TCP/IP for VSE/ESA 1.4

Performance considerations

Ingo Franzki
e-mail: ifranzki@de.ibm.com
VSE/ESA Development

VSE

e-business

Agenda

- TCP/IP for VSE/ESA
  - Startup
  - Multiple TCP/IP partitions
- Telnet Performance results
- FTP Performance results
- News with 1.4
- OSA-Express and QDIO
- Secure Socket Layer
  - Performance results

VSE

e-business

TCP/IP Startup job

- VSE partition size
  - Start with 20-30M partition size
- SETPFIX LIMIT
  - Start with LIMIT=900K
- Type of VSE partition
  - Can be static of dynamic
- Parameters to EXEC IPNET
  - SIZE=IPNET
  - IPINIT0x contains all TCP/IP parameters
  - DSPACE=3M max. size of dspace used by VTAM

VSE

e-business

TCP/IP Dispatch Priority

- Select PRTY sequence (low to high)
  - Batch, CICS, TCP/IP, VTAM, POWER
- A second TCP/IP partition is recommended
  - If besides Telnet ...
    - Concurrent FTP activity (not FTPBATCH)
    - Concurrent LPR/LPD activities
  - High concurrent FTP (or LPR/LPD) activity may/will impact e.g. Telnet response times
  - Or use FTPBATCH

VSE

e-business

Multiple TCP/IP Partitions

- Each TCP/IP Partition should have
  - A separate IP address
  - A separate host name
  - Its own set of adapters
  - Its own setup of startup parameters

- Functional reasons
  - Separation of workload
  - Separation of production and test
  - Separation of production workload
  - Separation of network (e.g. security)

VSE

e-business

Multiple TCP/IP Partitions - continued

- Performance reasons
  - Exploit more than 1 engine for TCP/IP
    - Only one engine per partition
  - Need of more virtual storage below the line
    - e.g. Telnet (VTAM) buffers
  - Individual customization
  - Separation of TELNET and FTP/LPR activities

- INET link has no performance benefits
  - Recomendation: let each partition have its own network link

VSE

e-business

TCP/IP's Access to VSE Data

- VSAM
  - VSAM macros and VSAM code in SVA
- POWER
  - POWER SAS (XPCC)
- LIBR
  - LIBRM macro
- ICCF
  - SLI (Read only) and DTSIPWR in SVA
- SD
  - DTFSD macro (BAM)

VSE

e-business

Batch FTP From a Separate Partition

- // EXEC FTP
  - Only FTP initialization is done from a batch partition
  - No performance related benefits
- // EXEC FTPBATCH
  - Potential exploitation of >1 engine of an n-way
  - Separate File-I/O routine used per FTP
  - Control of FTP batch CPU dispatch priority
  - More overhead for data transfer between batch and TCP/IP partition
  - Move of data between batch and TCP/IP partition using access registers

VSE

e-business

TN3270 Measurement Environment

- VSE/ESA 2.3
- TCP/IP 1.3 (E/G/J/K) and also 1.4
- Turbo Dispatcher, single engine
- DSW online workload
- 2 CICS/VSE partitions (F4 and F5)
- TCP/IP for VSE/ESA (F7)
- F4 and F5 balanced with F7
- 125 active terminals per CICS partition
  - ► driven by TPNS
- POOL=YES and TELNETD_BUFFERS=20

VSE

e-business

TN3270 Measurement Results

Overhead in terms of CPU-time per transaction

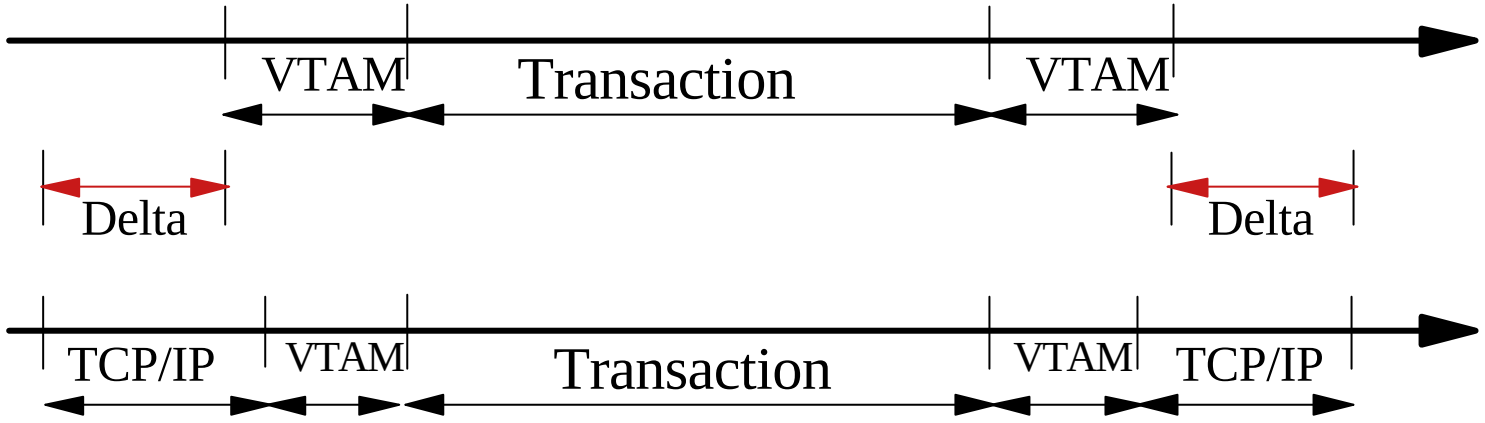

Delta/txn = msec/txn(TCP/IP) - msec/txn(VTAM)

Varies here between 20.0 and 24.9 msec

TN3270 Measurement Results - continued

Expected rel.-CPU-time and ITR-ratio vs SNA

$$\text{ITRR} = \text{ITR ratio} = \frac{\text{msec/txn (VTAM)}}{\text{msec/txn (TCP/IP)}}$$

In the measured cases, average overall (VTAM based) CPU-time of a transaction was about 20 msec (~280KI)
TCP/IP overhead was between 280KI and 350KI

Type/CPU-Heaviness of workload	Rel. CPU-time with TCP/IP	ITRR
DWS, measured (280KI)	2.0	0.5
Medium customer transaction (560KI)	1.5	0.67
Heavier customer transaction (840KI)	1.33	0.75
Heavy customer transaction (1000KI)	1.28	0.78

VSE

e-business

TN3270 Measurement Results - continued

- TCP/IP for VSE/ESA 1.4 vs. 1.3
  - 7-13 % less TCP/IP CPU-time overhead
- Response time impact is small
- TN3270 overhead
  - VSE/ESA native
    - Online utilization increases from 50% to 62%
  - VM/VSE Guest
    - Online utilization increases from 60% to 74%

TN3270 Measurement Results - continued

- TN3270 Virtual Storage capacity

	125 daemons		per daemon	
	-24	-31	-24	-31
TCP/IP GETVIS	476K	600K	3.8K	4.8K
VTAM GETVIS	0K	52K	0K	0.4K
SVA	20K	524K	0.16K	4.2K

Rough estimate for TN3270 VS-Capacity:

Max. #TN daemons = (remaining GETVIS-24) / 4K

Example:
A remaining GETVIS-24 of about 10M, gives
about 2500 Telnet daemons

VSE

e-business

FTP Measurement Results

- EDR = effective Data Rate (KB/sec)
- It is irrelevant who initiated an FTP transfer
- Transfer of a file from A to B may differ in EDR from transferring the identical file from B to A
  - Speed of physical HDD
  - Read/write caching
  - Blocksizes used (KB/IO)
- The higher the EDR of an FTP transfer, the higher is the required CPU utilization
- EDRs displayed by TCP/IP for VSE/ESA
  - Transfer sends
  - File I/O seconds

VSE

e-business

FTP Measurement Results - continued

Parameters	FTP speeds Source	Target	Network	CPUT/KB
Network speed and load			X	
TCP/IP parameters	X	X	X	X
FTP parameters	X	X	X	X
DASD speed (READ/WRITE)	X	X		
Local file definition				
- type	X	X		X
- log record length (NFS)	X	X		
- blocksize on disk	X	X		X
- I/O blocking (KB/IO)	X	X		X
- ASCII/EBCDIC/BINARY	X	X		X
size of files	X	X		X
Processor speed	X	X		X
other concurrent activities	X	X	X	
TCP4VSE PTF level	X	X	X	X

VSE

e-business

FTP Measurement Results - continued

EDR ranges (KB/sec) observed (1.3)

	FTP to VSE	FTP from VSE	Major impact
LIBR	340	470	DASD, network speed
POWER	115	290	DBLK
VSAM ESDS		460 360 160	to S/390 (CTCA) to RS/6000 CLAW Via CLAW & T/R

CPU resources (KI/KB) required (1.3)

	FTP to VSE	FTP from VSE	Dependencies
LIBR	18.9 - 20.1	11.9 - 13.3	
POWER	85	45	
VSAM ESDS		7.6 - 9.2	Convertion

VSE

e-business

FTP Measurement Results - continued

- TCP/IP for VSE/ESA 1.4 (ServPack A)
  - EDRs increased by 10% to 30%
  - CPU-time consumption decreased by about 25%

- Virtual Storage Capacity

	10 daemons		per daemon	
	-24	-31	-24	-31
TCP/IP GETVIS	3104K	40K	310K	4K

Max #FTP daemons = (remaining GETVIS-24) / 310K

Example:
A remaining GETVIS-24 of about 10M, gives about 32 FTP daemons

FTP with Batch Measurement Results - continued

Data rate comparison (VSAM)

	Overall EDR Transfer (KB/sec)		File I/O (KB/sec)	
	Real 9345	Virt. Disk	Real 9345	Virt. Disk
Interactive FTP	639	930	682	1462
Batch FTP	639	930	682	1462
FTPBATCH	511		682	

- Same rates as for Interactive FTP
  - Except transfer rate seen by FTPBATCH
- Overall EDRs for (single) FTPBATCH are about 15% lower here than from Batch FTP

VSe

e-business

FTP with BatchMeasurement Results - continued

- FTPBATCH with slightly higher CPU-time and with lower EDR
- FTPBATCH file transfers
  - Can be better workload balanced (controlled)
    - Via PRTY
  - Can run concurrently and thus achieve a higher sum of FTP EDRs
  - Allow to exploit >1 processor engines

VSE

e-business

News with TCP/IP for VSE/ESA 1.4

- Mainly functional enhancements
  - Implementation of slow start mechanism (outbound TCP)
  - CHECKSUM calculation also in H/W
- PQ40278 (ServPack A)
  - Many fixes and slight enhancements
    - better algorithm to reduce re-transmissions
    - TRACERT and DISCOVER commands

VSE

e-business

News with TCP/IP for VSE/ESA 1.4 (continued)

- PQ45314 (ServPack B)
  - Multi thread event processing
  - Entire redesign of EMAIL client
  - TELNETD: allocation of all buffers during initialization of TELNETDs (for POOL=NO only)
- PQ52348 (ServPack C)
  - Includes SSL for VSE
  - OSA Express Support (VSE/ESA 2.6 only)
- PQ55591 (ServPack D)
  - Removed limitation of 64K per send()

VSE

e-business

Multi Thread Event Processing

- More than 1 events can be processed at a time
  - Seperate TCP/IP internal task is assigned to any event (printout)
- Possible problem
  - Unpredictable order of events from independent jobs
  - Some printers only accepts one connection at a time
- SET SINGLEDEST=ON
  - Only one open connection possible to one destination

VSE

e-business

New: OSA-Express

- VSE/ESA 2.6
- Available for G5 and above
- Exploits Queued Direct I/O

	Gigabit Ethernet	Fast Ethernet 100 Mbps	ATM-LE 155 Mbps	Tokenring 4/16/100 Mbps
CHIPID TYPE=OSE (non-QDIO)	no	yes	yes	yes
CHPID TYPE=OSD (QDIO)	yes	yes	yes	yes

OSA-Express for IBM eServer zSeries and S/390, G221-9110-01, 11/2001

VSE

e-business

New: OSA-Express - continued

- Queued Direct I/O
  - Designed for very efficient exchange of data
  - Uses the QDIO Hardware Facility, without traditional S/390 I/O instructions
  - Without interrupts (in general)
  - Use of internal queues
  - With pre-defined buffers in memory for asynchronous use

VSE

e-business

OSA-Express Measurements

- Environment
  - VSE/ESA 2.6 on G6 native (LPAR)
    - TCP/IP 1.4 ServPak C
  - Linux on Netfinity
- Network attachment
  - Gigabit Ethernet QDIO (MTU=1500)
  - Fast Ethernet QDIO (MTU=1500)
  - Fast Ethernet Non-QDIO (MTU=1500)
- Workload
  - GET = VSE to Linux, 100MB $NULL file
  - PUT = Linux to VSE, 100MB $NULL file

OSA-Express Measurements - continued

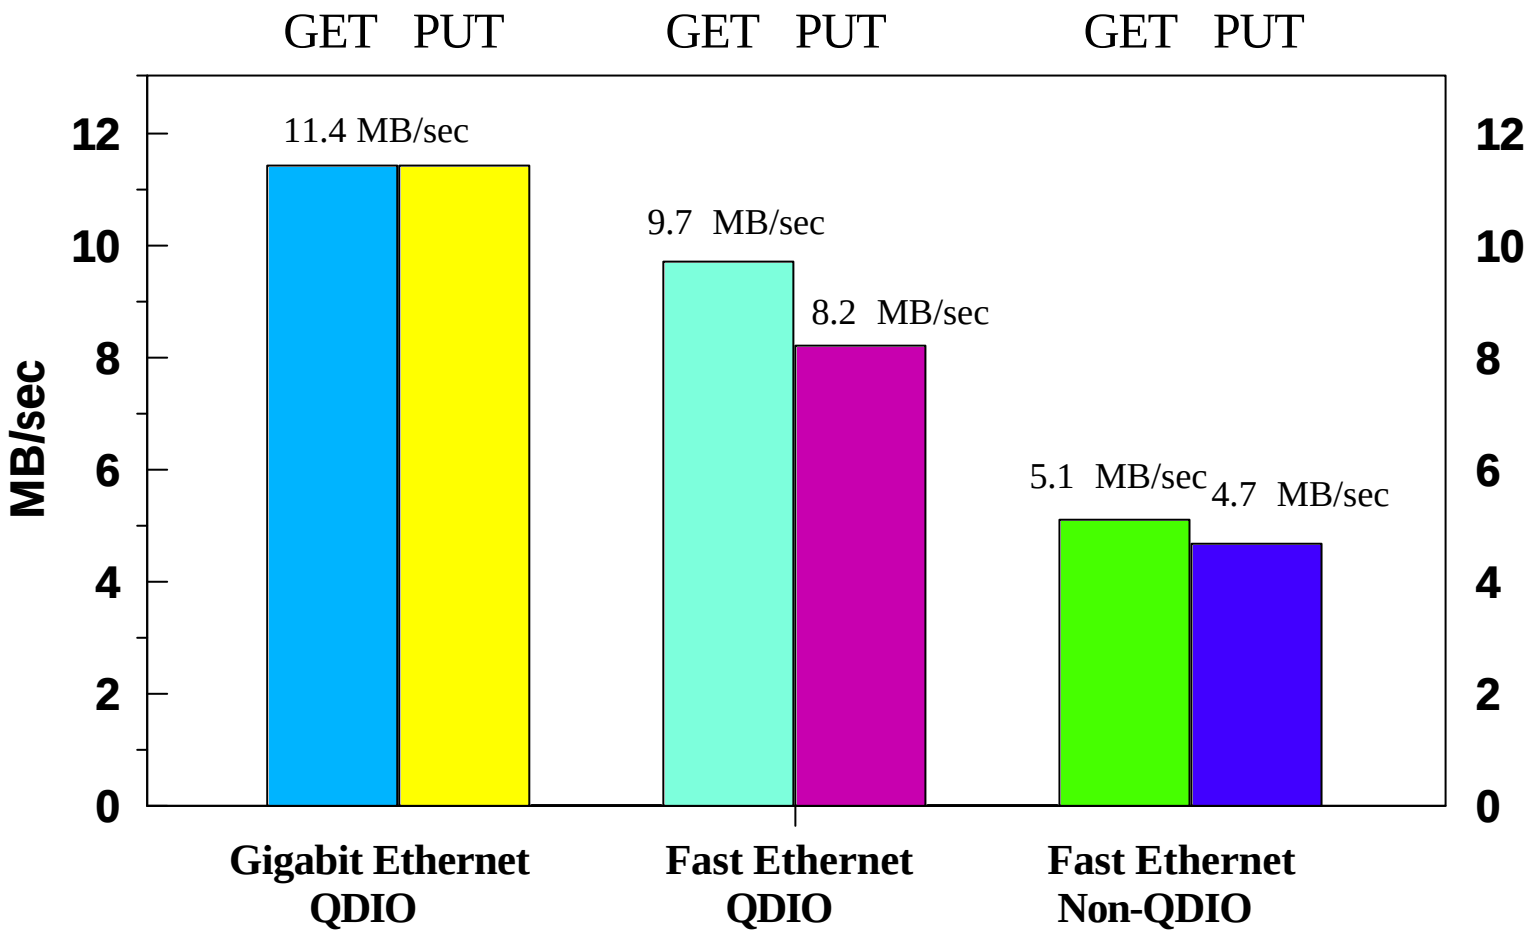

MTU = 1500

VSE

e-business

Gigabit Ethernet Measurements

- Environment
  - VSE/ESA 2.6 on G6 native (LPAR)
    - TCP/IP 1.4 ServPak C
  - Linux on G6 native (LPAR)
- Network attachment
  - Gigabit Ethernet QDIO
  - MTU = 576...9000
- Workload
  - GET = VSE to Linux, 100MB $NULL file
  - PUT = Linux to VSE, 100MB $NULL file

VSE

e-business

Gigabit Ethernet Measurements - continued

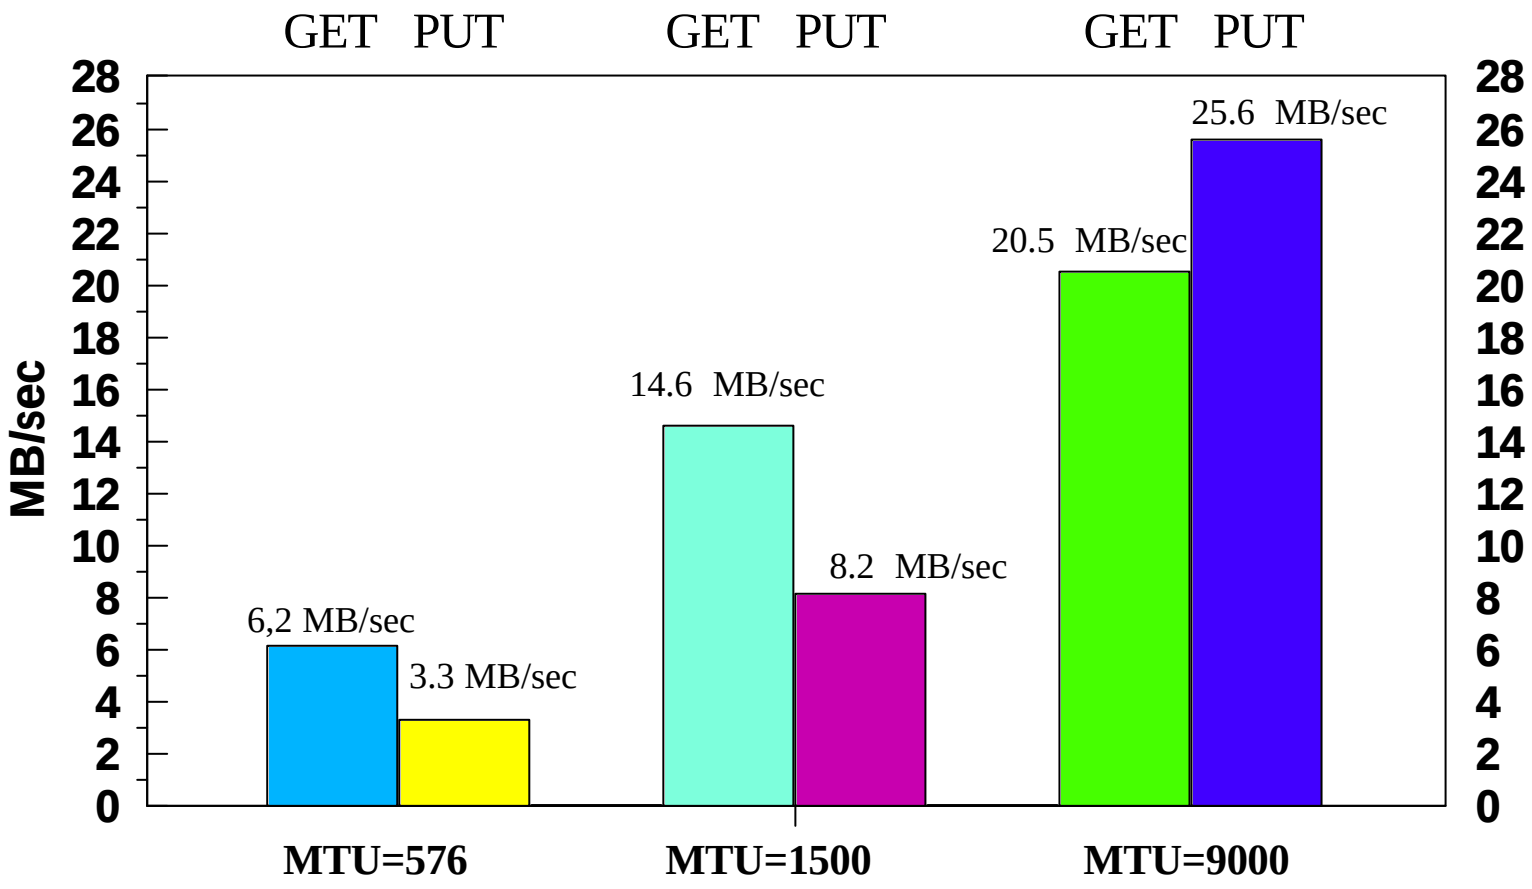

VSE

e-business

SSL Overview

- SSL for VSE is part of the TCP/IP base (ServPack C)
- Enabled with the Application Pak
- Integrated into TCP/IP for VSE/ESA
- Supports SSL 3.0 and TLS 1.0
- Key exchange: RSA
- Data Encryption: DES and Triple DES
- Hash algorithm: MD5, SHA
- Supports X.509v3 PKI Certificates
- SSL daemon implementation for HTTPS, Telnet
- SSL API compatible with the OS/390 SSL API

VSE

e-business

Key Management

- Keys and certificates are stored in a "keyring file"
  - In a VSE library
- SSL for VSE uses 3 VSE library members:
  - keyname.PRIV - the private key
  - keyname.CERT - the certificate
  - keyname.ROOT - the root certificate
- Stored in library CRYPTO.KEYRING per default
- Utilities available for key management and creation
  - CIALPRVK, CIALCERT, CIALROOT
- $SOCKOPT.PHASE defines the SSL parameters

VSE

e-business

SSL Daemon (SSLD)

- Define a SSL daemon for each TCP port that you want to secure:

  - DEFINE TLSD,ID=MYSSLD,

PORT=443,	HTTPS port
PASSPORT=443,	
CIPHER=0A096208,	Cipher suites
CERTLIB=CRYPTO,	library name
CERTSUB=KEYRING,	sublibrary name
CERTMEM=MYKEY,	member name
TYPE=1,	server application
MINVERS=0300,	SSL 3.0
DRIVER=SSLD	Driver phase name

VSE

e-business

Secure Socket Layer API

- Compatible to OS/390 SSL API
- Functions available for
  - Session initiating
  - Sending/receiving data
  - Ending a session
- SSL API is based on Socket API
- SSL API can be called from
  - LE-C programs
  - Assembler programs

VSE

e-business

CryptoVSE API

- Native cryptographic API (not available though LE)
- Provides cryptographic services:
  - Data encryption
    - DES
    - Triple DES
    - RSA PKCS #1
  - Message Digest
    - MD5
    - SHA-1
  - Digital Signatures
    - RSA PKCS #1 with SHA1 or MD5
  - Message Authentication
    - HMAC

VSe

e-business

Restrictions

- Cipher Suites supported:
  - 01 - RSA512_NULL_MD5
  - 02 - RSA512_NULL_SHA
  - 08 - RSA1024_DES40_CBC_SHA
  - 09 - RSA1024_DES_CBC_SHA
  - 0A - RSA1024_3DES_CBC_SHA
  - 62 - RSA1024_EXPORT_DES_CBC_SHA
- Only one Root certificate
- Certificate revocation lists not supported
- Keyring is not password protected
- Software encryption only

Performance Related Parameters

Parameters	Session initiating	Data exchange
Key exchange algorithm		
RSA512	X	-
RSA1024	X	-
Encryption Algorithm		
NULL	-	X
DES40CBC	-	X
EXPORT_DESCBC	-	X
DESCBC	-	X
3DESCBC	-	X
Hash Algorithm		
MD5	X	X
SHA	X	X
Session caching	X	-
Message Length	-	X

-Data exchange overhead is proportional to bytes/msg

-CPU-time overhead caused by SSL is in

-TCP/IP partition for SSL Daemon

-application partition for API usage

VSE

e-business

Measurement Environment

- ECHO Server
  - receives and decrypts a message (80 bytes)
  - encrypts and sends the message
  - Application/Transaction time is minimal
- VSE Connectors (SSL enabled)
  - download of a LIBR member to PC (1KB, 7KB)
    - mainly VSE outbound = encrypt
  - upload of a LIBR member to VSE (1KB, 7KB)
    - mainly VSE inbound = decrypt
  - Application time contains LIBR I/Os

VSE

e-business

Measurement Environment - continued

- Variations
  - SSL / non-SSL
  - Cipher Suites
    - 01 RSA512_NULL_MD5
    - 02 RSA512_NULL_SHA
    - 09 RSA1024_WITH_DES_CBC_SHA
    - 0A RSA1024_WITH_3DES_EDE_CBC_SHA
  - key length (512 / 1024 bit)

SSL Overhead

Non SSL

SSL

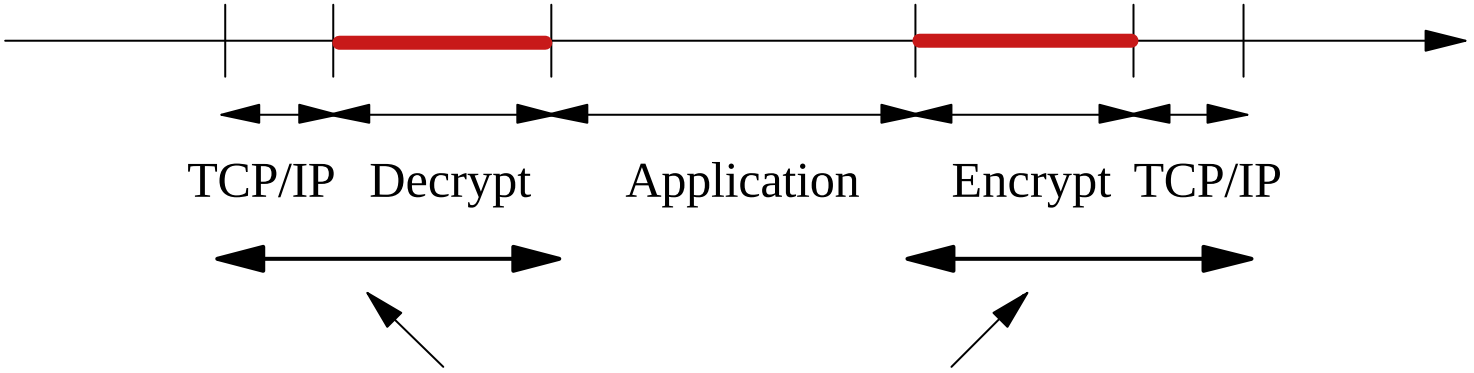

this has been measured

SSL Measurements Results

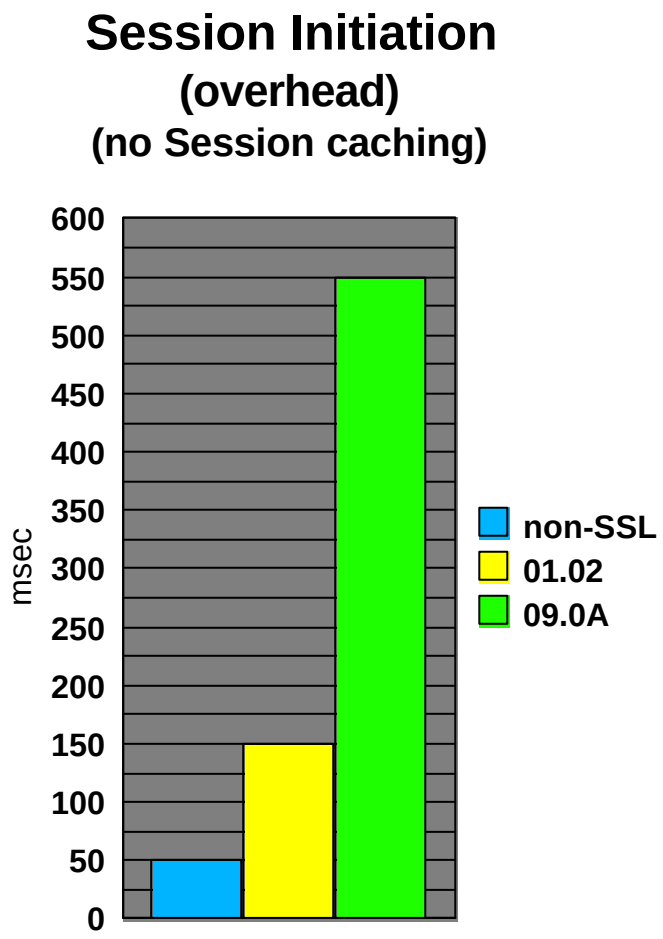

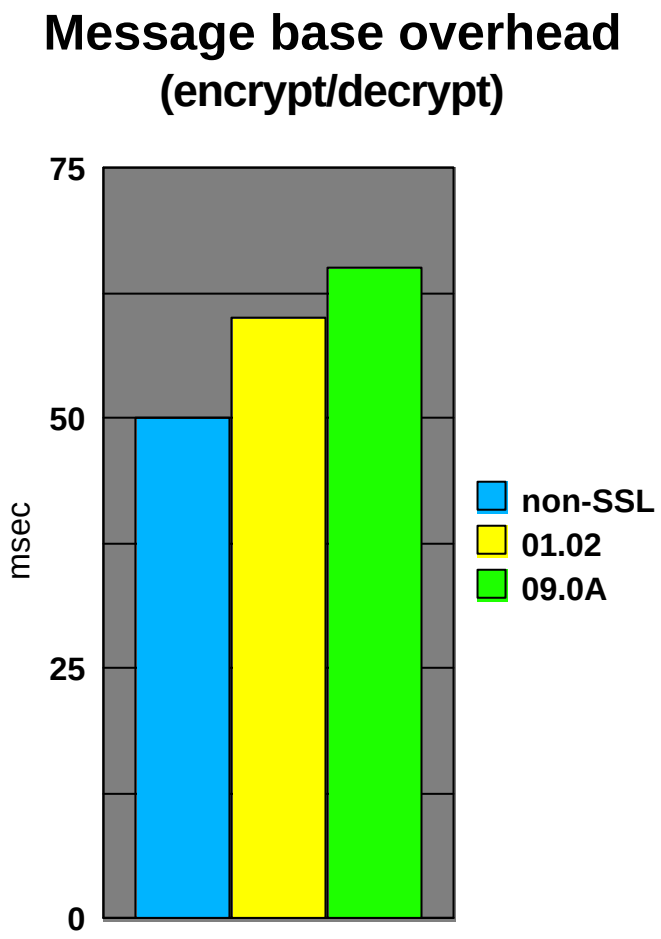

VSE

e-business

SSL Measurements Results - continued

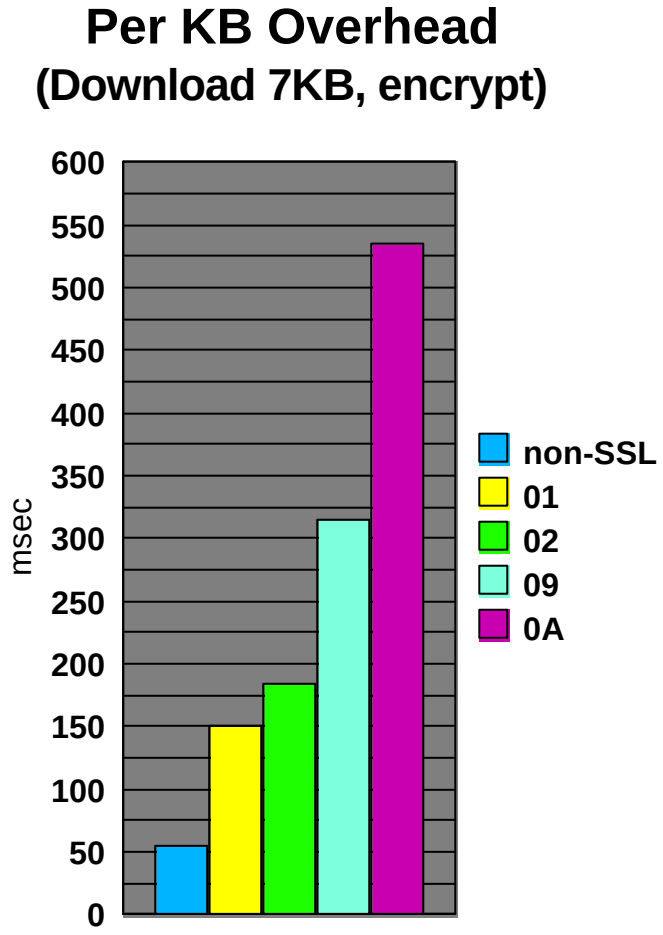

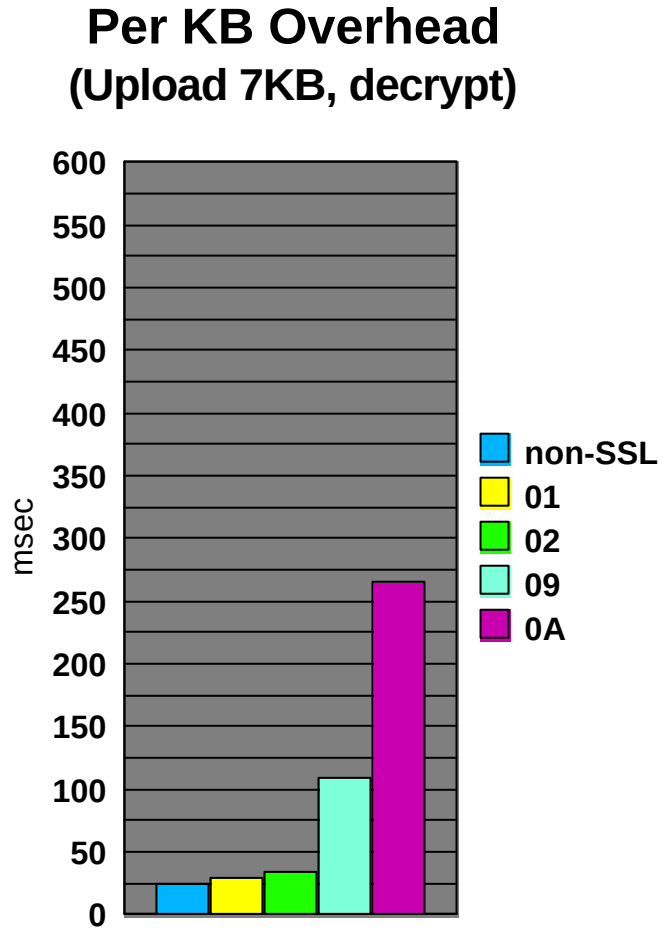

Note: These measurements includes application CPU-time for reading/writing LIBR member

SSL Measurements Results - continued

Per KB CPU-time
(Download, encrypt)
total = base + encrypt

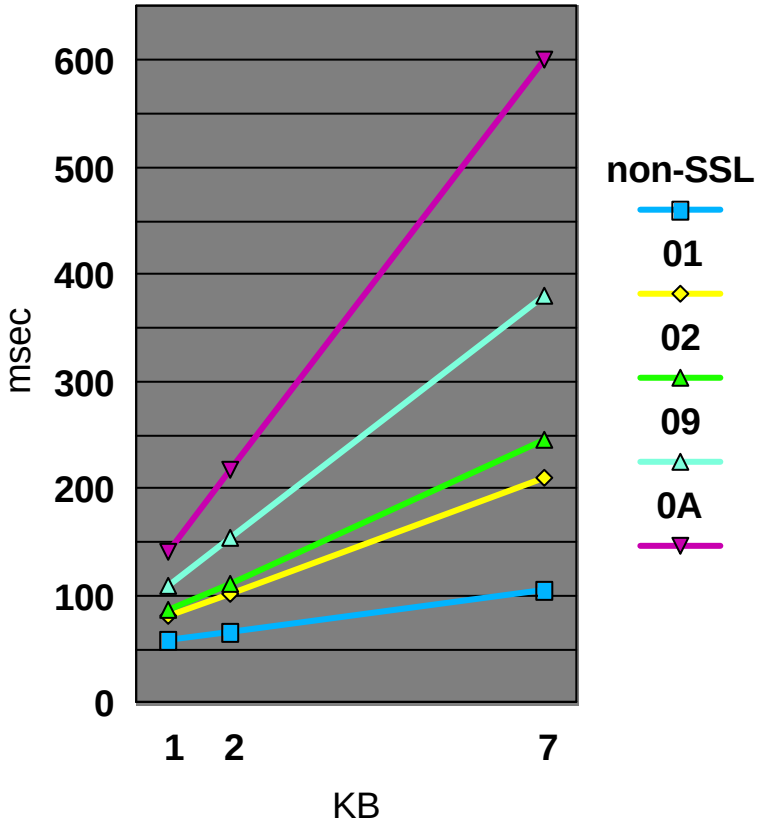

Per KB CPU-time
(Upload, decrypt)
total = base + decrypt

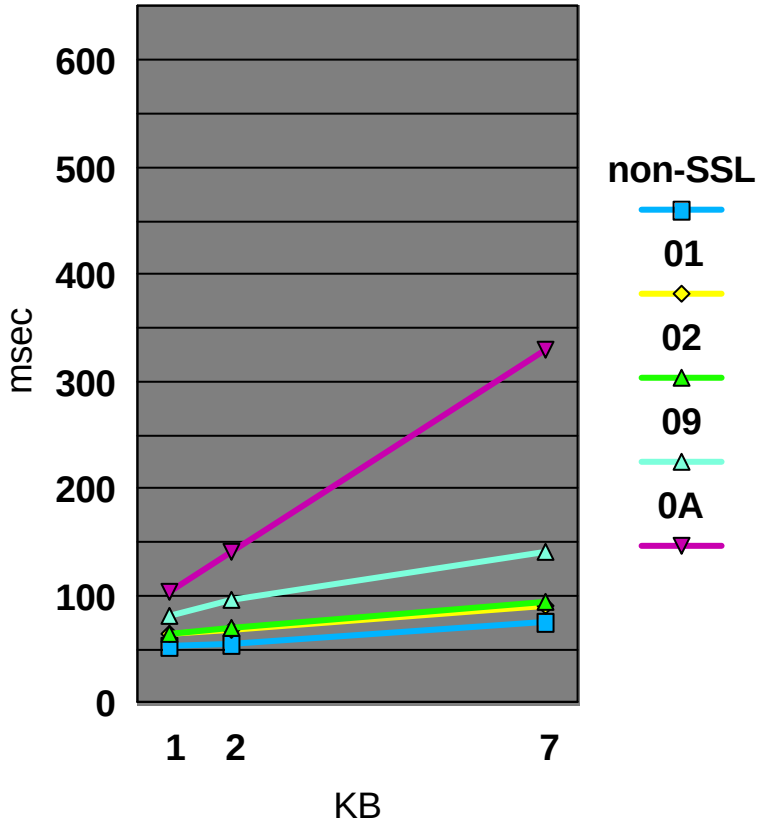

VSe

e-business

SSL Performance Recommendations

- Use SSL only if there is a need for
  - If at least one of the follwoing is required
    - Keeping secrets
    - Proving identity
    - Verifying information
- Cipher Suites 01 and 02 has less CPU-time consumption, but NO data encryption
  - RSA512_NULL_MD5, RSA512_NULL_SHA
- If data encryption is required
  - Use cipher suites 09 or 0A

VSE

e-business

Questions