OSA-Express QDIO performance enhanced in consolidated Linux under z/VM environments

A combination of enhancements to the Linux *qdio* and *qeth* drivers, z/VM, VM/ESA, and the OSA-Express microcode may increase the efficiency and performance of Linux servers using Gigabit Ethernet or Fast Ethernet OSA-Express running QDIO when consolidated under supported levels of VM. Gigabit Ethernet throughput (transactions per second) can potentially increase up to 40% while potentially reducing the z/VM CP cost (CPU cycles per transaction) by up to one half. To enable these enhancements, each of the support levels must be installed. Minimal effect will be noticeable until all three components are operating simultaneously at the described support levels.

Linux

Functional enhancement summary:

- Use an interrupt-driven model for outbound traffic to get more control over the OSA-Express adapter port
- Exploitation of several VM assists to reduce the number of SIE intercepts in inbound and outbound traffic

Required support level:

Linux *qeth* and *qdio* drivers for both 31-bit and 64-bit are provided on DeveloperWorks. (See: http://www10.software.ibm.com/developerworks/opensource/linux390/index.shtml)

- The noted performance improvements were included in *qdio* and *qeth* on May 31, 2002, for the kernel 2.4.17 "May 2002 stream"
- Retrofits of the enhanced *qdio* and *qeth* are targeted to be available on DeveloperWorks for the following existing distributions by June 17, 2002:
  - SuSE SLES7 31-bit kernel 2.4.7 (refresh available May 24, 2002)
  - SuSE SLES7 64-bit kernel 2.4.17 (available April 30, 2002)
  - Red Hat 7.2 31-bit kernel 2.4.9[-37] (available May 17, 2002)
- The distributors may pick up these enhanced *qdio and qeth* at some later point in time.
  - SuSE's next refresh is expected in 3Q 2002

VM/ESA and z/VM

Functional enhancement summary:

- Fewer SIGA calls, and improved QDIO handling
- Enhanced PCI interrupt handling for both inbound and outbound QDIO

Required support level:

VM/ESA V2.4.0

- APAR VM63036 Performance enhancement retrofit from VM62938
  PTF R240 UM30291
- APAR VM63034 for QDIO support
  PTF R240 UM30287

z/VM 3.1.0

- APAR VM63036 Performance enhancement retrofit from VM62938
  Applies to releases V3.1.0 and V4.1.0 only
  PTF R310 UM30292
- APAR VM63034 for QDIO support
  PTF R310 UM30288

z/VM 4.1.0

- APAR VM63036 Performance enhancement retrofit from VM62938
  Applies to releases V3.1.0 and V4.1.0 only
  PTF R410 UM30293
- APAR VM63034 for QDIO support
  PTF R410 UM30289

z/VM 4.2.0

- APAR VM62938 QDIO performance (APAR includes HiperSockets function enablement)
  Applies to 4.2.0 only
  PTF R420 UM30225
- APAR VM63034 for QDIO support
  PTF R420 UM30290

z/VM 4.3.0

- All enhancements included

OSA-Express QDIO mode on Gigabit Ethernet, Fast Ethernet, Token Ring and ATM Ethernet LAN Emulation,

Functional enhancement summary:

- Enhanced PCI interrupt handling for outbound QDIO

Required support level:

- zSeries 900 GA3 Driver 3G with OSA microcode level: 3.0A at MCLs: J11204.007 and J11204.008 (available May 03, 2002)
- zSeries 900 GA2 Driver 3C with OSA microcode level: 2.26 at MCLs: J10630.013 and J10630.014 (available May 20, 2002)
- zSeries 800 GA1 Driver 3G with OSA mircocode level: 3.0A at MCLs: J11204.007 and J11204.008 (available May 03, 2002)
- S/390 Parallel Enterprise Servers G5 and G6 with driver 26 with OSA microcode level: 4.25 at MCLs: F99904.032 and F99904.033 (available May 16, 2002)
- Note:
  There is a limit to the number of supported TCP/IP stacks per OSA-Express port when running in QDIO mode. This limit includes the sum of all operating system images or guests accessing/sharing the port at any one time (Linux, VM/ESA, z/VM, OS/390, z/OS, z/OS.e, TPF, VSE/ESA).

- Eighty (80) TCP/IP stacks per zSeries OSA-Express port (at driver 3C). OSA-Express on zSeries is a two port feature, with each port representing an individual adapter requiring a separate CHPID
- Fifteen (15) TCP/IP stacks per S/390 G5 or G6 OSA-Express port. OSA-Express on G5 and G6 is a one port feature, with the port representing a single adapter requiring a separate CHPID

For additional detail refer to
http://www-1.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/Flash10144

Recommendation:

It is highly recommended that customers planning to run multiple Linux virtual servers under VM/ESA or z/VM at relatively high traffic volumes running QDIO on OSA-Express Gigabit Ethernet or Fast Ethernet, should plan to install the suggested software levels and associated maintenance and MCL levels appropriate for their particular environment.

In addition, the Redbook, Linux on IBM eServer zSeries and S/390: ISP/ASP Solutions, SG24-6299-00, http://www.redbooks.ibm.com/abstracts/sg246299.html is now available and contains planning assistance for consolidating Linux under z/VM.

The OSA-Express Customer's Guide and Reference, SA22-7476-02 is available on IBM Resource Link(tm) at http://www.ibm.com/servers/resourcelink/

You are also encouraged to take advantage of the Solutions Assurance process available via w3.ibm.com/support/solassure.html including the zSeries Systems Assurance Checklists for Linux

For questions about this note, please contact
- Ann Jackson, zSeries WW Technical Support - Linux, at tie line 677-5975, outside 630-573-5975, or at Ann Jackson/Chicago/IBM@IBMUS
- Jim Goethals, zSeries Networking Offering Manager, at tie line 526-1367, outside 919-486-1367, or James Goethals/Raleigh/IBM@IBMUS