zSeries[7] Hardware Exploitation For V1R6

Topics

- z/OS V1R6
  - X-windows 64-bit virtual support
  - IPSec support of new zSeries hardware crypto functions
  - Direct OSA SNMP extended to LCS ports
- MTU size clarifications
- Overview of z/VM guest LAN support

e-business

zSeries hardware support in z/OS V1R6

IBM

Various zSeries hardware support items

- Enhanced X-Windows and Motif libraries to exploit 64-bit virtual storage
  - Driven by Java Virtual Machine support for 64 bit
  - Upgrade X Windows from X Windows System X11R6.1 to X11R6.6
    - Uses LE C/C++ sockets
  - Upgrade OSF/Motif 1.2.4 to Motif 2.1.30
  - z/OS provides the X Windows "client-side" function
  - Both X Windows and Motif are developed and distributed under Open Group (www.opengroup.org)
    - IBM is a member.
  - IPv6 support was not available at this time
- IPSec support of the zSeries synchronous crypto functions
  - SSL/TLS support picked up via System SSL library
- Extended OSA Direct SNMP management to include OSA-Express ports operating in LCS mode

X-Windows and Motif upgrade details

- In addition to many new functions available with the new releases, new libraries will be provided in the following manner:
  - Static Archive Libraries for 31-bit mode and 64-bit mode compiled with XPLINK
  - Dynamic Libraries (DLLs) for 31-bit mode and 64-bit mode compiled with XPLINK
- Since binary compatibility was not fully provided with the new offerings, z/OS will continue to provide the X Windows X11R6.1 and Motif 1.2.4 libraries and include files.
  - These libraries are compiled with 31-bit mode, NON XPLINK, and IBM Hexadecimal Floating Point.
- It is possible to link a non XPLINKed application with the new XPLINK libraries, but is not recommended
  - IBM recommends compiling and linking with the newest versions of the libraries
- Many new functions have been provided with the new libraries. To understand what has been provided,
  - go to: www.opengroup.org and follow the X windows and Motif links for product information.
- X11R6.6 Functions not supported by z/OS:
  - Low Bandwidth Extensions - provide compression / decompression support for improved performance over slow lines (new)
  - X3D-PEX - provides capabilities that allow each window on the display to act as a complete, independent virtual three-dimensional graphics workstation (part of old support)
  - XIE - X Imaging Extensions (deprecated)

IPSec support for zSeries 990 Crypto hardware

- New IBM eServer zSeries 990 or 890 servers became available 2003.
  - zSeries 990 GA 1, available 5/2003, does not have support for the cryptographic co-processor (CCF) that IPSEC accesses currently.
- Exploitation of IBM CP Assist for Cryptographic Functions provides IPSEC support for the Integrated Cryptographic Service Facility (ICSF) cryptographic instructions available in V1R5.
- The zSeries 990 and 890 server has new cryptographic instructions that improve symmetric encryption/decryption performance, as well as SHA1 performance.
- New ICSF encryption/decryption instructions available in V1R5 provides IPSEC access to the zSeries 990 and 890 cryptographic instructions.
- The CP Assist instructions are stated to provide excellent performance for DES, TDES, and SHA-1 services; and eliminate the requirement to execute crypto instructions on a CCF; thus achieving a more even workload across processors.

IPSec support for zSeries Crypto hardware

Encryption/ decryption method	IBM eServer zSeries 990 and 890 GA 1	IBM eServer zSeries 990 and 890 GA 2
crypto assist	IPSEC invokes crypto assist if present **(normal V1R5-and-later behavior**)	IPSEC invokes crypto assist if present **(normal V1R5-and-later behavior**)
hardware	not available	if crypto assist not present IPSEC invokes hardware **(normal pre-V1R5 behavior**)
software	if crypto assist not present or fails IPSEC invokes software **(normal pre-V1R5 behavior**)	if hardware not present or fails IPSEC invokes software

OSA Direct SNMP support for LCS

- For OSA SNMP management data, recommend using the OSA-Express Direct Subagent instead of the IBM MVS TCPIP Subagent.
- The OSA subagent communicates directly with the OSA adapters instead of requiring OSA/SF.
- Supports the OSA-Express Direct Enterprise-specific MIB, which provides much of the same data that's in the IBM TCPIP Enterprise-specific MIB in support of OSA
- The following OSA-Express MIB tables have been deprecated in the IBM MVS TCPIP Enterprise-specific MIB:
  - osaexpChannelTable
  - osaexpPerfTable
  - osaexpEthPortTable
  - osaexpEthSnaTable

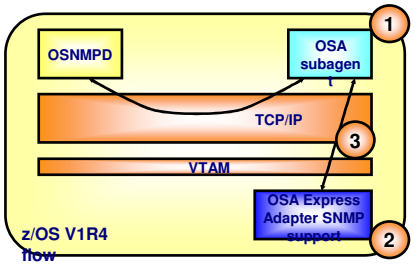

- The OSA subagent supports the same data from the EtherLike-MIB that the TCPIP subagent supports
- As of V1R6, the OSA-Express Direct Enterprise-specific MIB supports both QDIO and non-QDIO mode adapters
- The OSA subagent, IOBSNMP, is owned by OSA but shipped with z/OS Communications Server. Sample start JCL is IOBSNMP in hlq.SEZAINST.
- Refer to zSeries Open System Adapter-Express (OSA Express) Customer's Guide and Reference for more information on managing OSA.

e-business

MTU size clarifications

IBM

How to define MTU sizes for IPv4

e-business

Large MTU sizes have a very positive impact on file transfer type of workloads.

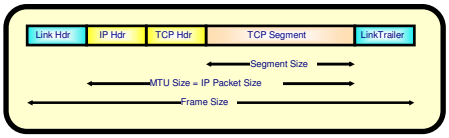

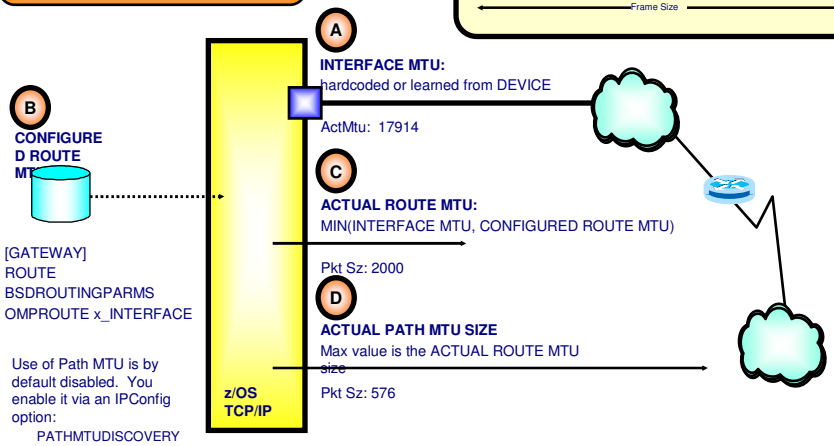

[GATEWAY]
ROUTE
BSDROUTINGPARMS
OMPROUTE x_INTERFACE

Use of Path MTU is by default disabled. You enable it via an IPConfig option:
PATHMTUDISCOVERY

MTU clarification

e-business

NOTES

- TCP/IP uses the MTU to determine the largest size frame to send. The MTU in effect for a given outbound send depends on several factors.

- interface_MTU
  - This is either a hardcoded size based on the physical device or it is a value learned from the device during activation. See <reference to the DEVICE/LINK summary chart and INTERFACE summary chart> for information about the interface MTU values for the various network interfaces types supported by TCP/IP. For an IPAQENET6 interface, you can configure a lower interface_MTU via the MTU keyword on the INTERFACE statement.
  - For an active link or interface, TCP/IP reports the interface_MTU on the ActMtu field on the Netstat DEVLINKS/-d command.

- configured_route_MTU
  - This is the MTU size configured for a route.
  - For a static route, you specify the configured_route_MTU on either a ROUTE statement in a BEGINROUTES block or on a GATEWAY statement in the TCP/IP profile.
  - For each IPv4 dynamic route added by OMPROUTE over an interface, the configured_route_MTU comes from the MTU value that you specify on the RIP_INTERFACE, OSPF_INTERFACE or INTERFACE statement for that interface in the OMPROUTE configuration file. If you do not specify an MTU for an interface, then OMPROUTE will use 576. For IPv6, OMPROUTE learns the interface_MTU value from TCP/IP and you cannot specify a configured_route_MTU in the OMPROUTE configuration file.
  - For each dynamic route added by OROUTED over an interface, the configured_route_MTU comes from the MTU value that you specify on the BSDROUTINGPARMS statement for that interface in the TCP/IP profile.

MTU clarification

e-business

NOTES

- actual_route_MTU
  - This is the minimum of the interface_MTU and the configured_route_MTU.
- path_MTU
  - This is the value determined by the path MTU discovery function. You can enable path MTU discovery for IPv4 via IPCONFIG PATHMTUDISCOVERY. Path MTU discovery is automatically enabled for IPv6. Path MTU discovery starts out by setting the path_MTU to the actual_route_MTU of the route. If packets would require fragmentation to get to the final destination, then path MTU discovery will determine the path_MTU by repeatedly decreasing this value until it can send packets to the final destination without fragmentation.
- Recommendations:
  - When using OSA-Express Gigabit Ethernet (which supports an interface MTU of 8892), be aware that not all routers or switches support a value of 8892. For example, the Cisco CAT65xx supports a maximum value of 4470. Either ensure that all routers or switches in your configuration support 8892 or specify a lower configured_route_MTU.
  - When using OMPROUTE, configure all nodes on a LAN to use the same MTU value. Otherwise, you may encounter problems such as OSPF adjacency errors.
  - When using OMPROUTE, specify a configured_route_MTU for each IPv4 interface.

e-business

Overview of z/VM guest LAN and virtual switch support

IBM

Using z/VM guest LANs and HiperSockets to implement multi-zone networking environments

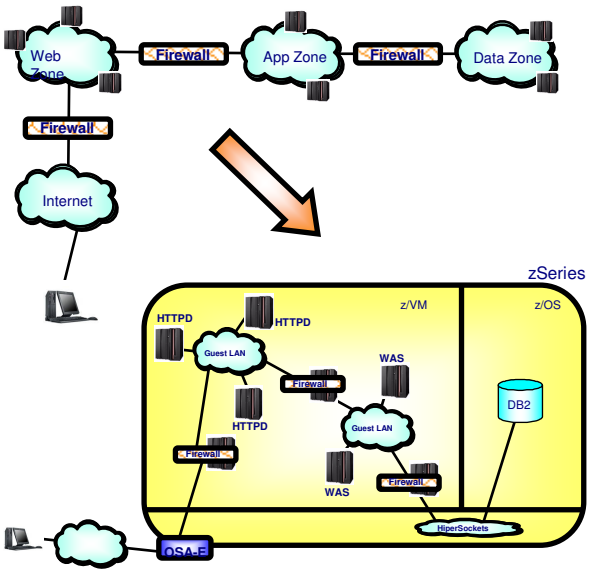

- A z/VM guest LAN is a simulated LAN
  - Ethernet IPv4 and IPv6 (QDIO)
  - HiperSockets IPv4 (iQDIO)
  - Unicast, Multicast, and Broadcast
  - No built-in connection to outside network
  - A separate IP subnet/prefix
  - Owned by z/VM "SYSTEM" or individual z/VM user ID
  - Created in SYSTEM CONFIG, directory, or via CP DEFINE LAN command
  - Guest LAN connectivity controlled at system level via access lists
  - Guests connect via virtual NICs (Network Interface Cards)
- You can create as many guest LANs in z/VM as you need
- Multiple guest LANs need to be interconnected via "routers" or "firewalls"
- Both z/OS and Linux can act as routers and/or firewalls
  - z/OS firewall technologies
  - Linux with native firewall functions or Stonegate firewall software

Locating the firewall functions outside zSeries

e-business

- If company security policies require use of a specific non-zSeries resident firewall technology, the individual security zones have to be extended outside zSeries and interconnected using firewalls that reside on non-zSeries platforms.
  - Various design options of which one is to continue using guest LANs, but interconnect the internal guest LANs via a router guest to the external LANs.
  - The guest LANs become part of the Web Zone and the App Zone in this example, but are separate IP subnets that are connected to the outside subnets via router functions in Linux, z/OS, or z/VM

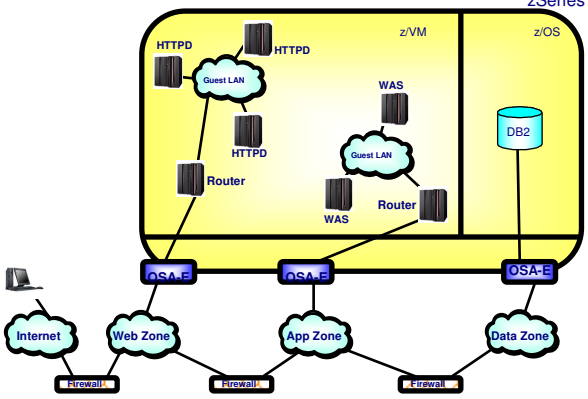

Locating the firewall functions outside zSeries using z/VM Virtual Switch

e-business

- By using the built-in z/VM layer-2 switch function known as the Virtual Switch, the guest LANs can be made part of the external LANs - part of the same subnet.
  - The routing function between the guest LANs and the external LANs can be removed and replaced by the virtual switch function
  - No dependency on Linux, z/OS, or z/VM routing

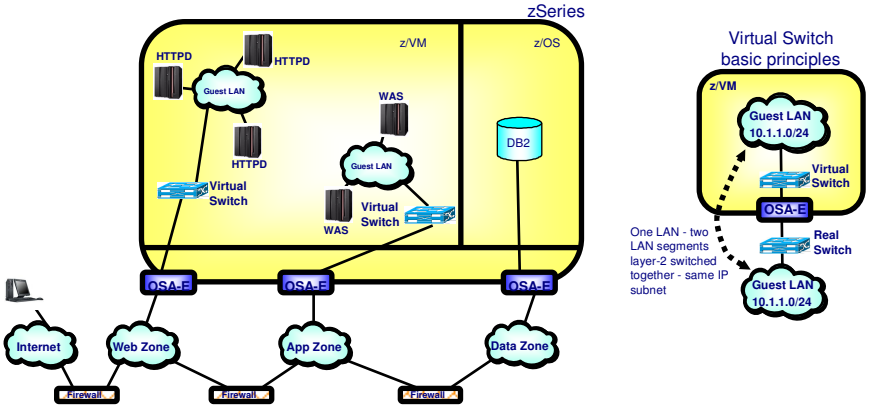

Combining outside firewalls, virtual switch, and VLAN technology

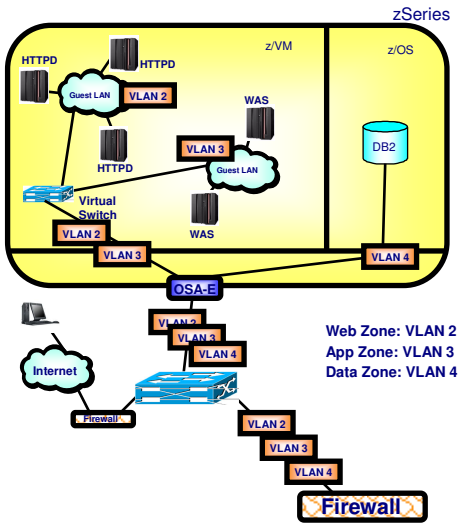

- z/VM's virtual switch is VLAN-aware and supports switching different VLAN IDs to different guest LANs.
- A single OSA adapter connected to a VLAN-aware external switch using a trunk mode connection can now serve all three security zones
- z/VM's virtual switch can be attached to up to three physical OSA adapters for scaling purposes if capacity requirements exceed those of a single network interface
- z/VM's virtual switch supports Ethernet IPv4 (QDIO)
  - Currently no support for other networking protocols, such as SNA, NETBIOS, or IPv6
- The virtual switch does depend on z/VM TCP/IP for initial QDIO control flows (known as the controller), but z/VM TCP/IP is not involved in data transfer via the virtual switch
- VLAN technology in combination with z/VM's virtual switch offers significantly simplified network connectivity options for a complex zSeries server environment

Overview of guest LAN technologies on zSeries

e-business

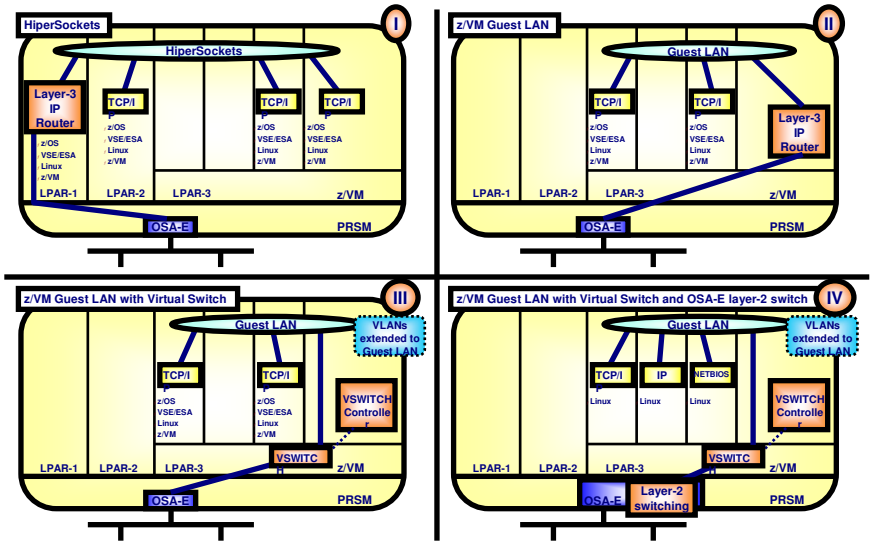

Comparison matrix

e-business

	HiperSockets	z/VM Guest LAN	z/VM Guest Lan with VSWITCH (z/VM 4.4)	z/VM Guest LAN with VSWITCH and OSA layer-2 switching (z/VM 5.1)
zSeries OS support	z/VM z/OS VSE/ESA Linux	z/VM z/OS VSE/ESA Linux	z/VM z/OS VSE/ESA Linux	Linux
Network protocol support	IPv4	IP4 and IPv6	IPv4	All protocols: IPv4, IPv6, IPX, NETBIOS, SNA, etc.
Device driver support	iQDIO	iQDIO or QDIO	QDIO	QDIO (For an OS to support other protocols than IP, its QDIO device driver must be able to handle protocols other than IP)
Need for mainframe layer-3 router to outside network	Required (z/OS and Linux support accelarated routing to/from HiperSockets)	Required	None (z/VM TCP/IP stack required for control, not layer-3 routing)	None (z/VM TCP/IP stack required for control, not layer-3 routing)
LPAR or z/VM guest support	Both	z/VM guests in a z/VM LPAR	z/VM guests in a z/VM LPAR	z/VM guests in a z/VM LPAR
Extend IEEE802.1q VLAN	No	No	Yes (for IPv4 frames)	Yes (for all protocol frames)

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Other company, product and service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.