

Instructions on Submitting a Claim for the Cohesity SmartFiles “Less-is-More” Guarantee

This document expands on the instructions in the Guarantee’s terms and conditions and answers some frequently asked questions.

Process for Submitting a Claim

1. Execute the test on both the other NAS system and Cohesity SmartFiles following the test conditions and steps outlined below.
2. Submit the claim evidence via email to less-is-more@cohesity.com with all the information outlined below.
3. Cohesity will validate the claim and respond with a decision. Validation may require the ability for a Cohesity employee to access the other NAS system and the Cohesity cluster to verify that the testing requirements were satisfied and the correct metrics were collected.

Test Conditions

1. Data
 - The same dataset must exist on both the competitor’s product and the Cohesity cluster and must be 10TB or greater.
 - The dataset may reside on multiple “volumes” (or equivalent storage container construct) on the competitive product and on multiple Views on the Cohesity cluster. The “less-is-more” guarantee will apply to storage efficiency of the entire dataset and not individual volumes or Views.
 - If multiple volumes and Views are used, no other data may reside on those volumes and Views.
2. System configuration - competitive product
 - If multiple volumes are used, the volumes under test cannot be sharing deduplication, compression, or other storage efficiency with volumes that are not under test.
 - The system can be enabled with any combination of deduplication, compression, and small file efficiency features available natively in the product.
3. System configuration - Cohesity SmartFiles
 - Either a physical cluster or VE can be used.
 - If multiple Views are used, the Views under test must be in a single Storage Domain. The Storage Domain cannot contain Views not under test.
 - The Storage Domain and the Views under test must have inline Dedupe and inline Compression enabled.
 - Enable Helios and RT (reverse tunnel) so Cohesity staff can examine the data remotely.

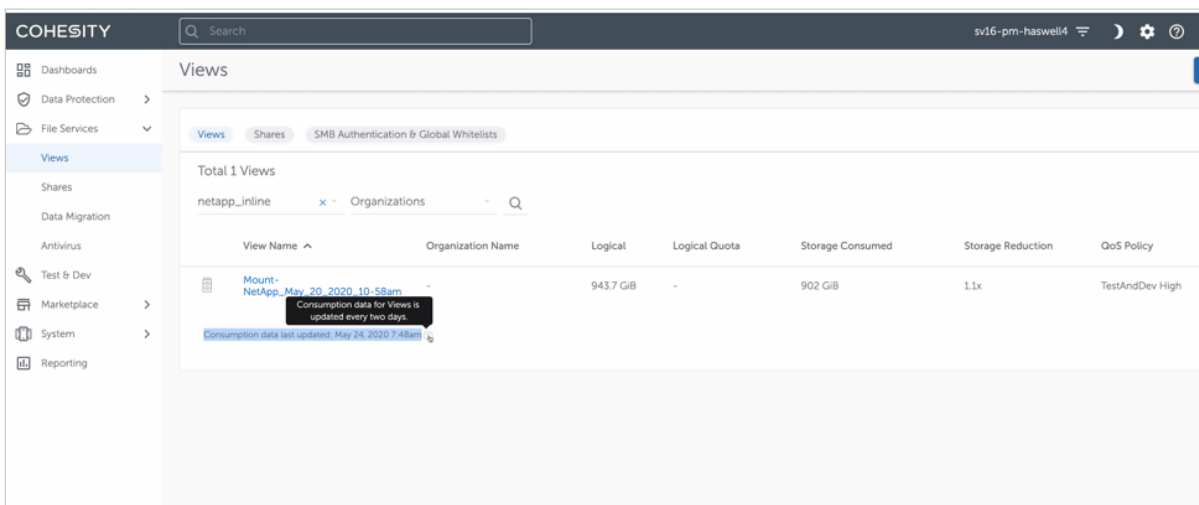
Test Process

Competitive Product

1. Create a new volume or volumes on the competitive product to host the test dataset.
2. Capture the usable storage capacity available on the volume(s) and system configuration settings.
 - Use the product's documented procedure and interfaces for capturing the usable capacity. This is after resilience (RAID/RP/EC).
 - Use the product's documented procedure for exporting the complete system configuration in human readable format such as a text file, PDF, or a screenshot of metrics from the product UI.
3. Copy or migrate the test dataset to the volume(s).
4. Wait until all storage efficiency operations have completed.
5. Capture the usable storage capacity available on the volume(s).

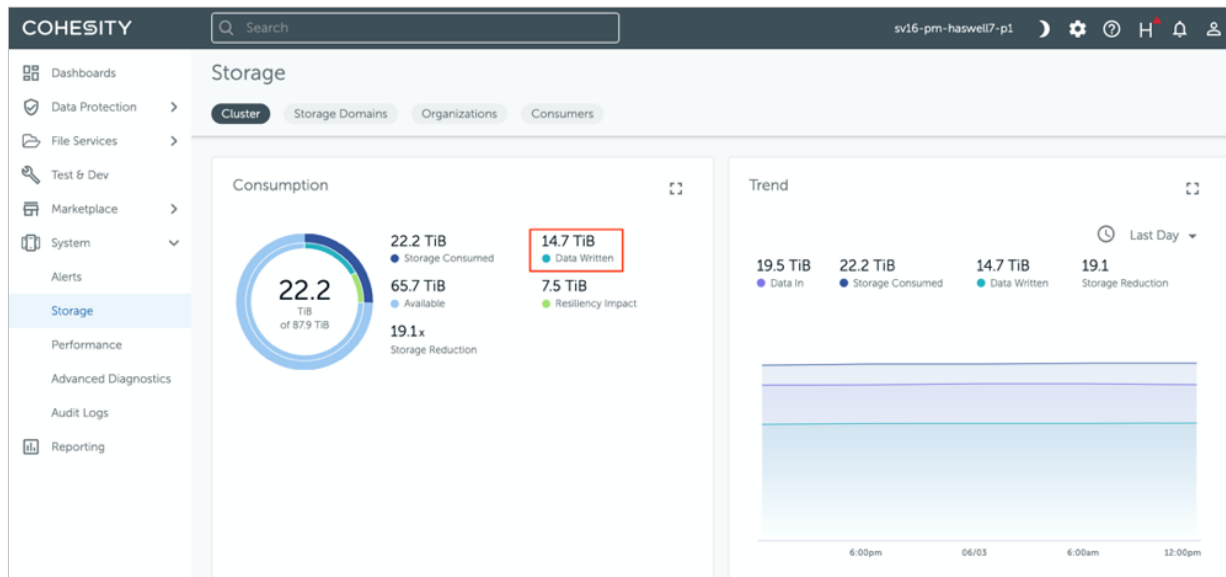
Cohesity

1. Create a new Storage Domain and new View(s) to host the test dataset.
 - Configure the Storage Domain as follows (all other settings may be left to default or set appropriately for your environment)
 - Inline Dedupe and Compression enabled
 - Configure Views
2. Copy or migrate the test dataset to the View(s).
 - Option 1: Backup and Clone (for NFS/SMB sources)
 - Use Cohesity's Data Protection to backup the competitive product's volume(s) under test to the Cohesity cluster
 - Once the backup is completed, clone the volume(s) into View(s).
 - Option 2: traditional copy (for NFS, SMB, or S3 sources)
 - Use any copy command or tool to copy the dataset to the Cohesity View(s).
3. Wait until all storage efficiency operations have completed. This can take a few days to complete the maximum data reduction.
 - To confirm that the storage efficiency operations have completed, check the Views page confirm that the "Consumption data last updated" shows a date and not "N/A".



4. Capture the usable storage capacity used in the Storage Domain.

- Log in to Cohesity DataPlatform.
- Navigate to Cluster > Summary
- Select Storage Domains -> Select Storage Domain where data is stored
- Select Storage Statistics
- Review Data Written

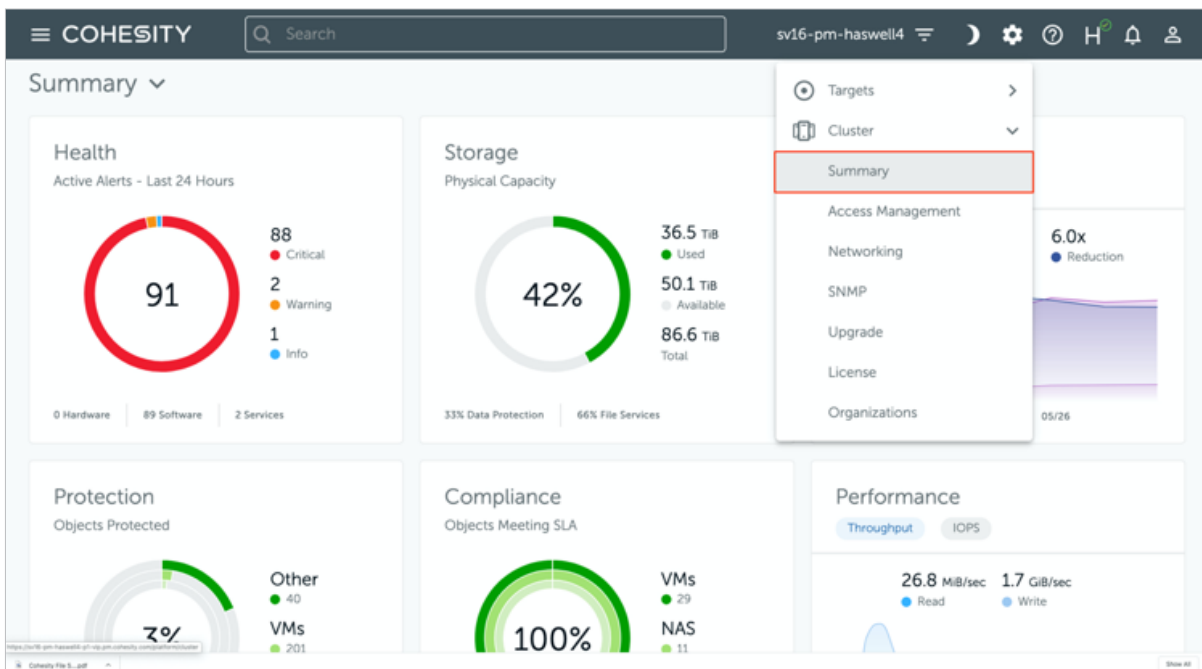


What to Submit in a Claim

In the email to less-is-more@cohesity.com, please provide the following information:

- Date(s) when the test was executed
- Information on the competitor's product:
 - product name
 - version of OS or storage software
 - Configuration of the product (names of all volumes, deduplication settings, compression settings)
 - Name of volume(s) storing the test dataset
- Information on the Cohesity product:
 - Version of DataPlatform (6.4.1 or later, Standard Edition)
 - Configuration of the Storage Domain and volume(s)
 - Name of volume(s) storing the test dataset
- Information on the dataset
 - Description of the dataset (e.g. user home directories)
 - Total logical capacity of the data

- Logical Capacity is the amount of data being written to the system. The data is physically reduced by efficiency technologies such as deduplication and compress.
 - From Windows, the File Explorer will show the logical size of the data, as the server is unaware of the storage deduplicating the data.
 - From Unix/Linux, du will show the logical size of the data.
- # of files
- MD5 hash (optional, this may be required)
 - Use find to list all the files in the directory then calculate the md5 hash for each file and pipe it sorted by filename to a file:
 - `find /dir1/ -type f -exec md5sum {} + | sort -k 2 > dir1.txt`
 - Do the same procedure to the another directory:
 - `find /dir2/ -type f -exec md5sum {} + | sort -k 2 > dir2.txt`
 - Then compare the result two files with diff:
 - `diff -u dir1.txt dir2.txt`
- Evidence of usable capacity on the competitor's product before writing the dataset and after writing the dataset. The total usable capacity consumed by the dataset is the difference between these two measurements.
- Evidence of usable capacity on the Cohesity SmartFiles View before and after writing the dataset. The total usable capacity consumed by the dataset is the difference between these two measurements.
- Cluster ID
 - 1. Log in to Cohesity DataPlatform.
 - 2. Navigate to Cluster > Summary




- 3. Summary page will show the Cluster ID

Summary Storage Domains Nodes Key Management System

Cluster Summary

86.6 TiB
Total Size



- Free 50.1 TiB
- Used 36.5 TiB

Cluster Name	sv16-pm-haswell4
Cluster ID	6165002496174814
Creation Date	May 16, 2019 12:57pm
Software	6.5.0a_release-20200423_b9986531
Hardware	C2500
Encryption	Disabled
Storage Domains	26
Nodes	4
Support Channel	Off ✎
Storage Capacity for Metadata	5 TiB
Storage Used for Metadata	23%

Configure Upgrade