

z/VM Service Critical Issues

Operating System Software Red Alerts

11/15/07
z/VM Service Manager
Contact: James McCormick
IBM Corporation
z/VM Development
1701 North Street
Endicott, NY 13760
mccormjp@us.ibm.com

November 15, 2007

VM64306 - PROBLEMS RUNNING LINUX ON FCP SCSI DASD THAT ARE ATTACHED WITH QIOASSIST

Customers using Linux distributions SLES10 and RHEL5 under z/VM versions 5.2 or 5.3 - in order to prevent data corruption or system stalls - need to install a PTF for z/VM APAR VM64306 if all of the following conditions are met:

- A Linux guest has guest memory above the 2GB boundary attached.
- The Linux guest uses an FCP I/O, or OSA networking in QDIO mode.
- FCP or OSA adapters are attached to the guest w/o using the NOQIOASSIST parameter in the device definition statement.

This is urgent information for customers running Linux guest systems under z/VM 5.3 and 5.2. Affected Linux distributions are SLES10 and RHEL5 which include V=V queue support.

When the guest memory exceeds 2GB, there is a severe issue, which can cause corrupted data to be transmitted to the OSA attached network or to/from FCP attached storage controllers. Potential impact is corrupted data in external storage or on the network. When using TCP/IP networking, excess package loss or retransmits may delay network traffic and cause application stalls. Exposure is greater when the guest system configuration contains multiple virtual CPUs.

Customers are urgently advised - as a workaround for this issue - to disable the V=V queue support (QIOASSIST) functions for all FCP and OSA adapters attached to affected guests by specifying the NOQIOASSIST parameter in the device definition statement or when attaching the adapter to the guest. It is not sufficient to set QIOASSIST=OFF for the guest or host system by issuing the SET QIOASSIST OFF command. Note, that enabling this workaround, will increase CPU consumption in the host system.

Customers running Linux distributions SLES8, SLES9 and RHEL4, as well as guest systems with a memory size smaller than 2GB are not affected by this issue. z/VM non-guest I/O as Virtual Switch and TCP/IP Stack as well as emulated FBA disks on FCP storage are also not affected by this issue.

z/VM APAR VM64287 is related to this problem, but does not completely resolve the issue. Open z/VM APAR VM64306 will fully address the problem, a fixtest is currently available from z/VM Level 2.

For further information on resolution of this issue, please watch QIOAssist Requirement Webpage (<http://www.vm.ibm.com/perf/aip.html>) which currently documents the z/VM, Linux and CP Millicode service levels required to run QIOAssist.

October 2007

VM64297 - PAGING SPIKE DEMAND SCAN SHUTS DOWN SYSTEM HANG

Large spike in paging causes apparent hang situation. DEMAND SCAN effectively shuts down while paging tries to catch up. This affects all z/VM R530 customers.

September 2007

VM64249 - PEVM63853 LPAR CHECK STOP DURING EDM H/W UPGRADE

During EDM (Enhanced Driver Maintenance) on zSeries hardware, the Linux-only LPAR running z/VM R520 stopped (a check-stop occurred on the partition). The HMC showed message RC=0C >> NOT OPERATING. The LPAR had to be inactivated and reactivated to be able to IPL z/VM. This problem only exists for z/VM R520 images that are running in LINUX-only logical partitions when a EDM occurs. It does not occur on z/VM R510 or R530 or other types of logical partitions.

There is a problem on z/VM R530 images in Linux only LPARs but it is much less severe. The hardware EDM is successful but the z/VM system doesn't correctly update its data about the CPU installed facilities. Also, notification of new CPU-related facilities are not presented to enabled guests. This problem is addressed in APAR VM64323 which is described below.

VM64323 - PEVM63853 NO LPAR MONITOR RECORDS AFTER EDM UPGRADE

During EDM (Enhanced Driver Maintenance) on IBM System z9, any LPAR running z/VM V520 or V530 displays message:

HCPSAS879I CPU INSTALLED FACILITIES HAVE CHANGED

After that, monitor records related to LPAR utilization are no longer generated because the PFXVLPAR bit was reset to zero.

After an EDM (Enhanced Driver Maintenance) on IBM System z9 hardware, an error during CPU Installed Facilities Change processing causes the PFXVLPAR bit in every CP prefix page to be reset to zero. Various CP functions use this bit to find out if CP is running in an LPAR. Because that bit is now incorrect, the following CP functions no longer function correctly.

- No CP Monitor records related to LPAR utilization will be generated.
- The response to DIAG x'00' will be incorrect. The bit that indicates that CP is running in an LPAR will be off. (Diagnose 0 DIAG0)
- In several places, CP will not inform LPAR that it wishes to give up its time slice. For example, CP uses this during spin lock processing. A symptom of this bug would be higher CP spin time trying to acquire locks that are contention-bound.
- The SET DYNAMIC_IO command will not work. This command enables or disables CP's ability to write a new IOCDS to the SE.
- It might take slightly longer to recover unresponsive CPUs after a SNAPDUMP.
- CP will just keep trying to find work rather than informing LPAR that it wishes to give up its time slice. This is known as Active Wait. This might cause elevated CPU consumption on VM LPARs.
- In certain cases, the QUERY LPARS command will show that CP is not running in an LPAR.