

VS@

e-business



VSE/ESA 2.6 and 2.7

Performance Considerations

Ingo Franzki

e-mail: ifranzki@de.ibm.com

VSE/ESA Development

VS@

e-business



Contents

- VSE/ESA 2.6
 - f* Performance Items (Reminder)
- VSE/ESA 2.7
 - f* Performance Items (Overview)
 - f* HiperSockets
 - f* Hardware Crypto support
- z800/z900 Remarks
- Dependencies for VSE/ESA Growth



e-business



VSE/ESA 2.6 Performance Items

- VSE/ESA 2.6 Base enhancements
 - f* Delete Label Function
 - saves >90% SVCs for sequential processing
 - f* LTA Offload for some AR commands
 - Less I/O by less FETCHes for LTA load
 - f* SVA-24 Phases moved above the line
 - \$IJBPTY (6KB)
 - f* Increased max number of SDL entries
 - Maximum value now 32765
 - f* SDL update from non-BG partitions
 - f* POWER Data file extension without reformat



e-business



VSE/ESA 2.6 Performance Items - continued

- VSE/ESA 2.6 Hardware Support
 - f* FICON Support (VSE/ESA 2.3 or higher)
 - f* New 2074 System Management Console
 - ESCON channel attached
 - Eliminates requirement for a non-SNA 3174 controller
 - f* OSA Express Adapter (e.g. Gigabit Ethernet)
 - Available for G5 and above
 - f* VSAM Support for large 3390-9 Disks (Shark)
 - f* Fastcopy Exploitation of ESS FlashCopy and RVA SnapShot



e-business



Hardware Support

- Queued Direct I/O
 - f* Designed for very efficient exchange of data
 - f* Uses the QDIO Hardware Facility, without traditional S/390 I/O instructions
 - f* Without interrupts (in general)
 - f* Use of internal queues
 - f* With pre-defined buffers in memory for asynchronous use
- Exploitation by TCP/IP for VSE/ESA
 - f* see TCP/IP Performance Considerations



e-business



VSAM SHROPT(4) Avoidance

- Connectors in VSE/ESA 2.5 require SHROPT(4) when updating VSAM files owned by CICS
 - New VSAM-via-CICS Service avoids SHROPT(4) by routing the VSAM requests to CICS
 - Communication between batch and CICS is XPCC
 - Naming convention for "VSAM-via-CICS files"
 - f* Each CICS is treated as "virtual" catalog
 - f* Files defined in CICS (via CEDA DEFINE FILE) are visible within this catalog
- #VSAM.#CICS.<applid>**
- indicates "virtual" CICS catalog
- APPLID of CICS region owning the files within this catalog

VS@

e-business



VSAM Redirector

- New connector with VSE/ESA 2.6
 - f* VSE is client
 - f* PC / workstation is server
- Exploits VSAM exit IKQVEX01
- Allows to redirect one or more VSAM files to a PC or workstation
- All VSAM requests of a particular file are redirected
 - f* Open / close
 - f* Get / put / point / delete / insert
- Transparent for applications
 - f* Usable from batch and CICS

VS@

e-business



VSAM Redirector - Performance Implications

- Is the file redirected ?
 - f* No: only at OPEN time (very small overhead)
 - f* Yes: at each request
- Network overhead ?
 - f* Yes, if file is redirected
 - f* Depends on
 - Number of VSAM requests
 - Size of records
- Data ownership
 - f* OWNER=REDIR
 - no VSAM I/O

VS@

e-business



VSE/ESA 2.7 Performance Items

- VSE/ESA 2.7 hardware support
 - f* z800/z900, Multiprice 3000, G5/G6
 - f* HiperSockets
 - f* Hardware Crypto Support
 - f* 32760 cylinder 3390 support
 - f* 3590 buffered tape mark
- VSE/ESA 2.7 enhancements
 - f* new TCP/IP for VSE/ESA release 1.5
 - f* \$IJBLBR above the line
 - f* II User Status Record above the line
 - f* VTAPE: removed DVCDN/DVCUP
 - f* POWER: reallocate queue file during warm start

VS@

e-business



VSE/ESA 2.7 Hardware support

- VSE/ESA 2.7 runs on the following machines
 - f* z800 (2066)
 - f* z900 (2064)
 - f* 9672 Parallel Enterprise Server (G5/G6)
 - f* Multiprice 3000 (7060)
 - f* equivalent emulators (Flex-ES)
- VSE/ESA 2.7 is based on the hardware instruction set described in the manual 'ESA/390 Principles of Operation' (SA22-7201).
- With VSE/ESA 2.7 it is assumed that all the ESA/390 instructions and facilities described in that manual can be used.

VS@

e-business



z800/z900 Remarks

- Prior to z800/z900 there is one cache for data and instructions
- z800/z900 has split data and instruction cache
- Performance implications:
 - f* If program variables and code that updates these program variables are in the same cache line (256 byte)
 - Update of program variable invalidates instruction cache
 - Performance decrease if update is done in a loop
 - f* See APAR PQ66981 for FORTRAN compiler

VS@

e-business



32760 cylinder 3390 support

- With announcement 101-341 at 11/13/2001 IBM announced the new 32760 cylinder 3390 volumes of the IBM TotalStorage Enterprise Storage Server (ESS).
 - f* This enhancement of the ESS F models was made available 11/30/2001.
- VSE/ESA 2.7 now supports these volumes.
 - f* helps relieve address constraints
 - f* improves the disk resource utilization
 - f* can be used to consolidate multiple disk volumes into a single address.

VS@

e-business



3590 Buffered Tape Mark support

- The 3590 control unit provides support for writing tape marks (TM) in buffered mode
- Writing TM's in "buffered" mode should enhance the performance
 - f* of all programs, which write many TM's as part of their file creation process (e.g. POFFLOAD)
- All the TM's written during OPEN/CLOSE (label processing) will remain to be written "UNbuffered"
 - f* all the programs which write TM's mainly or only during OPEN/CLOSE, will NOT benefit from this enhancement.

VS@

e-business



\$IJBLBR phase moved above the line

- The \$IJBLBR.PHASE has been split into two phases
 - f* \$IJBLBR.PHASE
 - f* \$IJBLB31.PHASE
- \$IJBLBR.PHASE will continue to reside in SVA-24
- \$IJBLB31.PHASE will reside in SVA-ANY (high SVA).
 - f* This will free about 180k in SVA-24.



e-business



II User status record above the line

- During Logon each II user gets besides others two storage areas allocated
 - f* User_Status_Record USR (904 bytes)
 - f* Panel_Hierarchy_List PHL (1352 bytes)
 - f* originally located in the CICS DSA (below)
- With VSE/ESA 2.7 the USR and PHL has been moved to ESDSA (shared above)
 - f* frees 2.3 KB in DSA below per user.
- ICCF TCTUALOC=ANY now supported
 - f* ICCF transaction programs has been changed to support a TCTUA (28 bytes) above the line



e-business



HiperSockets hardware elements (‘Network in a box’)

- Synchronous data movement between LPARs and virtual servers within a zSeries server
 - f* Provides up to 4 "internal LANs" HiperSockets accessible by all LPARs and virtual servers
 - f* Up to 1024 devices across all 4 HiperSockets
 - f* Up to 4000 IP addresses
 - f* Similar to cross-address-space memory move using memory bus
- Extends OSA-Express QDIO support
 - f* LAN media and IP layer functionality (internal QDIO = iQDIO)
 - f* Enhanced Signal Adapter (SIGA) instruction
 - No use of System Assist Processor (SAP)

VS@

e-business



HiperSockets hardware elements (**'Network in a box'**) - continued

- HiperSockets hardware I/O configuration with new CHPID type = IQD
 - f* Controlled like regular CHPID
 - f* Each CHPID has configurable Maximum Frame Size
- Works with both standard and IFL CPs
- No physical media constraint, no physical cabling, no priority queuing
- Secure connections

VS@

e-business



Measurement Environment

- z800 (2066-004)
 - f* 4 processors
- VSE/ESA 2.7 GA Driver in a LPAR (native)
 - f* 1 CPU active (~2066-001)
 - f* TCPIP00 (F7): OSA Express Fast Ethernet
 - f* TCPIP01 (F8): HiperSockets
- Linux for zSeries in a LPAR (native)
 - f* 3 CPUs active (shared)
 - f* eth0: OSA Express Fast Ethernet
 - f* hsi10: HiperSockets

VS@

e-business



Latency (Round trip time) - results


- Measurements has been done with PING command
 - Issued at Linux side
 - 10 Pings
 - PING sends a datagram to VSE
 - VSE sends a answer back to Linux
 - Time until answer arrives is measured
 - Round trip time

VS@

e-business



Latency (Round trip time) - results



HiperSockets is about 1.8 times faster in terms of latency

VS@

e-business



Throughput (MB/sec)

■ Measurements has been done with FTP

f Initiated at the Linux side

f Transferring 1GB (1000MB)

–without translation (binary)

–1 to 5 parallel streams

f PUT: send data to VSE

–VSE inbound

–sending a 1GB file to \$NULL file (in memory file)

–No file I/O is done by VSE/Linux

f GET: receive data from VSE

–VSE outbound

–receiving \$NULL file (in memory file) into /dev/null

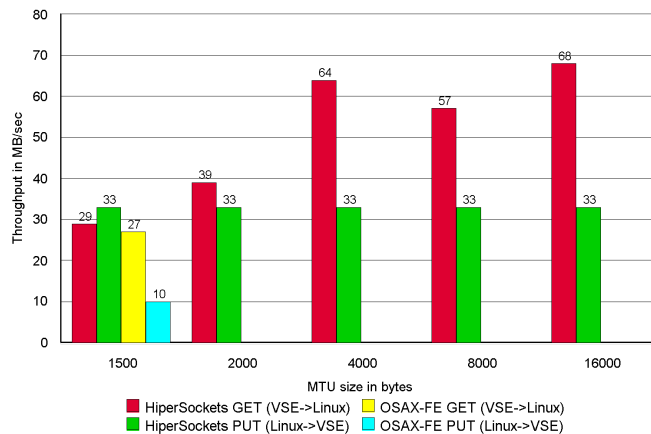
–No file I/O is done by VSE/Linux

VS@

e-business



Throughput (MB/sec) - results



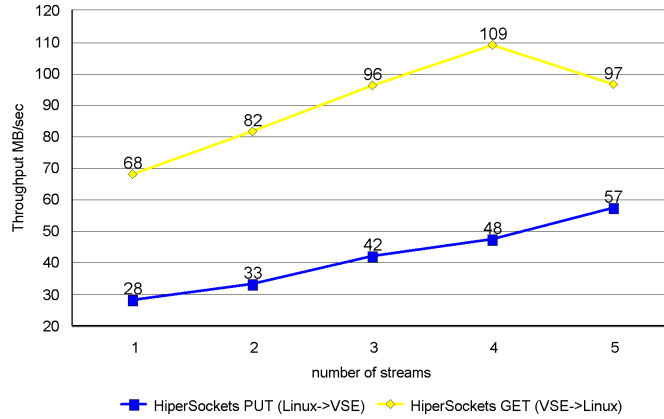
HiperSockets throughput is between 30-80 MB/sec

VS@

e-business



Throughput (MB/sec) - results (2)



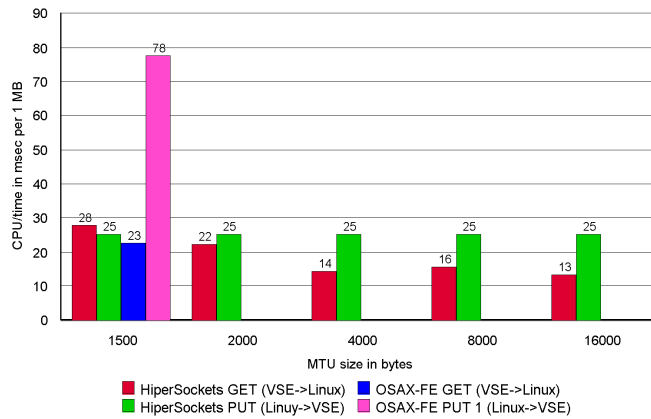
Maximum HiperSockets throughput of 109 MB/sec at 4 concurrent connections

VS@

e-business



CPU time per MB - results



About 15-30 msec CPU time per MB for HiperSockets (on a z800 2066-001)

VS@

e-business



Transaction per second


- Measurements has been done with an ECHO server
 - Client on Linux sends 100 bytes to server
 - Server on VSE echoes 100 bytes
 - Per TCP connection 10000 transactions are driven
 - Variations: Number of TCP connections
 - 1,2,3,4,5
 - 10,50,100
 - Measurements
 - Transactions per second
 - CPU time per transaction

VS@

e-business



Transactions per second - results



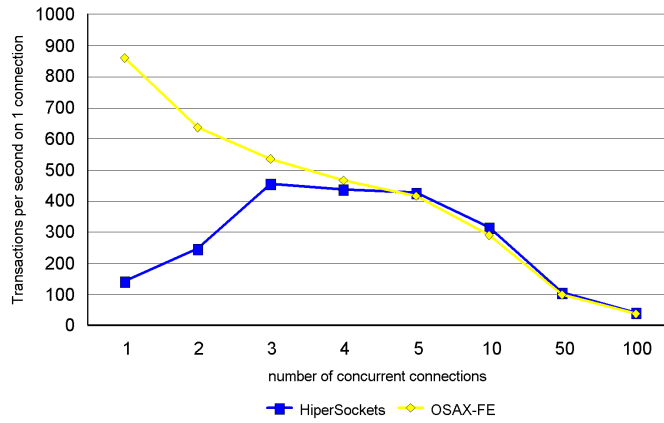
Maximum of 5200 transactions per second at 50 concurrent connections

VS@

e-business



Transactions per second on 1 connection - results



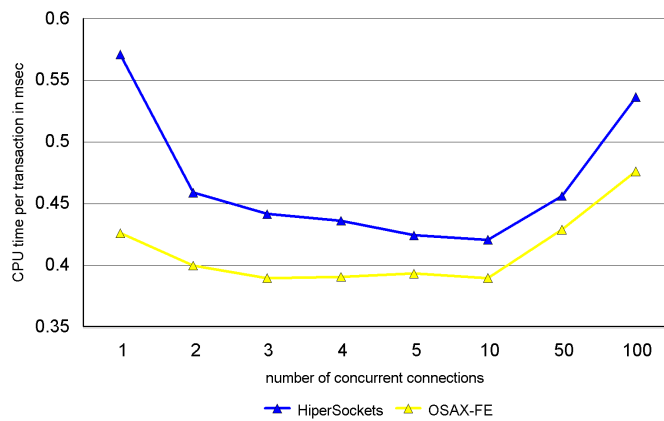
HiperSockets: Maximum of about 450 transactions per second on 1 connection (= about 2 msec response time)

VS@

e-business



CPU time per transaction



HiperSockets: About 0.45 msec CPU time per transaction for 2-50 connections

VS@

e-business



Measurement Results - conclusion

▪ HiperSockets

f Throughput

- Between 30-80 MB/sec
- Maximum throughput of 109 MB at 4 connections
- About 15-30 msec CPU time per MB

f Transactions per second

- Maximum of 5200 Transactions per second at 50 connections
- About 0.4-0.45 msec CPU time per transaction

VS@

e-business



Hardware Crypto Overview

▪ Requires VSE/ESA 2.7 and TCP/IP for VSE/ESA 1.5

▪ Supported crypto cards

f PCI Cryptographic Accelerator (PCICA)

- Feature code 0862
- Available for zSeries (z800, z900)

▪ The crypto card is plugged into the Adjunct Processor

▪ Currently only RSA (asymmetric) is supported

f Of benefit for Session initiation (SSL-Handshake)

▪ Also supported with

f z/VM 4.2 + APAR VM62905

f z/VM 4.3

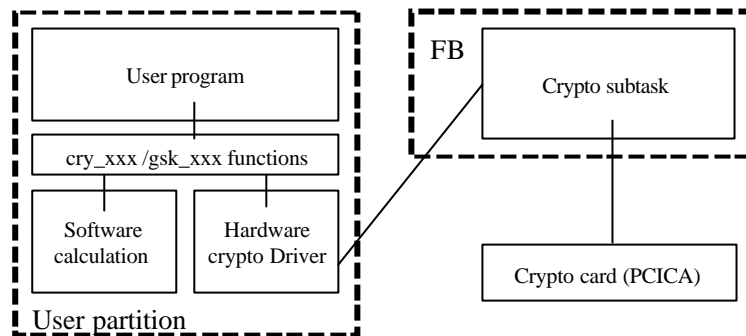


Hardware Crypto Overview - continued

- New crypto subtask in Security Server (SECSERV) running in FB

f Or as separate job if no SECSERV is running

f Crypto card is polled by crypto task



Measurement Environment

- VSE/ESA 2.7 running on a z900 (2064-109)

f on 1 processor (~2064-101)

f with a PCI Cryptographic Accelerator

- Testcase programs on VSE

f Crypto operations measurements

– calling cry_XXX functions (RSA, DES, SHA, MD5)

– each crypto operation is performed 10000 times

f Secured data transfer (SSL)

– performs SSL handshake

– performs encrypted data transfer

– counterpart program running on Windows (SSL-client)

- All RSA operations are measured

f with Hardware Crypto support

f with Software Crypto

– (support already available with TCP/IP 1.4/1.5 as shipped in VSE/ESA 2.6)



e-business



Measurement Environment - continued

Variations

f RSA encrypt/decrypt

-512 / 1024 bit key

f DES, DES CBC, 3DES CBC encrypt/decrypt

-software crypto only

-message length (128, 256, 512 bytes)

f SHA Hash, MD5 Hash, SHA HMAC, MD5 HMAC

-software crypto only

-message length (128, 256, 512, 1K, 2K bytes)

f SSL handshake/data transfer

-01 RSA512_NULL_MD5

-02 RSA512_NULL_SHA

-08 RSA512_DES40CBC_SHA

-09 RSA1024_DES_CBC_SHA


-0A RSA1024_3DES_EDE_CBC_SHA



e-business



Measurements Results - RSA



HW Crypto:

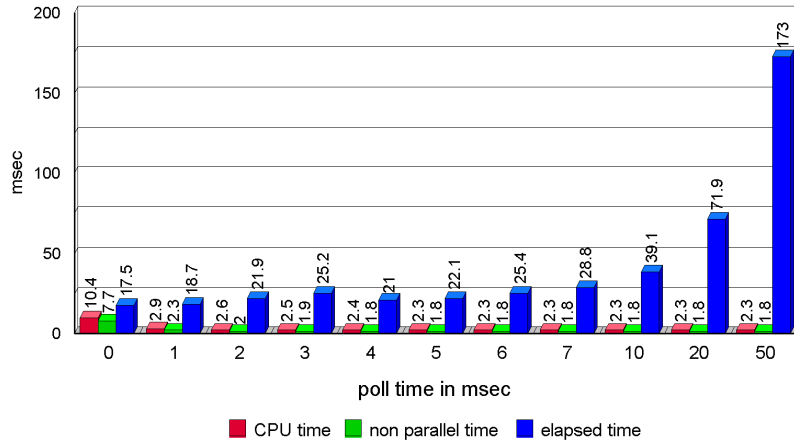
- CPU time and elapsed time is independent of operation / key length
- RSA operation takes about 2 msec CPU time and 28 msec elapsed time
- CPU time is always less than software crypto

VS@

e-business



Measurements Results - RSA polltime



Per default a polltime of 7 msec is used.

Can be changed with: MSG FB,DATA=WAITTIME=nn

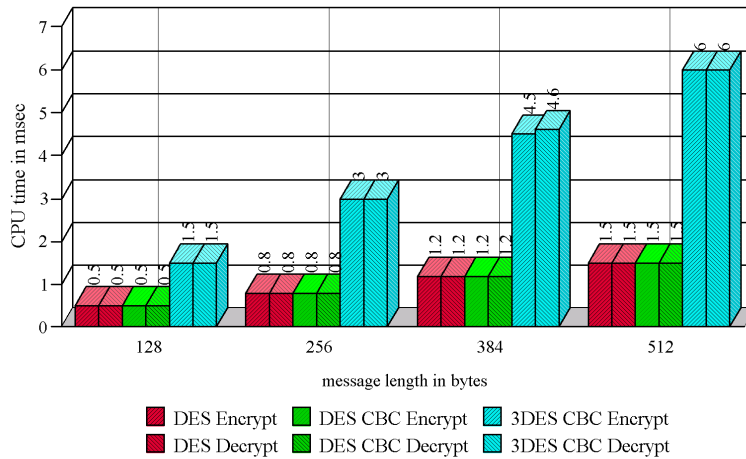
Smaller values increases CPU time, higher values increases elapsed time

VS@

e-business



Measurements Results - DES, DES CBC 3DES CBC (symmetric)



Software Crypto only!

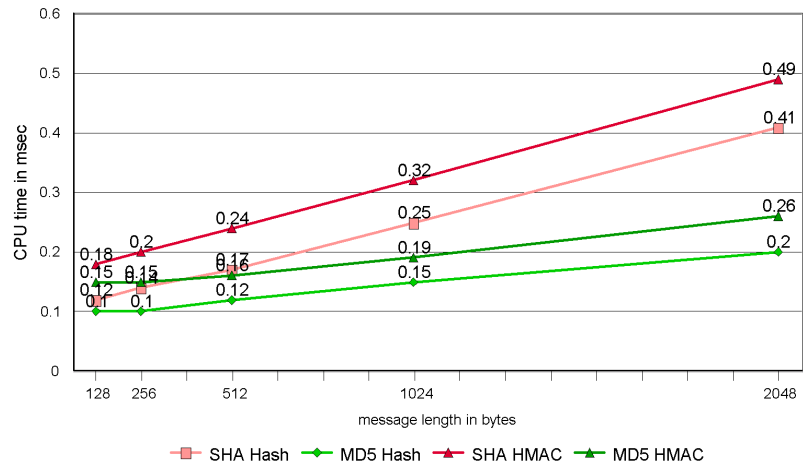
DES and DES CBC takes similar CPU times, 3DES CBC about 3.8 times



e-business



Measurements Results - SHA, MD5




SHA takes about 1.8 times more CPU time compared to MD5
Software Crypto only!



e-business



Measurements Results - SSL Handshake



HW Crypto:


- CPU time and elapsed time is independent of cipher suite used
- SSL handshake takes about 43-48 msec CPU time (connection establishment)



e-business



Measurements Results - SSL data transfer



CPU time depends on used hashing (SHA/MD5) and encryption algorithm (DES/3DES) Software Crypto only!




e-business




SSL data transfer overhead

Non SSL



SSL



this has been measured

VS@

e-business



Measurements Results - conclusion

- HW Crypto
 - f* Supports RSA operations only (e.g. used by SSL handshake)
 - f* CPU time/elapsed time is independent of operation and key length
 - f* Software RSA encryption is faster in terms of elapsed time (on large processors)
 - but hardware crypto saves CPU time
- SW Crypto
 - f* CPU time /elapsed time is very dependent on CPU speed and utilization

VS@

e-business



SSL Performance Recommendations

- Use SSL only if there is a need for
 - f* If at least one of the following is required
 - Keeping secrets
 - Proving identity
 - Verifying information
- Cipher Suites 01 and 02 has less CPU-time consumption, but NO data encryption
 - f* RSA512_NULL_MD5, RSA512_NULL_SHA
- If data encryption is required
 - f* Use cipher suites 08, 09 or 0A
 - f* 08 uses 512 bit keys, others 1024
 - f* 1024 bit RSA keylength is recommended (from a security point of view)

VS@

e-business



General Performance Hints for Connectors and TCP/IP

- Reduce amount of data transferred
 - f* Transfer only data that is needed
 - f* Issue only requests that are needed
- Use connection pooling
 - f* Reduce overhead of connection establishment
- Performance of connectors depends on
 - f* Network performance
 - f* Performance of "server"
 - f* Performance of "client" or middle tier

VS@

e-business



Dependencies for VSE/ESA Growth

- System dependencies
 - f* Many control-blocks etc.. still below the line
 - f* VTAM IOBUF areas in System GETVIS-24
 - f* Non-Parallel-Share limits n-way support
 - f* Number of tasks
 - Up to 255, 32 per partition, 208 subtasks in total
- Application dependencies
 - f* Integrated system concepts/functions
 - f* Functions/Applications dependencies
 - f* Number of users per TCP/IP partition



e-business



Dependencies for VSE/ESA Growth - continued

- Not being considered to be a limit
 - f* Number of partitions
 - 12 static + 150-200 dyn. partitions
 - f* Real storage (max. 2 GB)
 - f* Total virtual storage (max. 90 GB)
 - f* Total number of devices (3 digit CUU)
 - Max. 1024 devices (and 16 channels)
 - f* Total number of logical units
 - 255 per partition and 12x255=3060 in total
 - f* Label area
 - Max. about 9000 in total, and 712 in sub areas




e-business



Overhead Deltas for VSE Releases

-These numbers are for **average I/O intensive** workloads
 -More I/O intensive workloads can cause more overhead



VS@

e-business



Further Information

- VSE Homepage:
<http://www.ibm.com/servers/eserver/zseries/os/vse/>
- VSE Performance Homepage:
<http://www.ibm.com/servers/eserver/zseries/os/vse/library/vseperf.htm>
- Performance Documents from W. Kraemer
f available on the Performance Homepage
- VSE/ESA e-business Connectors User's Guide
<http://www.ibm.com/servers/eserver/zseries/os/vse/pdf/ieswue20.pdf>

VS@

e-business



Questions

