

caringo



Caringo Swarm and FileFly deployment guide for Milestone XProtect VMS

Table of Contents

INTRODUCTION	3
ARCHITECTURE	4
Hardware and software	4
Networking	5
WORKFLOW EXAMPLE	6
Caringo FileFly configuration	9
CONCLUSION	14

Introduction

The following document describes the installation and setup of Caringo Swarm and FileFly on a Milestone XProtect environment.

Caringo Swarm is a software defined object storage solution that provides a resilient and highly scalable pool of storage that can be used as a repository to keep massive amounts of information, including surveillance data.

Caringo FileFly simplifies the migration of secondary data from Windows machines and NetApp filers to Caringo Swarm. It brings complete and automated data lifecycle management of all the unstructured data, optimizing and consolidating it on a single and scale-out storage platform: Caringo Swarm.

FileFly is an agent that installs on a Windows machine, offloading data from a Windows drive letter to Swarm object storage platform, leaving a stub (pointer) on the primary storage. If an archived file gets accessed, FileFly will rehydrate the file automatically from Swarm back to the Windows server.

Milestone XProtect VMS requires a transparent mechanism to read and write the information as it stores the surveillance data in a drive letter on Windows. Caringo Swarm with FileFly has been proposed in order to meet this goal. Using this setup, an organization can have a limitless storage platform for a long-term retention of their surveillance assets and with quick and transparent access to them as there is no need to perform a classic restore from backup operation.

Architecture

This section of the document outlines architecture, hardware and software requirements.

Hardware and software

In regard to hardware, Caringo Swarm can be deployed either on virtual machines running on a VMware hypervisor, on standard x86 physical servers, or on a combination of both.

Please contact Caringo to spec out the hardware platform, as well as to define the full architecture of the solution: <https://www.caringo.com/contact>

We will work with you to understand your needs and determine the right storage strategy for your organization.

The software components required are:

- Caringo Swarm
- Caringo FileFly
- Milestone XProtect VMS

Milestone XProtect VMS can be deployed as a multi-server install or as single-server one. While multi-server is recommended to maximize performance, single-server deployments are supported as well.

In regard to Caringo components, FileFly migration software is installed on the Recording servers with Swarm running either on virtual machines or physical servers depending on sizing and scalability requirements.

You can find more information about Caringo products and their configuration in the additional documentation available both on Milestone Marketplace as well as on Caringo Connect: <https://connect.caringo.com>

Networking

Following the recommendations provided by Milestone, and best-practices from Caringo, the environment should be defined with at least three networks:

- Management network
- Feeds network
- Storage network

Each of these networks should have their own dedicated vNIC and vSwitch/VLAN on a virtual infrastructure, or a NIC and VLAN on a physical one.

Besides of the Swarm storage nodes, the Storage network is only present on the Milestone Recording server, and it is dedicated to migrating the files back and from the object storage platform.

FileFly can be deployed in many different ways: standalone, distributed, HA etc. In the example below, it was installed on the Recording server with its management running on a separate virtual machine (VM). Using this approach, the Recording server has a communications line dedicated to move data, from and to the object storage platform. Also, with this strategy, the deployment is simplified.

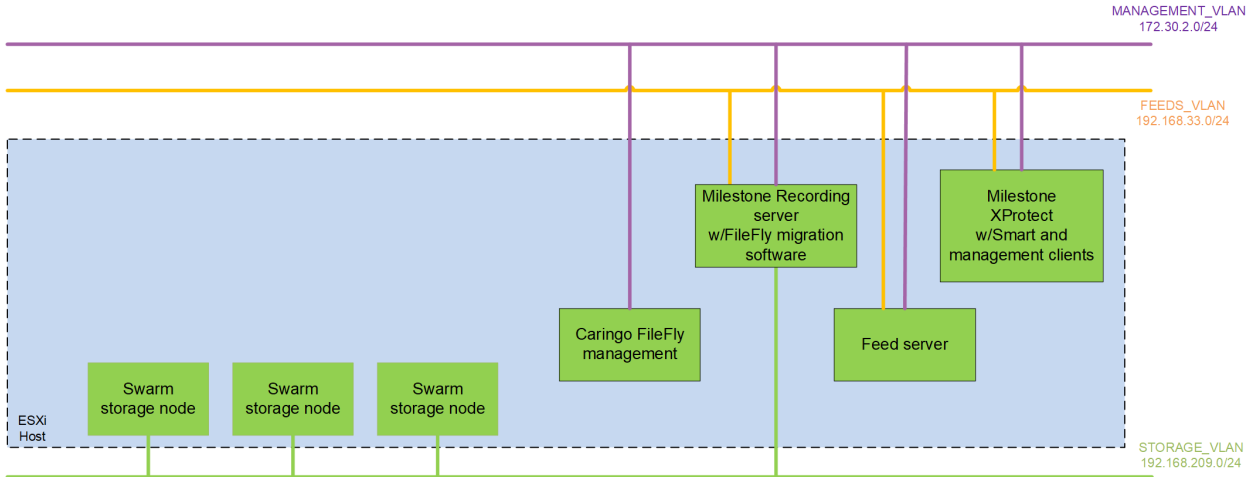


Fig.1. Network and components diagram example

Workflow example

In this section a full workflow example will be explained.

A simulated video surveillance fee provided by StableFPS have been configured for this workflow demo, with its remote path pointing to a RAM disk presented by the Feed server.

The Recording server has two disks attached. The drive letter **R:** is used for live recordings, while the unit **M:** is used to archive data.

In this example, all the video surveillance data that is more than 1 hour old is moved from **R:** to **M:** by XProtect VMS as it is shown in the following screenshot.

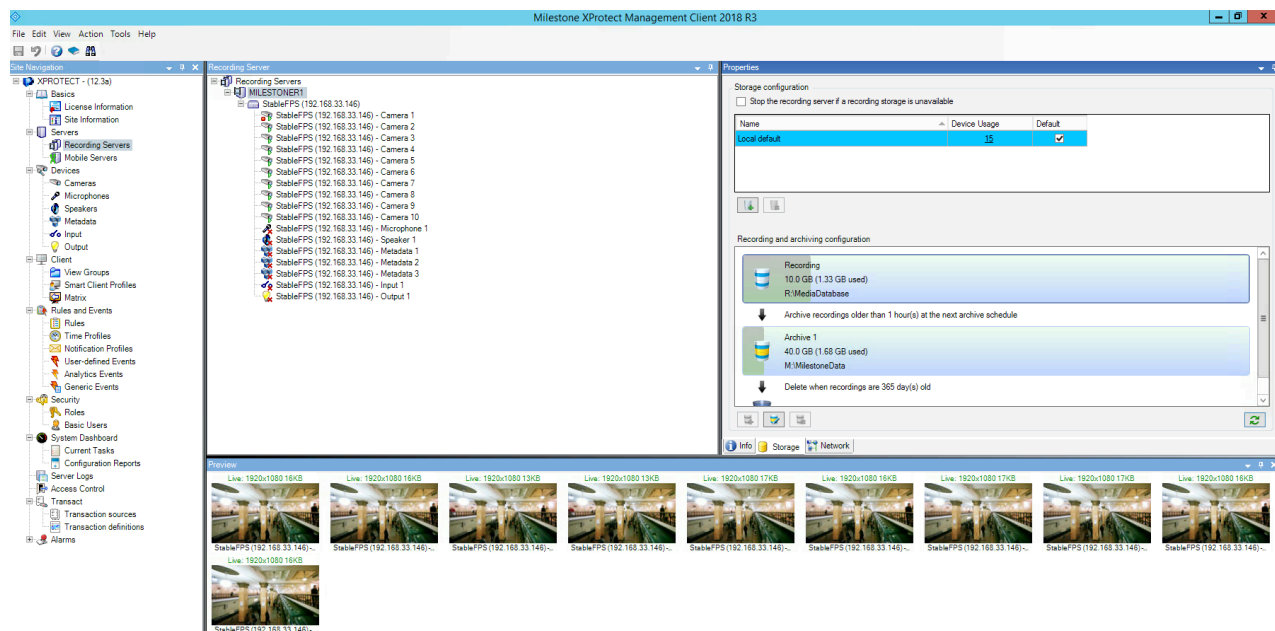


Fig.2. Recording server configuration

The next step is to migrate the media assets stored in the archive drive letter to the object storage platform. Based on rules and policies such as type of the file, size, last time modified etc. as well as source and destination, a migration task can be configured in FileFly.

In this example, the FileFly task will move the contents of the *.blk media files that are stored on **M:** to Caringo Swarm, leaving stubs in **M:**.

A stub is a 4KB file that acts as a pointer, the file content is no longer present in the Windows folder, as it has been moved to secondary storage: Swarm. When a user or application tries to access a file that has been migrated, a re-hydration process takes place automatically, bringing those contents back to primary storage.

The result is the end-user does not perceive any difference accessing the information while the surveillance data is consolidated on a limitless storage platform.

The FileFly tasks can be scheduled, so they will migrate back to Swarm the files that matches any given criteria, i.e. surveillance assets that are no longer needed on primary Windows storage as they have been watched already.

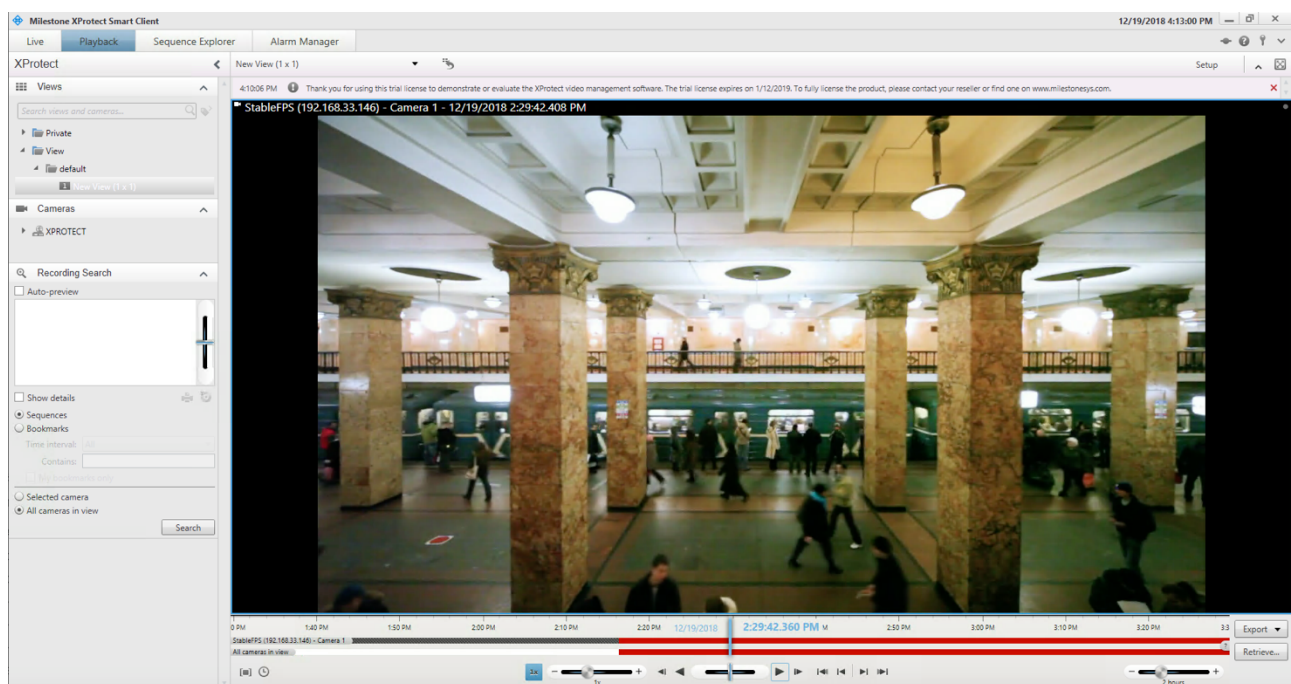


Fig.3. Watching archived footage with XProtect Smart client

The following screenshot show how the data is stored in the Windows server. The media files that have not been accessed remain migrated from the Windows server (*block39.blk*, *block40.blk* in the screenshot below) in an offline state with a size on disk of zero, while parts of the footage that are being watched are automatically rehydrated (*block38.blk* in the screenshot below).

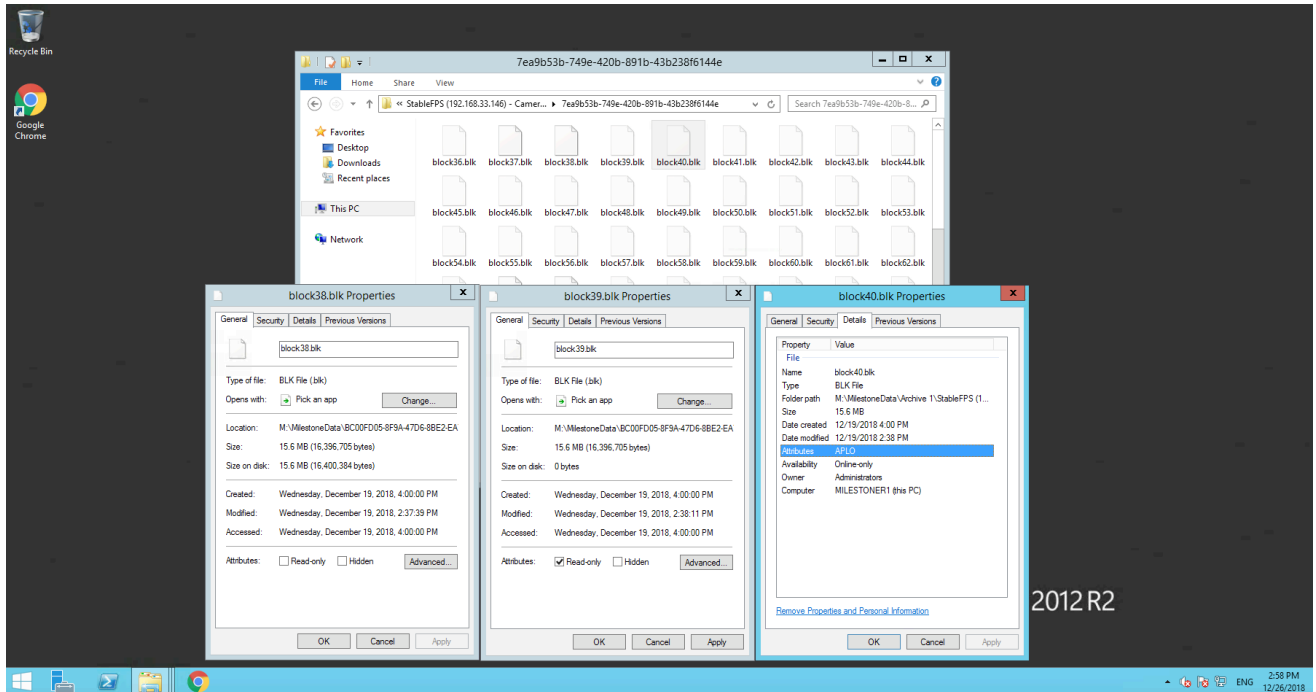


Fig.4. Recording server with contents archived in Caringo Swarm

The next one shows how the surveillance data is stored in Swarm. It includes rich metadata such as the name of the file, path, attributes etc.

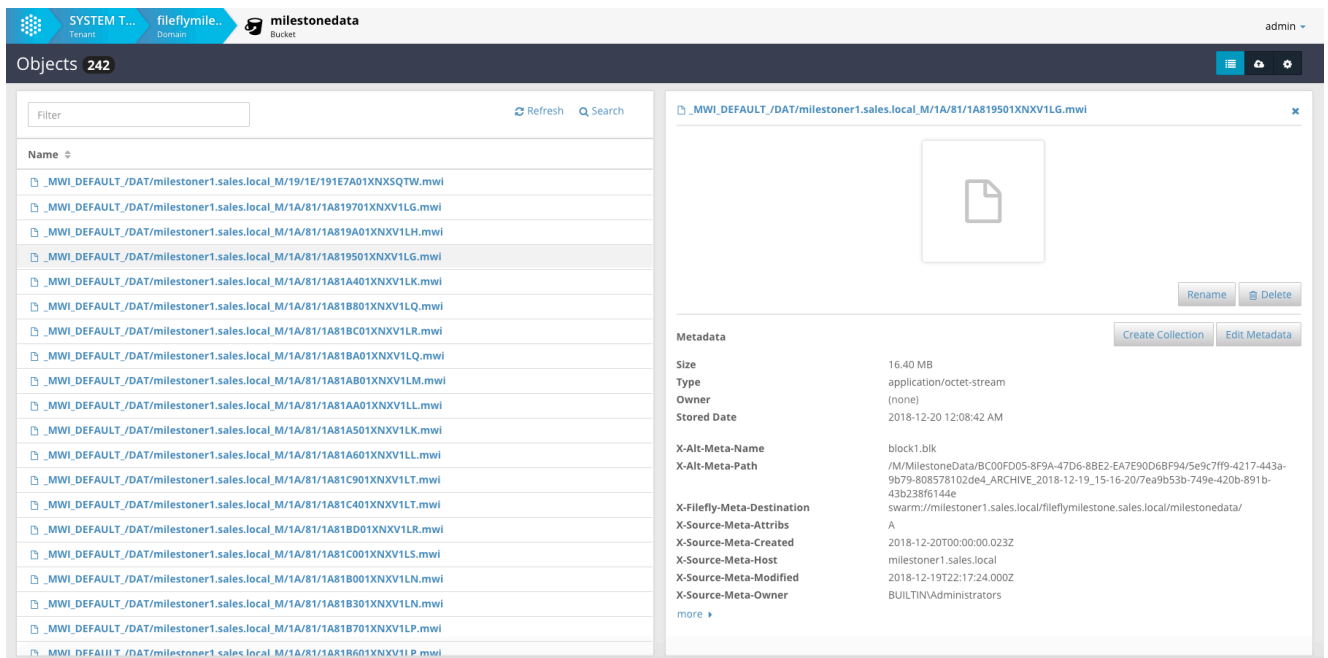


Fig.5. Milestone data stored in Caringo Swarm

Caringo FileFly configuration

The following screenshots show how FileFly is configured as part of this workflow example. Please, for more information, refer to the FileFly documentation available in both Caringo Connect website as well as Milestone Marketplace.

There are several steps:

- First, the FileFly migration software is installed in the recording server and it must be recognized by the FileFly admin portal (fig.6.)
- The source path must be defined. In this case **M:\MilestoneData** (fig.7)
- As well as the destination, which in this case is pointed to the storage domain *fileflymilestone.sales.local* and bucket *milestonedata*. The storage domain in Swarm is just an end-point, a logical entity that can be configured with its own data protection scheme, quota, permissions etc. if desired. (fig.8)
- The rules are next, for these tests only the media files (*.blk) are migrated, other small files such as indexes are not migrated as they are accessed all the time by Milestone XProtect VMS. (fig.9)
- The policies define the action FileFly will perform. In this case, migrate all *.blk files from the source to the destination. (fig.10)
- Finally, the task is created, it can be run manually or be scheduled using the calendar (fig.11)

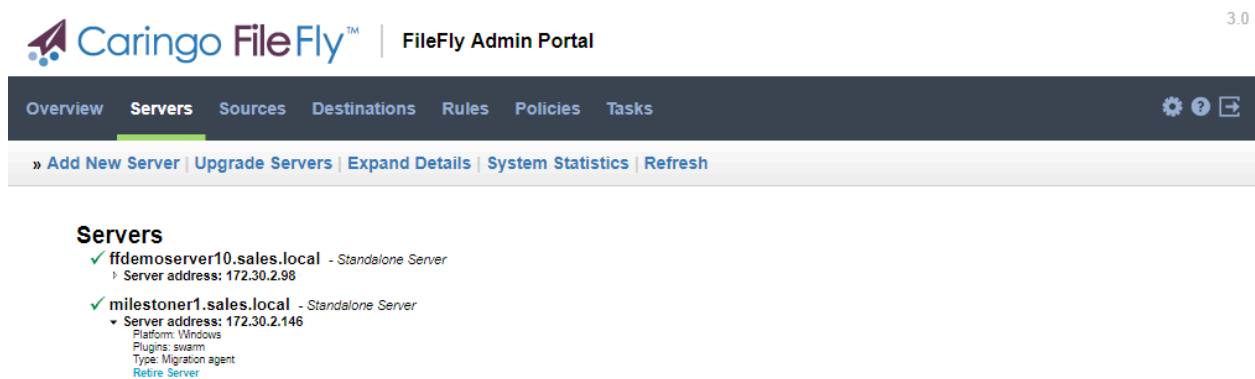


Fig.6. Recording server configured in FileFly

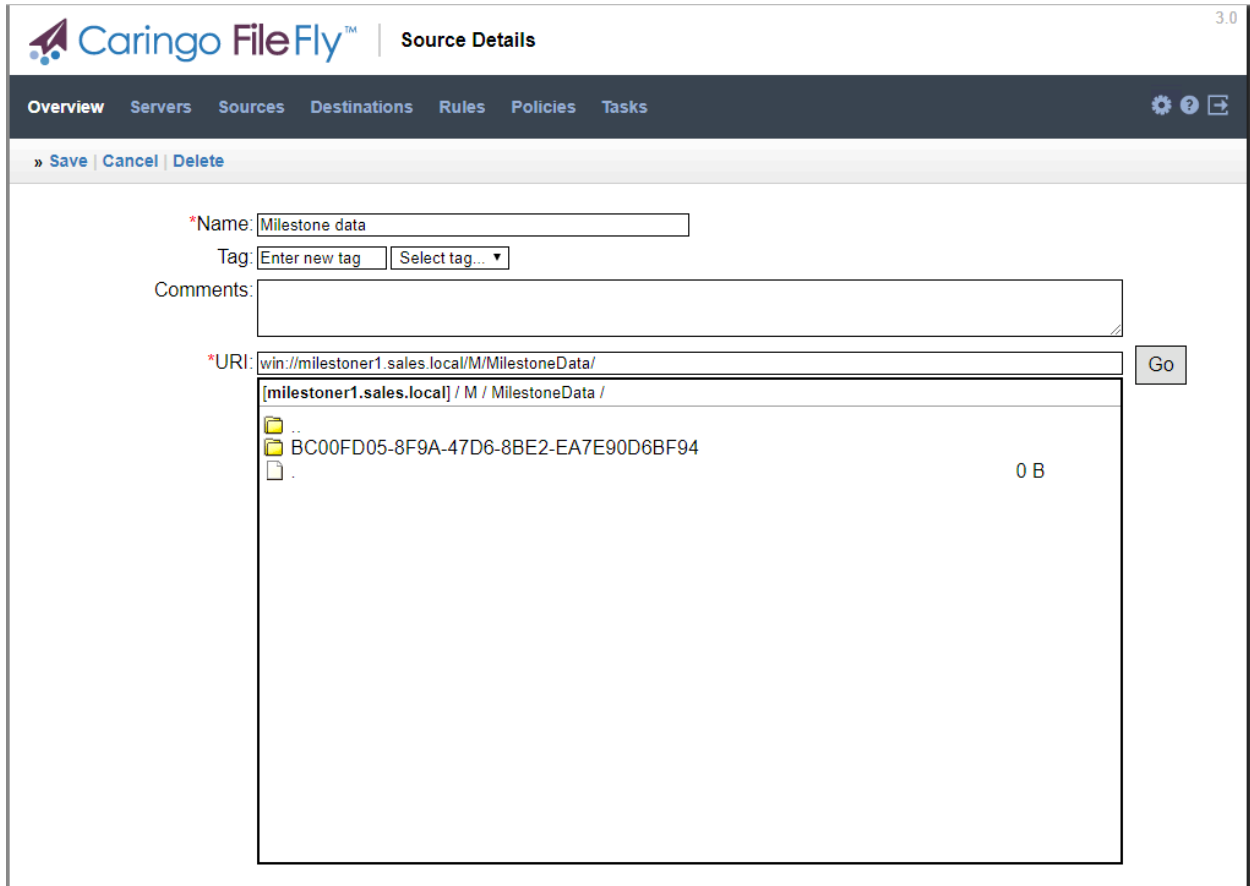


Fig.7. Source details

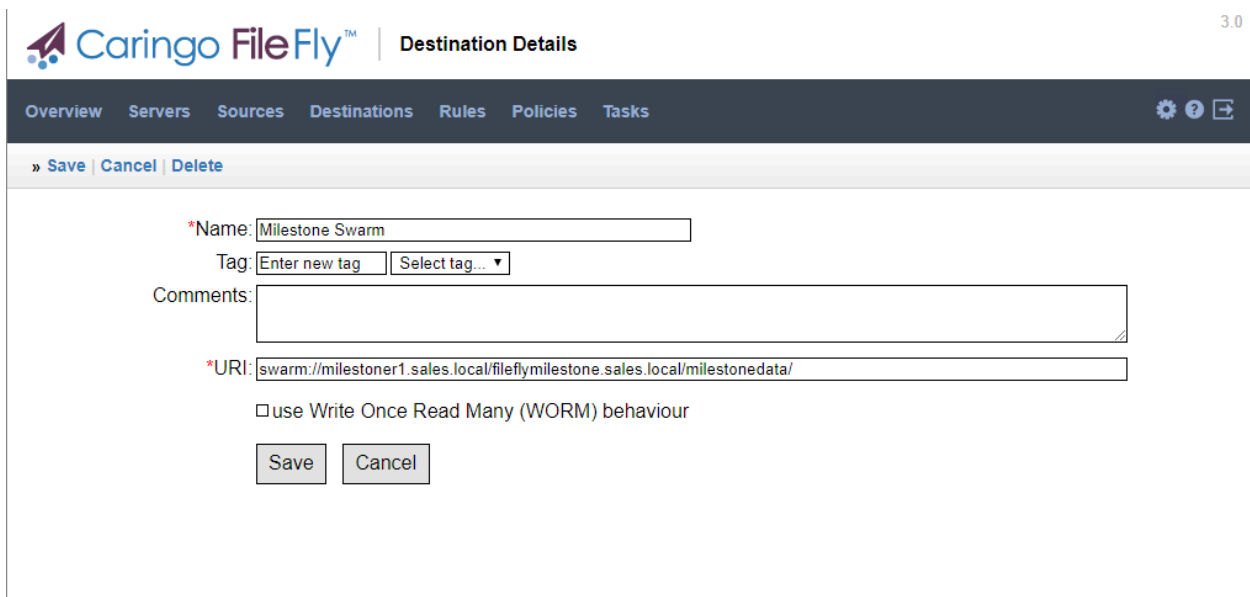


Fig.8. Destination details

Overview Servers Sources Destinations Rules Policies Tasks

» Save | Cancel | Delete

*Name:

Comments:

Negate:

File Matching

Patterns:

Use * to match all filenames
eg. *.doc, 2018-01-???.log (wildcard matching)
eg. /.*\.doc/, /2018-01-.*\.log/ (regular expression matching)

Size Matching

*Min size:

Max size:

Date Matching

After:

Before:

Age: old

Owner Matching

Patterns:

eg. fred.myorgunit.myorg, wilma.myorgunit.myorg
eg. MyDomain*, *.myorg (wildcard matching)

Fig.9. Rule details.

Overview Servers Sources Destinations Rules Policies Tasks
⚙️ ⓘ ↻

» Save | Cancel | Delete

*Name:

Comments:

Operation

*

Migrate file data from selected Sources(s) to Destination. Stub files remain at the Source location as placeholders until files are demigrated.

Demigrated files that have not been modified will be quick-remigrated to their original locations where possible.

skip *quick-remigrating* files demigrated less than day(s) ago

pause execution during work hours (see Settings page)

Rules

Available:

all files
 all files except databases
 not databases

Selected:

milestone blk files

Sources

Tag Filter:

Available:

[Untagged]:
 DemoShare
 img2-IIS
 Milestone data
 nfsexport

Selected:

[Untagged]:
 Milestone data

Hint: use shift/ctrl for multiple selection

Destination

Tag Filter:

Root:

Fig.10. Policy details

Policies

Available:

create drtool from source
drtool destination img2
drtool source img2
gather stats
GC scrub
migrate img2
Rehydrate demo
smart tier all files except databases

Selected:

Migrate BLK files MilestoneData

Hint: use shift/ctrl for multiple selection

Schedule

Enable:

Min	Hour	Day	Month	DoW
00	00	01	January	Sunday
05	01	02	February	Monday
10	02	03	March	Tuesday
15	03	04	April	Wednesday
20	04	05	May	Thursday
25	05	06	June	Friday
30	06	07	July	Saturday
35	07	08	August	
40	08	09	September	
45	09	10	October	
50	10	11	November	
55	11	12	December	

Hint: use shift/ctrl for multiple selection

Fig.11. Task details

Conclusion

Caringo Swarm is a software-defined object storage platform that can scale-out to hundreds of petabytes within the same cluster, providing a single pool of resources to consolidate information, with built-in data protection mechanisms and multi-tenancy capabilities.

FileFly is a transparent data migration solution based on policies that enables Windows applications to work seamlessly with Caringo Swarm.

These tests prove that Caringo Swarm with FileFly work consistently as an archive target for Milestone XProtect VMS surveillance data, since they provide a fast and straightforward access to these assets stored in a limitless storage platform avoiding the hassle of running manual restore operations.