



CXone Mpower EXPERT

磁盘延迟高、无HDD错误

https://kb-cn-stage.netapp.com/on-prem/ontap/Ontap_OS/OS-KBs/High_Latency_on_disks_with_no_H...

Updated: Wed, 22 Apr 2026 10:05:16 GMT

适用场景

- ONTAP 9
- FAS系统
- 配备硬盘驱动器(HDD)的Cloud Volumes ONTAP系统

问题描述

- 增加了最终NAS或SAN用户的延迟。
- `qos statistics volume latency show` 磁盘延迟为秒、但磁盘直方图上的统计信息显示的时间不超过8毫秒

```
cluster1::> qos statistics volume latency show -vserver vs0 -volume vs0_vol0 Workload ID Latency Network
Cluster Data Disk Qos Max Qos Min NVRAM Cloud FlexCache SM Sync VA -----
----- -total- - 455.00us 158.00us 0ms
297.00us 0ms 0ms 0ms 0ms 0ms 0ms 0ms 0ms vs0_vol0-wid1.. 15658 109.00ms 155.00us 0ms 273.00us 108.
2ms 0ms 0ms 0ms 0ms 0ms 0ms
```

'NetApp provides no representations or warranties regarding the accuracy or reliability or serviceability of any information or recommendations provided in this publication or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS and the use of this information or the implementation of any recommendations or techniques herein is a customers responsibility and depends on the customers ability to evaluate and integrate them into the customers operational environment. This document and the information

- 磁盘对象中仅显示一个或几个磁盘的io_QueueED高、大多数情况下通常仅显示一个磁盘
- 事件日志中没有单个磁盘的硬件错误、磁盘架或堆栈上也没有任何其他可能导致单个驱动器上的I/O排队的硬件问题
 - 驱动器发生故障可能会将I/O队列移至另一个磁盘

示例：`statistics` 命令 `disk` 对象在磁盘上显示高队列、而 `statt` 在磁盘 `0c.23.13` 上显示 100% 繁忙的磁盘 I/O

```
Cluster::> set -privilege diag Warning: These diagnostic commands are for use by NetApp personnel only. Do
you want to continue? {y|n}: y Cluster::*> statistics start -object disk -counter io_queued Statistics collection is
being started for sample-id: sample_148 Cluster::*> statistics show -filter "io_queued>100" Object: disk
Instance: 0d.23.13 Start-time: 12/5/2022 16:48:26 End-time: 12/5/2022 16:51:58 Elapsed-time: 212s Scope:
node1 Number of Constituents: 1 (complete_aggregation) Counter Value -----
----- io_queued 818 1 entry was displayed. Cluster::*> node run -node node1 -command statit -b
```

注：观察到在 2013 年 4 月 23 日排队 30 秒后、可以看到热磁盘正在使用 `statit`

```
Cluster::*> node run -node node1 -command statit -e ... disk ut% xfers ureads--chain-usecs writes--chain-usecs
cpreads-chain-usecs greads--chain-usecs gwrites-chain-usecs /aggr1_node1/plex0/rg2: 0d.23.18 25 69.19 0.01
2.00 2033 31.50 60.67 175 37.69 51.19 169 0.00 .... . 0.00 .... . 0a.21.16 24 69.77 0.01 2.00 6964 32.11 59.56
194 37.66 51.18 90 0.00 .... . 0.00 .... . 0d.22.17 57 231.26 133.56 5.45 2075 26.91 29.41 573 70.79 9.87 641 0.
00 .... . 0.00 .... . 0d.23.22 57 230.68 132.96 5.46 1845 26.83 29.56 646 70.90 9.74 604 0.00 .... . 0.00 .... . 0d.
23.13 95 295.63 198.16 4.10 5472 26.83 29.76 1371 70.63 9.91 1975 0.01 .... . 0.00 .... . 0d.22.18 57 231.26
133.55 5.38 2080 26.84 29.60 561 70.86 9.73 634 0.00 .... . 0.00 .... . 0a.20.18 57 231.69 133.54 5.42 1846 27.
00 29.46 647 71.15 9.78 608 0.00 .... . 0.00 .... . 0a.20.16 57 233.00 134.49 5.48 1879 27.08 30.09 634 71.43 9.
83 594 0.00 .... . 0.00 .... . 0d.22.19 57 231.98 134.18 5.41 2099 26.87 29.67 567 70.93 9.87 646 0.00 .... . 0.00 .
... . Cluster::*> set admin
```