

Cohesity Granular Recoveries for Virtualized Microsoft Exchange Environments



Table of Contents

About This Guide: 1

Importance of Granular Recovery v.s. Restore:

- Cohesity Storage Platform Overview: 2

The Cohesity Kroll Ontrack PowerControls Solution

- Overview: 2
- Architecture: 2
- Performing Granular Level Recoveries: 2
- Create Clone of Exchange Server: 3
- Make the Database Available to Kroll: 3
- Mount Recovered Exchange VMDK: 7
- Mount Recovered Database on Kroll Server: 9
- Recover Object from Backup: 11
- Unmount and Destroy Clone: 13

Conclusion

- Conclusion: 14

About This Guide

This paper details the steps and best practices to enable granular level recovery for Microsoft Exchange Servers using Kroll Ontrack PowerControls. Combining Kroll Ontrack PowerControls with Cohesity Data platform provides a comprehensive webscale approach to data protection and recovery for Exchange Server environments.

Audience

This paper is written for Exchange and Backup administrators familiar with administering and managing Exchange backup environments.

The Audience must also be familiar with:

- Kroll Ontrack PowerControls for Exchange
- Cohesity Data Platform
- Microsoft Exchange Server

Importance of Granular Recovery v.s. Restore

Historically, the only way to recover a single email or item in a public folder required a lot of work and a lot of time to complete. Administrators would have to scour boxes of tape or racks of appliances to find the exact backup or version of the database. Once located, the process of downloading or extracting the entire database was required in order to find that single item. The ability to recover individual objects from a larger data set, such as an Exchange Database, have enabled administrators to deliver much smaller RTOs to their customers. Removing the requirement to fully download large objects, in favor of extracting only what is needed while leaving the backup in place.

These minimal RTOs are now commonplace throughout the industry. The Cohesity Data Platform and Kroll OnTrack PowerControls for Exchange have combined to deliver best in class Granular Level Recoveries backed by the infinite scale of a fully integrated data protection platform for virtual machines.

Cohesity Data Platform Overview

Cohesity introduced the world's first scale-out data management platform to enable organizations to standardize secondary workflows on a completely unified and fully distributed solution. Cohesity's scale-out distributed file system was built from the ground up to ensure complete scalability to enable organizations to flexibly grow their environment by simply adding nodes to a cluster. With this scalability, organizations can eliminate the costs of data migrations and forklift upgrades, while benefiting from the simplicity of a homogenous solution.

In addition to ensuring complete scalability, Cohesity enables the delivery of all secondary workloads by providing a powerful, unified platform with industry-leading capacity and performance across the most common file system protocols. In a single 2U Block, Cohesity delivers 48TB of flash and HDD storage to efficiently balance performance, capacity, and price to give organizations the flexibility to offload secondary workloads from expensive primary storage arrays. Cohesity's integrated backup, archive, replication, disaster recovery, and public/private cloud support, combined with its inherent context awareness, eliminate the need for backup software and other ancillary backup infrastructure.

With Cohesity, organizations can benefit from stringent data protection policies by gaining valuable insight from their data. Cohesity's context-aware file system makes it easy to understand storage utilization trends, reclaim unused resources, and create summary and detailed reports with just the click of a button.

The Cohesity Kroll Ontrack PowerControls Solution

Overview:

Kroll Ontrack® PowerControls™ is a forensically sound email management tool that helps IT administrators granularly search and restore mailboxes, messages, attachments and other Microsoft® Office Outlook items - without restoring the entire database.

Architecture:

Setting up and configuring the Kroll Ontrack PowerControls suite of software is simple. For supported Microsoft Exchange environments, a Kroll Ontrack PowerControls for Exchange application, along with a Microsoft Outlook Client, needs to be available to perform granular level recoveries.

Cohesity recommends as a best practice that the Kroll Ontrack PowerControls software is installed on a stand alone virtual machine. The Microsoft Outlook client, on this virtual machine, should be configured to connect to the production Exchange Server with explicit Exchange Administrator privileges. This account should be enabled to recover objects to any Mailbox on the production Exchange environment.

Supported Configuration

- Kroll Ontrack PowerControls 11 for Exchange
- Cohesity Storage Platform version 2.5+
- Microsoft Exchange Server 2013+ Running on VMware

Performing Granular Level Recoveries:

Performing a Granular Level Recovery involves two steps.

1. Make the database available to the Kroll Virtual Machine
2. Open the Exchange Database in Kroll then search, select, and recover the object

Note: It is assumed that all Exchange Servers running on VMware are protected with Application Consistency enabled, at least one backup has completed successfully, and available for recovery.

Create Clone of Backup Exchange Server

Make the Database Available to Kroll:

Optimized for time to recovery, the advantages of granular level recoveries enables the administrator to pick and choose the individual objects to restore where traditionally the entire database would need to be recovered from backup only to pull a few individual objects. Kroll paired with Cohesity puts the time to recovery in the hands of the administrator, now able to easily adhere to strict Service Level Agreements.

To Make the Exchange Database available to Kroll:
Login to Cohesity, then navigate to the **Test & Dev** screen to initiate the clone.

The screenshot shows the Cohesity web interface. The top navigation bar includes 'COHESITY', 'Test & Dev' (highlighted with an orange box), 'Analytics', 'Protection', 'Platform', and 'Monitoring'. On the right, it says 'Welcome, admin' with links for 'Admin', 'Help', and 'Logout'. Below the navigation bar is the 'Manage Clones' section, which includes a 'Clone VMs' button (highlighted with an orange box) and a summary of clone tasks: 0 Running Tasks, 0 Scheduled, 2 Success, 0 Errors, 1 Clone Destroys, and 0 Destroy Errors. Below this is a table titled '3 Clone Tasks' with columns for Task Name, Objects Cloned, Start Time, Task Duration, and Status.

| Task Name | Objects Cloned | Start Time | Task Duration | Status |
|--|----------------|--------------------|---------------|-----------|
| Clone-VMs-2016-02-11_07-31PM | 1 object | 02/11/2016 7:35pm | 5 seconds | Success |
| Clone-VMs-2016-02-08_11-35AM | 1 object | 02/08/2016 11:36am | 5 seconds | Success |
| Clone-VMs-2016-02-08_11-30AM | 1 object | 02/08/2016 11:35am | < 1 second | Destroyed |

At the bottom of the interface, there is a footer with copyright information: © 2015 Cohesity, Support, Feedback, Help, cohesity.com.

Figure 1: Select Test & Dev to navigate to the managed clones view

Select Clone VMs and search for the VM that contains the Exchange Database.

When a full recovery of the Exchange Database is wanted, simply click the “Add to Cart” button then the “Continue” button from the main Cohesity Dashboard:

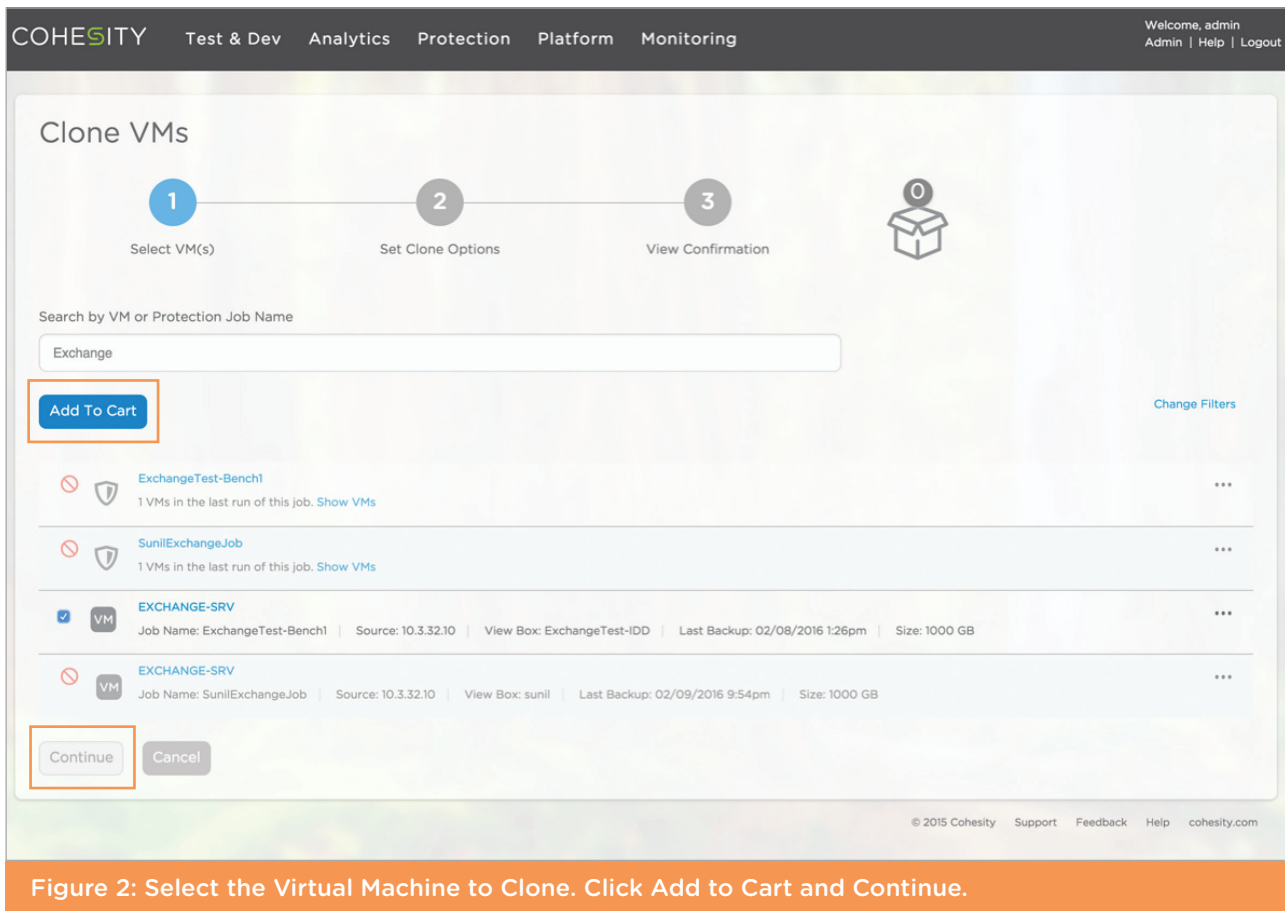


Figure 2: Select the Virtual Machine to Clone. Click Add to Cart and Continue.

Select the VM that matches the search. The VM has the “VM” icon vs the Cohesity Protection icon. Next, name the new Clone VM and configure clone to have the network detached and leave the VM powered off, as specified in the **Additional Options** section of the clone.

Clone Options

Rename Cloned VMs

YES

Add Prefix: Add Suffix:

Clone Location

* Source: * Resource Pool: * View:

Networking Options

Detach network
 Attach to a new network

Additional Options

YES Leave cloned VMs powered off
 NO Continue clone even if errors occur when cloning VMs

Selected Objects

| Selected Objects | Clone to |
|---|--|
| EXCHANGE-SRV <small>View Box: ExchangeTest-IDD Job Name: ExchangeTest-Bench1</small> | EDB-REC_EXCHANGE-SRV_New <small>Snapshot: 02/08/2016 1:26pm, 1000 GB (Latest Snapshot) </small> |

Figure 3: Configure the Clone settings. Ensure that the Clone is left off. To change the recovery point, select the icon in the Clone to column.

To choose a point in time copy that's different from the current backup, select the edit icon next to the VM listed in the "Clone to" column in the Selected Objects section and select the appropriate backup version.

1 Recover Points for EXCHANGE-SRV

| Recover Point | App-Consistent | Type |
|--|----------------|-------------|
| <input checked="" type="radio"/> 02/08/2016 1:26pm | Yes | Full Backup |

Figure 4: Select the Recovery Point, and then save to continue.

Select the **Save** button to confirm the selection of the point in time copy. From there, simply select the **Finish** button to complete the clone selection process.




Figure 5: One all Clone settings are configured. Click Finish to continue.

At the end of this process, Cohesity will automatically create an NFS datastore, mount the datastore to the resource pool specified, and register the cloned VM in vCenter.

Next, navigate to your vCenter console to mount the newly cloned VM's VMDK to the virtual machine running the Kroll Ontrack PowerControls for Exchange software.

To Start, select the Kroll virtual machine on vCenter, and **Edit** the Settings.

Mount Recovered Exchange VMDK

Then select Clone VMs and search for the VM that contains the Exchange Database.

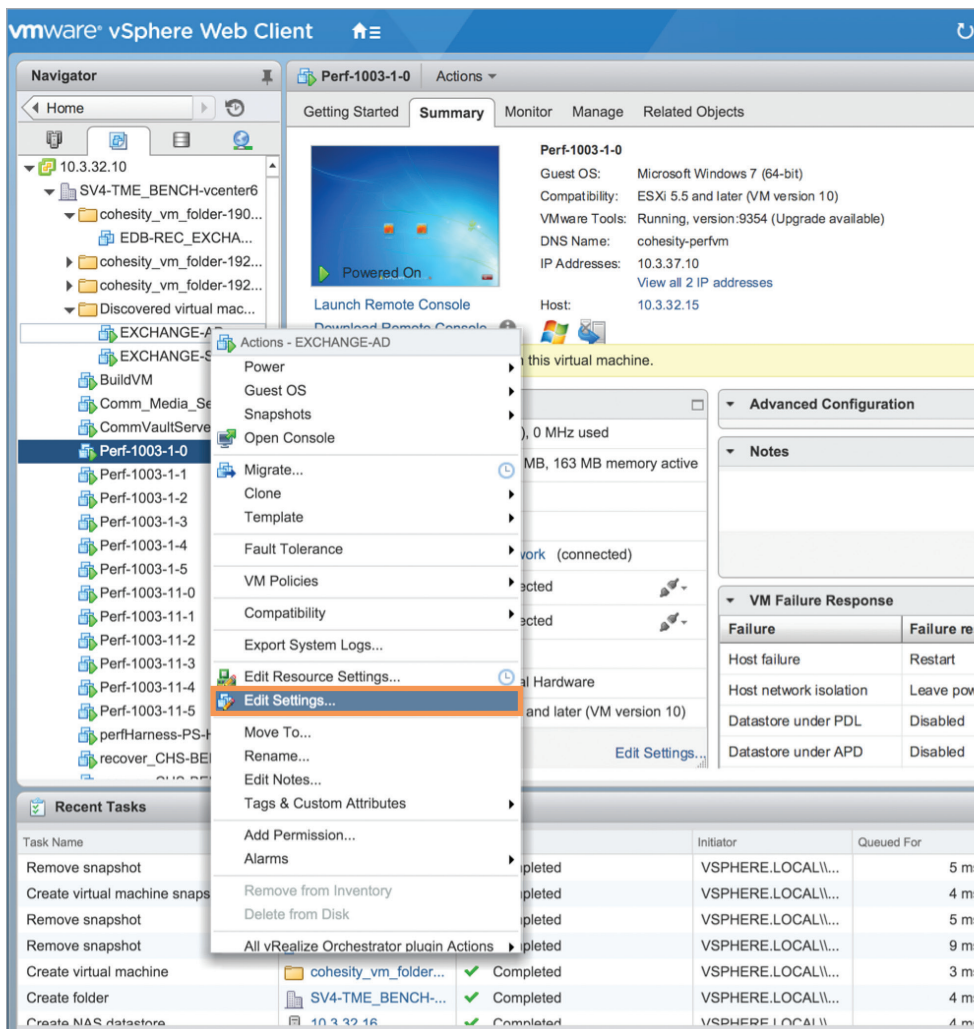


Figure 6: Right-Click the Kroll VM and select Edit Settings.

From the **Edit Settings** menu, use the dropdown on the bottom of the menu, and select Existing Hard Drive, and press the **Add** button.

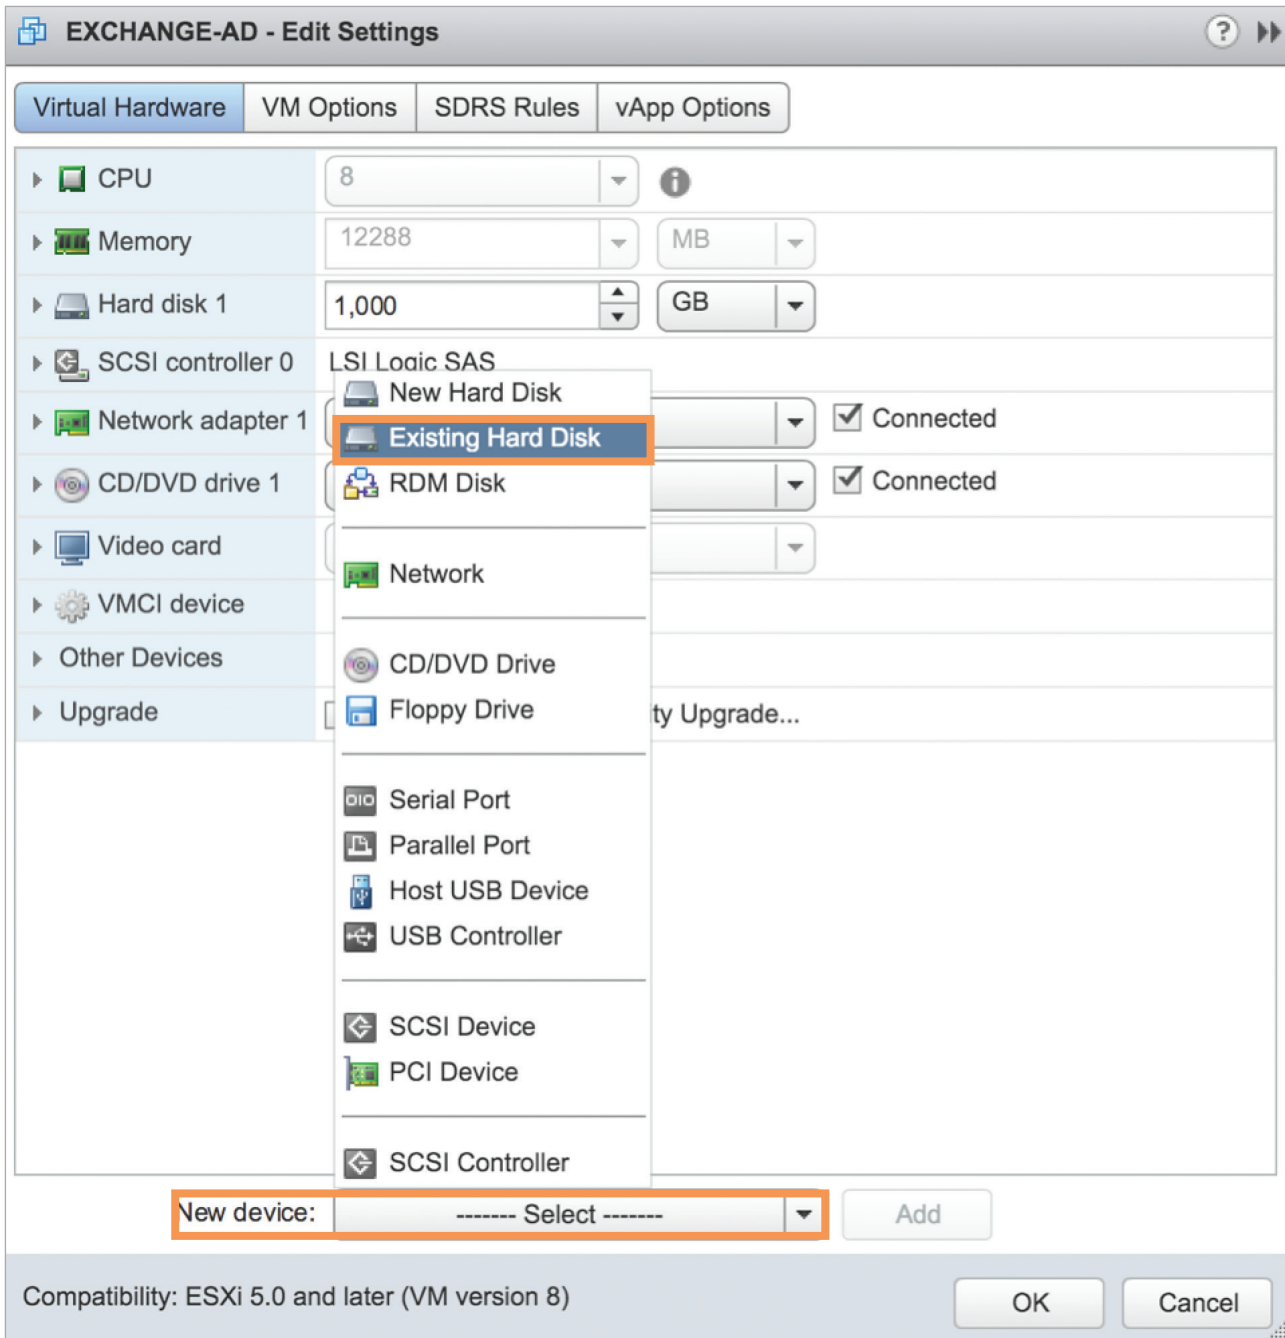


Figure 7: Select Existing Hard Disk in the New Device drop-down.

Then navigate to the Cohesity provisioned datastore to find the clone, and select all of the VMDKs which contains the Exchange Server Databases, Logs, and Temporary files.

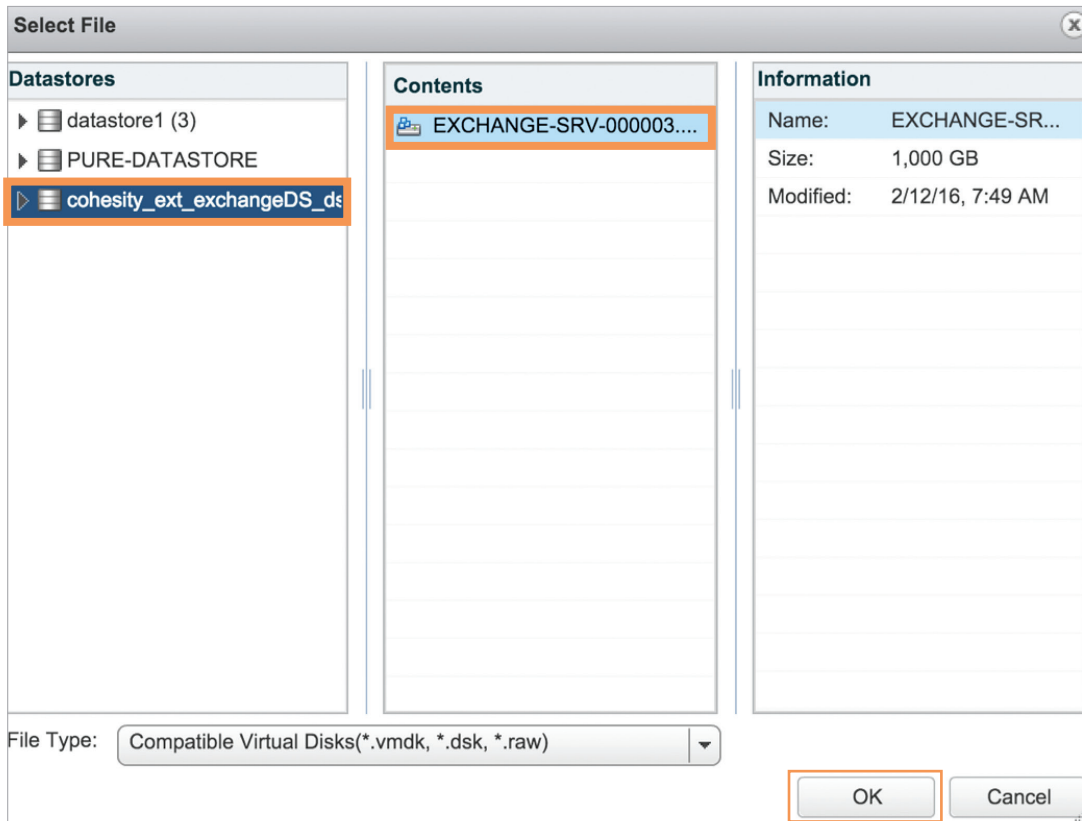


Figure 8: Navigate to the newly mounted Cohesity Datastore, and find the created Clone. Select the VMDK that contains the EDB.

Mount Recovered Database on Kroll Server

After the VMDK is mounted to the Kroll Ontrack PowerControls for Exchange system, log into the remote console of the server and navigate to the **Disk Management** control panel.




Figure 9: Navigate to the Disk Management MMC Snap-in

Next rescan the system by clicking on **More Actions** and selecting Rescan Disks

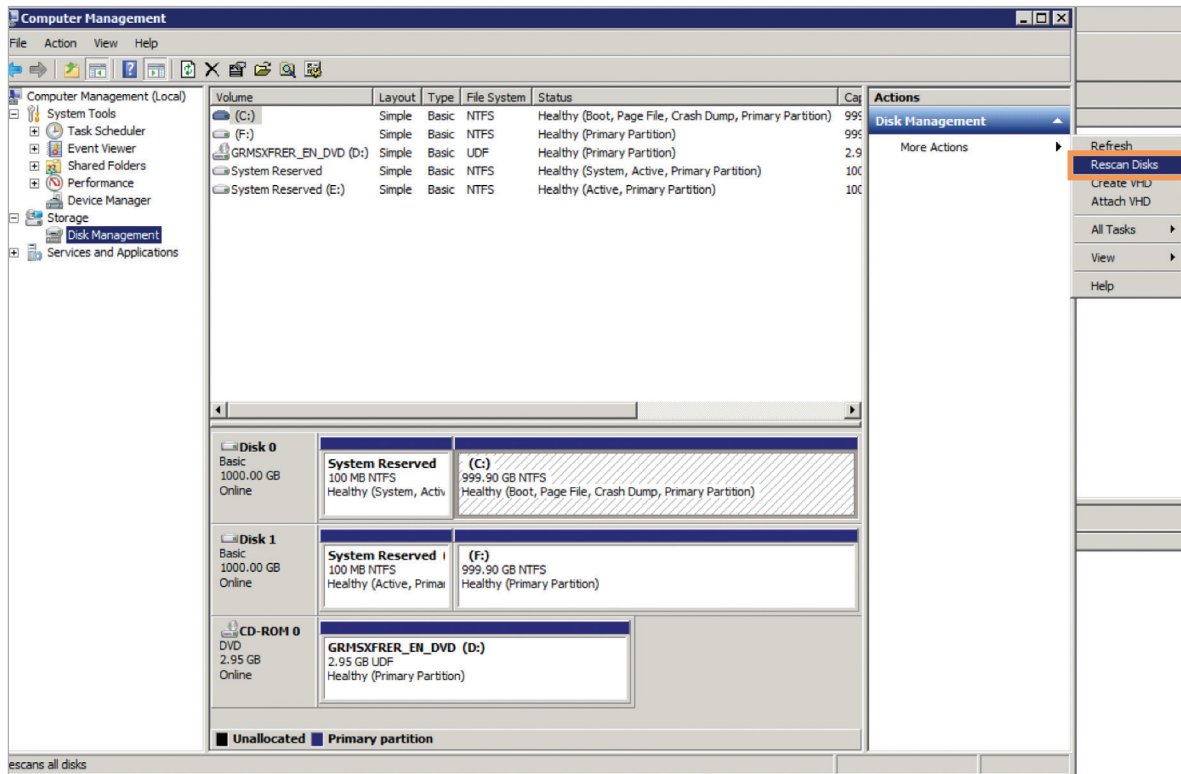


Figure 10: Select Rescan Disks to find the newly mounted VMDK file.

Finally, right-click on the newly found partition, and click **Online**. This will mount the volume to your system, and give the new drive a drive letter (F: in this example).

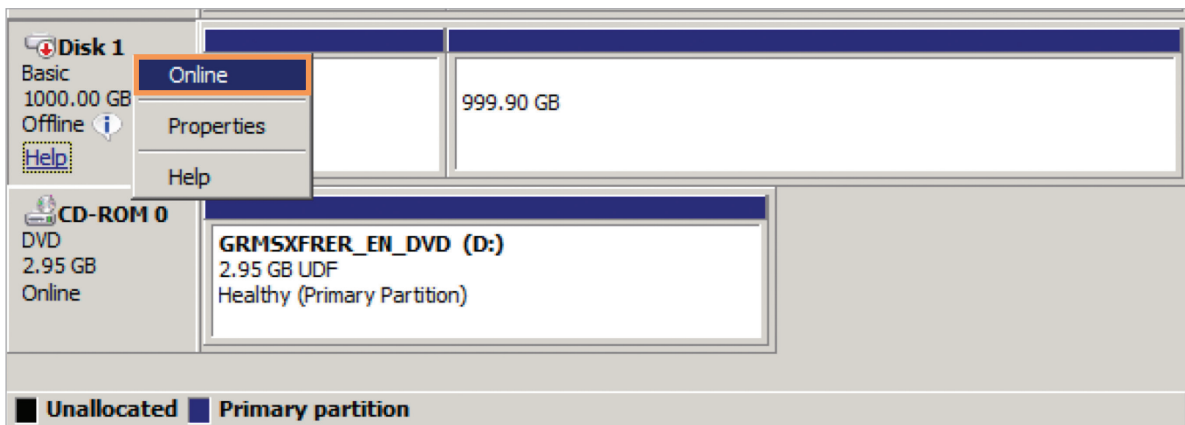


Figure 11: Bring the disk online.

Mounting the cloned disk to the Kroll system completes the first step in the recovery process. To start the second, and final, step in the process, launch the Kroll Ontrack PowerControl for Exchange software and select **Open Source**.

Recover Object from Backup

To perform the Granular Level Recovery from the Exchange Database, select the **EDB File** option, and select the location of the Exchange Database. If the Log Files and Temporary Files locations are not automatically populated, specify the locations by clicking on the **Browse** button and locating them on the attached disks. Finally, click “Finish”. This will allow Kroll to open the Database File and populate the UI with all Mailboxes in a tree view.

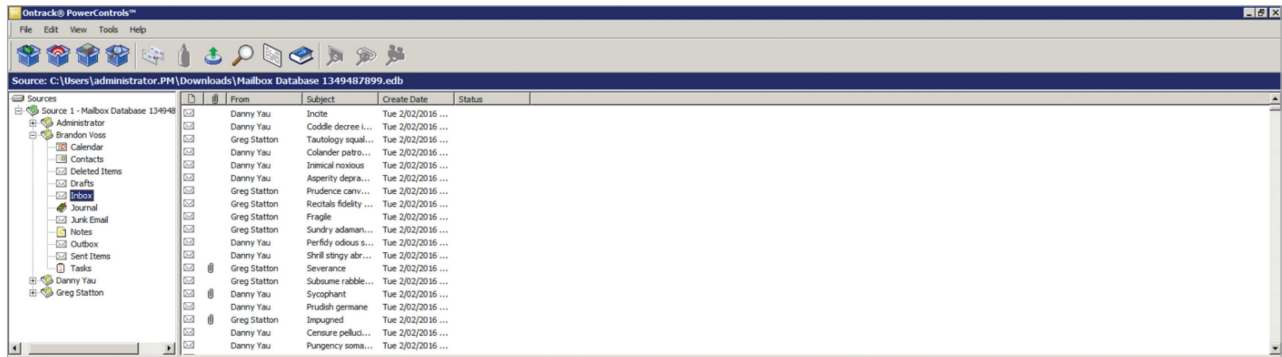


Figure 12: Load the EDB file in the Kroll management interface.

It is recommended that the production target Exchange server be specified before the search and select objects process happens. To do this first select the **Open Target Exchange Server** button on the top menu bar, and provide the connection information. If your Exchange environment supports Autodiscover, Kroll will find the Exchange Server Name automatically.

Data Wizard

Target Selection

Please select the type of target you want to open and its location. If you do not want to open a target at this time, click Skip.

Select Target:

- Single Mailbox (on-premise or Office 365)
- All Mailboxes (on-premise)

Microsoft Exchange Server

Domain Controller: EXCHANGE-AD

Server Name: exchange-srv.pm.cohesity.com

- Connect To Public Folders
- Connect using Autodiscover
- Connect using Outlook Anywhere

< Back Next > Skip >> Cancel

Figure 13: Connect to the Production Exchange Infrastructure to begin the object recovery.

Note: By selecting All Mailboxes Kroll will load **all mailboxes** found in the selected Exchange Database. If administrators wish, **Single Mailbox** can be selected to reduce the load time from the Exchange Database.

The Target Exchange Database will appear in the bottom pane of the UI with all the available Mailboxes in a similar tree view as the backup view above it. Select the object from the backup Exchange Database that you want to restore then drag it to the appropriate Mailbox

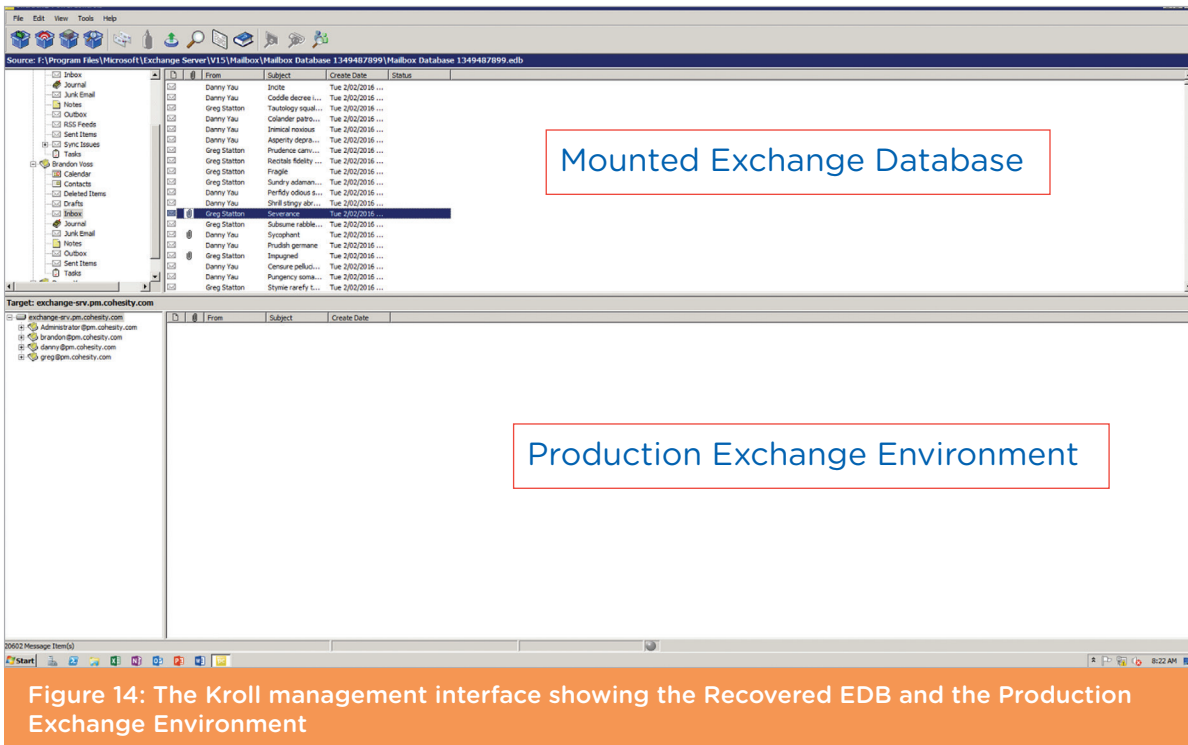


Figure 14: The Krill management interface showing the Recovered EDB and the Production Exchange Environment

Unmount and Destroy Clone

After you have performed all of the Granular Level Recoveries successfully, as a best practice Cohesity recommends that the VMDKs be detached and the Cohesity Clone be destroyed. To detach the VMDKs navigate back to the vCenter administration interface, right-click on the Krill virtual machine, and select **Edit Settings**. From the **Settings** window, remove the mounted VMDKs by clicking the 'X' next to the extra **Hard Drives**.

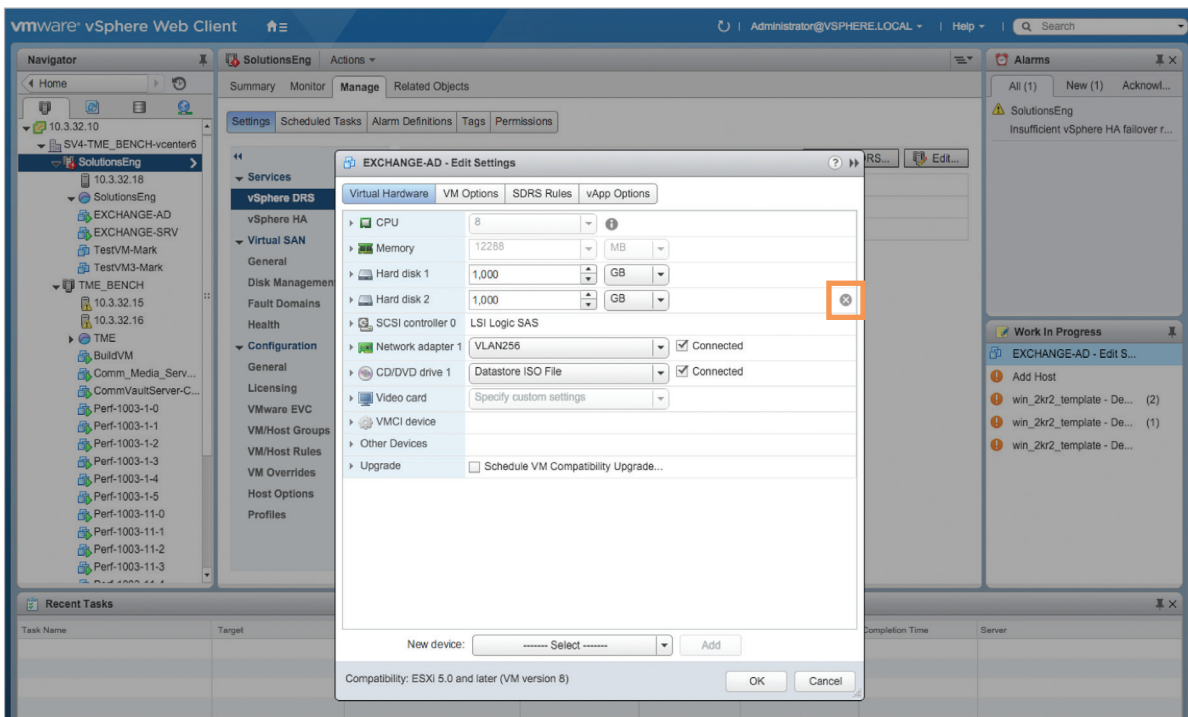


Figure 15: Unmount the Clone's VMDK by removing the added Hard Disk. Press the "X" on the right

Once the VMDKs have been removed, navigate to the Cohesity Web UI and proceed to the **Test & Dev** tab to perform the **Tear Down Clone** process.

COHESITY Test & Dev Analytics Protection Platform Monitoring Welcome, admin
Admin | Help | Logout

Clone-VMs-2016-02-17_03-10PM [Back to Clones](#)

Tear Down Clone

Cloned by admin | Cloned to 10.3.32.10:TME

Success
Status

1
Object

02/17/2016 3:11pm
Start Time

4 seconds
Duration

Clone Options

- VM Prefix: "Greg_"
- VM Suffix: "_1"
- Parent Source: 10.3.32.10
- Resource Pool: TME
- View Name: cohesity_ext_ExchangeRecovery
- Continue on Error: Yes
- Power State: Off

1 Cloned Object

| Object | Snapshot Cloned | Status | Message |
|----------------|-------------------|---------|---------|
| > EXCHANGE-SRV | 02/08/2016 1:26pm | Success | |

© 2015 Cohesity Support Feedback Help cohesity.com

Figure 16: To complete the tear down process, navigate back to Test & Dev, select the Clone job, and Tea Down Clone

Conclusion

By combining the power of the Cohesity Data Platform and the granular level recovery capabilities of Kroll PowerControls for Exchange, a truly webscale approach to Data Protection can be achieved. Benefiting both the administrators with reducing the complexities of management, and their customers by enabling near-zero RTOs for single objects on any virtualized Exchange environment.