



The Zenoss Enablement Series:

Zenoss Best Practices for Post-Installation Control Center Storage Management

Document Version 500-P1

Zenoss, Inc.

www.zenoss.com

Copyright © 2016 Zenoss, Inc 11305 Four Points Drive, Bldg 1 - Suite 300, Austin, Texas 78726, U.S.A.
All rights reserved.

Zenoss and the Zenoss logo are trademarks or registered trademarks of Zenoss, Inc. in the United States and other countries. All other trademarks, logos, and service marks are the property of Zenoss or other third parties. Use of these marks is prohibited without the express written consent of Zenoss, Inc. or the third-party owner.

Cisco, Cisco UCS, Cisco Unified Computing System, Cisco Catalyst, and Cisco Nexus are trademarks or registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

Flash is a registered trademark of Adobe Systems Incorporated.

Oracle, the Oracle logo, Java, and MySQL are registered trademarks of the Oracle Corporation and/or its affiliates. Linux is a registered trademark of Linus Torvalds.

SNMP Informant is a trademark of Garth K. Williams (Informant Systems, Inc.). Sybase is a registered trademark of Sybase, Inc.

Tomcat is a trademark of the Apache Software Foundation.

vSphere is a trademark of VMware, Inc. in the United States and/or other jurisdictions.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries. All other companies and products mentioned are trademarks and property of their respective owners.

Table of Contents

Introduction	1
Theory	1
Procedures	2
Storage is full. Which Storage is it?	2
Master host.....	2
Resource pool host.....	2
Special Notes on Application Data and Snapshots	3
Monitoring Storage	3
DFS Capacity	4
Expanding available storage	4
Adding additional disks to PV/VG/LV	4
Resizing the DM Base Device	5
Docker	5
ISVCS and Backup Volume Capacities.....	5
Increase Available Storage Space with fstrim	6
Backups	6
Adding Storage for Application Data or Snapshots	7
Additional Resources	9

Introduction

If Control Center is installed, it is very important to regularly monitor the storage requirements of each entity within and including Zenoss Resource Manager. It is important to develop and implement a management strategy. Closely and proactively monitoring storage use and attending to potential issues before they become bigger problems is essential. If Zenoss, or one of its components, runs low on storage space, the system can fail, collected/stored data can become corrupted, and numerous potentially difficult to diagnose errors and messages can be generated. This KB provides a set of best practices to manage storage for Zenoss Resource Manager.

This document applies to **Control Center 1.1.x – Zenoss Resource Manager 5.1.x**

Theory

Zenoss and its components have a strong requirement for stable and ample storage. Storage can be controlled in one of several ways:

- Moving data to another location to free storage space
- Adding addition storage space
- Limiting the data that is collected/stored

The types of issues caused by storage shortage include:

Location	Cause	Fix
General	Storage is full	Determine which storage and make adjustments
General	System/Data corruption	<ul style="list-style-type: none"> i. Run repair commands ii. Restore backups iii. Request Zenoss' support assistance
DFS	Thin pool exhaustion (capacity inadequate)	Expand storage: <ul style="list-style-type: none"> i. Add new disks to PV/VG/LV ii. Expand existing disks and expand VG/LV
	DM Base device full (capacity inadequate)	resize base device: <ul style="list-style-type: none"> • serviced-storage resize command
Docker	Master & Resource Pool Host too small (capacity inadequate)	<ul style="list-style-type: none"> • adjust capacities
	ISVCS Containers too small (capacity inadequate)	Free up temporary space. Containers are long-running and might benefit from fstrim to free unused space.

Location	Cause	Fix
ISVCS	Capacity inadequate	i. add additional space ii. decrease use
Backups	Backup storage too crowded (capacity inadequate)	i. Exclude performance data from backups (CC 1.1.8) ii. add additional space iii. decrease use iv. move off-box

Procedures

The following procedures are recommended best practices for managing storage in Zenoss.

Storage is full. Which Storage is it?

An installation of Control Center requires multiple storage areas, each with its own specific usage. This storage differs between a host configured as a Control Center master and hosts set up as a resource pool. The following describes each of the various required storage areas:

Master host

The Master host requires five storage areas:

- **Docker data storage:**
The Docker data storage area contains images and snapshots of the containers it manages.
- **Control Center internal services data:**
The Control Center internal services data storage area contains a variety of run-time data, including the ZooKeeper data.
- **Application data:**
The Control Center application data storage area provides space for Resource Manager databases and application snapshots.
- **Application data backups:**
The Control Center backups storage area can be located on a local partition or a remote file server. The Control Center browser interface stores backup files in this storage area, and uses a subdirectory for temporary files that are created during a restoration.
- **Control Center metadata (high-availability systems only):**
The Control Center metadata storage area is required to mirror the data between the master nodes in high availability deployments only.

Resource pool host

The Resource pool host uses two storage areas:

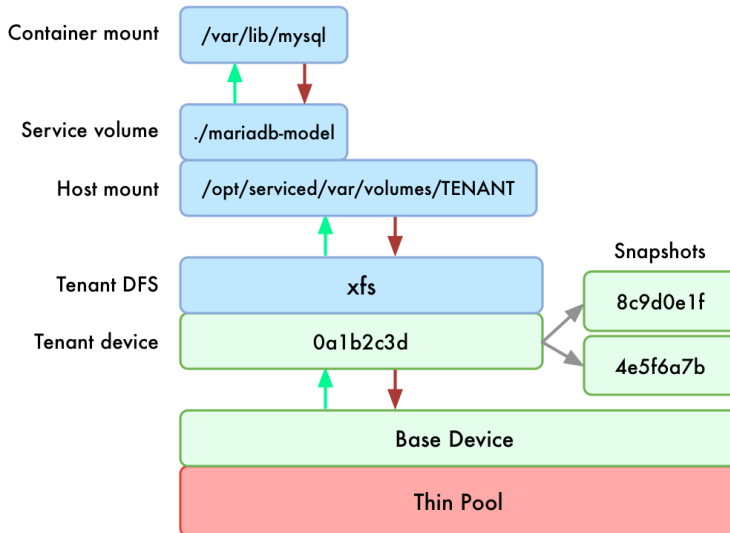
- **Docker data storage:**
The Docker data storage area contains the images and snapshots of the containers it manages.
- **Zookeeper cluster storage:**

If the host is part of a Zookeeper cluster it must be configured with storage for the Zookeeper data.

Special Notes on Application Data and Snapshots

Control Center is able to take snapshots of the Application Data volume. These snapshots are copy on write (COW) that track changes to the file system over time. Initially, a snapshot consumes very little space, but slowly expands incrementally as changes are made to the application data. Note that although removing a 20 Meg file from the application reduces the space used in the application data volume, it increases the size of the snapshot because it tracks the change.

The following diagram shows the file system architecture:



Because application data and snapshots share the same base device, it is extremely important to ensure that snapshots do not fill up the base device storage. For this reason, snapshots should be considered as short term save points only. This means snapshots that require long term saving should instead be made with full backups.

Each of the various storage areas used by Control Center must be appropriately sized for each deployment. Visit the [Zenoss Documentation](https://www.zenoss.com/resources/documentation) site (<https://www.zenoss.com/resources/documentation>) to read the *Resource Manager Planning Guide* for considerations and guidance.

Monitoring Storage

New functionality included in Resource Manager 5.1.5 and Control Center 1.1.7 provides the ability to monitor some storage used by the Control Center. Included with this monitoring are the following four default thresholds:

- Thin pool free space low at 80% used
- Thin pool free space is very low at 90% used
- Tenant Volume free space is low at 80% used
- Tenant Volume free space is very low at 90% used

NOTE: To enable email notifications when these thresholds are breached, add one or more administrator email accounts to these thresholds. Visit the [Zenoss Wiki](http://wiki.zenoss.org/Notify_Me_of_Important_Events) (http://wiki.zenoss.org/Notify_Me_of_Important_Events) for

specific instructions about how to be notified about specific events.

As of Control Center 1.1.7, the command `serviced volume status` can be run on the master host to provide direct information about the storage. For example:

```
[root@ip-10-111-23-200 ~]# serviced volume status
Status for volume /opt/serviced/var/volumes:

Driver:          devicemapper
Driver Type:     direct-lvm
Volume Path:     /opt/serviced/var/volumes

Thin Pool
-----
Logical Volume:          serviced-serviced--pool
Metadata (total/used/avail): 104 MiB / 2.145 MiB (2.1%) / 101.9 MiB (98%)
Data (total/used/avail):    89.01 GiB / 3.137 GiB (3.5%) / 85.87 GiB (96%)

230bew5ab7cmes2tw9husxpq1 Application Data
-----
Volume Mount Point:      /opt/serviced/var/volumes/230bew5ab7cmes2tw9husxpq1
Filesystem (total/used/avail): 100 GiB / 1.753 GiB (1.8%) / 98.25 GiB (98%)
Virtual device size:     100 GiB
```

DFS Capacity

Issues that result from DFS capacity challenges can require actions that include:

- Expanding the available storage due to thin pool exhaustion
- Adding additional disks to expand available storage
- Resizing the DM Base Device

Expanding available storage

To display information about the base device storage, use the `df` and `serviced volume status` commands. If the returned information shows that the DFS is nearing the capacity of the base device, it is necessary to expand the device mapper device beyond its initial base size. To expand the base size, use the `serviced-storage resize` utility.

NOTE: To successfully increase the capacity of the base device, there must be sufficient free space available on the thin pool. Increasing the capacity might additionally require expanding the volume. See the [Expand the Volumes](#) section for or additional information.

Adding additional disks to PV/VG/LV

For information on this procedure, see section titled Adding Storage for Application Data or Snapshots [Adding Storage for Application Data or Snapshots](#).

Resizing the DM Base Device

To display the base device current information, use the `lsblk` command:

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE
```

Example result showing the current 100G base device:

```
...
sdf                               8:80      0    100G      0
disk
├─serviced-serviced--pool_tmeta   253:2      0    104M      0 lvm
│ └─serviced-serviced--pool       253:6      0    89.9G     0 lvm
│   └─docker-8:33-138780768-3H7mNdL7PeuOryy3YB9b5Q 253:15     0    100G     0 dm
...

```

Example

To resize the base device from 100G to 150G:

```
serviced-storage resize -o dm.thinpooldev=/dev/mapper/serviced-serviced-pool \
-d /opt/serviced/var/volumes 150G
```

The `lsblk` command shows the changes after resizing:

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNT
...
sdf                               8:80      0    100G     0 disk
├─serviced-serviced--pool_tmeta   253:2      0    104M     0 lvm
│ └─serviced-serviced--pool       253:6      0    89.9G     0 lvm
│   └─docker-8:33-138780768-3H7mNdL7PeuOryy3YB9b5Q 253:15     0    150G     0 dm
/exports/serviced_volumes_v2/aqnv12s7s3x11ex3x5wz0unpw
...

```

The `df` command shows the new size, for example:

```
df -h /opt/serviced/var/volumes/aqnv12s7s3x11ex3x5wz0unpw/
Filesystem                                Size  Used Avail Use%
Mounted on
/dev/mapper/docker-8:33-138780768-3H7mNdL7PeuOryy3YB9b5Q 148G   19G  123G  14%
/opt/serviced/var/volume
```

Docker

ISVCS and Backup Volume Capacities

ISVCS can run short of space. To ensure this does not occur, the space must be monitored. If the space becomes limited, there are two main options:

- Expand the volume(s)
- Decrease required space usage - accomplished by controlling how long data is kept, selecting which data is kept and restoring unused space to a thin pool volume via the `fstrim` command.

Expand the Volumes

Because the ISVCS and Backups volumes are handled as standard XFS filesystems, the process to expand either of them follows standard Linux procedures for expanding storage. For example:

RedHat Storage Administration Guide - Increasing the size of an xfs file system, located at:
https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Storage_Administration_Guide/xfsgrow.html

Decrease Storage Space Usage

The process for decreasing space usage depends on the system, ISVCS or Backups volumes:

- **ISVCS:**

For ISVCS, storage space usage can be controlled by adjusting two parameters: the *size permitted* or *time to age out*. These values affect the Control Center Logstash data. The values are configurable within the `/etc/default/serviced` file.

The defaults are **10GB** (maximum permitted size) or **14 days** (time to age out), whichever limit is hit first. Searching the `serviced` file for the following lines displays the current settings that can be adjusted:

```
SERVICED_LOGSTASH_MAX_DAYS=14
SERVICED_LOGSTASH_MAX_SIZE=10
```

- **Backups:**

By default, Control Center does not delete backup files. When automating backups through `cron` or other methods, it is a best practice to implement a step to remove backups older than a certain age.

Increase Available Storage Space with `fstrim`

Thin pool storage enables the Linux operating system to provide storage to a device dynamically only as it is required. For example, if a 50 MB file is written to thinly provisioned storage, the operating system must provide additional storage from the pool to that device. If the 50 MB file is deleted, in theory the now unnecessary 50 MB space should be returned to the pool as available space. In practice, the unused space is not returned because returning this space to the pool is not automatic. Instead, the unused storage space is only returned to the pool when the `fstrim` command is run. For example:

```
fstrim -v [/mnt/X]
```

Where `[/mnt/X]` is the mount point for your thin pool volume.

For this reason, `cron` is configured by default to automatically run `fstrim` on the application data device to return unused storage to the pool.

NOTE: If there are major changes to the device it is useful to run `fstrim` manually to free the space sooner without waiting for the next `cron` cycle.

By default, `fstrim` discards all unused blocks in the filesystem. Options can be used to modify this behavior based on range or size. See the `fstrim` man page (<http://man7.org/linux/man-pages/man8/fstrim.8.html>) for additional information.

Backups

It is very important that backups be performed on a regular basis no matter whether HA or DR is in use. Backups should exist for the production environment and potentially the system of record from the initial deployment. Having accurate and tested backups can mitigate problems caused by software or hardware issues.

Control Center has the ability to both backup and restore the entirety of Control Center and the Zenoss Resource Manager installation (excluding Analytics, see below). Note that a backup can be restored either to the system it was taken on or to an alternate system. The backup contains all necessary information to facilitate restoration.

Restoration onto the same system reverts the Control Center resource host, its service configuration and the Resource Manager state to the point in time when the backup was initiated. Backups include:

- All Zenoss Resource Manager modeling (what devices are being monitoring and how they are being monitored).
- Event and monitored performance data.
NOTE: This assumes no external HBase cluster configuration for the performance data - see additional information, below.

Backups should be stored on a machine other than the Control Center master. Backups should be copied or migrated to an off-system location for safe keeping and to help regulate storage space usage on the master.

Backup restoration must be successfully tested on a regular basis. Zenoss recommends performing backups on a regular basis to non active resource targets such as separate DR systems, or a DEV, STAGING or QA environment. This restoration from backup test is necessary to demonstrate that the backup restores successfully. It also demonstrates that the restored system is an accurate representation of the state of the deployment when the backup was performed. This means that data, including Device Classes, performance data, etc. is representative and up to date for when the backup was taken.

If possible, it is recommended to have the system you restore any backup on able to mirror at least one device from your production system.

In addition, ensure there is a recent backup that successfully restores before undertaking any of the following:

- Upgrading - perform a daily Zenoss backup and take a VM snapshot if necessary/possible
- Testing - weekly restore testing and confirmation that data is recent
- Using a temporary VM or DR site.

Adding Storage for Application Data or Snapshots

To perform this procedure, an unformatted primary partition is required on the master host. This procedure prepares a partition as a device mapper thin pool on a master host and adds it to the thin pool for applications. The storage can be used for application data or for snapshots.

1. Log in to the master as *root* or as a user with superuser privileges.
2. Identify the current size of the thin pool for applications.

```
lvs -o lv_size serviced
```

Example result:

```
LSize  
150.04g
```

In this example, the current size is 150GB. Record the size for use in a subsequent step.

3. Identify the partition to prepare, and record its size. For example:

```
lsblk -p --output=NAME,SIZE,TYPE Device
```

Example output:

```
NAME          SIZE TYPE
/dev/sdc      150G disk
|--/dev/sdc1  150G part
```

In this example, the partition to prepare is `/dev/sdc1`, and its size is 150GB. Record the size for use in a subsequent step.

4. Create a new physical volume.

```
pvcreate Partition
```

NOTE: Replace *Partition* with the partition identified previously, for example:

```
pvcreate /dev/sdc1
```

On success, the `pvcreate` command displays the name of the new physical volume. Note that the name is the same as the partition.

5. Extend the serviced volume group used for applications:

```
vgextend serviced Physical-Volume
```

NOTE: Replace *Physical-Volume* with the name of the physical volume created in the previous step, for example:

```
vgextend serviced /dev/sdc1
```

6. Add space to the thin pool for applications with the updated volume group.

```
lvextend -L+New-SpaceG serviced/serviced-pool
```

NOTES:

- Replace *New-Space* with the size of the partition, identified in the second step.
- It is necessary to include the units, in this case "G" for gigabytes.
- The `lvextend` command extends both the data and metadata areas of the thin pool.

```
lvextend -L+150G serviced/serviced-pool
```

7. Add storage space for application data.

The preceding steps expand the amount of storage available for snapshots. To use the additional storage for application data, perform this step.

- a) Add the size of the application data storage area, identified in the first step, to the size of the new partition, identified in the second step.

- b) Display the Control Center *tenant ID*.

```
ls /opt/serviced/var/volumes
```

The following output shows an example Control Center tenant id:

```
dovygh6m2gz4zgdzxb2ud70g9
```

- c) Update the Control Center.

```
serviced-storage resize -d /opt/serviced/var/volumes \
-o dm.thinpooldev=/dev/mapper/serviced-serviced--storage \
Tenant-ID Total-SpaceG
```

NOTES:

- Replace Tenant-ID with the identifier displayed in the previous substep, and replace Total-Space with the total amount of application data storage space.
- It is necessary to include the units, in this case "G" for gigabytes.

For example:

```
serviced-storage resize -d /opt/serviced/var/volumes \  
-o dm.thinpooldev=/dev/mapper/serviced-serviced--storage \  
dovygh6m2gz4zgdzxb2ud70g9 150G
```

Additional Resources

The following resources contain additional information concerning storage and management for Zenoss:

- [About Resource Manager Version 5.0 \("Europa"\) Data Stores](https://support.zenoss.com/hc/en-us/articles/202434009)
(https://support.zenoss.com/hc/en-us/articles/202434009)
- [Error: "no space left on device" When Backing Up with Control Center](https://support.zenoss.com/hc/en-us/articles/208578106)
(https://support.zenoss.com/hc/en-us/articles/208578106)
- [How to Configure a Threshold to Generate an Event](https://support.zenoss.com/hc/en-us/articles/202191959) (https://support.zenoss.com/hc/en-us/articles/202191959)
- [How to Remove Unwanted Docker Images](https://support.zenoss.com/hc/en-us/articles/208468616)
(https://support.zenoss.com/hc/en-us/articles/208468616)
- [How to Use Zenoss Enterprise Backup \(ZEB\) to Back Up the Resource Manager Databases](https://support.zenoss.com/hc/en-us/articles/202394999)
(https://support.zenoss.com/hc/en-us/articles/202394999)
- [How to Recreate a Control Center 1.1x Thin Pool](https://zenoss.zendesk.com/knowledge/articles/210876686)
(https://zenoss.zendesk.com/knowledge/articles/210876686)