

[Subscribe](#)[Share ▼](#)[Past Issues](#)[Translate ▼](#)[View this email in your browser](#)

Welcome to the first round of proactive notifications from Zenoss Support. You're receiving this email because you have an account at [support.zenoss.com](http://support.zenoss.com). The goal of this newsletter is to provide you with relevant news and information about product releases, important KB articles, and other support-related information.

## Recent Product Releases

- [Resource Manager 5.1.5 / Control Center 1.1.7 Release Notes](#)
- [Service Impact Release Notes 5.1.3](#)
- [Zenoss Quarterly Zenpack Update for 2016-Q2](#): Significant work has been done on a variety of Zenpacks.

## News for all supported versions (4.2.4, 4.2.5, 5.x)

- [Zenoss Resource Manager cannot monitor targets with upgraded OpenSSH server](#): Zenoss Resource Manager cannot currently monitor targets with upgraded SSH packages of version 1.6.9 and higher (OpenSSH 6.9+). Zenoss is testing an upgrade of our communication framework version to resolve this incompatibility. Releases for Resource Manager versions 4.2.4, 4.2.5 and 5.1.x are targeted for September dependent on testing results.

## News for 5.1.x

[Subscribe](#)[Share ▼](#)[Past Issues](#)[Translate ▼](#)

results in storage issues for Zenoss Resource Manager, including running out of storage space. When `fstrim` is run on an LVM thin volume it fails with the following log message: `fstrim: /mnt: the discard operation is not supported`. This means that instead of performing the expected `fstrim` procedure and reducing the data usage reported by the `lvs` command, the storage use increases and can fill up the disk. You should update the kernel on your systems to address the `fstrim` issue.

- [How To Set The HBASE TTL \(Time To Live\) Value For Resource Manager 5.1.1](#): Upgrading to Zenoss Resource Manager 5.1.x or higher from 5.0.x can result in the TTL (time to live) HBASE data retention policy being set incorrectly. The TTL setting controls how long performance data is retained for targets within Resource Manager. If this is set incorrectly, you may be at risk of running low on space or not retaining performance data for your desired length of time.
- In CC 1.1.6 we introduced the *serviced volume status* command. Due to a Linux kernel bug, there is a potential to run out of memory when running that command. In CC 1.1.7 we have changed *serviced volume status* to not report `storage.device.allocated` and `storage.snapshot.allocated` statistics to work around this issue. These functions will be restored once an updated Linux kernel is released. If you are running CC 1.1.6, you should upgrade to CC 1.1.7, or add the following option to prevent hitting the kernel bug:  
`"SERVICED_STORAGE_STATS_UPDATE_INTERVAL=999999999"`  
to `"/etc/default/serviced"`.

As always, be sure to check out our Knowledge Center for additional helpful information: [Zenoss Knowledge Center](#).

Best Regards,

Your Zenoss Support Team

[Subscribe](#)

[Share ▼](#)

[Past Issues](#)

[Translate ▼](#)

OWN IT.



---

*Copyright © 2016 Zenoss Support, All rights reserved.*