

BEST PRACTICES GUIDE FOR MAINTAINING ENOUGH FREE SPACE ON ISILON CLUSTERS AND POOLS

October 31, 2016

Abstract

Consuming all available storage space on an Isilon cluster or pool can have a major impact on your workflow and the cluster's data protection capabilities. However, it is one of the easiest problems to avoid. This document describes the issues that can arise from out-of-space situations, and provides guidelines to help you avoid them.

Table of Contents

Implications of having a full cluster or pool.....	3
Error messages indicating cluster or pool is full or nearly full	3
What to do if you have a full or nearly full cluster or pool	4
Best practices to maintain enough free space on your cluster or pool	4
Maintain sufficient free space	4
Maintain appropriate protection levels.....	4
Monitor cluster capacity.....	4
Manage your data	5
Provision additional capacity	5
Manage snapshots	6
Make sure all nodes are supported	7
Make sure all nodes in a pool are compatible.....	7
Enable Virtual Hot Spare	7
Enable Spillover.....	7

Implications of having a full cluster or pool

When a cluster or pool is more than 90% full, the system can experience slower performance and possible workflow interruptions in high-transaction or write-speed-critical operations.

Further, when a cluster or pool approaches full capacity (over 98% full), the following issues can occur:

- Substantially slower performance
- Workflow disruptions (failed file operations)
- Inability to write data
- Inability to make configuration changes
- Inability to run commands that would ordinarily be used to free up space
- Inability to delete data
- Data unavailability
- Potential failure of client operations or authentication (to connect/mount and navigate)
- Potential for data loss

Allowing a cluster or pool to fill can put the cluster into a non-operational state that can take significant time (hours or even days) to correct. Therefore, it is important to keep your cluster or pool from becoming full.

Error messages indicating cluster or pool is full or nearly full

You might get any of the following error messages when attempting to write to a full, or nearly full, cluster or pool:

Error	Where the error appears
The operation can't be completed because the disk "share name" is full.	OneFS web administration interface, or the command-line interface on a Mac with an NFS mount.
No space left on device.	OneFS web administration interface, or the command-line interface on a Mac client.
No available space.	OneFS web administration interface, or the command-line interface of a Windows client.
ENOSPC (error code)	On the cluster in <code>/var/log/messages</code> . This error code will be embedded in another message.
Failed to satisfy layout preference.	On the cluster in <code>/var/log/messages</code> .
Disk Quota Exceeded.	Cluster command-line interface, or an NFS client when you encounter a Snapshot Reserve limitation.

What to do if you have a full or nearly full cluster or pool

Follow the best practices listed below to keep your cluster or pools from getting too full. If your cluster or pool is full, or nearly full, and you are unable to maintain enough free space on it, contact Isilon Technical Support.

Best practices to maintain enough free space on your cluster or pool

Follow these guidelines to ensure your cluster or pool always has enough free space to keep your workflows active and protect your data.

For more information about adding additional features or capacity to your system, contact your sales representative.

For instructions on how to use the features described below, see the [OneFS Administration Guide](#). Alternatively, consult the [OneFS Online Help](#).

Maintain sufficient free space

Having enough free space allows both you and Isilon Technical Support time to react appropriately in case of a node failure.

For optimal performance in any size cluster, EMC recommends that you maintain 10% free space in each pool of a cluster.

To better protect smaller clusters (containing 3 to 7 nodes) we recommend that you maintain 15 to 20% free space. A full smartfail of a node in smaller clusters may require more than one node's worth of space. Maintaining 15 to 20% free space can allow the cluster to continue to operate while EMC helps with recovery plans.

Plan for contingencies: having a fully updated backup of your data can limit the risk of data loss if a node fails.

Maintain appropriate protection levels

Make sure that your cluster and pools are protected at the appropriate level. Every time you add nodes, re-evaluate protection levels. For more information about protection levels, see the [OneFS Administration Guide](#) and *How to determine if an Isilon cluster is in a window of risk for data loss*, [article 336424](#).

OneFS 7.2 and later includes a function that calculates a suggested protection level based on cluster configuration, and alerts you if the cluster falls below the suggested level. For more information about this feature, see the “Suggested protection” section of the [OneFS Administration Guide](#).

Monitor cluster capacity

- **Configure alerts.** Set up event notification rules so that you will be notified when the cluster begins to reach capacity thresholds. Make sure to enter a current email address in order to receive the notifications. For instructions, see the “Events and alerts” section of the [OneFS Administration Guide](#).
- **Pay attention to alerts.** The cluster sends notifications when it has reached 95 percent and 99 percent capacity. On some clusters, 5 percent (or even 1 percent) capacity remaining might mean that a lot of space is still available, so you might be inclined to ignore these notifications. However, it is best to pay attention to the alerts, closely monitor the cluster, and have a plan in place to take action when necessary.

- **Pay attention to ingest rate.** It is very important to understand the rate at which data is coming in to the cluster or pool. Some ways to do this are:
 - SNMP
 - SmartQuotas
 - FSAnalyze
- **Use SmartQuotas** to monitor and enforce administrator-defined storage limits. SmartQuotas manages storage use, monitors disk storage, and issues alerts when disk storage limits are exceeded. Although it does not provide the same detail of the file system that FSAnalyze does, SmartQuotas maintains a real-time view of space utilization so that you can quickly obtain the information you need. A SmartQuotas license is required to use this feature. Consult with your account team and Isilon Technical Support to determine a strategy for your workflow and operational needs. If SmartQuotas is not licensed on the cluster, you can use FSAnalyze.
- **Run FSAnalyze jobs.** FSAnalyze is a job-engine job that the system runs to create data for InsightIQ's file system analytics tools. FSAnalyze provides details about data properties and space usage within the `/ifs` directory. Unlike SmartQuotas, FSAnalyze updates its views only when the FSAnalyze job runs. Since FSAnalyze is a fairly low-priority job, it can sometimes be preempted by higher-priority jobs and therefore take a long time to gather all of the data. An InsightIQ license is required to run an FSAnalyze job.

Manage your data

- Regularly archive data that is rarely accessed.
- Regularly delete unused data.
- Ensure that pools do not become too full by setting up file pool policies to move data to other tiers and pools.

Provision additional capacity

To ensure that your cluster or pools do not run out of space, you can create more clusters, replace smaller-capacity nodes with larger-capacity nodes, or add new nodes to existing clusters or pools. If you decide to add new nodes to an existing cluster or pool, contact your sales representative to order the nodes long before the cluster or pool runs out of space. EMC recommends that you begin the ordering process when the cluster or pool reaches 80% used capacity. This will allow enough time to receive and install the new equipment and still maintain enough free space.

Suggested timeline:

Used space	Action
80%	Begin planning to purchase additional storage.
85%	Receive delivery of the new storage.
90%	Install the new storage.
95%	There is a possibility of data unavailability if capacity is not increased.

Manage snapshots

Sometimes a cluster has many old snapshots that take up a lot of space. Reasons for this include:

- The cluster is configured to take a lot of snapshots, but is not configured to have the snapshots expire as quickly as they are created.
- There is a device that is down or smartfailed on the cluster (in other words, the cluster is in a “degraded protection” state). You cannot run a SnapshotDelete job (or any job other than FlexProtect or FlexProtectLin) while the cluster is in this state. This could cause the number of snapshots to increase without being noticed.
- There is no SnapshotIQ license on the cluster, so you cannot delete old snapshots that still exist.

Best practices for managing snapshots:

- Configure the cluster to take fewer snapshots, and for the snapshots to expire more quickly, so that less space will be consumed by old snapshots. Take only as many snapshots as you need, and keep them active for only as long as you need them.
- If you need to delete snapshots and there are down or smartfailed devices on the cluster, or if the cluster is in an otherwise “degraded protection” state, contact Isilon Technical Support for assistance.
- Avoid taking snapshots of `/ifs`, `/ifs/data`, and `/ifs/home` in favor of more specific targets when possible. In particular, avoid taking nested snapshots, redundant snapshots, or overly scoped snapshots. For example, if you schedule snapshots of `/ifs/data` and `/ifs/data/foo` and `/ifs/data/foo/bar`, consider taking snapshots of only the intermediate or most granularly scoped part (`/ifs/data/foo` or `/ifs/data/foo/bar`).
- If you have a lot of old, unneeded snapshots taking up space, and you do not have a SnapshotIQ license, contact Isilon Technical Support to discuss your options and provide assistance with deleting the old snapshots.
- Do not delete SyncIQ snapshots (snapshots with names that start with SIQ), unless the only remaining snapshots on the cluster are SyncIQ snapshots, and the only way to free up space is to delete those SyncIQ snapshots. Deleting SyncIQ snapshots resets the SyncIQ policy state, which requires a reset of the policy and potentially a full (initial) or differential (target-aware initial) sync. A full or differential sync could take many times longer than a regular snapshot-based incremental sync.
- Delete snapshots in order, beginning with the oldest. Do not delete snapshots from the middle of a time range. Newer snapshots are mostly pointers to older snapshots, and they look larger than they really are. Deleting the newer snapshots will not free up much space. Deleting the oldest snapshot ensures you will actually free up the space. You can determine the snapshot order (if not by name or date) by running the following command:

```
isi snapshot snapshots list
```

The snapshot IDs (first column) are non-conserved, serial values.

NOTE: If the oldest snapshot is a SyncIQ snapshot and you do not want to delete it, you can delete the next-oldest snapshot. Even if the next-oldest snapshot is in the same path, you will still free up some amount of data.

Make sure all nodes are supported

Each version of OneFS supports only certain nodes. Refer to the “OneFS and node compatibility” section of the [Isilon Supportability and Compatibility Guide](#) for a list of which nodes are compatible with each version of OneFS. When upgrading OneFS, make sure that the new version supports your existing nodes. If it does not, you might need to replace the nodes.

Make sure all nodes in a pool are compatible

Space and performance are optimized when all nodes in a pool are compatible. When you add new nodes to a cluster, OneFS automatically provisions nodes into pools with other nodes of compatible type, hard drive capacity, SSD capacity, and RAM. Occasionally, however, the system might put a node into an unexpected location. If you believe that a node has been placed into a pool incorrectly, contact Isilon Technical Support for assistance.

Different versions of OneFS have different rules regarding what makes nodes compatible. For more information on compatible nodes (also called equivalent nodes), see the following resources:

- The “Storage Pools” section of the [OneFS Administration Guide](#) for your version of OneFS
- The “Node class compatibility” section of the [Isilon Supportability and Compatibility Guide](#)
- [Node Compatibility \(Equivalency\) and SSD Compatibility in OneFS 7.2 and Later - Isilon Info Hub](#)

Enable Virtual Hot Spare

The purpose of Virtual Hot Spare (VHS) is to keep space in reserve in case you need to smartfail drives when the cluster gets close to capacity. Enabling VHS will not give you more free space, but it will help protect your data in the event that space becomes scarce.

VHS is enabled by default. Isilon strongly recommends that you *do not* disable VHS unless directed by a Support engineer. If you disable VHS in order to free some space, the space you just freed will probably fill up again very quickly with new writes. At that point, if a drive were to fail, you might not have enough space to smartfail the drive and re-protect its data, potentially leading to data loss.

If VHS is disabled and you upgrade OneFS, VHS will remain disabled. If VHS is disabled on your cluster, first check to make sure the cluster has enough free space to safely enable VHS, and then enable it by following the steps in *OneFS: How to enable and configure Virtual Hot Spare (VHS)*, [article 471814](#).

Enable Spillover

Spillover allows data that is being sent to a full pool to be diverted to an alternate pool. Spillover is enabled by default on clusters that have more than one pool. If you have a SmartPools license on the cluster, you can disable Spillover. However, it is recommended that you keep Spillover enabled. If a pool is full and Spillover is disabled, you might get a “no space available” error but still have a large amount of space left on the cluster.

Copyright © 2016 EMC Corporation. All rights reserved. Published in USA.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>). For documentation on EMC Data Domain products, go to the EMC Data Domain Support Portal (<https://my.datadomain.com>).