

Linux on System z

Programming Specifications for

LS1213

Query Host Access to Volume support

Version 5 from 2015-03-10 for distribution¹

Line Item Owner: Stefan Haberland
Department 3303 / 7103-03
Schoenaicher Str. 220
71032 Boeblingen, Germany
Phone +49-7031-16-1760
LotusNotes ID: Stefan Haberland/Germany/IBM@IBMDE
e-mail: stefan.haberland@de.ibm.com

¹ git commit 6547eb22e072ea985ee872674dbab5aeb90e72 from 2016-07-06 16:26:03 +0200

1 Query Host Access to Volume support

Owner: Stefan Haberland

Plan Candidate ID: LS1213

1.1 Review

The following people have reviewed and signed-off this specification:

Reviewer	Role	sign-off date
Oliver Stark	<i>Test</i>	2015-03-09
Dorothea Matthaeus	<i>ID</i>	2015-02-24
Mustafa Mesanovic	<i>Performance</i>	2015-02-20
Peter Oberparleiter	<i>Architect</i>	2015-02-26

The status of this specification is: final

1.2 Abstract

With this feature, applications can query if a DASD volume is online to another operating system instances by checking the online status of all attached hosts from the storage server.

The lsdsd command can display this information, and the zdsfs command as well as fdasd and dasdfmt are able to evaluate it.

1.3 Description

The storage server is the only instance knowing about the online status of the device on all attached operating system instances in an LPAR or so called hosts. The DS8000 provides an interface to query this information. It is used to check if a DASD device is set online and therefore potentially is used on another operating system instance. This information can help to reduce the chance for outages or possible data corruption due to concurrent access to DASD volumes from different operating system instances.

The zdsfs tool will use this information to prevent concurrent access from Linux and z/OS. Also the tools fdasd and dasdfmt will check the online count and present either a warning or abort according to the options given.

1.4 Configuring, building and installing

None.

1.5 Setting up the environment

None.

1.6 Using this feature

A new sysfs attribute `host_access_count` is added to the DASD sysfs directory which returns the number of operating system instances that have the corresponding device set online. It will show a number between 1 and 256.

This sysfs entry is usable only for online devices. It will return error ENODEV if used with offline devices and EOPNOTSUPP if the storage server does not support this feature.

```
# cat /sys/bus/ccw/devices/0.0.bf45/host_access_count
13
```

Additionally a new debugfs entry `host_access_list` is introduced showing the details for each host connected to the DASD device. This includes the following information:

pgid

The path group ID is a 22 digit hexadecimal internal number assigned by the operating system when setting the DASD online. This ID uniquely identifies the operating system against the storage server.

status_flags

A 2 digit hexadecimal mask containing status flags. reserved - 0x20, online - 0x50

sysplex_name

A 8 character field containing the name of the SYSPLEX in EBCDIC

supported_cylinder

An integer value describing the maximum number of cylinders per volume supported by the host

timestamp

device reserve timestamp in seconds since January 1, 1970

The information shown in this interface are intended to be parsable by userspace tools.

1.6.1 lsdasd

lsdasd is extended to show the information in a more human readable way. Therefore lsdasd is extended with the option `-H` or `--host-access-list`.

```
# lsdasd -H bf45
Host information for 0.0.bf45
Path-Group-ID      FL  Status  Sysplex      Max_Cyls      Time
=====
88000d29e72964ce8570b8  50  ON      TRX1LNK1     268434453      0
88000e29e72964ce8570c3  50  ON              268434453      0
88000f29e72964ce8570d1  50  ON              268434453      0
88011d29e72964ce8570d4  50  ON              268434453      0
88011e29e72964ce8570d9  50  ON              268434453      0
88011f29e72964ce8570e3  50  ON              268434453      0
88022d29e72964ce8570e6  50  ON              268434453      0
88022e29e72964ce8570ea  50  ON              268434453      0
88022f29e72964ce8570f1  50  ON              268434453      0
88033d29e72964ce8570f7  50  ON              268434453      0
88033e29e72964ce8570fe  50  ON              268434453      0
88033f29e72964ce85710e  50  ON              268434453      0
80004229e72964ce7dce74  00  OFF              65520           0
80004a29e72964ce7db60d  00  OFF              65520           0
80003c29e72964ce8481a6  00  OFF              65520           0
80004629e72964ce7f1c13  70  ON-RSV          65520  1424174863
```

The fields are the same as for the debugfs output except that the flags value is interpreted by the 'STATUS' field in a human readable way. Status values are:

- ON - device online
- OFF - device offline
- ON-RSV - device online and reserved
- OFF-RSV - device offline and reserved

1.6.2 zdsfs, lsdasd, dasdfmt

The zdsfs tool is extended to query the `host_access_count` field to check the number operating system instances using this device. If this number exceeds 1 the device is used by more than the current system and zdsfs prints a warning message when mounting the device. To prevent the user from mounting a device concurrently used by other operating system instances zdsfs is extended with a `-o check_host_count` option.

The tools fdasd and dasdfmt will print a warning message if the host access open count is greater than 1 as well. They will get the option `-C` or `--check_host_count` to secure the operation and cancel work if the open count exceeds 1.

1.7 Upgrading from and coexisting with a previous version

None.

1.8 Using online documentation

The man page of lsdasd is extended by this paragraph:

```
-H|--host-access-list
    Show information about all operating system instances that use
    this device.
```

The man page of zdsfs is extended by this paragraph:

```
zdsfs options:

[...]
-o check_host_count
    Stop processing if the device is used by another operating
    system instance.
```

The man page of dasdfmt is extended by this paragraph:

```
-C or --check_host_count
    Stop processing if the device is used by another operating
    system instance.
```

The man page of fdasd is extended by this paragraph:

```
-C or --check_host_count
    Stop processing if the device is used by another operating
    system instance.
```

1.9 Intra-kernel interfaces

None.

1.10 Dependencies

1.10.1 Dependencies on IBM software

None.

1.10.2 Dependencies on Open Source Software

None.

1.10.3 Dependencies on distributors

None.

1.10.4 Dependencies on hardware

The storage server should run a firmware version including support for the query host access information feature otherwise reading from the sysfs and debugfs attribute will fail and return -EOPNOTSUPP.

1.11 ID documents

The Device Drivers, Features, and Commands document should be updated to include the changes by this item. Especially the chapter “Working with DASDs” should contain the changes made to the lsdasd, dasdfmt, fdasd and zdsfs commands.

1.12 Security considerations

The debugfs attribute is created with file permissions that only allow read access for user root. This prevents the PGID information from being misused to avoid reserve states for example.

1.13 Service and health check considerations

A new health check might be useful to indicate that DASD devices are possibly used by other operating system instances.

1.14 Performance considerations

None.

1.15 Test considerations

The tools lsdasd, fdasd, dasdfmt and zdsfs have to be tested as well as the sysfs and debugfs interfaces. Please note that the storage server has to support the feature. The item was developed and tested on a DS8870 FW level 87.21.5.0.

End of document