

Linux on System z Programming Specifications for

KRN1405

PCI: Function-Measurement-Block Enhancements for z13

Version 5 from 2016-07-06 for distribution¹

Line Item Owner: Sebastian Ott
Department 3303 / 7103-03
Schoenaicher Str. 220
71032 Boeblingen, Germany
Phone +49-7031-16-1085
LotusNotes ID: Sebastian XB Ott/Germany/IBM@IBMDE
e-mail: sebott@de.ibm.com

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

1 PCI: Function-Measurement-Block Enhancements for z13

Owner: Sebastian Ott

Plan Candidate ID: KRN1405

1.1 Review

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

Reviewer	Role	sign-off date
Arnd Schneider	<i>Test</i>	2015-06-17
Dorothea Matthaeus	<i>ID</i>	2015-05-22
Horst Hartmann	<i>Performance</i>	2015-05-21
Martin Schwidefsky	<i>Architect</i>	2015-05-21

The status of this specification is: final

1.2 Abstract

This feature adds support for PCI function-type specific measurement data.

1.3 Description

This line item enhances the `statistics` interface to display PCI function-type specific measurement data.

1.4 Configuring, building and installing

None.

1.5 Setting up the environment

None.

1.6 Using this feature

PCI function measurement is enabled by default for any PCI function. Measurement data for a PCI function can be accessed by reading from the `statistics` file associated with the PCI function in the debug file system (usually mounted at `/sys/kernel/debug/`). To read measurement data for a (RoCE) function named `0000:00:00.0` use:

```
cat /sys/kernel/debug/pci/0000:00:00.0/statistics
FMB @ 0000000078cd8000
Update interval: 4000 ms
Samples: 14373
Last update TOD: cefa44fa50006378
      Load operations:      1002780
      Store operations:     1950622
Store block operations:    0
      Refresh operations:    0
      Received bytes:       0
      Received packets:     0
```

```
Transmitted bytes:    0
Transmitted packets: 0
Allocated pages:     9104
Mapped pages:        16633
Unmapped pages:      2337
```

Generating measurement data for a PCI function can be disabled at runtime by writing “0” to the `statistics` file (write “1” to `reenable`).

Machines prior to z13 will show the generic (PCI function type independent measurement data).

1.7 Upgrading from and coexisting with a previous version

None.

1.8 Using online documentation

None.

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

Available with z13.

1.11 ID documents

The `statistics` interface is currently not documented but an example could be added to the *PCI Express support* chapter of the device drivers book. But only the interface in general should be described not the individual counters.

1.12 Security considerations

None.

1.13 Service and health check considerations

None.

1.14 Performance considerations

None.

1.15 Test considerations

Nothing special - just verify that it works as expected and doesn't break anything.

End of document