

Understanding z/OS Communications Server storage use

Alfred B Christensen – alfredch@us.ibm.com
IBM Raleigh, NC

Thursday 5-Aug-2010 - 3:00 PM to 4:00 PM



SHARE in Boston

Understanding z/OS Communications Server storage use

Session number:	???
Date and time:	Thursday August 5, 2010 - 3:00 PM - 4:00 PM
Location:	Room 109 (Hynes Convention Center)
Program:	Communications Infrastructure
Project:	Communications Server
Track:	SNA/IP Integration and z/OS Systems Programming
Classification:	Technical
Speaker:	Alfred B Christensen, IBM
Abstract:	<p>If you have ever wondered how the z/OS Communications Server uses storage in your z/OS LPARs, this session is for you. The session will provide an overview and explanation of the z/OS Communications Server storage concepts and model - how TCP/IP and VTAM uses private storage, CSA, ECSA, CSM storage (ECSA, fixed, data spaces), 64-bit common, etc. The session will provide hints and tips about how to monitor Communications Server storage usage using display command and RMF reports.</p>



Trademarks, notices, and disclaimers

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

- | | | | | |
|--|---|---|--|--|
| <ul style="list-style-type: none"> • Advanced Peer-to-Peer Networking® • AIX® • alphaWorks® • AnyNet® • AS/400® • BladeCenter® • Candle® • CICS® • DataPower® • DB2 Connect • DB2® • DRDA® • e-business on demand® • e-business (logo) • e business (logo)® • ESCON® • FICON® | <ul style="list-style-type: none"> • GDDM® • GDPS® • Geographically Dispersed Parallel Sysplex • HiperSockets • HPR Channel Connectivity • HyperSwap • i5/OS (logo) • i5/OS® • IBM eServer • IBM (logo)® • IBM® • IBM zEnterprise™ System • IMS • InfiniBand® • IP PrintWay • IPDS • iSeries • LANDP® | <ul style="list-style-type: none"> • Language Environment® • MQSeries® • MVS • NetView® • OMEGAMON® • Open Power • OpenPower • Operating System/2® • Operating System/400® • OS/2® • OS/390® • OS/400® • Parallel Sysplex® • POWER® • POWER7® • PowerVM • PR/SM • pSeries® • RACF® | <ul style="list-style-type: none"> • Rational Suite® • Rational® • Redbooks • Redbooks (logo) • Sysplex Timer® • System i5 • System p5 • System x® • System z® • System z9® • System z10 • Tivoli (logo)® • Tivoli® • VTAM® • WebSphere® • xSeries® • z9® • z10 BC • z10 EC | <ul style="list-style-type: none"> • zEnterprise • zSeries® • z/Architecture • z/OS® • z/VM® • z/VSE |
|--|---|---|--|--|
- * All other products may be trademarks or registered trademarks of their respective companies.

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

- Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
- 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 trademarks of Microsoft Corporation in the United States, other countries, or both.
- InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
- IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

Notes:

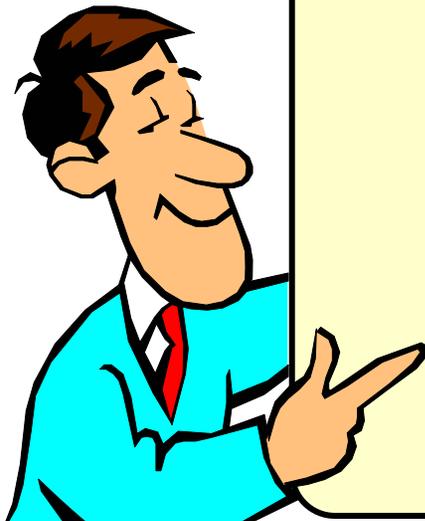
- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job 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 improvements equivalent to the performance ratios stated here.
- IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Refer to www.ibm.com/legal/us for further legal information.

Agenda



- Background information
- TCP/IP storage use
- Communications Storage Manager (CSM) storage use
- VTAM storage use
- Use of RMF to monitor z/OS CS storage use
- Summary



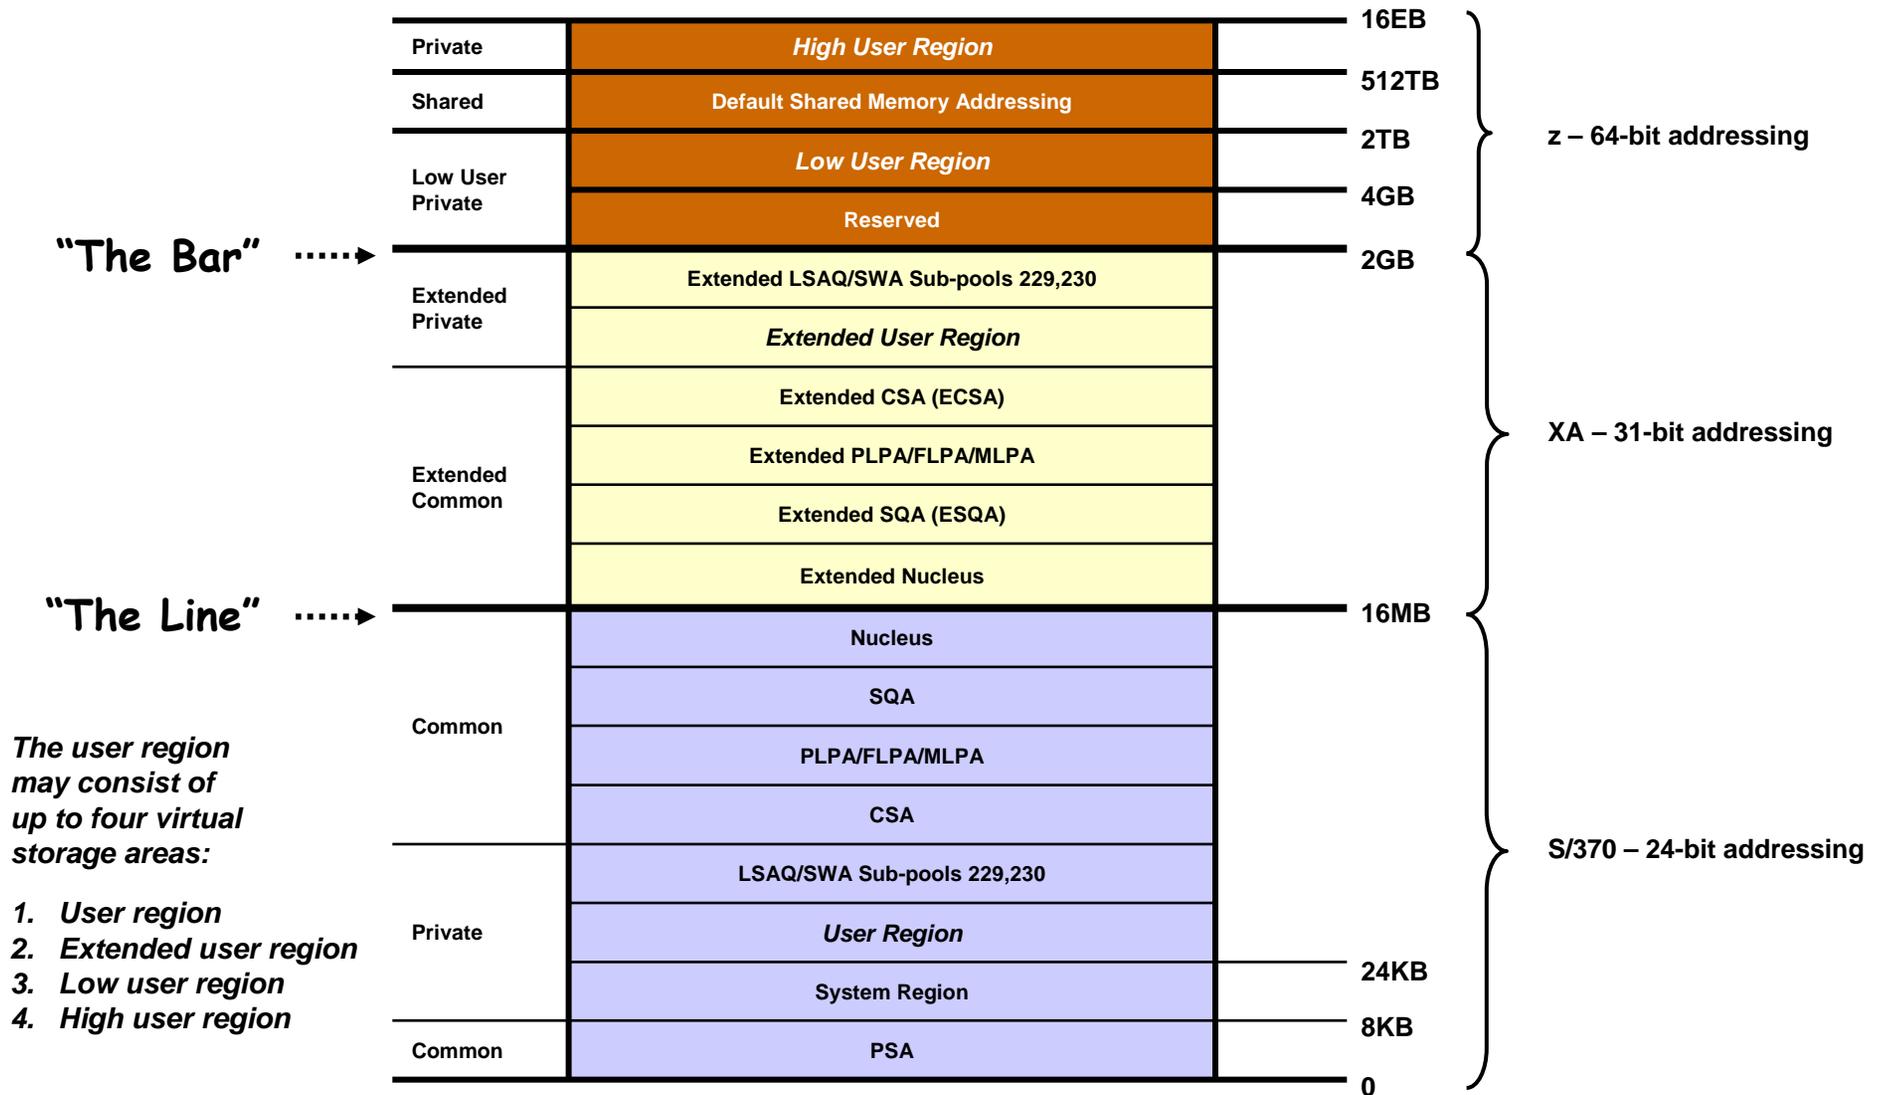
Disclaimer: All statements regarding IBM future direction or intent, including current product plans, are subject to change or withdrawal without notice and represent goals and objectives only. All information is provided for informational purposes only, on an “as is” basis, without warranty of any kind.

Understanding z/OS Communications Server Storage Usage

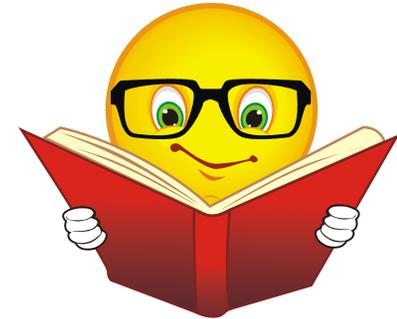
Background information



z/OS virtual storage map



KB, MB, GB, TB, PB, EB, ZB – and the lot ..



Name	Short	Size in bytes	2 to the order of
1 kilobyte	KB	1,024	10
1 megabyte	MB	1,048,576	20
1 gigabyte	GB	1,073,741,824	30
1 terabyte	TB	1,099,511,627,776	40
1 petabyte	PB	1,125,899,906,842,624	50
1 exabyte	EB	1,152,921,504,606,846,976	60
1 zettabyte	ZB	1,180,591,620,717,411,303,424	70
1 yottabyte	YB	1,208,925,819,614,629,174,706,176	80

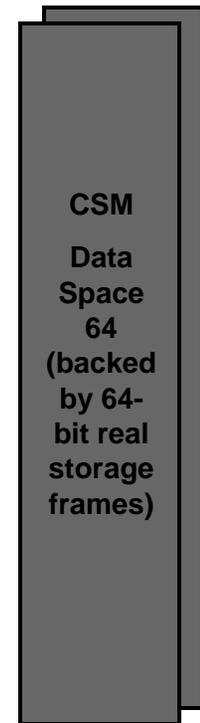
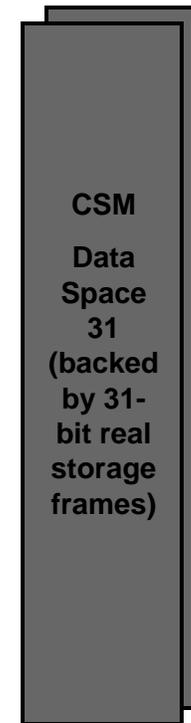
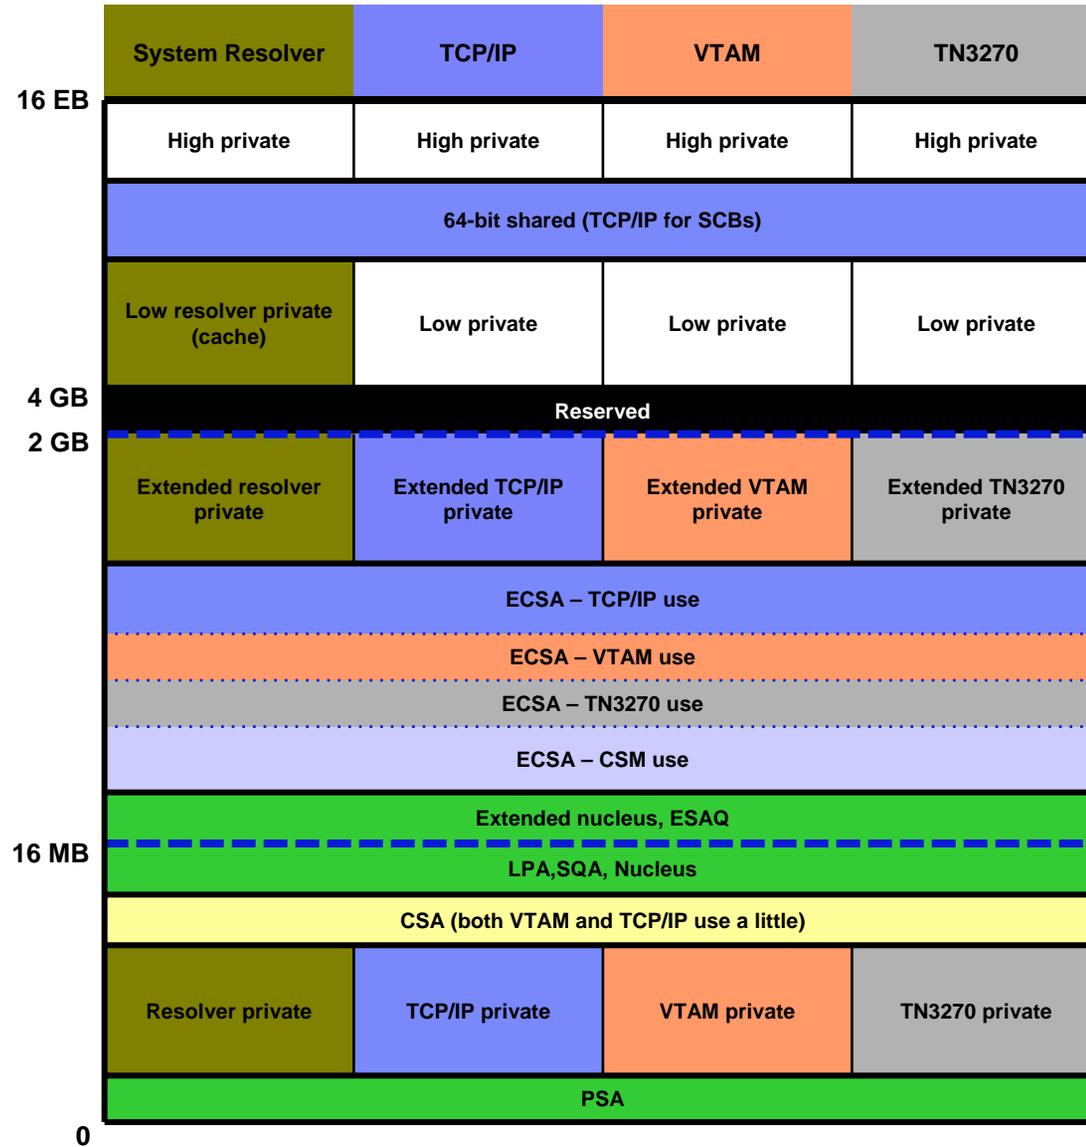
A perspective of address space size over time

- A 31-bit XA address space is 128 times larger than a 24-bit S/370 address space
- A 64-bit z/Architecture address space is 8 billion times larger than a 31-bit XA address space and 1024 billion times larger than a S/370 address space

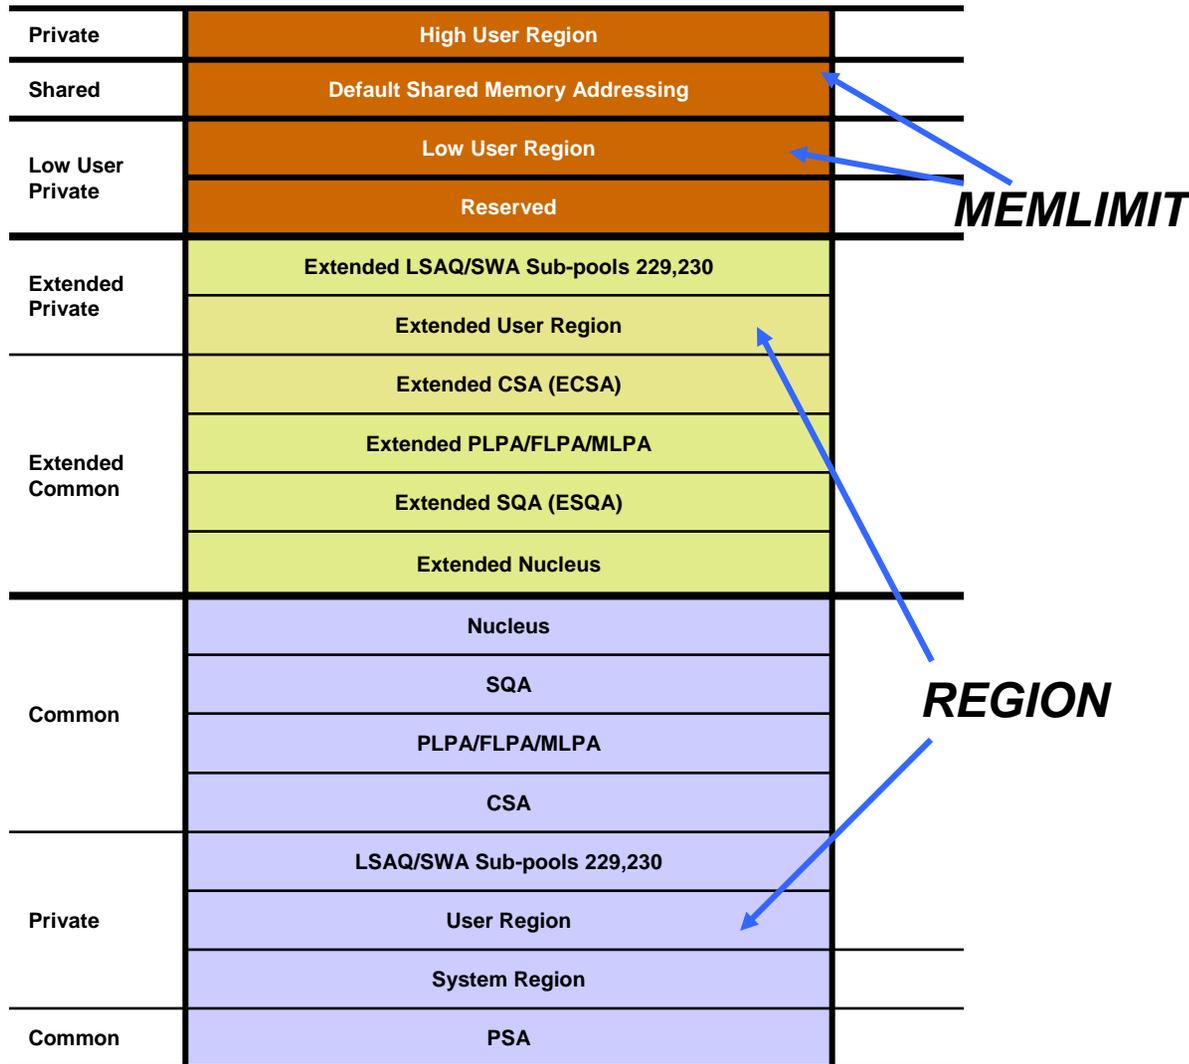
Architecture	Size	Relative size	Comparison – metric	Comparison – non-metric
24-bit S/370 address space	16 MB	1	7 cm	2.8 inches
31-bit XA address space	2 GB	128	9 m	30 feet
64-bit z/Architecture address space	16 EB	1 099 511 627 776	Distance between Earth and Mars 78 341 212 Km	Distance between Earth and Mars 48 678 972 miles

Note: distance between earth and Mars varies between 66 million Km and 402 million Km. 78 341 212 Km is the average distance.

Communications Server virtual storage usage overview



Limiting the amount of virtual storage in the user region(s)



- **The REGION JCL keyword limits amount of space in the user region and the extended user region**

 - Includes user region below the line and between the line and the bar, but not above the bar
 - Does not include space in any of the common areas
 - REGION=0M gives all available storage
- **The MEMLIMIT JCL keyword limits the amount of space in the low and high user regions**

 - Above the bar
- **Installation defaults can be defined in JES and SMF**

 - To be used as default when no REGION or MEMLIMIT keywords are specified in the JCL

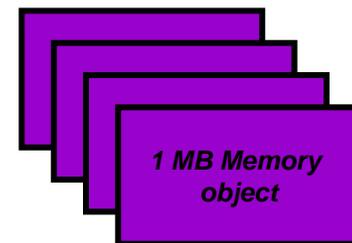
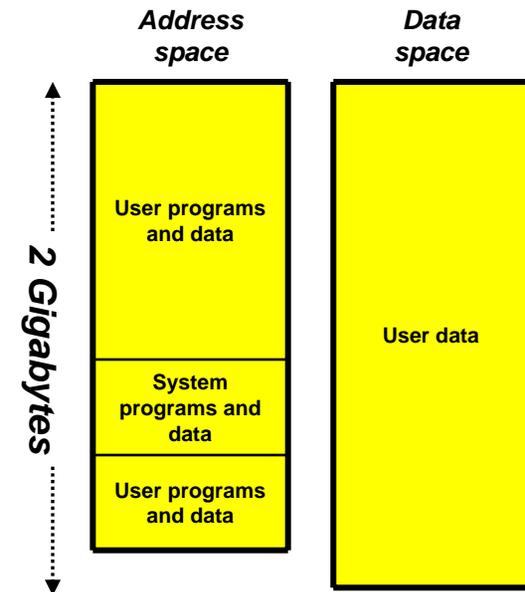
Installation exits may override JCL keywords

- **Installation exits can enforce installation standards – will override any specifications found in JCL**
 - **IEALIMIT**
 - Is no longer a recommended exit routine for this purpose
 - Can only enforce standards for user region below the 16 MB line and is linked into the MVS nucleus
 - **IEFUSI**
 - Is the preferred exit routine to enforce these limits
 - Resides in the LPA
 - Can enforce standards for:
 - The user region below the 16 MB line
 - The extended user region between the 16 MB line and the 2 GB bar
 - The user region above the 2 GB bar (MEMLIMIT JCL keyword)
 - PPT attribute NOHONORIEFUSIREGION can be used for selected programs to ignore IEFUSI standards – defined in the SCHEDxx SYS1.PARMLIB member

Data Spaces and memory objects

- **Data Space**
 - A data space is max 2 GB in size
 - 31-bit addressing is used to access data in a data space
 - A data space may be backed by 64-bit real storage frames (if the system has more than 2 GB of real storage)
 - What CSM means with the term ‘DATA SPACE 64’:
 - This is not 64-bit virtual storage, this is a normal 2 GB data space that is allowed to be backed by 64-bit real storage frames
 - CSM uses so-called common data spaces
 - System-wide number of common data spaces is limited by the MAXCAD keyword in IEASYSxx
 - Default is 50

- **Memory object**
 - Virtual storage above the bar may be allocated and made part of an address space’s virtual storage addressing range
 - Obtained in multiples of 1 MB (a so-called memory object)
 - A memory object may be a common memory object or a shared memory object – residing in the area between 2 TB and 512 TB
 - A memory object may also be part of the private addressing range within an address space – part of the user region



System wide storage-related definitions

- SYS1.PARMLIB – IEASYSnn member
 - **CSA=(a,b)**
 - "a" specifies the size of the CSA, located below 16MB
 - "b" specifies the size of the extended CSA, located above 16MB
 - Make sure this number is big enough to accommodate TCP/IP's, VTAM's, and CSM's ECSA requirements
 - Plus the requirements of whatever else you run on that system
 - Example:
 - CSA=(4M,256M)
 - ❑ 4MB CSA below the 16 MB line
 - ❑ 256MB ECSA above the 16 MB line
 - **HVCOMMON=a**
 - "a" specifies the size of the 64-bit common area (xxG or xxT)
 - The 64-bit common area will be placed below the 4T line.
 - The value you specify will be rounded up to a 2 gigabyte boundary
 - Default is 64 GB, minimum is 2 GB, maximum is 1 TB

Understanding z/OS Communications Server Storage Usage

TCP/IP storage use



TCP/IP storage use overview

- **TCP/IP user region**
 - Primarily extended user region
 - Very little, if any, in the user region below the line
 - Nothing in low and high user regions above the bar
 - TCP/IP refers to this storage as POOL storage

- **Extended Common Service Area (ECSA)**
 - The part of ECSA that is obtained and managed by TCP/IP
 - Not part of CSM storage
 - During TCP/IP initialization, a few modules are loaded into ECSA using dynamic LPA functions
 - Prior to z/OS V1R11 this was done by TCP/IP itself using directed load
 - Main module is EZBTIINI, which in the current z/OS V1R11 distribution is 7,117,424 bytes in size
 - Otherwise ECSA storage is primarily used for common control blocks
 - ECSA storage is obtained in either key 0 or key 6
 - Most workload-related storage is obtained in key 6
 - ECSA storage is obtained with TCP/IP or the “System” as owner
 - Almost all TCP/IP’s ECSA is obtained with the “System” as owner

- **Common 64-bit memory**
 - In z/OS V1R11, TCP/IP uses common 64-bit memory objects for Socket Control Blocks (SCBs)

D TCPIP,[stackname],STOR

- This does not include CSM storage
- Limits can optionally be configured in the TCP/IP Profile for POOL and ECSA storage

- **POOL** storage
 - The amount of storage TCP/IP uses in its user regions

- **ECSA** storage
 - In releases prior to z/OS V1R12, ECSA includes the size of load module EZBTIINI, which is loaded into ECSA during initialization (in z/OS V1R11 using dynamic LPA functions).
 - EZBTIINI resides in storage key 0.
 - The total size of this module is roughly 6,950K in z/OS V1R11, which is up about 660K from z/OS V1R10.
 - In z/OS V1R12, ECSA does not include any load modules that have been loaded into common storage
 - The size of load modules in ECSA is reported separately in the output from this command
 - Most of the ECSA storage is from an RMF perspective reported under the MVS master address space (RMF job name of %MVS).
 - The ECSA storage is obtained with OWNER(SYSTEM)
 - You cannot get an RMF Common Storage report to match the number above

- **64-BIT common**
 - This storage is obtained in 1 MB objects, but only backed by real storage or page data set space to the extent TCP/IP uses it

- **CSA Modules**
 - This is common storage used by TCP/IP load modules. Reported in z/OS V1R12.



D TCPIP,[stackname],STOR

z/OS V1R11

```

13.59.24  d tcpip,tcpcs,stor
13.59.26  EZZ8453I TCPIP STORAGE
EZZ8454I TCPCS      STORAGE      CURRENT    MAXIMUM    LIMIT
EZZ8455I TCPCS      ECSA        9645K     10087K    NOLIMIT
EZZ8455I TCPCS      POOL        14017K    14171K    NOLIMIT
EZZ8455I TCPCS      64-BIT COMMON    1M        1M        NOLIMIT
EZZ8459I DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY
    
```

z/OS V1R12

```

09.19.47  d tcpip,,stor
09.19.47  EZZ8453I TCPIP STORAGE
EZZ8454I TCPCS      STORAGE      CURRENT    MAXIMUM    LIMIT
EZZ8455I TCPCS      ECSA        2789K     3245K     NOLIMIT ---
EZZ8455I TCPCS      POOL        14212K    14212K    NOLIMIT
EZZ8455I TCPCS      64-BIT COMMON    1M        1M        NOLIMIT
EZZ8455I TCPCS      CSA MODULES   7423K     7423K     NOLIMIT ---
EZZ8459I DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY
    
```

ECSA storage is now reported as workload-related storage use only. Load modules in ECSA are reported separately in the CSA Modules line.



D TCPIP, TN-Server-Name, STOR

- **Each TN3270 server address space on your system also support the D TCPIP,,STOR command**
 - Output will reflect that TN3270 server address space’s use of storage
 - Enter the command for each TN3270 server address space

```

13.30.24  d tcpip,tn3270a,stor
13.30.24  EZZ8453I TELNET STORAGE
EZZ8454I TN3270A  STORAGE          CURRENT    MAXIMUM    LIMIT
EZZ8455I TN3270A  ECSA              97K        117K       NOLIMIT
EZZ8455I TN3270A  POOL              6636K     7392K      NOLIMIT
EZZ8455I TN3270A  64-BIT COMMON    0M         0M         NOLIMIT
EZZ8455I TN3270A  CTRACE           262372K   262372K   262372K
EZZ8459I DISPLAY TELNET STOR COMPLETED SUCCESSFULLY
    
```

- There are no limits that can be configured for a TN3270 server address space
- The storage options include a CTRACE storage type
 - The CTRACE storage is in the TN3270 server’s private area (not common and not data space), but is not included in the POOL storage type line
 - It is currently set to 256 MB
 - The limit here is not configurable
 - The size reflects how much storage is obtained for tracing, not what currently is being used for tracing

How do you limit TCP/IP's storage use?

- **POOL storage**
 - Via the REGION keyword in TCP/IP's start up JCL
 - User and extended user region size
 - Can be overridden by installation exits
 - Primarily IEFUSI
 - If used, should be accompanied by limitations in TCP/IP's profile
 - TCP/IP Profile
 - GLOBALCONFIG POOLLIMIT xxK or xxM
 - 0 means no limit and is the default

- **ECSA storage**
 - TCP/IP Profile
 - GLOBALCONFIG ECSALIMIT xxK or xxM
 - 0 means no limit and is the default

- **The GLOBALCONFIG limits will be enforced by TCP/IP itself**
 - Warning messages at 80%, 90%, and 98% utilization
 - EZZ4360I, EZZ4361I, and EZZ4362I for ECSA
 - EZZ4364I, EZZ4365I, and EZZ4366I for POOL
 - Storage relieved message issued when under 75% again
 - EZZ4363I for ECSA and EZZ4367I for POOL

Considerations for choosing POOL and ECSA limits

- **Monitor TCP/IP storage use during your acceptable peak periods**

- **Add a reasonable fudge factor to the observed maximum usage values**
 - The ECSALIMIT ensures that TCP/IP does not overuse the z/OS system's common storage.
 - It is intended to improve system reliability by limiting TCP/IP's common storage usage.
 - Accommodate for temporary application “hang” conditions, where TCP/IP must buffer large amounts of inbound or outbound data.

- **If you choose to limit POOL storage, make sure you use a value that is lower than or equal to what your installation exit (IEFUSI) enforces**
 - The benefit of specifying POOL limit is that you will receive warning messages before storage obtain calls start failing with not enough storage available to satisfy the requests

- **The values can be changed via OBEYFILE command processing**

Other considerations for TCP/IP storage controls - UDP

- **UDP applications:**
 - Control queue size of UDP messages
 - UDPCONFIG UDPQUEUELIMIT
 - When UDPQUEUELIMIT is ON then receive queues for UDP ports are limited to 2000 messages or 2800K
 - To further control UDP queue sizes, enable IDS policies with traffic regulation and UDP message queue size limits per UDP application
 - Does not apply to EE UDP traffic !!



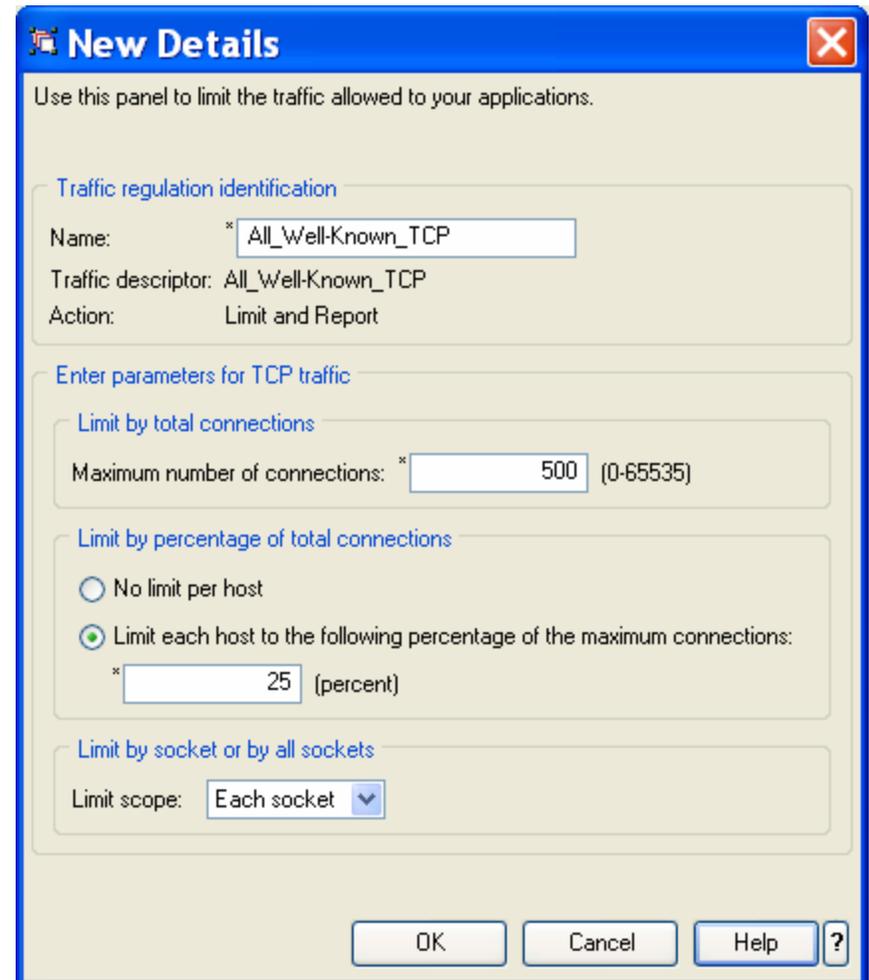
This will also benefit use of CSM storage

Other considerations for TCP/IP storage controls - TCP

▪ TCP applications

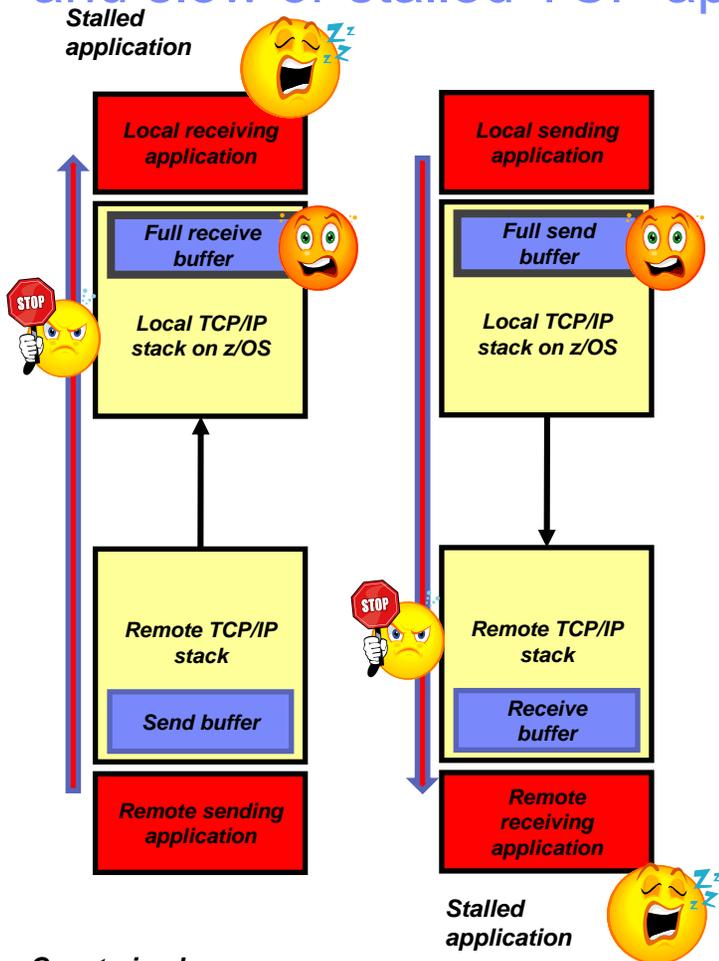
- Control number of concurrent connections with TCP Listeners to reduce impacts of flooding attacks:
 - IDS policy with traffic regulation and connection limits

- Monitor “misbehaving” applications that build up large send or receive buffers:
 - Monitor new z/OS V1R11 syslogd messages about applications with long send or receive queues
 - Watch for EZZ8661I to EZZ8664I messages
 - These messages will be logged through TRMD and syslogd even if you have no IDS policy enabled



This will also benefit use of CSM storage

z/OS V1R11 storage improvements for storage shortages situations and slow or stalled TCP applications



Constrained:

- Queue full & data at least 30 sec old
- Any data at least 60 sec old

Entry alert when 90% of constrained threshold reached, exit alert when down at 80%

- Data in a send buffer is page fixed – awaiting IO operations to be initiated
 - When application is not making progress or fixed storage is constrained
 - All new data added to TCP send queue is marked as page-able
 - When storage becomes constrained, all unsent data on send queues for all non-local TCP connections is marked as page-able
 - Before data is sent to remote stack it is changed back to fixed, as required by the DLC

- It was very difficult to identify which local applications caused excessive amounts of space to be used on the send or receive queues
 - Alerts issued to indicate TCP queue in constrained state
 - Indicate old data on send or receive queue
 - Identify connection (connection id, job name, addresses, ports)
 - Constrained state entry and exit indicated
 - Issued to syslogd using TRMD

Understanding z/OS Communications Server Storage Usage

Communications Storage Manager (CSM) storage use

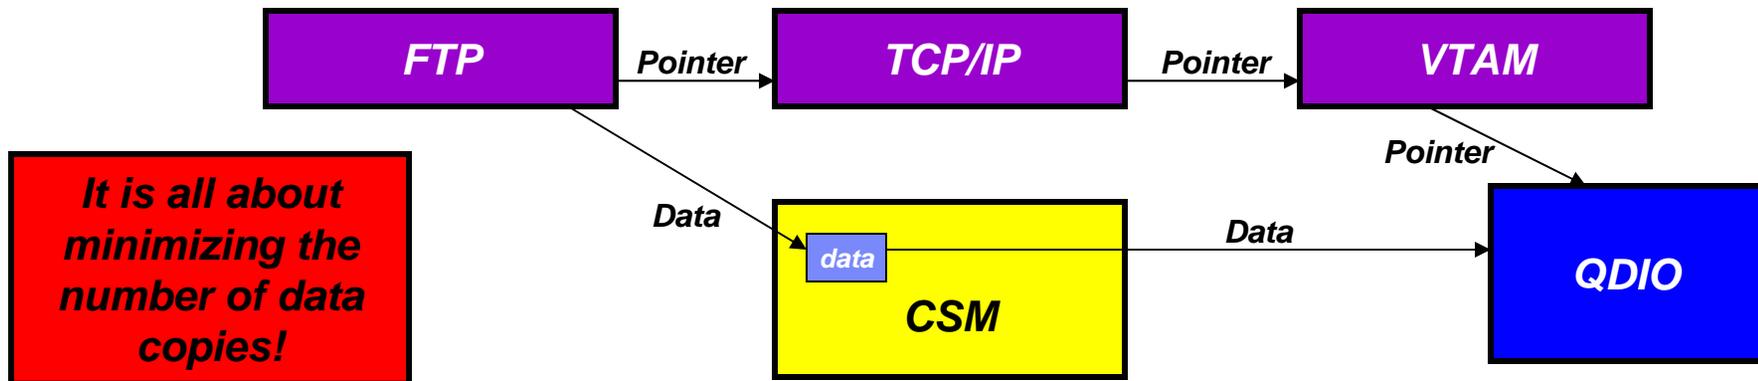


Communications Storage Manager (CSM)

- **The communications storage manager (CSM) is a VTAM component that allows authorized host applications to share data with VTAM, TCP/IP and other CSM users without the need to physically copy the data.**

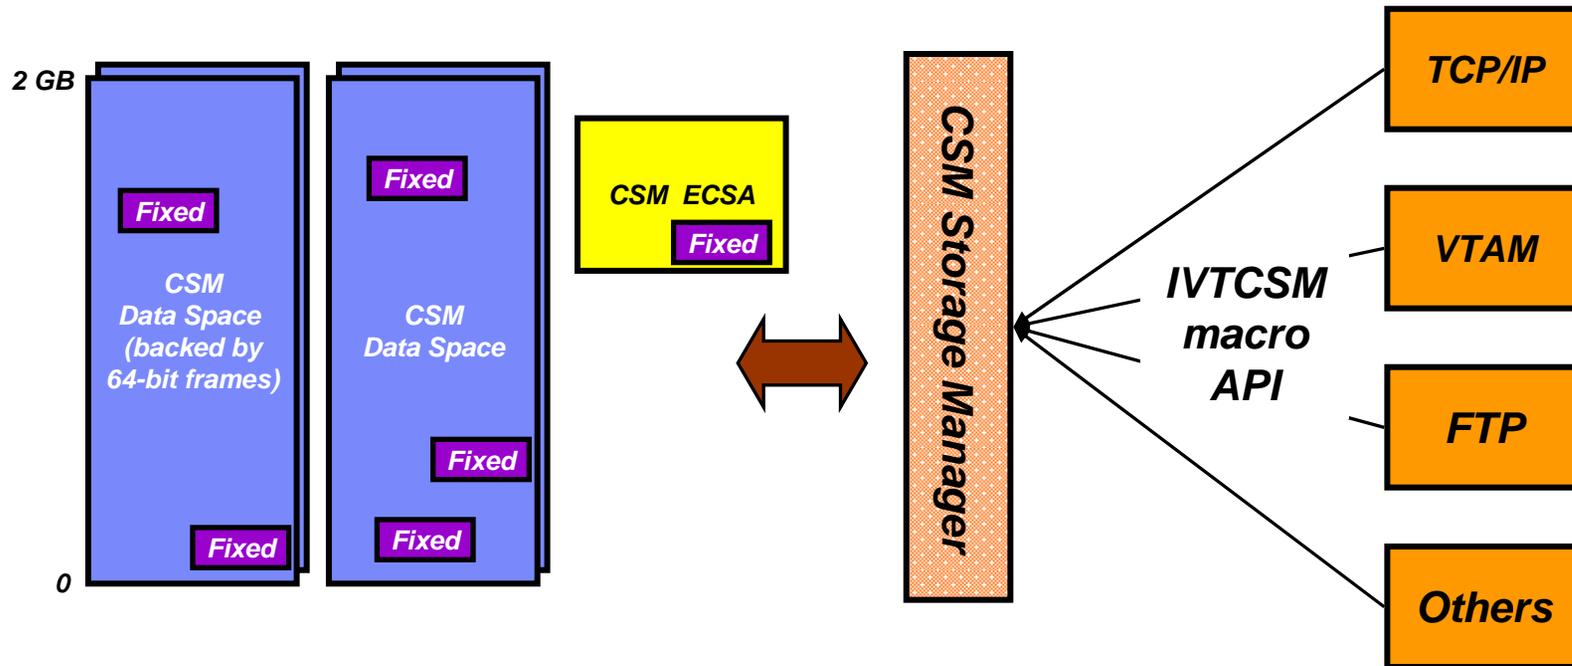
- **CSM includes a public application programming interface (API) that provides a way to:**
 - Obtain and return CSM buffers
 - Change ownership of buffers
 - Copy buffers
 - Manage CSM buffers

- **The storage key for CSM buffers is key 6**
 - Most CSM storage is obtained with the “System” as owner

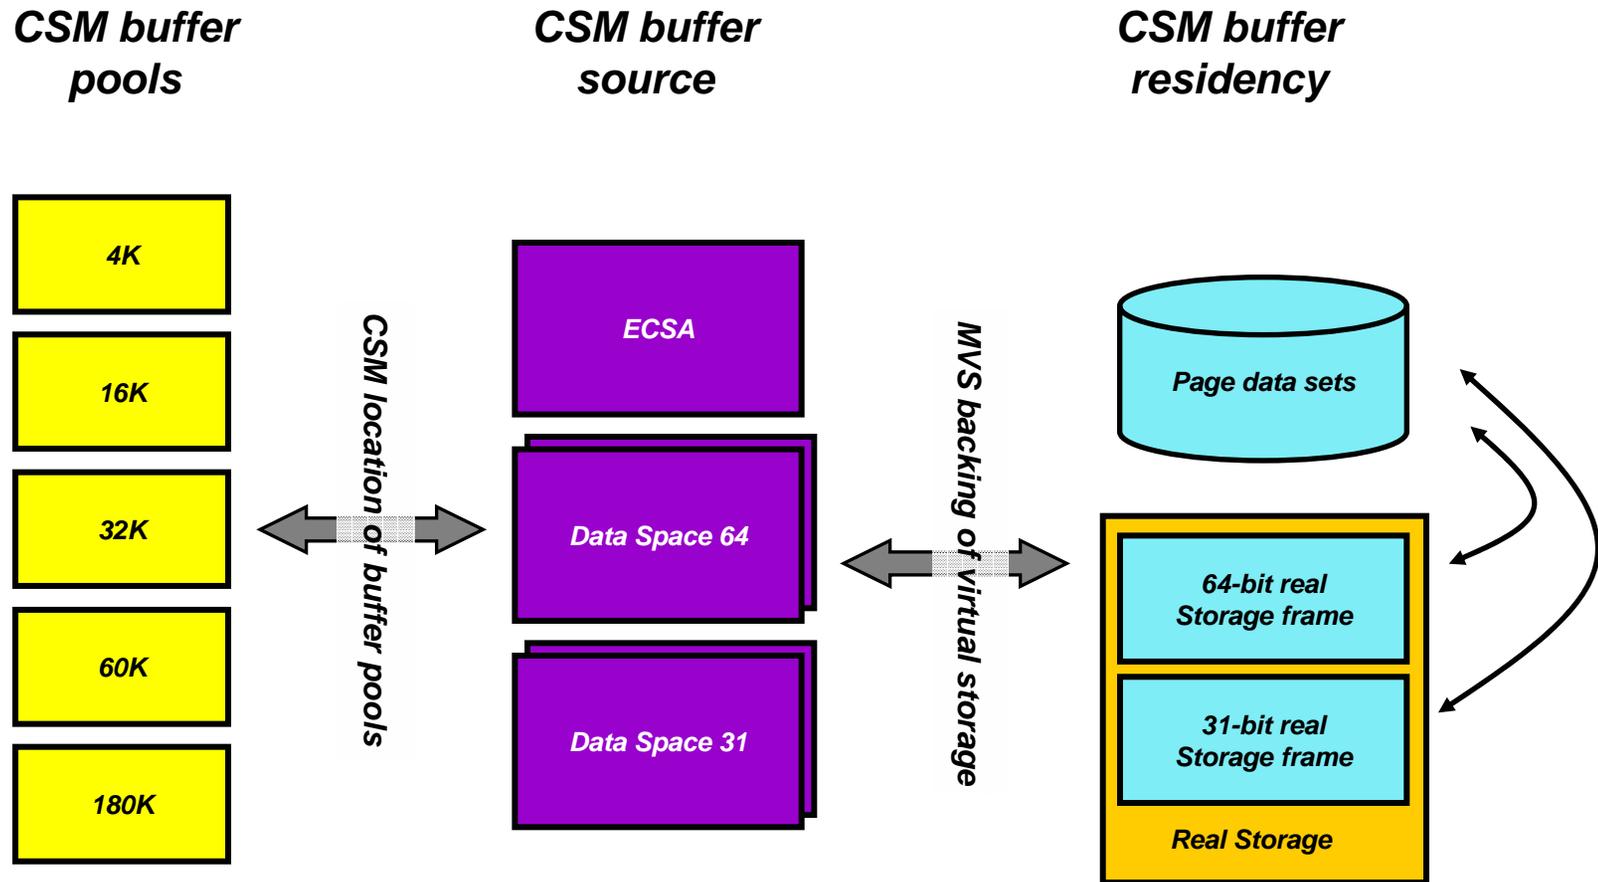


CSM virtual storage overview

- Users of CSM storage use the IVTCSM macro to interact with the CSM storage manager
- CSM storage is organized into buffer pools of fixed sizes (4K, 16K, 32K, 60K, and 180K)
- CSM storage is either ECSA or data space storage
 - CSM may create up to a total of five data spaces if needed
- CSM storage may be fixed or pageable
 - CSM storage may be page fixed in order to support an IO operation that requires fixed storage



CSM buffer pools, source, and residency



How to control CSM storage

- **IVTPRM00 member of SYS1.PARMLIB – Note: Only the 00 suffix is allowed – IVTPRM00.**
 - Read by VTAM during initialization
 - Can be modified and read via a MODIFY CSM without any parameters

- **You can control:**
 - The maximum amount of CSM storage that can be fixed at any point in time
 - Includes both ECSA and data space fixed storage
 - The maximum amount of ECSA storage, CSM can use at any point in time
 - How each of the five buffer pools is to be managed per buffer source
 - The default initial number of buffers to create the first time a user of CSM services request a pool to be created of a given size in one of the three CSM buffer sources
 - If initial number of buffers is set in IVTPRM00, it will always be used – otherwise the initial number of buffers depend on what the first requester asks for
 - The minimum number of free buffers in the pool before expansion must be done
 - The number of buffers the pool is to be expanded by when expansion occurs

- **Each of the data spaces CSM uses is 2 GB in size and no further controls are available for those**

Max fixed	100M
Max ECSA	100M

A health check suggests these two should be set to 120M as an initial value.

Buffer size	4K	16K	32K	60K	180K
Initial	64	16	16	16	2
Minimum free	8	4	2	2	1
Expansion units	16	8	4	4	2



How to monitor overall CSM storage use

Comments

- A “D NET,CSM” command will provide a quick overview of how much storage has been allocated by CSM, and how much of it is in-use or free for use by a CSM user.

- Three main locations:
 - ECSA
 - Data Space 31
 - Data Space 64

- Sample IVTPRM00 for the system used for this D CSM command:
 - FIXED MAX(120M)
 - ECSA MAX(120M)

IVT5530I	BUFFER	BUFFER						
IVT5531I	SIZE	SOURCE		INUSE	FREE	TOTAL	#buf	Exp?
IVT5532I	-----							
IVT5533I	4K	ECSA		200K	312K	512K	128	yes
IVT5533I	16K	ECSA		0M	256K	256K	16	no
IVT5533I	32K	ECSA		0M	512K	512K	16	no
IVT5533I	60K	ECSA		0M	0M	0M	0	
IVT5533I	180K	ECSA		0M	360K	360K	2	no
IVT5535I	TOTAL	ECSA		200K	1440K	1640K		
IVT5532I	-----							
IVT5533I	4K	DATA SPACE	31	0M	256K	256K	64	no
IVT5533I	16K	DATA SPACE	31	0M	0M	0M	0	
IVT5533I	32K	DATA SPACE	31	0M	0M	0M	0	
IVT5533I	60K	DATA SPACE	31	0M	0M	0M	0	
IVT5533I	180K	DATA SPACE	31	0M	0M	0M	0	
IVT5535I	TOTAL	DATA SPACE	31	0M	256K	256K		
IVT5532I	-----							
IVT5533I	4K	DATA SPACE	64	4372K	236K	4608K	1152	yes
IVT5533I	16K	DATA SPACE	64	0M	256K	256K	16	no
IVT5533I	32K	DATA SPACE	64	64K	448K	512K	16	no
IVT5533I	60K	DATA SPACE	64	0M	0M	0M	0	
IVT5533I	180K	DATA SPACE	64	0M	360K	360K	2	no
IVT5535I	TOTAL	DATA SPACE	64	4436K	1300K	5736K		
IVT5532I	-----							
IVT5535I	TOTAL	DATA SPACE		4436K	1556K	5992K		
IVT5532I	-----							
IVT5536I	TOTAL	ALL SOURCES		4636K	2996K	7632K		
IVT5538I	FIXED	MAXIMUM =	120M	FIXED	CURRENT =	7161K		
IVT5541I	FIXED	MAXIMUM USED =		7161K	SINCE LAST DISPLAY	CSM		
IVT5594I	FIXED	MAXIMUM USED =		7161K	SINCE IPL			
IVT5539I	ECSA	MAXIMUM =	120M	ECSA	CURRENT =	1955K		
IVT5541I	ECSA	MAXIMUM USED =		1955K	SINCE LAST DISPLAY	CSM		
IVT5594I	ECSA	MAXIMUM USED =		1955K	SINCE IPL			
IVT5559I	CSM DATA SPACE	1 NAME:	CSM64001					
IVT5559I	CSM DATA SPACE	2 NAME:	CSM31002					

CSM buffer pool usage for user of CSM storage

- To see how much CSM storage each of the CSM 'users' currently are using, issue a D NET,CSM,OWNERID=ALL command
 - This command can also be issued with a specific ownerID
- In this example, only VTAM (VTAMCS) and a TCP/IP stack (TCPCS) are using CSM storage

```

D NET,CSM,OWNERID=ALL command
IVT5508I DISPLAY ACCEPTED
IVT5549I PROCESSING DISPLAY CSM COMMAND - OWNERID SPECIFIED
IVT5530I BUFFER BUFFER
IVT5551I SIZE SOURCE STORAGE ALLOCATED TO OWNER
IVT5532I -----
IVT5553I 4K ECSA 48K
IVT5553I 32K ECSA 64K
IVT5554I TOTAL ECSA 112K
IVT5532I -----
IVT5553I 4K DATA SPACE 64 20K
IVT5554I TOTAL DATA SPACE 64 20K
IVT5532I -----
IVT5554I TOTAL DATA SPACE 20K
IVT5532I -----
IVT5556I TOTAL FOR OWNERID 132K
IVT5557I OWNERID: ASID = 002A JOBNAME = VTAMCS
IVT5532I -----
IVT5530I BUFFER BUFFER
IVT5551I SIZE SOURCE STORAGE ALLOCATED TO OWNER
IVT5532I -----
IVT5553I 4K ECSA 128K
IVT5553I 32K ECSA 64K
IVT5554I TOTAL ECSA 192K
IVT5532I -----
IVT5553I 4K DATA SPACE 64 4324K
IVT5553I 16K DATA SPACE 64 16K
IVT5554I TOTAL DATA SPACE 64 4340K
IVT5532I -----
IVT5554I TOTAL DATA SPACE 4340K
IVT5532I -----
IVT5556I TOTAL FOR OWNERID 4532K
IVT5557I OWNERID: ASID = 004D JOBNAME = TCPCS
IVT5599I END

```



Detailed CSM buffer pool usage overview

- Per each of the 15 CSM buffer pools, a detailed usage display can be used
 - D NET,CSMUSE,POOL=pool-name

F NET,CSM,

- MONITOR=YES
- MONITOR=NO
- MONITOR=DYNAMIC
 - CSM dynamically activates CSM buffer monitoring when CSM storage usage approaches the constrained level

- The pool names are:

- 4KECSA
- 16KECSA
- 32KECSA
- 60KECSA
- 180KECSA
- 4KDS
- 16KDS
- 32KDS
- 60KDS
- 180KDS
- 4KDS64
- 16KDS64
- 32KDS64
- 60KDS64
- 180KDS64

D NET,CSM,MONITOR

```

D NET,CSMUSE,POOL=4KECSA
IVT5508I DISPLAY ACCEPTED
IVT5574I PROCESSING DISPLAY CSMUSE COMMAND - POOL SPECIFIED 654
IVT5584I USAGE DETAILS - 4KECSA POOL - POOL TOTAL = 164K
IVT5532I -----
IVT5576I AMOUNT MONITOR ID OWNERID JOBNAME
IVT5532I -----
IVT5577I 80K 21 003A TCPCS
IVT5579I BUFFER USE FOR 21 : USECNT USERDATA MONITOR HISTORY
IVT5580I 20 F2C5F0F2 00000021
IVT5532I -----
IVT5577I 52K 21 002A VTAMCS
IVT5579I BUFFER USE FOR 21 : USECNT USERDATA MONITOR HISTORY
IVT5580I 5 F0C4F0F0 00000021
IVT5580I 4 F2C5F0F1 00000021
IVT5580I 4 F2C5F0F0 00000021
IVT5585I DETAIL TOTAL FOR 4KECSA POOL = 132K
IVT5599I END
    
```

Monitor IDs are documented in z/OS Communications Server: IP and SNA Codes Chapter 4

CSM monitor IDs – lots of good hints about where CSM storage is being used

```

IVT5577I          80K      21          003A      TCPCS
IVT5579I BUFFER USE FOR 21 :   USECNT  USERDATA  MONITOR HISTORY
IVT5580I          20     F2C5F0F2          00000021
IVT5532I -----
IVT5577I          52K      21          002A      VTAMCS
IVT5579I BUFFER USE FOR 21 :   USECNT  USERDATA  MONITOR HISTORY
IVT5580I          5      F0C4F0F0          00000021
IVT5580I          4      F2C5F0F1          00000021
IVT5580I          4      F2C5F0F0          00000021
    
```

Range	Description
X'00' – X'1F'	CSM Monitor IDs
X'20' – X'2F'	DLC Monitor IDs
X'30' – X'8F'	VTAM Unique Monitor IDs
X'90' – X'97'	TCP/IP IF Layer Monitor IDs
X'98' – X'9F'	TCP/IP IP Layer Monitor IDs
X'A0' – X'AF'	TCP/IP Transport Layer Monitor IDs
X'B0' – X'FF'	TCP/IP Misc Monitor IDs

For monitor ID 21 (DLC Read Operation):

•VTAM

- VTAM owns 52K of 4KECSA CSM buffers used for read operations at the DLC layer
- The read operations are associated with devices 0D00 (CTC), 2E01 (QDIO write), and 2E00 (QDIO read)

•TCP/IP

- TCP/IP owns 80K of 4KECSA CSM buffers used for read operations at the DLC layer
- TCP/IP's read operation is associated with device 2E02 (QDIO data path)

- For monitor ID X'21', the user field will provide the EBCDIC sub-channel read device unit address associated with this device. Some devices can configure how much read storage is used (for example, QDIO devices).
- Other monitor IDs use the user data field for other types of information.



Understanding z/OS Communications Server Storage Usage

VTAM storage use



VTAM storage use overview

- **VTAM user region**
 - Primarily extended user region
 - Some in the user region below the line

- **Common Service Area (CSA) – referred to by VTAM as CSA24**
 - VTAM uses CSA below the line for a few modules and control blocks

- **Extended Common Service Area (ECSA)**
 - The part of ECSA that is obtained and managed by VTAM
 - Not part of CSM storage
 - VTAM has a few modules loaded into ECSA
 - Otherwise ECSA storage is primarily used for common control blocks
 - ECSA storage is obtained in either key 0 or key 6
 - Most of VTAM’s ECSA storage is obtained in key 6
 - ECSA storage is obtained with VTAM or the “System” as owner
 - Workload related ECSA storage is generally obtained in key 6 with VTAM as owner

How do you limit VTAM's storage use?

- **VTAM's private storage (VTAM's POOL storage)**
 - Via the REGION keyword in VTAM's start up JCL
 - Private and extended private region size
 - Can be overridden by installation exits - primarily by IEFUSI

- **CSA+ECSA storage**
 - Via the CSALIMIT VTAM start option (ATCSTRxx), which covers both CSA and ECSA
 - CSALIMIT=0 - Default, limit set to 90% of system limit (IEASYSnn CSA+ECSA limit)
 - CSALIMIT=nn - Warning when reached, but will continue until only 25% of system limit is available
 - CSALIMIT=(nn,F) - Will not go above the limit specified
 - Can be changed dynamically via a MODIFY VTAMOPTS or MODIFY CSALIMIT command

- **CSA storage below the line**
 - Via the CSA24 start option, which specifically covers 24-bit addressable CSA storage (CSA24 is a subset of the CSALIMIT value)
 - CSA24=a
 - CSA24=0 is the default and means no limit
 - Can be changed dynamically via a MODIFY VTAMOPTS or MODIFY CSALIMIT command



How do you monitor VTAM's storage use?

- **Via a D NET,BFRUSE command**

- Towards the end of the output from that command is a breakdown of VTAM's current storage use (limit, current, high-water mark)
 - CSA
 - CSA24
 - Private

In this example, this is the default limit: 90% of IEASYSnn CSA+ECSA

IST449I CSALIMIT = 240012K, CURRENT = 2562K, MAXIMUM = 2578K	CSA+ECSA
IST790I MAXIMUM CSA USED = 2578K	
IST1667I SYSTEM CSA LIMIT = 266680K	
IST1831I 91% OF SYSTEM CSA STORAGE REMAINING = 243320K	
IST449I CSA24 LIMIT = NOLIMIT, CURRENT = 60K, MAXIMUM = 63K	CSA
IST790I MAXIMUM CSA24 USED = 63K	
IST595I IRNLIMIT = NOLIMIT, CURRENT = 0K, MAXIMUM = 0K	
IST981I VTAM PRIVATE: CURRENT = 1030K, MAXIMUM USED = 1183K	Private
IST924I -----	
IST1565I CSA MODULES = 1756K	Modules in CSA24 and ECSA
IST1565I CSA24 MODULES = 40K	
IST1565I PRIVATE MODULES = 7497K	

Understanding z/OS Communications Server Storage Usage

Use of RMF to monitor z/OS CS storage use





Communications Server storage is in Key 6

- The RMF Monitor II Virtual Storage Activity report reports the amount of Key 6 storage in CSA and ECSA
 - The Communications Server uses key 6 storage

VIRTUAL STORAGE ACTIVITY											
z/OS V1R11			SYSTEM ID 3090			DATE 06/16/2009			INTERVAL 09.44.920		
			RPT VERSION V1R11 RMF			TIME 11.05.15			CYCLE 1.000 SECONDS		
COMMON STORAGE SUMMARY											
NUMBER OF SAMPLES 50											
STATIC STORAGE MAP			ALLOCATED CSA/SQA								
AREA	ADDRESS	SIZE	----- BELOW 16M -----			----- EXTENDED (ABOVE 16M) -----					
			MIN	MAX	AVG	MIN	MAX	AVG	MIN	MAX	AVG
EPVT	16E00000	1682M									
ECSA	6D71000	257M	SQA	568K 11.05.16	568K 11.05.16	568K	14.5M 11.13.15	14.7M 11.08.55	14.6M		
EMLPA	6D70000	4K	CSA	372K 11.05.16	372K 11.05.16	372K	23.0M 11.05.16	23.2M 11.06.15	23.1M		
EFLPA	6D6D000	12K									
EPLPA	2CD0000	64.6M	ALLOCATED CSA BY KEY								
ESQA	1A0C000	18.8M	0	168K 11.05.16	168K 11.05.16	168K	12.5M 11.05.16	12.5M 11.05.16	12.5M		
ENUC	1000000	10.0M	1	84K 11.05.16	84K 11.05.16	84K	832K 11.05.16	832K 11.05.16	832K		
----- 16 MEG BOUNDARY -----			2	36K 11.05.16	36K 11.05.16	36K	16K 11.05.16	16K 11.05.16	16K		
NUCLEUS	FD6000	168K	3	0K 11.05.16	0K	0K	0K 11.05.16	0K	0K		
SQA	E5A000	1520K	4	0K 11.05.16	0K	0K	4K 11.05.16	4K 11.05.16	4K		
PLPA	BF3000	2460K	5	4K 11.05.16	4K 11.05.16	4K	668K 11.05.16	668K 11.05.16	668K		
FLPA	BF2000	4K	6	80K 11.05.16	80K 11.05.16	80K	8964K 11.05.16	9160K 11.06.15	9048K		
MLPA	BE5000	52K	7	0K 11.05.16	0K	0K	228K 11.05.16	228K 11.05.16	228K		
CSA	800000	3988K	8-F	0K 11.05.16	0K	0K	0K 11.05.16	0K	0K		
PRIVATE	2000	8184K	SQA EXPANSION INTO CSA								
PSA	0	8K		0K 11.05.16	0K	0K	0K 11.05.16	0K	0K		



RMF Monitor III Common Storage report

- Not too useful since most of the common storage that is obtained by the Communications Server components is obtained with the "SYSTEM" as owner (%MVS as jobname).

```

RMF V1R11 Common Storage                               Line 1 of 61
Command ===>                                           Scroll ===> CSR

Samples: 90      System: 3090  Date: 07/06/09  Time: 13.05.00  Range: 90  Sec

          ----- Percent -----      ----- Amount -----
System Information      CSA  ECSA  SQA  ESQA      CSA  ECSA  SQA  ESQA
IPL Definitions                3980K 257M 1520K 19M
Peak Allocation Values        16    9   80   69      639K  23M 1211K 13M
Average CSA to SQA Conversion  0    0                0    0
Average Use Summary           8    9   35   69      324K  23M  531K 13M
Available at End of Range     92   91   65   31      3656K 234M  989K 5903K

Unalloc Common Area: 4564K

          Service      ELAP  -- Percent Used -  ----- Amount Used -----
Jobname  Act C Class    ASID  Time  CSA  ECSA  SQA  ESQA      CSA  ECSA  SQA  ESQA
%MVS                3    6  31  38      127K 15M  478K 7342K
%REMAIN            0    0   0   0        424 105K  128  6240
VTAMCS           S SYSSTC  0041  3.0H  1    1   0   0      22032 2485K  0   696
TCPCS           S SYSSTC  0058  2.9H  0    0   0   0        136 47272  0   991
TN3270A         S SYSSTC  0061  2.9H  0    0   0   0        136 24432  0   632
ABCRESO         S SYSSTC  0040  2.9H  0    0   0   0         0   312  0  1280
OMPROUTE        S SYSSTC  0064  2.9H  0    0   0   0         0   952  0  1200
    
```

VTAM's modules in ECSA, CSM's, and TCP/IP's ECSA storage is accounted for under %MVS

VTAM's workload-related ECSA storage is accounted for under VTAM itself

TCP/IP components have minor ECSA amounts accounted under their own address spaces



RMF Monitor III Storage Memory Objects (64-bit storage)

- Both the resolver address space and TCP/IP use 64-bit virtual storage since z/OS V1R11
 - Resolver uses 64-bit user region storage
 - TCP/IP uses 64-bit common storage

```

RMF V1R11  Storage Memory Objects                Line 1 of 8
Command ===>                                     Scroll ===> CSR

Samples: 90      System: 3090  Date: 07/06/09  Time: 13.05.00  Range: 90  Sec

----- System Summary -----
-- Memory Objects --      ----- Frames -----      --- Area Used % ---
Common Shared Large      Common Fixed Shared  1 MB      Common Shared  1 MB
      6      0              800      0      0              0.0      0.0

-----
Jobname      Service
C Class      ASID      ----- Memory Objects -----  Frames  ----- Bytes -----
Total  Comm  Shr Large  1 MB  Total  Comm  Shr
SMSPDSE  S SYSTEM  0008      12      0      0              76.0M      0      0
TRACE    S SYSTEM  0004      8      0      0              8192K      0      0
GRS      S SYSTEM  0007      4      0      0              140G      0      0
ABCRESO  S SYSSTC  0040      4      0      0              4096K      0      0
JESEAUX  S SYSSTC  0027      3      3      0              3072K 3072K      0
ZFS      S SYSSTC  0049      2      0      0              22.0M      0      0
*MASTER* S SYSTEM  0001      1      1      0              1024K 1024K      0
TCPCS    S SYSSTC  0058      1      1      0              1024K 1024K      0
    
```

Resolver uses 64-bit private for the DNS cache

TCP/IP uses 64-bit common memory for its sockets control blocks

Understanding z/OS Communications Server Storage Usage

Summary



Summary

- **Limiting CS common storage use should be done to protect other subsystem's access to common storage**
 - Severe network spikes (normal or the results of an attack) can require large amounts of common storage for a period of time
 - To avoid CS monopolizing all common storage on a system, limits should be enforced
 - There are no general values that will work for everyone
 - Monitor your system for a period of normal and acceptable peak workloads
 - Communications Server components have mechanisms built in to deal with storage constraints where the limits are approached
 - The intended objective of those mechanisms is to keep the system running in a controllable state, reducing the amount of network traffic to/from the system for a brief period of time
 - If other subsystems are unsuccessful in obtaining any more common storage, they may or may not be able to “survive”

Summary

- **To get a complete picture of Communications Server common storage use, you need to add up the following items:**
 - ECSA use from the D TCPIP,,STOR command for all stacks
 - For z/OS V1R12: Also the CSA modules from D TCPIP,,STOR command
 - ECSA use from the D TCPIP,TN-Server,STOR command for all TN3270 server address spaces
 - CSM total ECSA use from the D NET,CSM command
 - VTAM total ECSA use from the D NET,BFRUSE command

- **For all items, you can collect**
 - Current use
 - High-water mark
 - Limit (if one has been configured)



And as usual,
a little REXX
program can
do wonders
to bringing
order into all
that data ..

```

-----
System: MVS098 Sysplex: LOCAL Date: 18 Mar 2010 Time: 16:48:02
CPU: 002097-E56 LPAR/VM-name: MVS098
z/OS 01.12.00 5694-A01 CPU-id: FF01E15E2097
-----
ASName      Stor-type      Current HighWater      Limit      In-use
-----
System      CSA              338432    654152      3968K      8.33%
System      ECSA             24067K    24916K      262720K    9.16%
System      SQA              540088    1216K       1520K      34.70%
System      ESQA             13673K    14360K      20240K     67.56%
TCPCS      ECSA              2752K     3243K        0          N/A
TCPCS      POOL              8892K     8925K        0          N/A
TCPCS      64-bit Commo    1024K     1024K        0          N/A
TCPCS      ECSA Modules    7428K     7428K        0          N/A
VTAMCS     ECSA              2838K     2844K      240019K    1.18%
VTAMCS     POOL              1034K     1172K        0          N/A
VTAMCS     ECSA Modules    1776K
VTAMCS     CSA Modules     40960
VTAMCS     PRIV Modules    7595K
TN3270A    ECSA             119808    119808        0          N/A
TN3270A    POOL              7826K     7827K        0          N/A
TN3270A    64-bit Commo      0          0            0          N/A
TN3270A    CTRACE           262372K   262372K     262372K   100.00%
CSM        ECSA              1554K     0            122880K    1.26%
CSM        Dataspace 31      0          262144       2048M     0.00%
CSM        Dataspace 64     8660K     9960K        2048M     0.41%
CSM        Fixed             11001K    0            122880K    8.95%
-----
(E)CSA     Workload          7261K
(E)CSA     Modules           9244K
    
```



For more information



URL	Content
http://www.twitter.com/IBM_Commserver 	IBM Communications Server Twitter Feed
http://www.facebook.com/IBMCommserver 	IBM Communications Server Facebook Fan Page
http://www.ibm.com/systems/z/	IBM System z in general
http://www.ibm.com/systems/z/hardware/networking/	IBM Mainframe System z networking
http://www.ibm.com/software/network/commserver/	IBM Software Communications Server products
http://www.ibm.com/software/network/commserver/zos/	IBM z/OS Communications Server
http://www.ibm.com/software/network/commserver/z_lin/	IBM Communications Server for Linux on System z
http://www.ibm.com/software/network/ccl/	IBM Communication Controller for Linux on System z
http://www.ibm.com/software/network/commserver/library/	IBM Communications Server library
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/zos/support/	IBM z/OS Communications Server technical Support – including TechNotes from service
http://www.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs	Technical support documentation from Washington Systems Center (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)
http://www.ibm.com/systems/z/os/zos/bkserv/	IBM z/OS Internet library – PDF files of all z/OS manuals including Communications Server

For pleasant reading