity//Protect snapshots. Purity//Protect snapshots are fast and deliver superior space efficiency and scalability.

Replication – This option is used when data has to be replicated to another Pure Storage array. Pure Storage array-based replication is based on snapshot technology which does async replication between two FlashArrays. In Pure Storage array-based replication, users can define the replication target, frequency, storage quotas, and bandwidth throttling options.

VMCopy – The above two options are driven at the storage level. VMCopy drives VMware-based data transfer using VADP (VMware vStorage API for Data Protection). This option is used to back up the VMs from the vSphere hosted on any storage to a Pure Storage FlashArray volume.

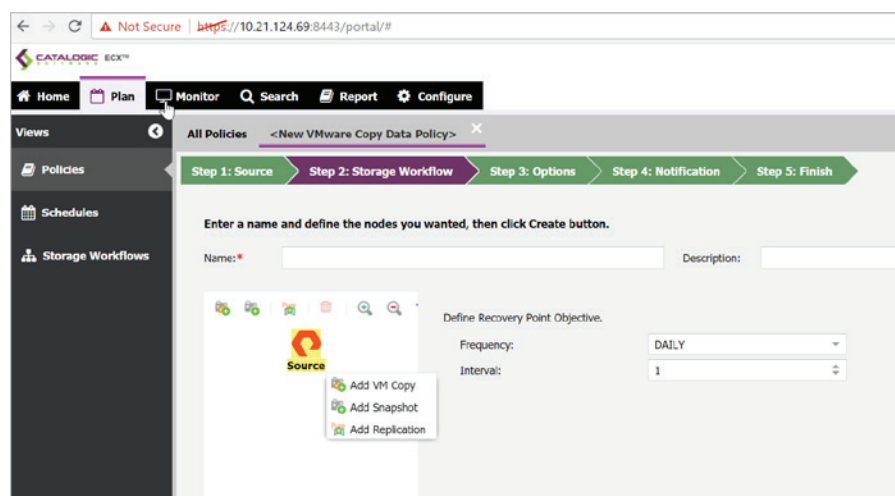


Figure 13. Storage Workflows within Copy Data policy

Use Data policies leverage these snapshot copies of production data and make them available instantly for a multitude of use-cases, including disaster recovery, dev-test automation, and improved DevOps. Use Data policies can be created for individual virtual machines or multiple virtual machines and assigned to end users based on Roles Based Access Control (RBAC). Figure 14 shows the VMware Use Data Policy.

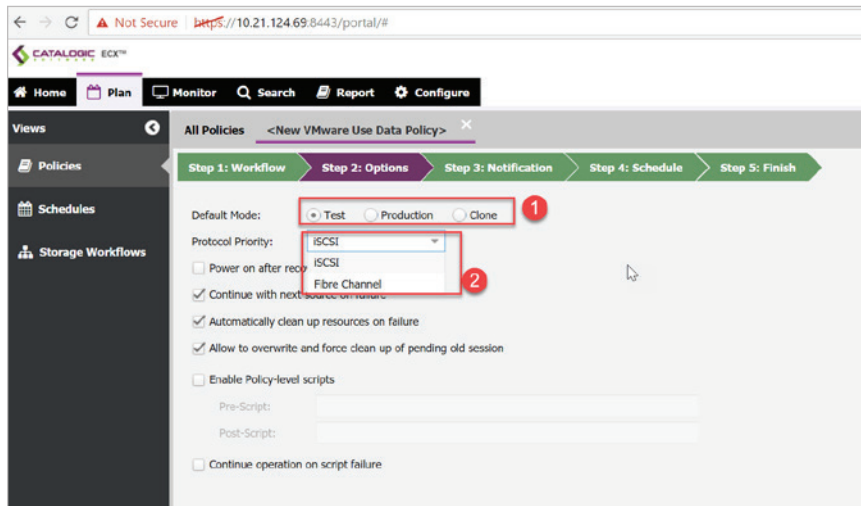


Figure 14. Creating a Use Data Policy

Users can recover their Virtual Machines in different modes depending on their requirements. The three modes are as follows:

TEST

Test Mode creates virtual machines for development/testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. In Test Mode, the Virtual machines run from a snapshot volume. Users can test and verify changes and have the option to either promote the virtual machine to Production or cleanup after testing and verification is completed. Through fenced networking, you can establish a safe environment to test policies without interfering with production. Virtual machines created through Test Mode are also given unique names and identifiers to avoid conflicts within the production environment.

PRODUCTION

Production Mode is used for disaster recovery scenarios at the local site from primary storage or a remote disaster recovery site, replacing original virtual machines. As part of the recovery process, Production mode retains all configurations of the virtual machine, including names and unique identifiers. All copy data policies associated with the virtual machine continue to run as per the original schedule associated with it.

CLONE

Clone Mode creates copies of virtual machines for scenarios requiring permanent or long-running copies of a test environment in a fenced network. Virtual machines created using Clone mode are also given unique names and identifiers to avoid conflicts with the production environment. With clone mode, users need to ensure that they have enough storage capacity, since clone mode copies data from a snapshot to a vSphere datastore.

USE CASE I – ARRAY-BASED SNAPSHOT OF VIRTUAL MACHINES

In this scenario, we will snapshot and replicate virtual machines using ECX integration with Pure Storage FlashArray. You can snapshot VMware data, including virtual machines, datastores, folders, vApps, and datacenters, with snapshots using a VMware Copy Data policy. For this use case, we will use two virtual machines (1 Windows and 1 CentOS) in vCenter_primary.

To create a VMware Copy Data Policy, as shown in Figure 15.

1. Click the **Plan** tab. On the **Views** pane, select **Policies**.
2. Click the **All Policies** tab. Click **New**, then select **VMware** in the **Copy Data** column, as shown in Figure 16. The VMware Copy Data Policy editor opens, as shown in Figure 17.

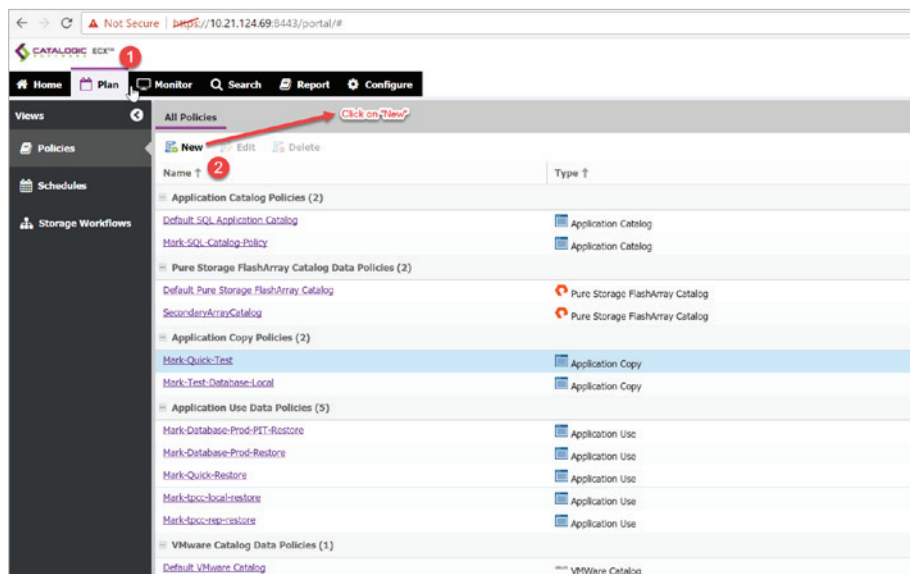


Figure 15. Creating Copy Data Policy

VMware Copy Data Policy editor opens as shown in Figure 16.

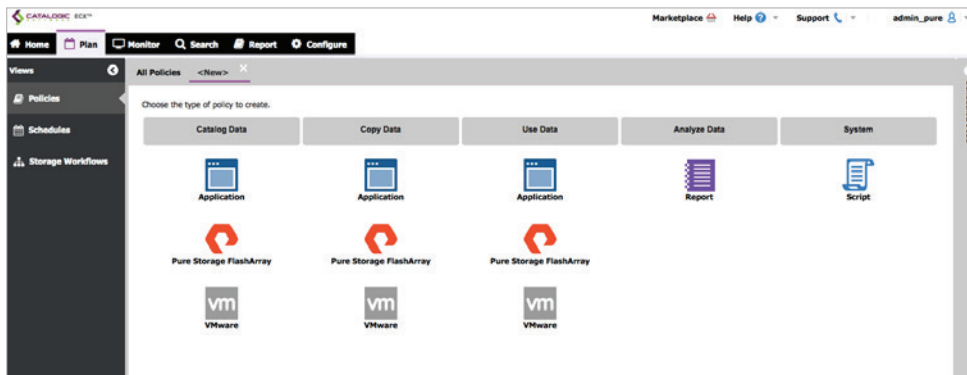


Figure 16. Creating Copy Data Policy

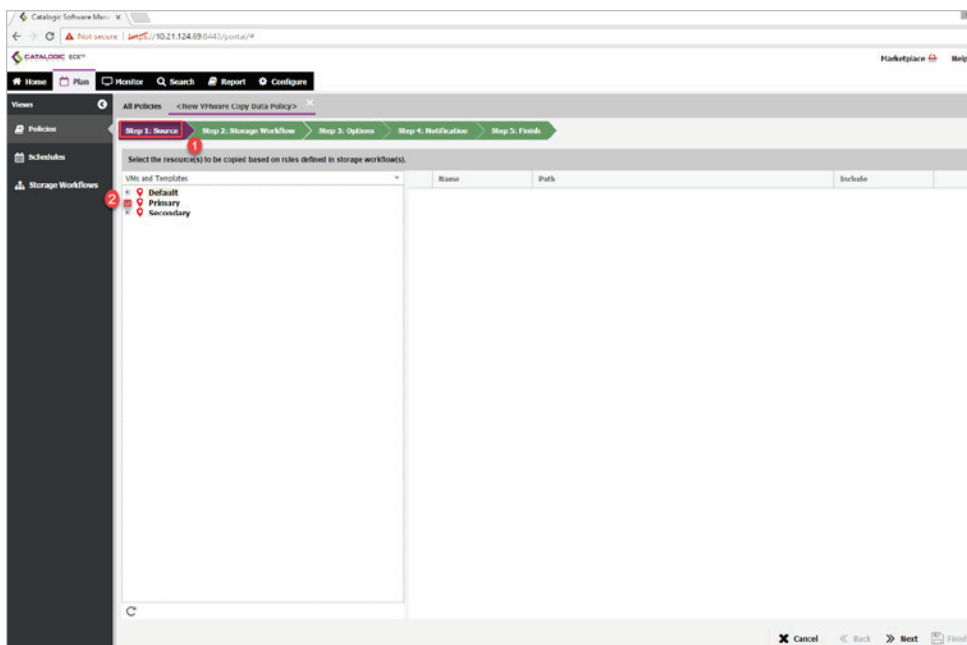


Figure 17. Creating Copy Data Policy

3. Click the **Source** tab, as shown in Figure 17, and expand the resource under **Primary**. From the drop-down menu, select **VMs** or **Storage**. From the list of sites, expand the site **Primary**, browse **vCenter_primary Site**, select the two VMs **Prod1** and **Prod2**, and then click **Next**, as shown in Figure 18.

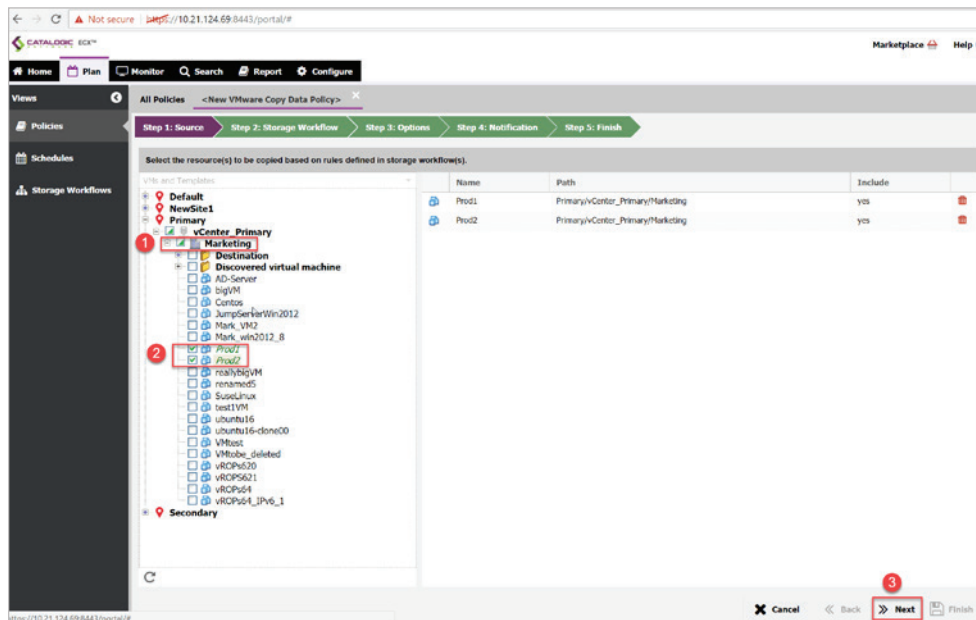


Figure 18. Creating Copy Data Policy

- Click the **Storage Workflow** tab and then click on **Create a new workflow**, as shown in Figure 19. Storage admins can create multiple templates based on SLA requirements. Click on **Create a New Storage Workflow**.

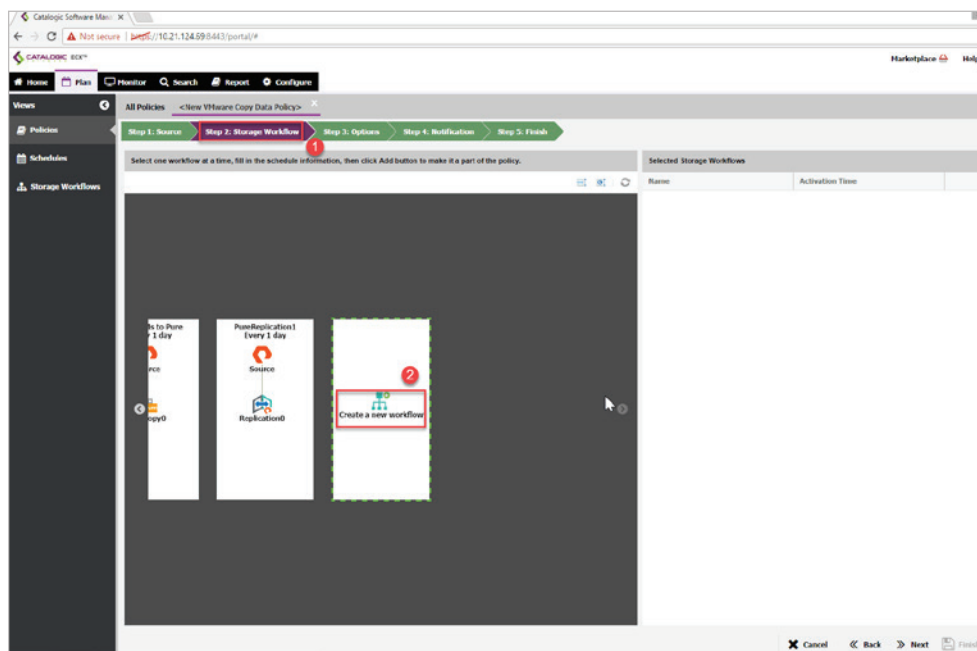


Figure 19. Creating Copy Data Policy

You will be directed to the **New VMware Copy Policy** tab, as shown in Figure 20.

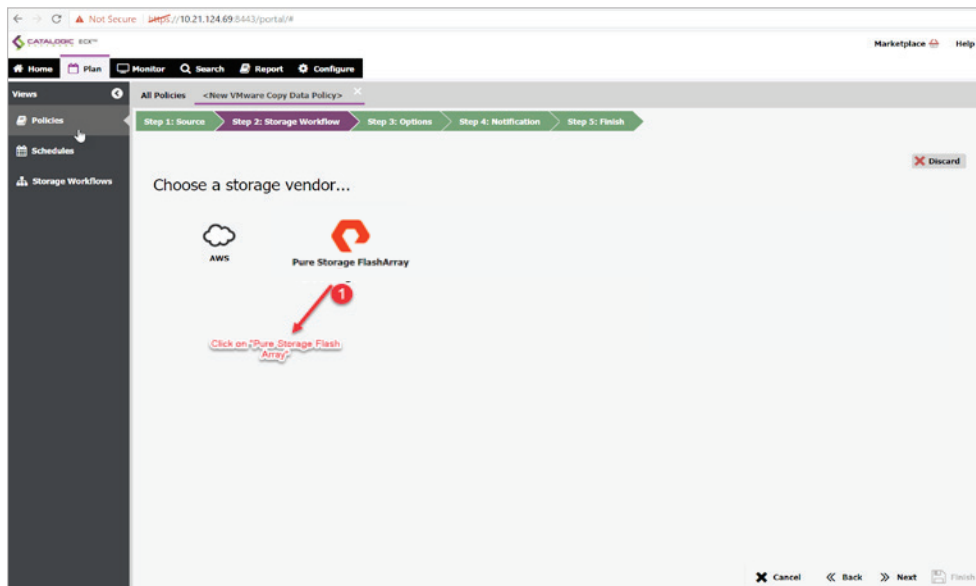


Figure 20. Creating Copy Data Policy

Click on the **Pure Storage FlashArray** icon, as shown on Figure 20. You will be directed to the page, as shown in Figure 21. This gives you three options: **Add Snapshot**, **Add Replication**, and **VM Copy**, also shown in Figure 21.

Click on **Add Snapshot** since we will be creating the Storage workflow using Add Snapshot functionality. You can define snapshot retention based on the service level agreements of the organization.

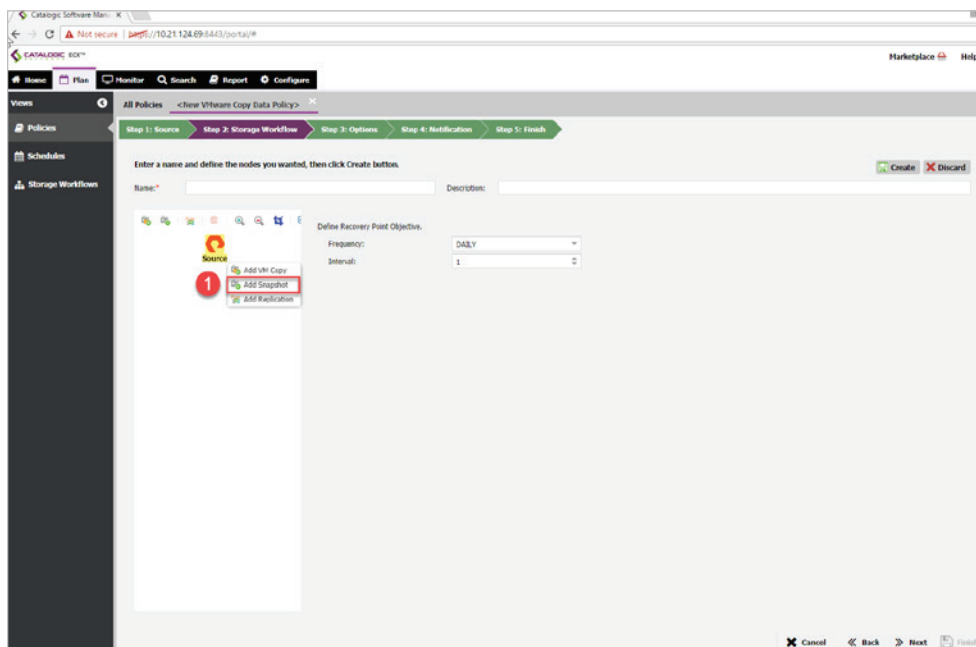


Figure 21. Creating Copy Data Policy

You will be directed to the page, as shown in Figure 22. In this page, enter the name of the **Storage Workflow**, “PureSnapshot”, as shown in Step 1, Figure 22, and an optional Snapshot Prefix label, as shown in Step 2. Click the **Create** button, as shown in Step 3.

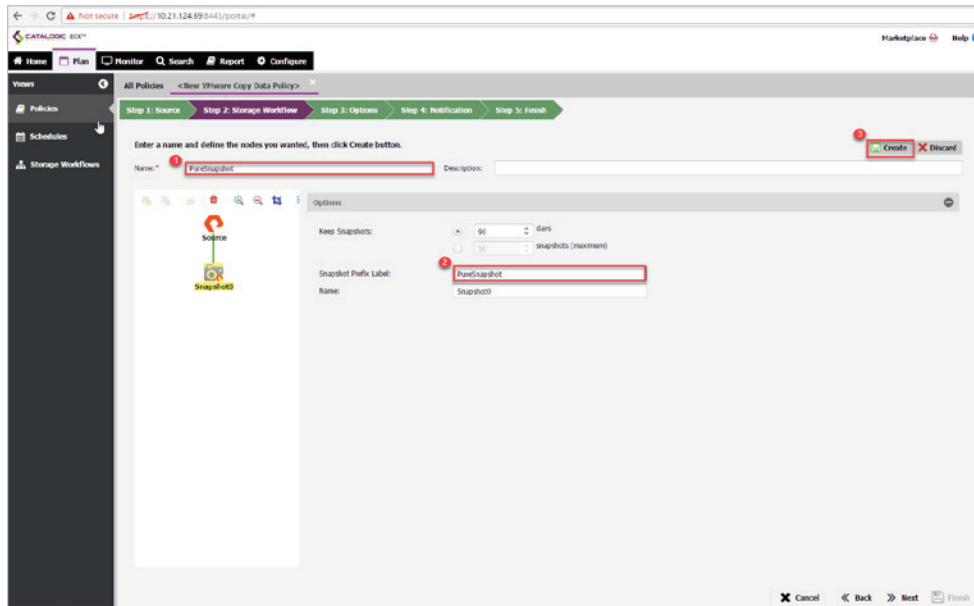


Figure 22. Creating Copy Data Policy

Similarly, we will create two more storage workflows, namely, **Add VMCopy** and **Add Replication**.

To create a VMCopy workflow, select **Add VMcopy**, as shown in Figure 23.

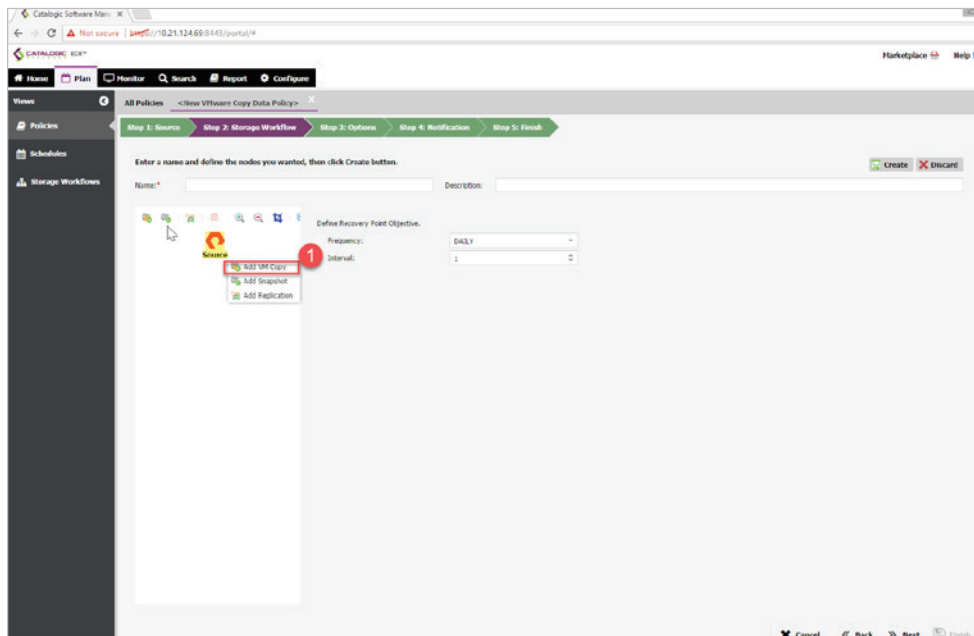


Figure 23. Creating Storage Workflow “Copy VMs to Pure”

Enter the name “Copy VMs to Pure”, as shown in Figure 24. Browse the **Secondary** and select the FlashArray registered in the site.

As shown in Figure 24, you can also change the retention period of snapshots, and optionally add **Target Volume Prefix Label** and **Snapshot Prefix Label**, as shown in Figure 24, Steps 4 and 5, respectively. Select the appropriate protocol in the Storage Workflow – which is mandatory and important. In this case, we selected **Fibre Channel** from the drop-down menu, as shown in Step 6, Figure 24. For this paper, we use the **Clone Full Copy Method** which will take advantage of the faster Pure Storage array-based replication as compared to **VADP**, which is VMware-based replication, as shown in Step 7. Lastly, click on **Create**, as shown in Step 8.

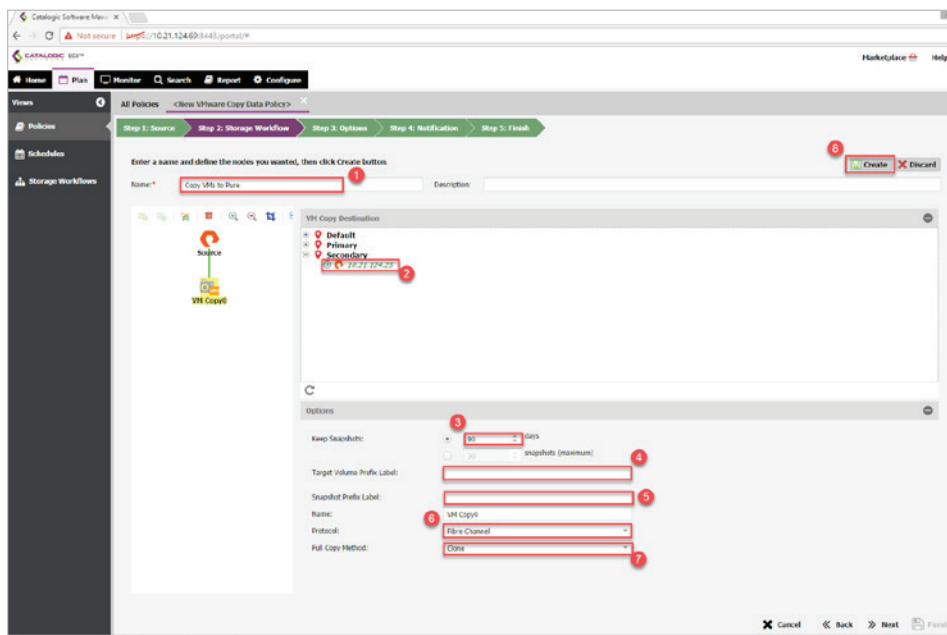


Figure 24. Creating Storage Workflow “Copy VMs to Pure”

To create a Replication Workflow, select **Add Replication**, as shown in Figure 25.

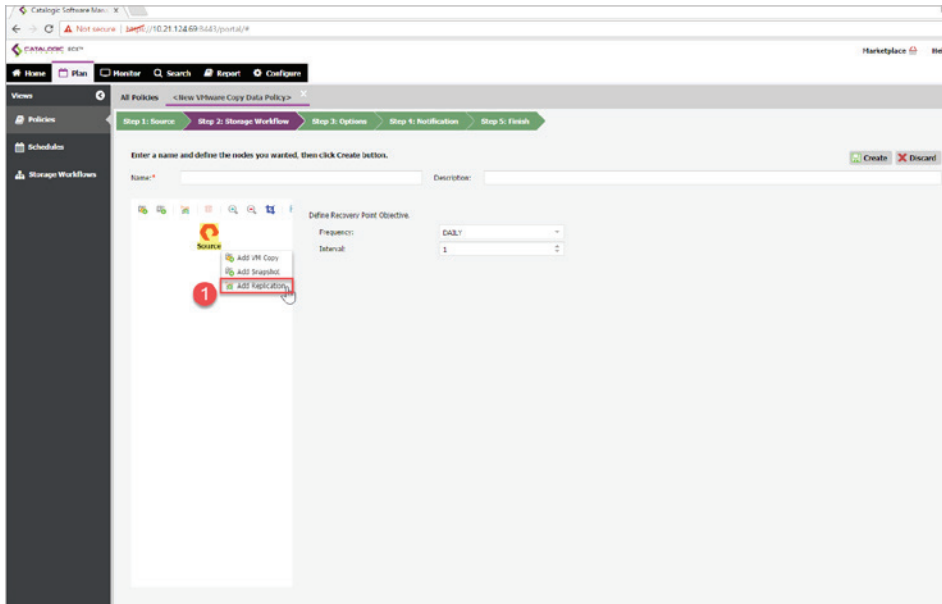


Figure 25. Creating Storage Workflow “PureReplication”

You will be directed to the **Storage Workflow** page. Name the workflow “PureReplication”, as shown in Figure 26. Browse the Site **Secondary** and select the FlashArray registered as the replication destination. If desired, set the retention policy for the snapshot, as shown in Step 2. Here we will use the default of 90 days, also shown in Step 2. We can also set the destination retention time for snapshots, as shown in Step 3. Lastly, we can optionally add the **Name** for snapshots and snapshot prefix label, as shown in Step 4 and Step 5. Click on **Create**, as shown in Step 6, to add the storage workflow to your copy data policy.

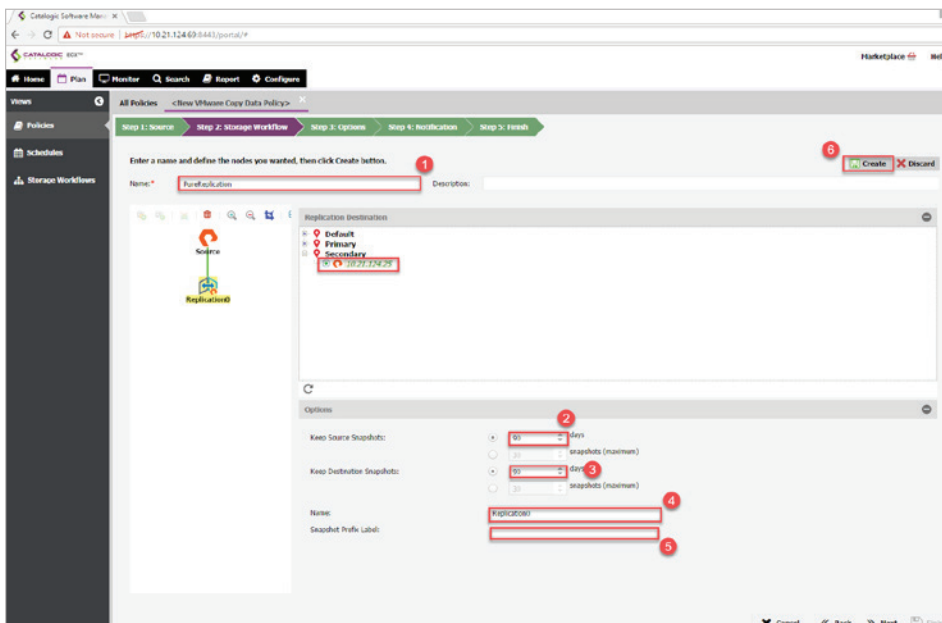


Figure 26. Creating Storage Workflow “PureReplication”

Once the storage workflows are created, click **Add** on the workflow to apply them to the policy, as shown in Step 1, Figure 27.

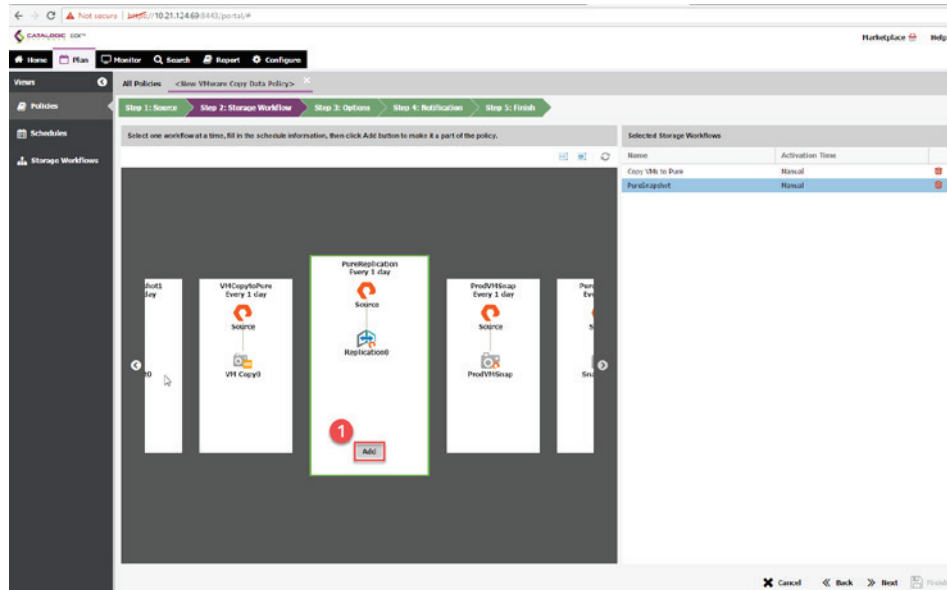


Figure 27. Adding Storage Workflow to Copy Data Policy

You can set a time to run the Storage Workflow with the **Specify Activation Time** trigger, as shown in Step 1, Figure 28. Enter the day and time to activate the Storage Workflow, as shown in Step 2 and Step 3. You can optionally select **Activate Manually** to invoke the Storage Policy to run at any time you want. If configuring more than one Storage Workflow in a policy, use the **Same as workflow** option to trigger multiple Storage Workflows to run at the same time.

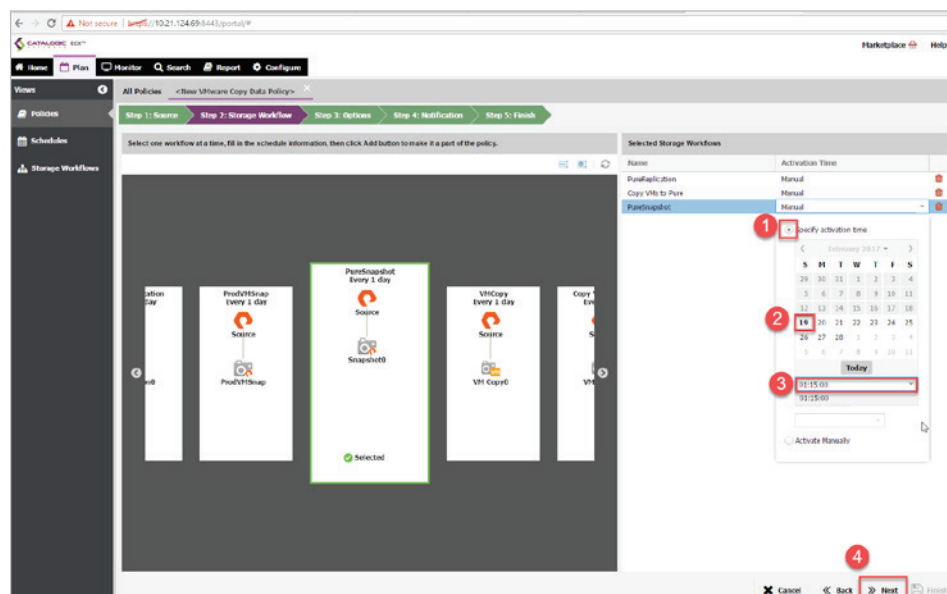


Figure 28. Adding activation time to Storage Workflow "PureSnapshot"

*Note: Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow. You can see the activation time of the **PureReplication Workflow** in Figure 29. Once the Storage Workflow and its activation time are selected, as shown in Figure 29, click on **Next** to proceed to the next step.*

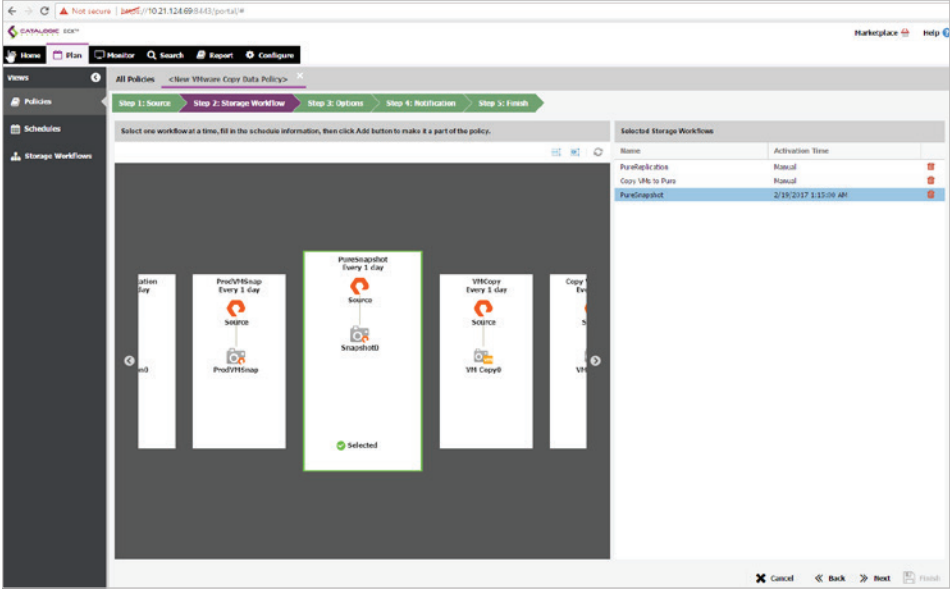


Figure 29. Adding activation time to Storage Workflow “PureSnapshot”

Click the **Options** tab. Set the policy options. We will set the following:

- a. **Create VM snapshots for all VMs**, as shown in Step 1, Figure 30 – This option is to configure virtual machine snapshots. For this use case, we will select **Make all virtual machines application/file system consistent**, as shown in Step 2. Application-consistent copy data captures data in memory and transactions in process.
- b. Click on **Next** to go to the next step.

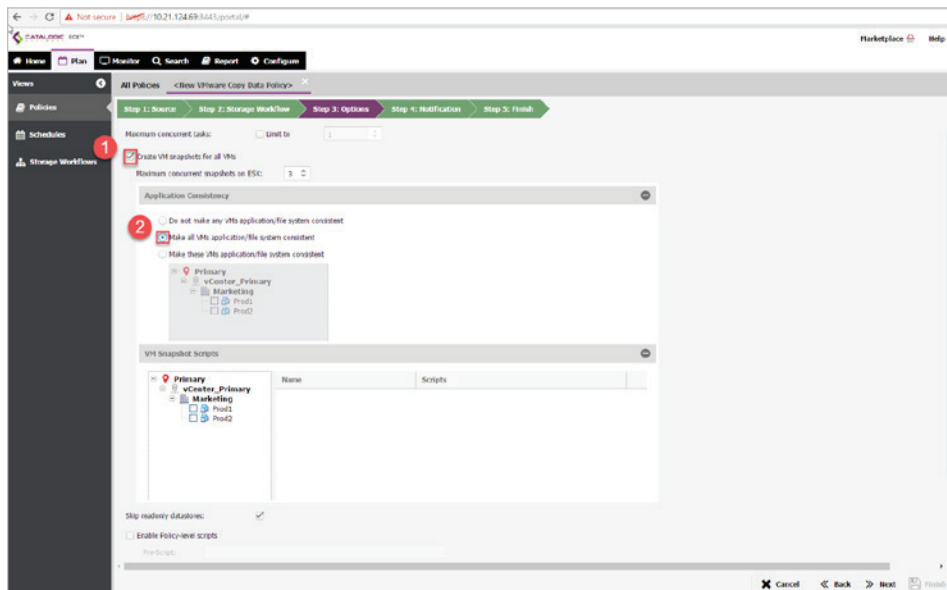


Figure 30. Options in Copy Data Policy

Click the **Notification** tab, as shown in Figure 31. Users can select an SMTP Server from a list of available SMTP resources. Since we have not configured any SMTP server, we will skip this step. Click on **Next**.

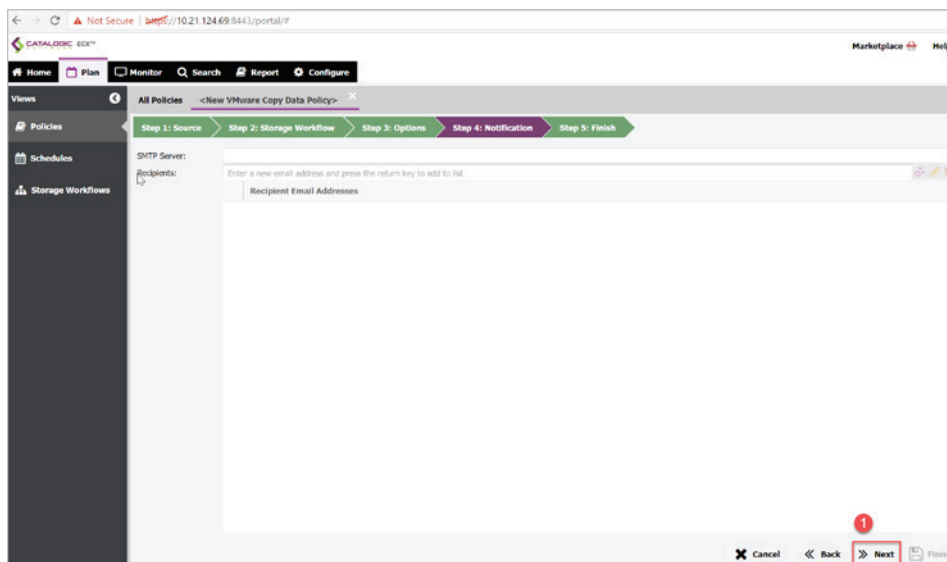


Figure 31. Notifications in Copy Data Policy

Click the **Finish** tab, as shown in Figure 32. Enter a name for your policy, as shown in Step 1, and optionally a meaningful description. We will use the name “ProdVMCopy” for this policy. Click **Finish**, as shown in Step 2.

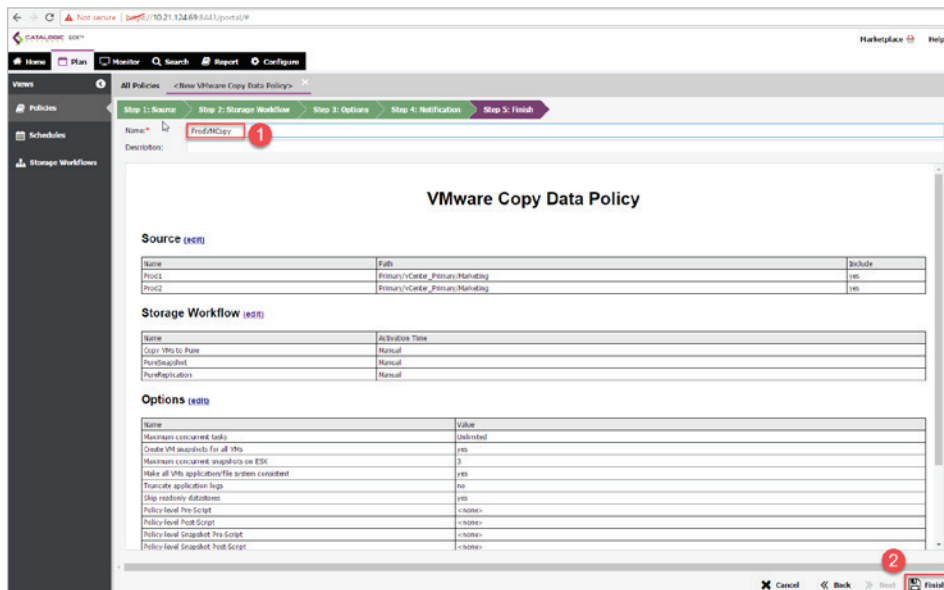


Figure 32. Copy Data Policy “ProdVMCopy”

The policy runs as defined by your triggers, or can be run manually from the **Monitor** tab. We can also add or remove the virtual machines in the future if we want, since this policy is a template which can be modified at any time.

Follow the steps to execute the Copy Data Policy “ProdVMCopy”, as shown in Figure 39.

1. Click on the **Monitor** tab, as shown in Step 1, Figure 33.
2. In the **Views** pane, select **Jobs**, as shown Step 2.
3. Select the job **ProdVMCopy** to run by clicking in the row containing the job name, as shown in Step 3.
4. Click **Start**.
5. You will be prompted to select a Storage Workflow which will be used in this Copy Data Policy. As shown in Figure 34, we can select any Storage Workflow to be executed with the Copy Data Policy. We selected the **Pure Snapshot1** Storage Workflow as shown in Step 1. Then click **OK**.

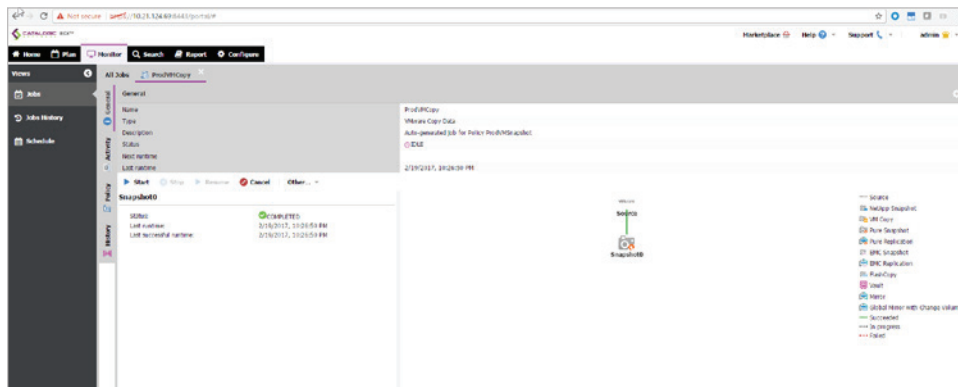


Figure 36. Run Copy Data Policy

If you would like to see the detailed activity of the execution of your policy, click on **Activity**, as shown in Figure 37.

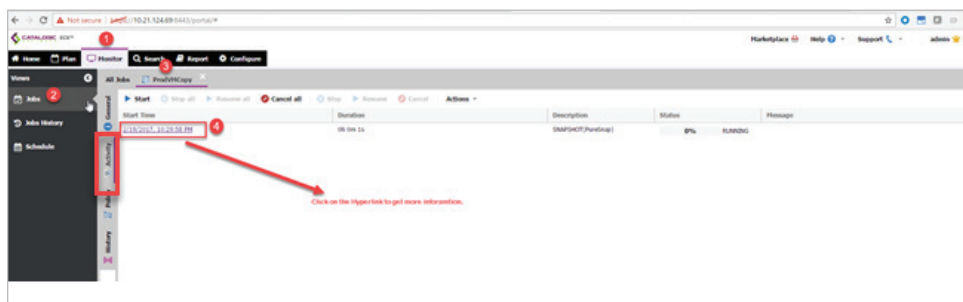


Figure 37. Run Copy Data Policy

USECASE II – INSTANT RECOVERY OF VIRTUAL MACHINES TO ORIGINAL PRODUCTION ENVIRONMENT

In ECX, recovery workflows can use a local copy or a replicated copy to spin up virtual machines in any VMware infrastructure. These workflows allow business lines to access data copies for multiple use cases. In this use case, we will go through an example of restoring the two virtual machines (**Prod1** and **Prod2**) that were copied in the previous section to the original vCenter (**vCenter_primary**). To perform this restore, we will create a Use Data policy.

To create a Use Data policy

1. Click the **Plan** tab, as shown in Step 1, Figure 38.

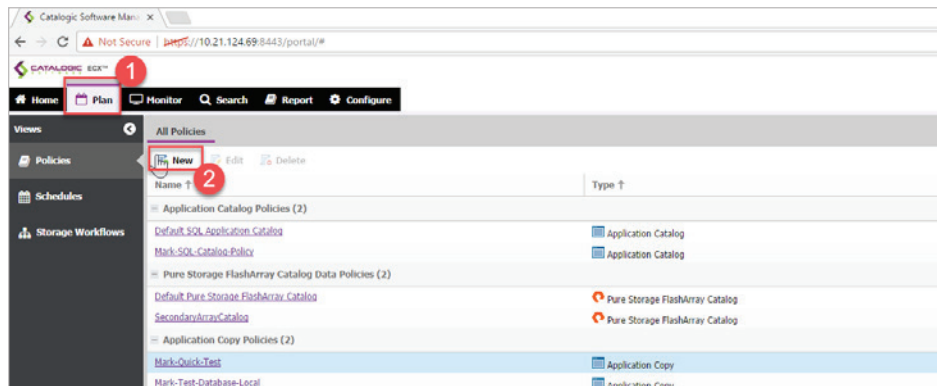


Figure 38. Run Copy Data Policy

2. In the **New** tab, click on **VMware** in the Use Data column, as shown in Step 1, Figure 39. You will be directed to the page, as shown in Figure 41 .

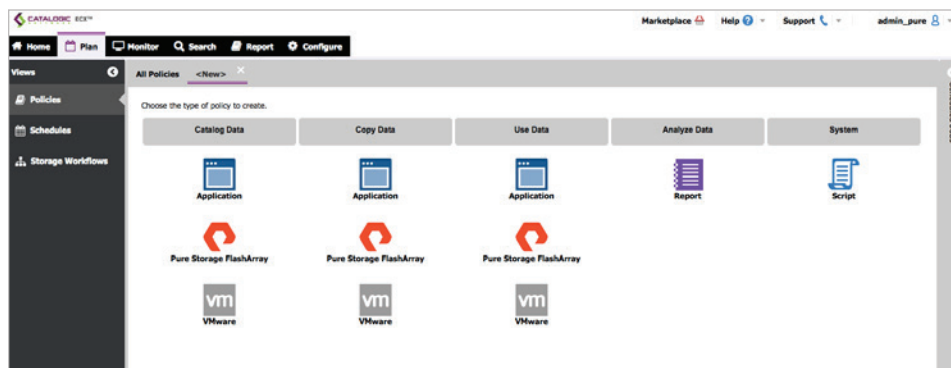


Figure 39. Run Copy Data Policy

In the **Workflow** tab, click the **Instant Virtualization** template, as shown in Step 1, Figure 40.

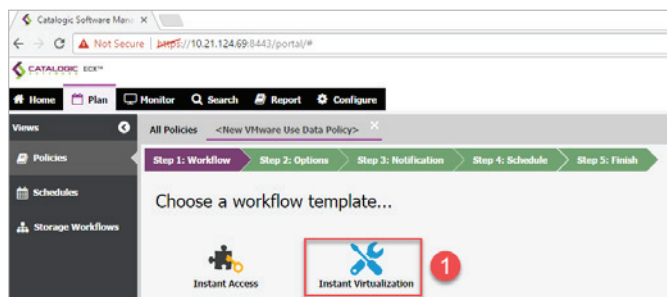


Figure 40. Create Use Data Policy

- Click **Source**. Browse the **Primary** site, as shown in Step 2, Figure 41, and select the VMs **Prod1** and **Prod2** in **vCenter_Primary**, as shown in Step 3. You can change the order in which the resources are recovered by dragging and dropping the resources in the grid. Click on **Copy**, as shown in Step 4.

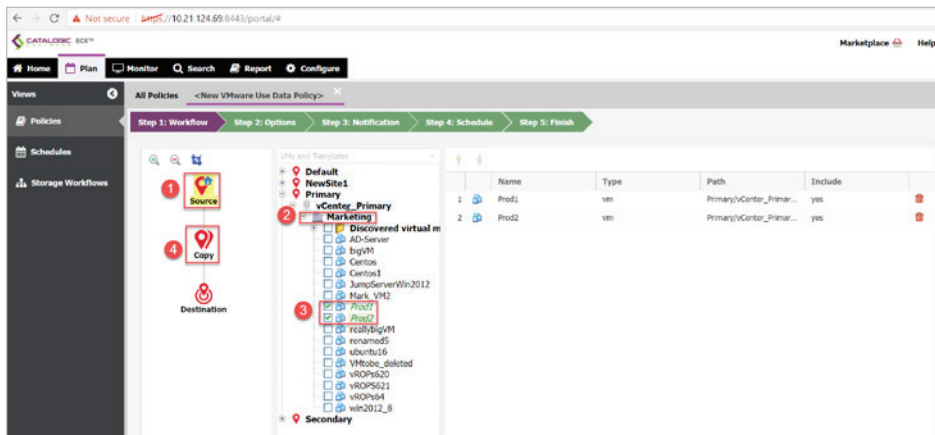


Figure 41. Create Use Data Policy

- Figure 41 shows all the sites that have copies of those virtual machines. We will select the site **Secondary**, as shown in Step 1, Figure 42. ECX catalog automatically tracks the lineage and location of the snapshots which were created by the Copy Data policy.

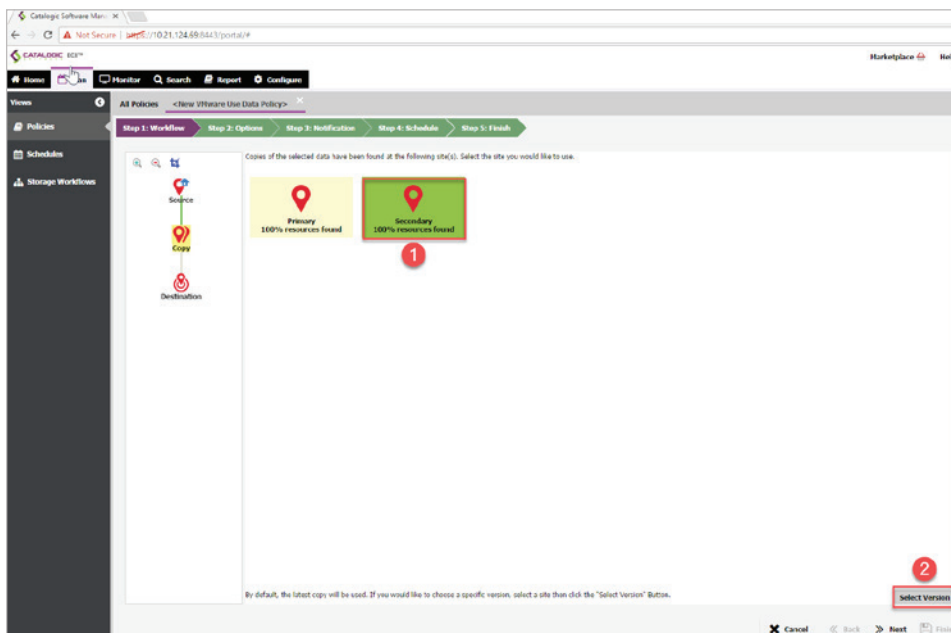


Figure 42. Create Use Data Policy

Figure 43. Create Use Data Policy

- Click **Destination**, as shown in Step 1, Figure 44. To restore to the original host or cluster, we select **Use original host or cluster** with system defined IP configuration, as shown in Step 2. You can also choose **Use alternative host or Cluster**.

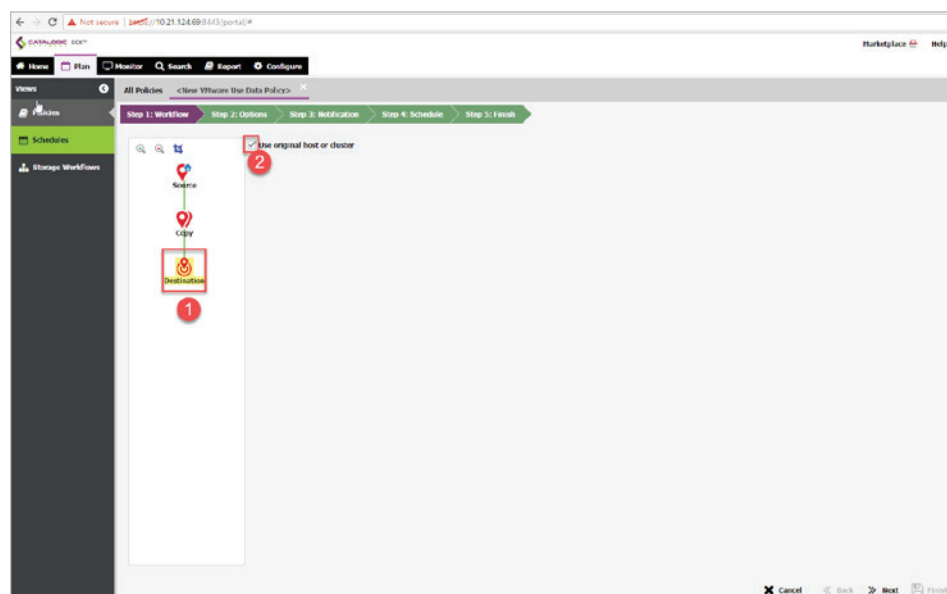


Figure 44. Create Use Data Policy

If you choose to use **Use alternative host or Cluster** by unchecking **Use original host or cluster** with system defined IP configuration, then you will have to fill in **Virtual Networks, Datastores, Subnets** and **VM Folders**, as shown in Figure 45.

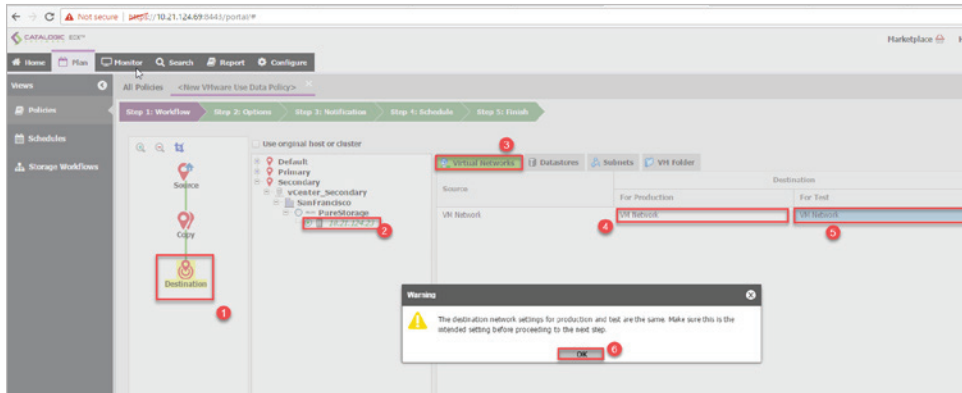


Figure 45. Create Use Data Policy

- Click the **Options** tab, as shown in Step 1, Figure 46. Select policy options, as shown in Step 1 and Step 2. **Default Mode:** Set the VMware Use Data policy to run in **Production** mode by default. Once the policy is created, it can be run in **Test**, **Production**, or **Clone** mode through the **Monitor** tab. We will retain the other options as **Production**.

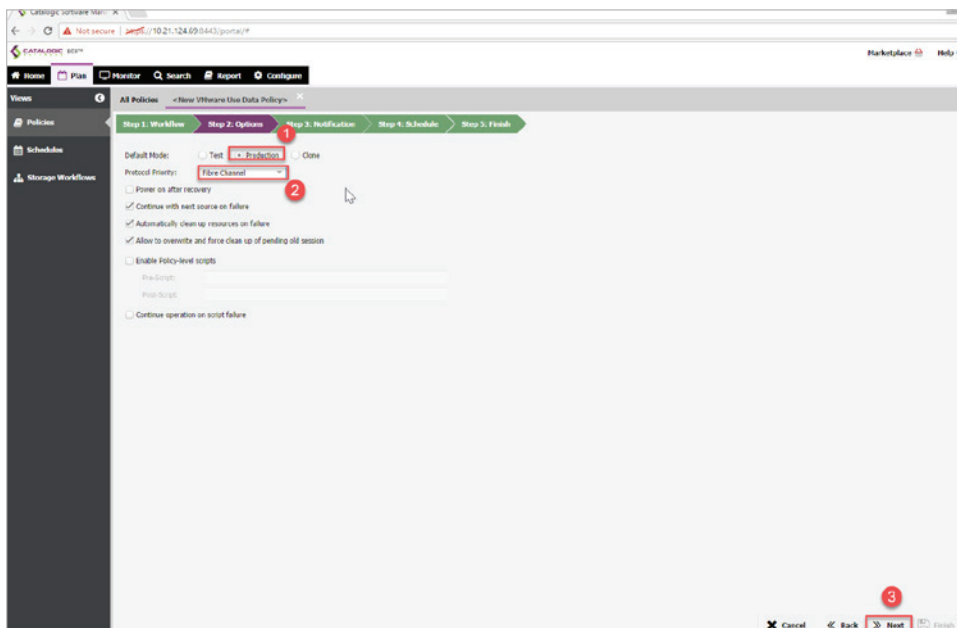


Figure 46. Create Use Data Policy "Production Mode"

- Click the **Schedule** tab, as shown in Step 1, Figure 47. Select **Start job now**, as shown in Step 2, which will run the job immediately, as defined by the order of the recovery

sources in the **1: Workflow** tab. Optionally, select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.




Figure 47. Create Use Data Policy

Note: To create and select a new schedule, click on the **Plan** tab and then click on **Views**, as shown in Step 1, Figure 48. Next select **Schedules**, as shown in Step 2, and then click **New**, as shown in Step 3.




Figure 48. Create Schedule for Use Data Policy

Create a schedule, as shown in Figure 49, then return to the Use Data policy editor, refresh the **Available Schedules** pane, and select the new schedule.

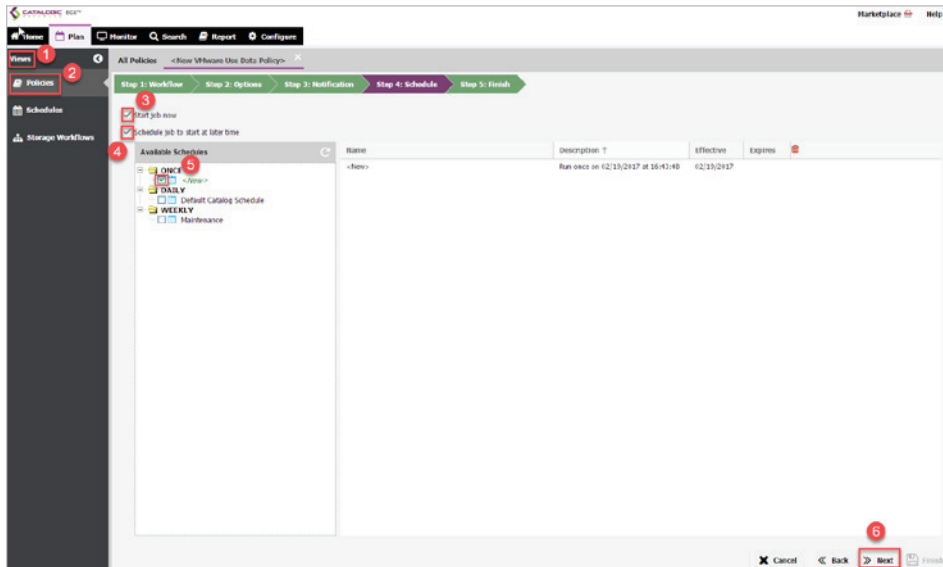


Figure 49. Create Schedule for Use Data Policy

8. Click the **Finish** tab, as shown in Figure 50. Enter a name for your policy and a meaningful description. We named the policy “ProdVMRestore”, as shown in Step 2. Click **Finish**. The policy runs as defined by your triggers and can be monitored from the **Monitor** tab.

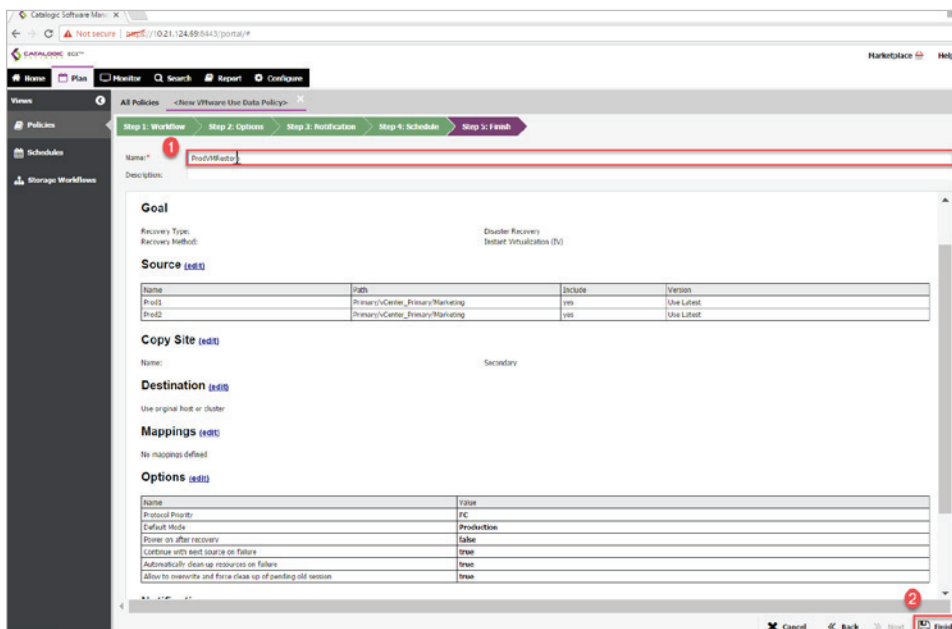


Figure 50. Create Use Data Policy

As described earlier, Production mode replaces the original machine images with images from the latest snapshot that was selected in the policy. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data policies associated with the virtual machine continue to run, as shown in Figure 50.

USECASE III – INSTANT RECOVERY OF VIRTUAL MACHINES TO DEV-TEST ENVIRONMENT

In this use case, we will go through an example of restoring the two virtual machines (**Prod1** and **Prod2**), from a replicated snapshot copy to an alternate vCenter (**vCenter_secondary**) using **Test** mode. To perform this restore, we can either create a new Use Data policy or edit the last used Use Data Policy, **ProdVMRestore**.

1. Edit the last created Use Data Policy called **ProdVMRestore**.




Figure 51. Edit Use Data Policy

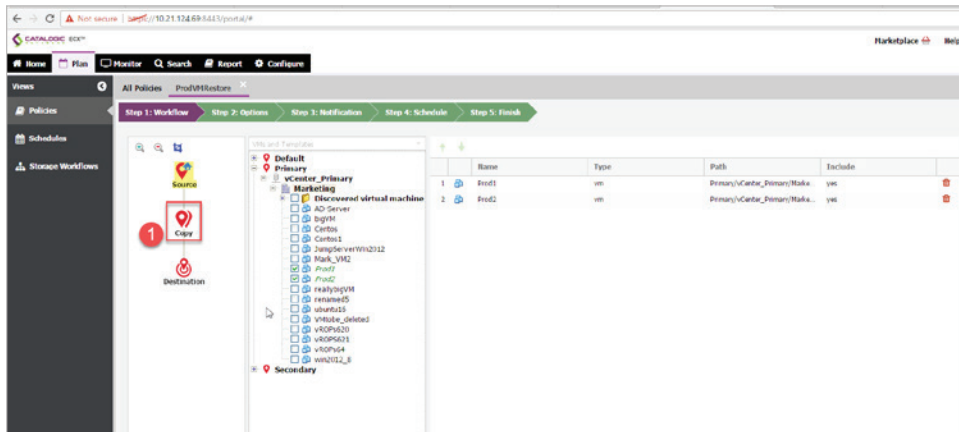


Figure 52. Edit Use Data Policy

2. In **Copy** selection, choose the site **Secondary**, as shown in Step 2, Figure 53. Click on **Next**.

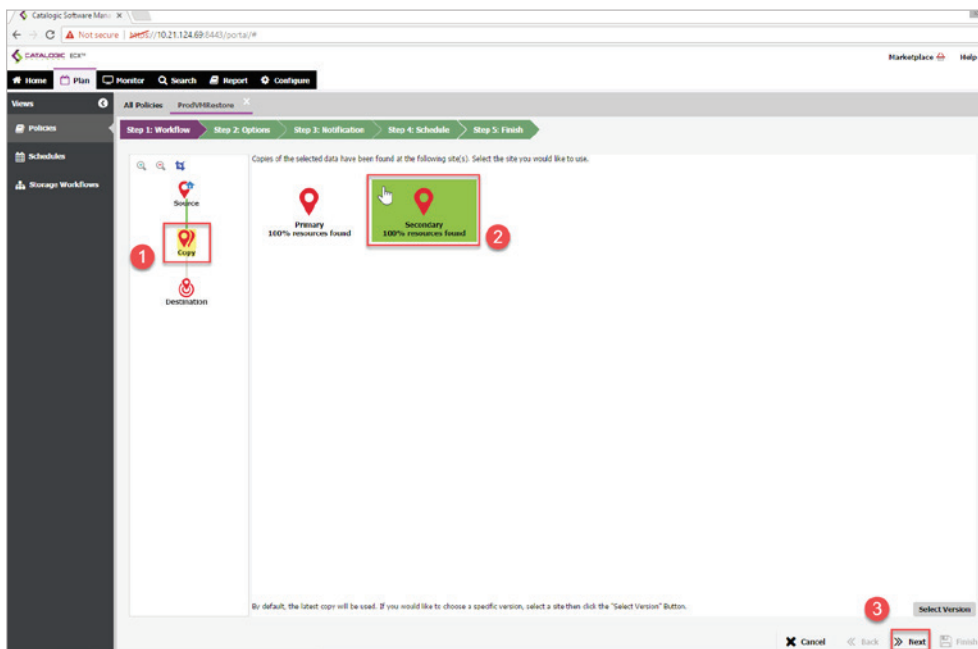


Figure 53. Edit Use Data Policy

3. Click **Destination**, as shown in Step 1, Figure 54. Uncheck the option **Use original host or cluster** if you would like to restore it on another vCenter. A destination panel opens with all the sites and corresponding vCenter resources. For this scenario, we will browse the **Secondary** site and select **vCenter_secondary** and the ESXi Host **10.21.124.23**, as shown in Figure 54.

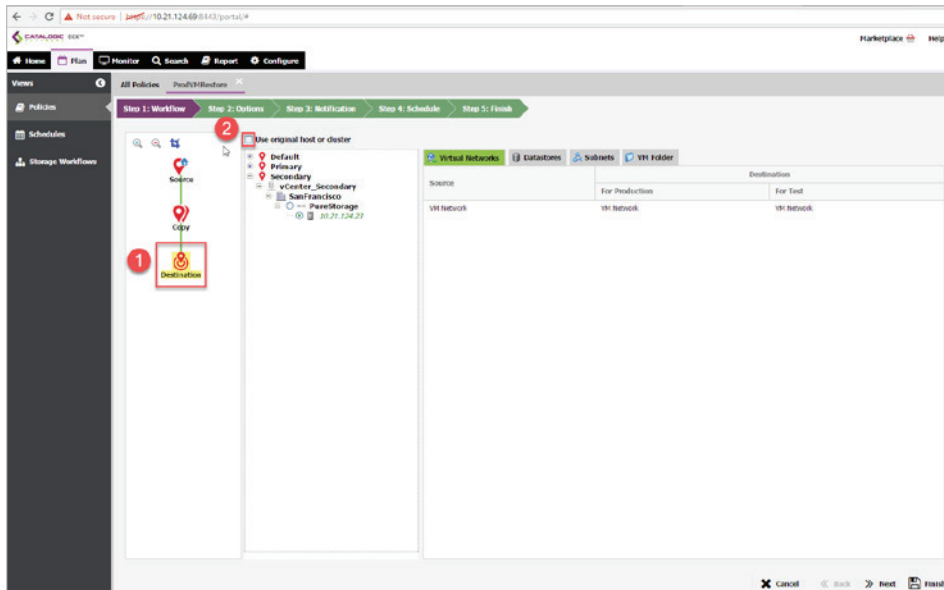


Figure 54. Edit Use Data Policy

4. Select the virtual network and datastores mapping. The **Virtual Networks** pane displays all of the virtual networks associated with your VMware Use Data policy sources. New virtual networks must be selected for use at the recovery site, as well as new datastores in the **Datastores** pane. Select a production and test network in the **Virtual Networks** tab, and a destination datastore in the **Datastore** tab.
 - a. **Virtual Networks** – Set virtual networks for production and test recovery jobs.

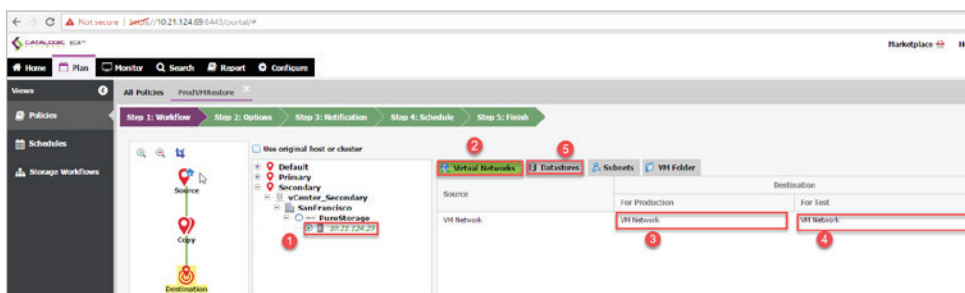


Figure 55. Edit Use Data Policy

- b. **Datastores** – Set the destination datastore, **DUMP3**. You can choose any datastore that you want, as shown in Figure 56.

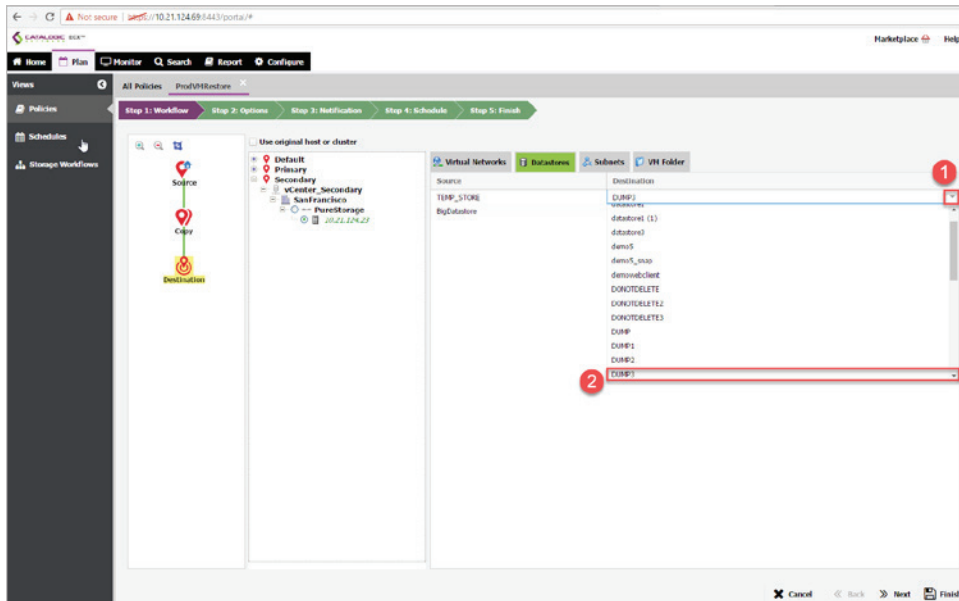


Figure 56. Edit Use Data Policy

- c. **Subnet** – Set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported. We will skip the step to retain the original subnet and IP address, as shown in Figure 57.

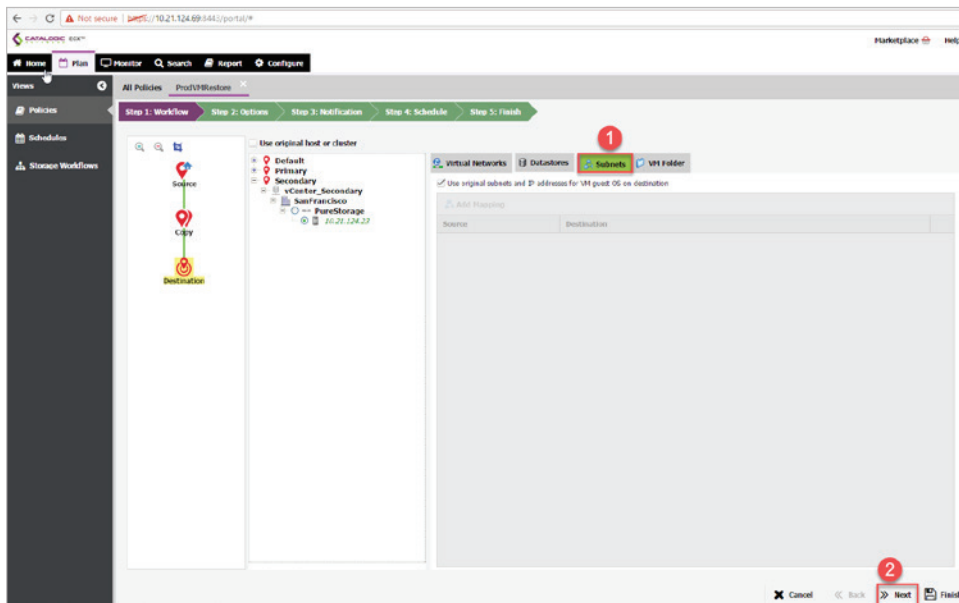


Figure 57. Edit Use Data Policy

- Click the **Options** tab. Select the policy options. **Default Mode:** Set the VMware Use Data policy to run in **Test** mode, as shown in Figure 58. **Test** mode will create a virtual machine with the suffix “_test”; this virtual machine will be running off of a snapshot volume. Once the policy is created, it can be run in **Test**, **Production**, or **Clone** mode through the **Monitor** tab. We will retain the other options as is.

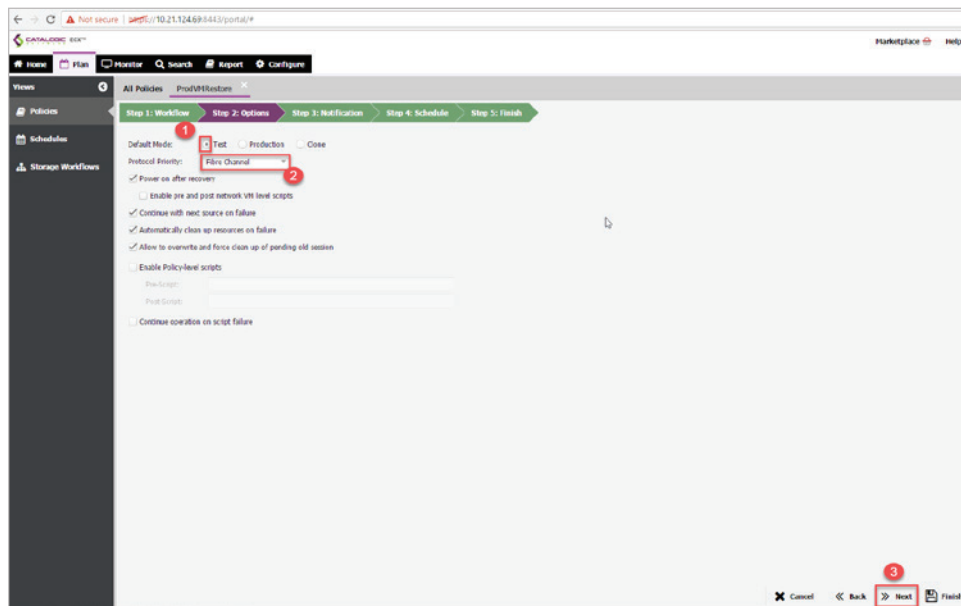


Figure 58. Edit Use Data Policy

- We will skip the steps **Schedule** and **Notifications** and manually initiate the policy from the **Monitor** tab.
- Click the **Finish** tab. Enter a name for your policy and a meaningful description. We kept the name of the policy **ProdVMRestore**. Click **Finish**, as shown in Figure 59.

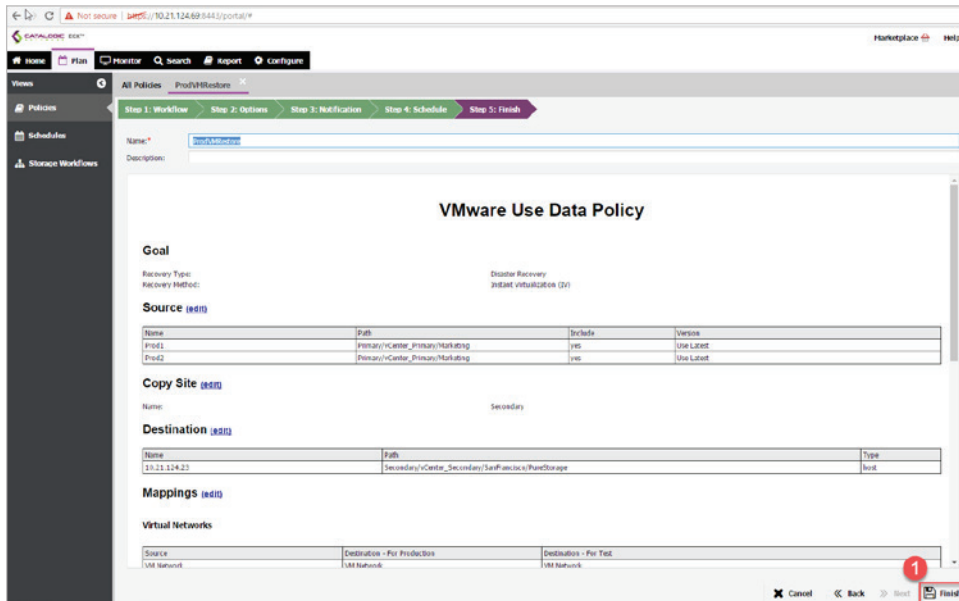


Figure 59. Edit Use Data Policy

- Click the **Monitor** tab, highlight the policy **ProdVMRestore**, and click the **Start** button. You will be prompted to select a **Start** option. We will select **Test** and click the **OK** button, as shown in Figure 60.

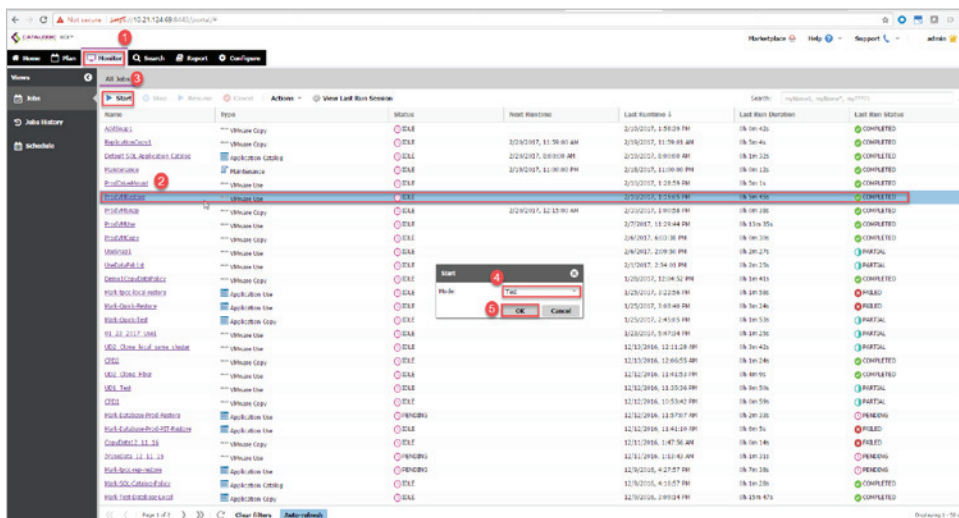


Figure 60. Monitor Use Data Policy

Monitor the execution of the **ProdVMRestore** policy here, as shown in Figure 61.

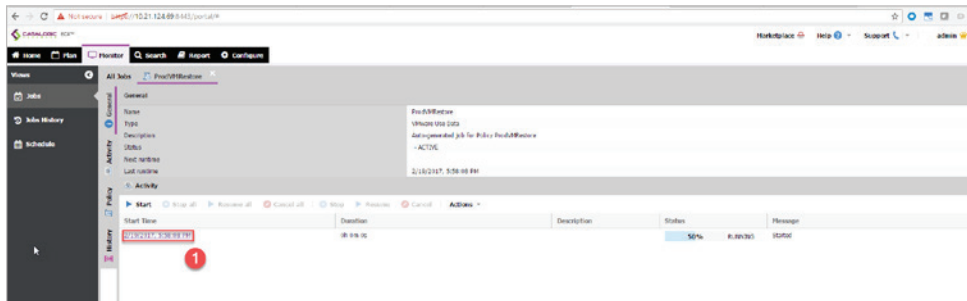


Figure 61. Monitor Use Data Policy

You can view the detailed log of all the steps being executed by clicking on the **Log** tab, as shown in Step 1, Figure 62.

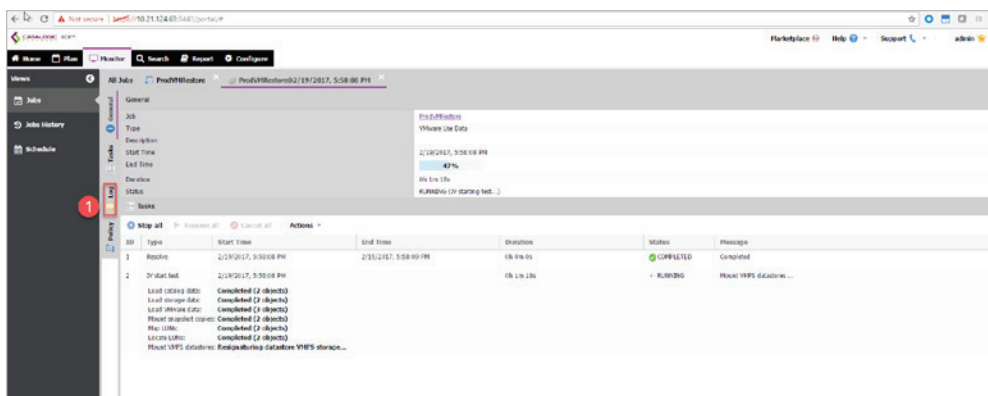


Figure 62. Monitor Use Data Policy

- In the vCenter web console, we can see the virtual machines **Prod1-test00** and **Prod2-test00** were successfully created and powered on, as shown in Figure 63. The VMs are created with a suffix '-test##' to indicate the mode of operation.



Figure 63. Check virtual machines on vCenter (Secondary Site).

10. You can also validate the virtual machines **Prod1** and **Prod2** on the Primary Site as shown in Figure 64.

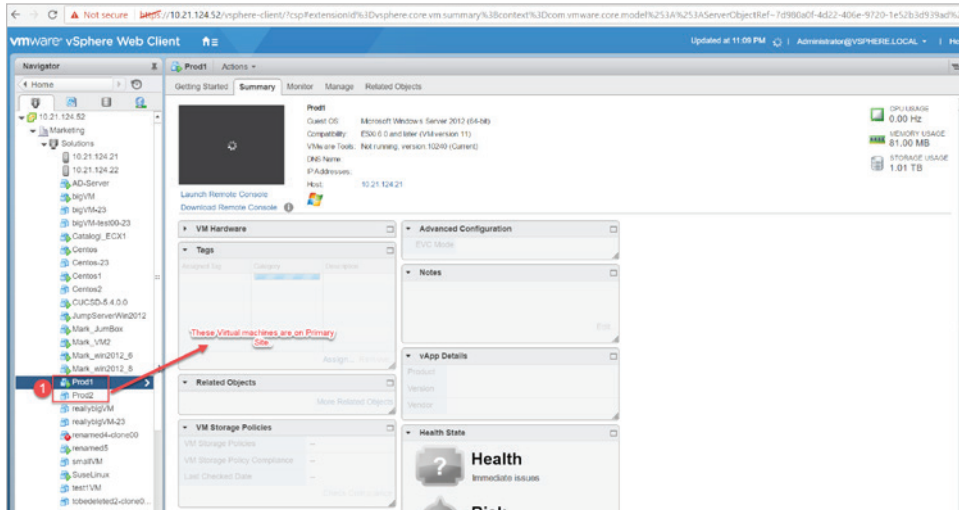


Figure 64. Virtual machines on Primary Site

11. Once the policy completes successfully, select one of the following options from the **Actions** menu on the **General** tab of the job session on the **Monitor** tab, as shown in Step2, Figure 65: **End IV** (Cleanup), **RRP** (vMotion), or **Clone** (vMotion).
- End IV (Cleanup)** destroys the virtual machine and cleans up all associated resources. Since this is a temporary/testing virtual machine, all data is lost when the virtual machine is destroyed.
 - RRP (vMotion)** is equivalent to using the **Production** selection in the policy **Options** screen. This option migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the “For Production” Network.
 - Clone (vMotion)** is equivalent to using the **Clone** selection in the policy **Options** screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the “For Test” network.

We will select **RRP** to make the VM permanent. **Prod1-test00** and **Prod2-test00** will get renamed to **Prod1** and **Prod2** on the secondary site. And now they will be running off of a permanent datastore, as compared to the snapshot volume previously.

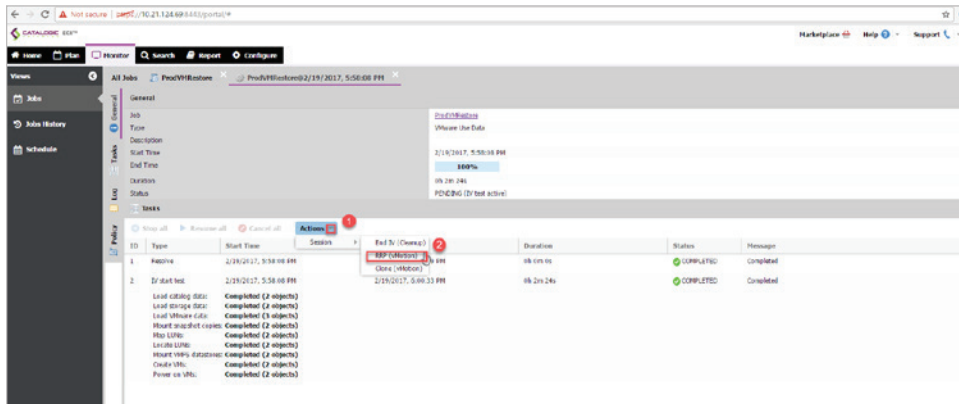


Figure 65. Virtual machines on Primary Site

Once the virtual machines become permanent using the Rapid return to Production (RRP) method, **Prod1-test00** and **Prod2-test00** are renamed **Prod1** and **Prod2** on the secondary site, as shown in Figure 66.

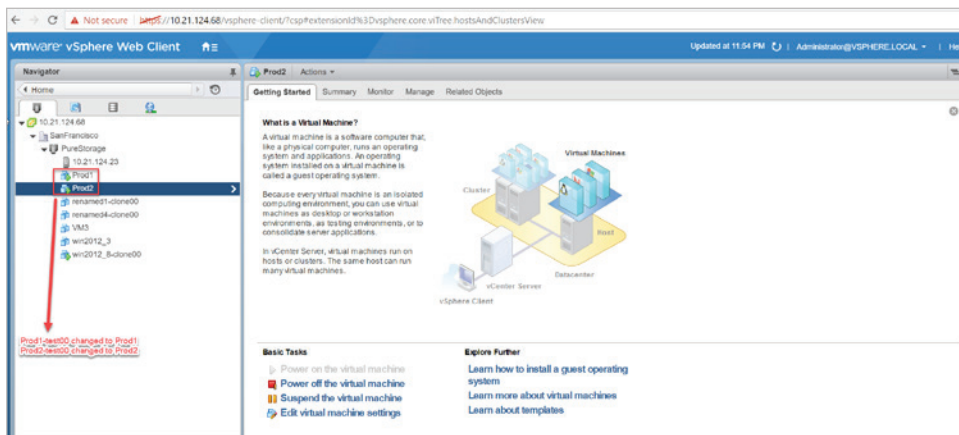


Figure 66. Virtual machines on Secondary vCenter changed from Prod1-test00 and Prod2-test00 to Prod1 and Prod2

SUMMARY

In this whitepaper, we have demonstrated how to manage, orchestrate, and analyze Copy Data in an IT environment deployed on VMware and Pure Storage FlashArray infrastructure using Catalogic Software's intelligent Copy Data management platform, ECX.

This document covers three major use cases in today's IT environment, and these can be used for Automated DR, Dev-Test or DevOps, and Business Analytics. 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 of Pure Storage FlashArray-based snapshot and replication, without adding complexity or requiring added expertise.

REFERENCES

1. Catalogic ECX 2.5.1 User Guide

https://mysupport.catalogicsoftware.com/product.php/ECX251/ECX_UsersGuide.pdf

2. Catalogic ECX 2.5.1 Quick Start Guide PDF

https://mysupport.catalogicsoftware.com/product.php/ECX251/ECX_UsersGuide.pdf

ABOUT THE AUTHOR

Saurabh Talwar (Sunny) is a PureStorage Solutions Architect for VMware. Sunny has extensive experience in defining and developing technical product solutions & marketing strategy for enterprise products. Sunny has worked in release engineering, marketing, and Quality Engineering departments, and has worked closely with Engineering to understand Product Architecture. He came to Pure from HGST (Subsidiary of Western Digital), where he was designing solutions for SSDs. Sunny was doing performance analysis with competitor SSDs in order to do comparative analysis and write reports for the product marketing team. Sunny was also developing demos with real world applications like MySQL and Aerospike, showing the performance of PCIe Cards, and answering any questions about product at conferences like VMworld and FlashMemory Summit. Apart from work, Sunny likes to swim and do Crossfit.



© 2017 Pure Storage, Inc. All rights reserved.

Pure Storage, Pure1, Evergreen, and the "P" Logo are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Catalogic ECX is a registered trademark of Catalogic Software in the U.S. and other countries. VMware vCenter and VMware vSphere are registered trademarks of VMware in the U.S. and other countries. Oracle and Oracle Enterprise Linux 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.

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.

ps_wp_catalogic-ecx-integration-for-vmware_01



Pure Storage, Inc.

Twitter: [@purestorage](https://twitter.com/purestorage)

www.purestorage.com

650 Castro Street, Suite #260
Mountain View, CA 94041

T: 650-290-6088

F: 650-625-9667

Sales: sales@purestorage.com