

CICS TS V4.2 Scalability

Chris Baker
CICS Strategy and Planning
IBM United Kingdom

Wednesday 14th March 2012
Session 10287





© IBM Corporation 2012. All Rights Reserved.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: ibm.com/legal/copytrade.shtmlAIX, CICS, CICSplex, DataPower, DB2, DB2 Universal Database, i5/OS, IBM, the IBM logo, IMS/ESA, Power Systems, Lotus, OMEGAMON, OS/390, Parallel Sysplex, pureXML, Rational, Redbooks, Sametime, SMART SOA, System z , Tivoli, WebSphere, and z/OS.

A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml.

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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce

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

Microsoft and Windows are trademarks of Microsoft Corporation 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

Intel 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.

Agenda

- Quick summary of previous CICS release scalability items
- CICS TS V4.2 delivers new and enhanced capabilities in the area of scalability ...
 - New **concurrency** option for greater OTE exploitation
 - Threadsafe Mirrors for Function Shipping over IPIC connections
 - Threadsafe CICS-DBCTL interface when connected to IMS V12
 - More threadsafe EXEC CICS API and SPI commands
 - New 64-bit infrastructure and exploitation
 - More VSAM LSR Pools
- Summary

Notes

The newest CICS Transaction Server V4.2 release is focused on helping our customers to deliver a smarter transaction processing experience. New and enhanced capabilities have been made in 5 key areas. This session will cover one of those areas - Scalability - in detail.

The Scalability theme of CICS Transaction Server for z/OS V4.2 is concerned with being able to do more work, more quickly in a single CICS region.

We will start with a summary of the scalability capabilities in previous releases before moving on to cover the new and enhanced Scalability capabilities delivered in CICS Transaction Server for z/OS V4.2 enhancements which centre on Open Transaction Environment (OTE) enhancements but also cover 64-bit exploitation and an enhancement to File Control to allow use of more VSAM LSR pools.

Scalability Items in Previous Releases

- CICS TS V3.2
 - CICS-MQ exploitation of OTE
 - Threadsafe File Control (Post GA)
 - 64-bit container support
 - Extended ESDS support
 - Shared Data Tables > 2GB using multiple data spaces
 - XCF group limit changes
- CICS TS V4.1
 - Exploitation of z/OS XLMSS parser
 - IPv6 support
 - 64-bit positioning – extended MVS linkage support
 - Java positioning – JVMSERVER for dynamic scripting support

Notes

Previous releases of CICS TS had differently named themes to group together enhancements but if we revisit CICS TS V3.2 and CICS TS V4.1 there are indeed a number of enhancements that can be reclassified under a scalability theme. This demonstrates an ongoing roadmap of enhancing CICS to allow more work to run, more quickly, in a single CICS region.

OTE Settings in CICS TS V4.1 and below

- Up until now you can define a program as ...
 - **concurrency(quasirent)** meaning it must run on QR TCB
 - **concurrency(threadsafe)** meaning it can run on an open TCB or QR TCB whichever is in use at the time, it doesn't care which one is used
- There is no **concurrency** attribute to tell CICS that the application must run on an open TCB from the very start
- Instead, the only way you could do this was to define it with attribute **api(openapi)**
 - This has disadvantages (see next slide)

Notes

A CONCURRENCY(QUASIRENT) program always runs on the QR TCB.

A CONCURRENCY(THREADSAFE) program is a program that has been coded to threadsafe standards and contains threadsafe logic. It is capable of running on either the QR TCB or an open TCB. It starts off running on the QR TCB. If processing, such as a DB2 request, causes a switch to an open TCB, then on return from the resource manager to the program the program continues on the open TCB.

Prior to CICS TS V4.2, there is no concurrency option that specifies that the program should run on an open TCB from the start. Instead, this can be achieved by defining it as API(OPENAPI) but this has disadvantages which are explained on the next slide.

OTE Settings in CICS TS V4.1 and below

- Defining a program with attribute **api(openapi)**
 - Tells CICS the application will use non-CICS APIs
 - Has the side effect of running the application on an open TCB
 - Has the drawback of forcing CICS to match execution key and TCB key so that non-CICS APIs will work!
 - This is not required for CICS APIs
 - Means CICS has to use L9 open TCBs for user key programs
 - Not good for DB2, MQ, or CICS Sockets applications
 - *Accessing a Resource Manager always uses an L8 open TCB*

Notes

Defining a program with attribute API(OPENAPI) has the side effect of causing the program to run from the start on an open TCB. However, primarily it is telling CICS that the program will issue non CICS supported API commands, i.e. something other than EXEC CICS, EXEC SQL, WebSphere MQ commands, etc, for example, an MVS command.

Unlike CICS commands, an MVS command ,for example, will only work correctly when the key of the TCB matches the execution key. So when running EXECKEY(USER) which is key 9, an OPENAPI program is given a key 9 TCB called an L9 TCB.

Contrast this with a program that only issues CICS supported commands. In this case CICS does not reference the key of the TCB, so a program can be running in User key or CICS key and can run under a key 8 TCB such as the QR or an L8 TCB.

Users who define programs as OPENAPI get the advantage of starting on an open TCB but suffer the disadvantage that an L9 will be used (assuming storage protection is active). This is a disadvantage because DB2 requires an L8 TCB and so we switch from L9 to L8 and back again for every DB2 request. This is why OPENAPI should not be used for DB2 programs or MQ programs.

OTE Settings in CICS TS V4.2

- In CICS TS V4.2 ...
 - We have separated out whether an application must run on an open TCB from what type of APIs it uses
- **CONCURRENCY(REQUIRED)**
 - Application must be coded to threadsafe standards
 - States that the application **MUST** run on an open TCB
 - Application starts on an open TCB
 - If a TCB switch to the QR TCB is made for an EXEC CICS command, a switch back to the open TCB is made when returning to the application
- Existing API keyword defines what APIs are used
 - This defines what type of CICS TCB is used

Notes

CICS TS V4.2 provides a new CONCURRENCY(REQUIRED) setting which specifies that the program requires to run on an open TCB. It will run on an open TCB from the start, and if CICS has to switch to the QR TCB to process a non-threadsafe CICS command, it will return to the open TCB when it returns to the application program. Now the user can define that the program must start on an open TCB, independently of defining what APIs it uses.

The API parameter determines what type of open TCB is used.

OTE Settings in CICS TS V4.2 ...

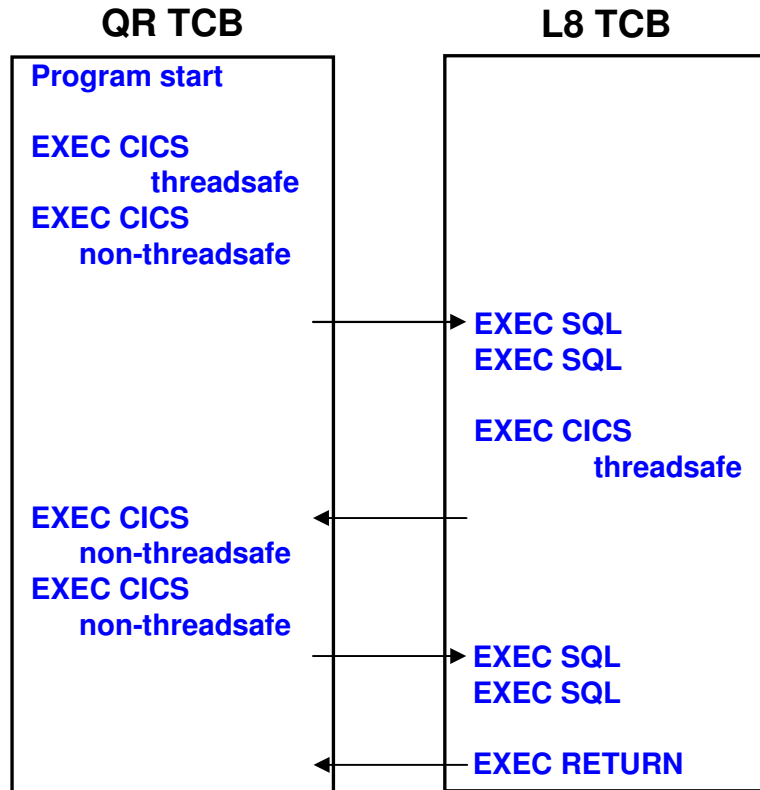
- **CONCURRENCY(REQUIRED) and API(CICSAPI)**
 - The application will run on an open TCB from the start
 - It only uses CICS supported APIs (including DB2, IMS, and MQ)
 - CICS will always use an L8 TCB in this instance irrespective of the execution key, as CICS APIs do not rely on the key of the TCB
 - Great for applications that are to use Resource Managers
 - Such as DB2 and MQ as the same L8 TCB is used!
- **CONCURRENCY(REQUIRED) and API(OPENAPI)**
 - The application will run on an open TCB from the start
 - As it will use non-CICS APIs ...
 - It will run on an L8 or an L9 TCB depending on the execution key
 - *This is the same as with CICS TS V4.1*
 - Only use **OPENAPI** when non-CICS supported APIs are to be used!

Notes

For CONCURRENCY(REQUIRED) programs the type of open TCB used depends on what APIs the program is to use:

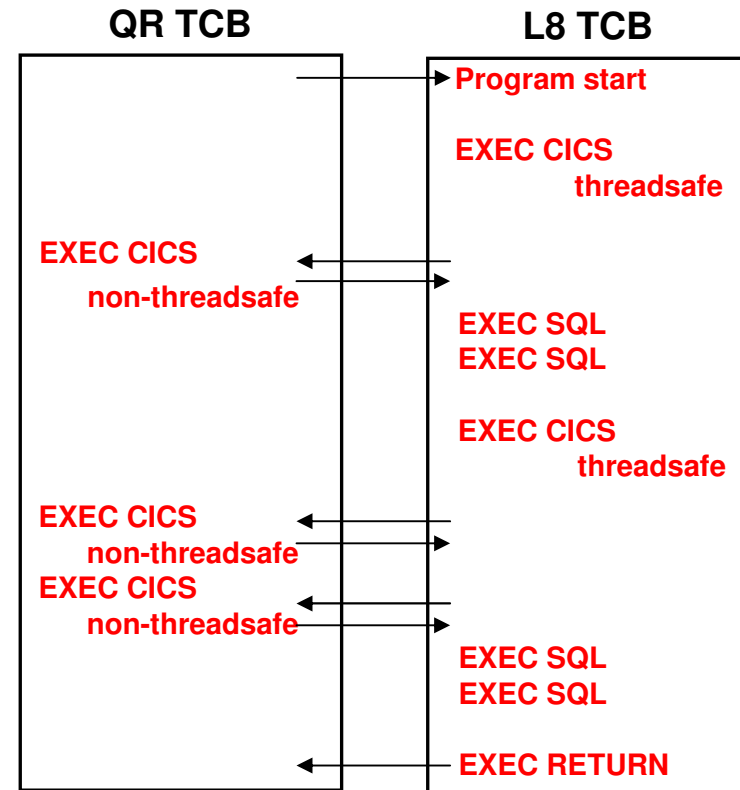
- If the program uses only CICS supported APIs (including access to external resource managers, such as DB2, IMS, and WebSphere MQ) then it should be defined with the program attribute API(CICSAPI). In this case CICS always uses an L8 open TCB, irrespective of the execution key of the program, because CICS commands do not rely on the key of the TCB.
- If the program is to use other non-CICS APIs then it must be defined with the program attribute API(OPENAPI). In this case CICS uses an L9 TCB or an L8 TCB depending on the execution key of the program. This is to allow the non-CICS APIs to operate correctly. This OPENAPI behaviour is the same as previous releases.

CONCURRENCY(THREADSAFE) versus CONCURRENCY(REQUIRED)



The program for this transaction is defined
CONCURRENCY(THREADSAFE)
 API (CICSAPI)

EXECCKEY(USER) or EXECCKEY(CICS)



The program for this transaction is defined
CONCURRENCY(REQUIRED)
 API (CICSAPI)

EXECCKEY(USER) or EXECCKEY(CICS)

Notes

Existing threadsafe CICS-DB2 applications, which have taken advantage of the performance gains of being able to run on the same TCB as the DB2 call by being defined as CONCURRENCY(THREADSAFE) API(CICSAPI), can be further enhanced by defining them as CONCURRENCY(REQUIRED) API(CICSAPI). This definition means that the programs can run on an L8 open TCB (irrespective of their execution key as CICS can for example run the program in key 9 on an key 8 TCB) from the start without waiting for the first DB2 call to move them on to the open TCB.

Achieving additional benefit depends on how many, if any, non-threadsafe CICS commands the application executes

Threadsafe Mirror

- DFHMIRS is now Threadsafe
 - Supplied definition now specifies CONCURRENCY(THREADSAFE)
- IPIC transformers are now threadsafe
 - Non IPIC code remains non threadsafe!
- Only requests Function Shipped over IPIC will run on open TCB
 - File Control (FC) and Temporary Storage (TS)
 - Distributed Program Link (DPL)
 - If the target application program is defined as threadsafe and the Mirror is already on an open TCB
- Review your DFHSIT specification for FCQRONLY if using IPIC
 - Specify FCQRONLY=NO as there is no longer any need to turn off threadsafety in the FOR

Notes

The CICS-supplied mirror program, DFHMIRS, which is used by all mirror transactions, is now defined as threadsafe. In addition, the IPIC transformers have been made threadsafe. For IPIC connections only, CICS runs the mirror program on an L8 open TCB whenever possible. For threadsafe applications that function ship commands to other CICS regions using IPIC, the resulting reduction in TCB switching improves the performance of the application compared to other intercommunication methods. To gain the performance improvement, you must specify the system initialization parameter FCQRONLY=NO in the file-owning region.

File Control requests that are Function Shipped using IPIC connectivity provide threadsafe file control with significant potential throughput improvements over LU6.2 in CICS regions with multiple processors available.

Temporary Storage requests that are Function Shipped using IPIC connectivity are threadsafe and no longer need to switch to the QR TCB before being Function Shipped.

Threadsafe Function Shipping over IPIC

- Remote FC and TS requests over IPIC are now threadsafe
 - No TCB switch to QR in the AOR
 - Transformers can run on open TCB in the AOR and ship the request
- File Control and Temporary Storage requests in the remote region can run on an open TCB
 - Mirror switches to an open TCB for 1st FC or TS request
 - FC request will run on an open TCB provided ...
 - FCQRONLY=NO is set
 - It's a local VSAM LSR file or VSAM RLS file
 - *i.e. existing File Control threadsafe support*
 - TS request will run on an open TCB
- MIRRORLIFE on the IPCONN definition controls lifetime of mirror
 - UOW or TASK allows the mirror to continue?

Notes

For remote File Control (FC) or Temporary Storage (TS) requests Function Shipped over IPIC connections, CICS will no longer force a switch to the QR TCB if it is running currently on an open TCB. The requests will be Function Shipped running on the open TCB.

In the FOR or QOR, the mirror decides when to switch to an open TCB. It does so for the first File Control or Temporary Storage request received over an IPIC connection. The idea is for long running mirrors to keep the mirror transaction running on an open TCB.

A new option MIRRORLIFE has been added to the IPCONN attributes for Function Shipped File Control (FC), Transient Data (TD), and Temporary Storage (TS) requests using an IPIC connection. MIRRORLIFE improves efficiency and provides performance benefits by specifying the lifetime of mirror tasks and the amount of time a session is held.

Threadsafe CICS-DBCTL

- CICS-DBCTL interface will use OTE when connect to IMS V12
 - At connect time ...
 - CICS and IMS determine if each other can support OTE
 - With IMS V10 and V11
 - CICS-DBCTL TRUE enabled as QUASIRENT
 - Toleration APAR PM31730 (IMS V10), PM31729 (IMS V11)
 - With IMS V12
 - CICS-DBCTL TRUE enabled as OPENAPI
 - Exploitation APARs PM31420 + PM47327, PM45414 are required

Notes

CICS provides a CICS IMS Database control (CICS-DBCTL) interface for IMS to satisfy DL/I requests that are issued by applications running in a CICS region. In CICS TS V4.2 the CICS-DBCTL interface has been defined as threadsafe and CICS can run the CICS-DBCTL task-related user exit (TRUE) on an L8 open task control block (open TCB).

The Open Transaction Environment (OTE) is supported from IMS Version 12 with PTFs for APARs PM31420 + PM47327, and PM45414.

applied. IMS indicates to CICS during the connection process that the OTE is supported and consequently CICS defines the CICS-DBCTL TRUE as an open API TRUE.

An open API TRUE is run on an L8 open TCB, which is dedicated for use by the calling CICS task. Running an application on an open TCB improves throughput and performance by reducing the use of the QR TCB. Threadsafe CICS applications that run on an L8 open TCB and use threadsafe CICS-DBCTL commands now avoid up to four TCB switches for each call to IMS. For more information about CICS IMS applications and the OTE, see Enabling CICS IMS applications to use the open transaction environment (OTE) through threadsafe programming.

If your IMS version does not support the OTE, CICS runs the CICS-DBCTL TRUE on the QR TCB.

Threadsafe CICS-DBCTL ...

- TCB switching is controlled by the IMS Database Resource Adapter (DRA)
 - IMS V12 will not switch TCBs and will run the IMS request on the calling L8 TCB
 - Prior to IMS V12 DRA will continue to be called on the QR TCB and will switch to a DRA thread TCB
- CICS code implementing CICS-DBCTL interface made threadsafe
- CALL DLI and EXEC DLI APIs are now threadsafe when run with IMS V12

Notes

The IMS DRA is the code that prior to CICS TS V4.2 will switch from the QR TCB to an IMS thread TCB and back again for each IMS request. With CICS TS V4.2 and IMS V12 the DRA will not switch TCBs but will run the request on the calling TCB which will be an L8 TCB. Hence the CALL DLI and EXEC DLI commands are threadsafe when running with IMS V12 and will run on the L8 TCB.

The CICS side of the CICS-DBCTL interface has been made threadsafe so that the processing can run on an L8 TCB.

Threadsafe CICS-DBCTL ...

- For non threadsafe CICS applications there is no change to the amount of TCB switching
 - Switch to L8 TCB to invoke DFHDBAT and IMS
 - Switch back to QR TCB when returning to the application
- For threadsafe CICS applications save TCB switches
 - Remain on L8 TCB after return from IMS request
 - Expect same benefits as seen for threadsafe CICS-DB2 applications
 - Reduced CPU
 - More throughput

Notes

For a non threadsafe application there is no reduction in the amount of switching. Instead of switching from the QR TCB to an IMS thread TCB and back again for each IMS request, it switches from the QR TCB to an L8 TCB and back again.

For a threadsafe application, If it is running on the QR TCB it will switch to an L8 TCB and then stay on an L8 TCB when returning to the application.

For a threadsafe application that is already running on an L8 TCB, or for a CONCURRENCY(REQUIRED) application running on an L8 TCB then no TCB switching occurs for the IMS request.

Threadsafe CICS-DBCTL ...

- CICS-IMS applications when using IMS V12 will use L8 TCBs
 - Adjust MAXOPENTCBS as necessary
- DRA startup table DFSPZPxx still controls how many IMS threads can be used
 - MINTHRDs – the number of threads that remain ‘signed on’
 - DRA keeps IMS threads but dissociated from the TCB
 - MAXTHRDs – will still limit the number of IMS threads from this CICS region
 - Even though an L8 TCB may be available, IMS limits how many can be used for threads
 - *If exceeded, MAXTHRDs will queue in the IMS DRA as it does today (pre OTE)*

Notes

MINTHRD=xxx

This parameter specifies the number of threads for this CICS system that, once initialized, remain created while the DRA is active. These threads remain allocated until this CICS system is disconnected from DBCTL, except if a thread is stopped by a /STOP command or by a thread failure. Additional threads are created, up to the number specified in MAXTHRD, or the number specified in MAXREGN, or the maximum of 999, whichever of these values is the lowest. These additional threads (not the MINTHRDs) are released when there is not enough system activity to require them. The maximum value you can specify for MINTHRD is 999, and the default is 1. For information about specifying values for MINTHRD, see Specifying numbers of threads. See also MAXREGN in IMS system generation macros used by DBCTL.

MAXTHRD=xxx

This parameter specifies the maximum number of transactions for which this CICS region can have PSBs scheduled in DBCTL. Any schedule requests that are over this limit are queued in the DRA. You can balance the load sent to a single DBCTL from multiple CICS regions by specifying appropriate values for MAXTHRD in each CICS region. The maximum value you can specify for MAXTHRD is 999 (but it should not exceed the value specified for MAXREGN) and the default is 1, or the value you specified in MINTHRD. For information about specifying values for MAXTHRD, see Specifying numbers of threads. See also MAXREGN in IMS system generation macros used by DBCTL.

Threadsafe Syncpoint

- Syncpoint Commands now made threadsafe*
 - EXEC CICS SYNCPOINT
 - EXEC CICS SYNCPOINT ROLLBACK
 - EXEC CICS RESYNC
 - * Some TCB switching may still occur, but heavily reduced
- EXEC CICS API no longer switches to QR TCB for syncpoint requests
 - Recovery Manager (RM) domain now makes the decision if a TCB switch is required
 - All RM domain clients register at startup and tell RM if they are threadsafe
 - RM domain keeps task on the open TCB whenever possible
- End of task syncpoint can run on an open TCB before CICS switches to the QR TCB to terminate the task

Notes

The Recovery Manager (RM) processes this Syncpoint command on an open TCB wherever possible to minimize TCB switching. Syncpoint processing can take place on an open TCB for all resource types declared as threadsafe that were accessed in the unit of work. If resource types not declared as threadsafe were accessed in the unit of work, the Recovery Manager switches to the QR TCB for those resource types. A CICS resource type declares itself to the Recovery Manager as threadsafe if the EXEC CICS commands relating to the resource type are threadsafe.

Prior to CICS TS V4.2, CICS would switch to QR prior to end of task syncpoint. In CICS TS V4.2 it remains on an open TCB, if it is running on one, until end of task syncpoint has been called. Afterwards it switches to the QR TCB for the task detach logic.

Threadsafe Syncpoint ...

- Non-threadsafe Recovery Manager clients will force a TCB switch to the QR TCB
 - Examples ...
 - Local Terminal, MRO, Transient Data, ...
- Recovery Manager clients may themselves switch to the QR TCB
 - Examples ...
 - File Control for BDAM files
 - RMI for a quasi-reentrant TRUE
- Example ...
 - Threadsafe application running on an open TCB that has updated DB2 and MQ issues a syncpoint ...
 - In CICS TS V4.1 – 9 TCB switches occur
 - In CICS TS V4.2 – 0 or 1 TCB switch occurs

Notes

Terminal driven transactions involve CICS terminal control logic which is not threadsafe, so a switch to QR will occur during phase 2 syncpoint processing. This is an example of Recovery Manager clients that will cause a switch to QR.

Other RM clients will themselves switch to the QR TCB if necessary. For example, File Control is threadsafe when processing local VSAM files, but it will switch to the QR TCB for BDAM files.

Prior to CICS TS V4.2 a threadsafe application running on an L8 TCB that had updated DB2 and MQ and then issued a syncpoint would suffer 9 TCB switches.

- A switch to QR would be made at the start of syncpoint.
- Switches to L8 and back to QR would occur when calling DB2 for PREPARE
- Switches to L8 and back to QR would occur when calling MQ for PREPARE
- Switches to L8 and back to QR would occur when calling DB2 for COMMIT
- Switches to L8 and back to QR would occur when calling MQ for COMMIT

In CICS TS V4.2, if it was a terminal driven transaction one TCB switch to the QR TCB would occur. For a non terminal driven transaction (and assuming no other non threadsafe resources had been touched) then no TCB switches will occur.

Threadsafe EXEC CICS API Commands

- QUERY SECURITY
- SIGNON, SIGNOFF
- VERIFY PASSWORD, VERIFY PHRASE
- CHANGE PASSWORD, CHANGE PHRASE

- EXTRACT TCPIP, EXTRACT CERTIFICATE

- All Call and EXEC Level Named Counter Server commands
- Built In Functions (BIF) for DIGEST and DEEDIT

Notes

The listed API commands have been made threadsafe in CICS TS V4.2 and can run on an open TCB.

Threadsafe EXEC CICS SPI Commands

- New EXEC CICS SPI commands that are threadsafe ...
 - INQUIRE CAPDATAPRED, INQUIRE CAPINFOSRCE, INQUIRE CAPOPTPRED
 - INQUIRE EPADAPTER, SET EPADAPTER
 - INQUIRE OSGIBUNDLE, INQUIRE OSGISERVICE
 - INQUIRE TEMPSTORAGE, SET TEMPSTORAGE
- Existing EXEC CICS SPI commands made threadsafe ...
 - INQUIRE, SET, PERFORM CLASSCACHE
 - INQUIRE, SET, PERFORM JVMPOOL
 - INQUIRE JVM
 - INQUIRE JVMPROFILE

Notes

The listed SPI commands have been made threadsafe in CICS TS V4.2 and can run on an open TCB.

CICS and 64-bit Support – Overview

- Enabled the CICS domain architecture to run in, and exploit 64-bit addressing mode
- Allowed CICS domains to use stack storage, domain anchor, and all associated domain control blocks in virtual storage above the bar
- Provided below the bar VSCR ...
 - Single System Scaling ...
 - More concurrent tasks, larger applications, ...
 - Increasing pressure on storage usage above the 16MB line!
 - Both in z/OS storage and in EDSA storage!

Notes

The aim of the CICS 64-bit Re-Architecture is to provide a CICS domain architecture environment that exploits the underlying z/Architecture for 64-bit addressing and to provide the infrastructure for the future so that CICS applications will be able to utilize and exploit 64-bit addressing mode. This will enable CICS to remove some of the previous limitations that affect scalability and availability by delivering large address spaces with the exploitation of the 64-bit addressing provided by the z/Architecture. With z/OS 64-bit virtual storage, CICS can make use of this large 64-bit virtual storage to increase capacity by supporting a larger number of concurrent users and concurrent transactions as well as keeping up with the virtual storage demands of increased workload of existing applications and the larger memory requirements of new applications and new technologies.

CICS and 64-bit Support – Overview ...

- CICS Transaction Server for z/OS V4.2 contains significant changes to the CICS domain architecture that exploits the underlying z/Architecture for 64-bit addressing and provides the infrastructure for CICS domains to utilize and exploit 64-bit addressing mode
- These changes affect a number of CICS internal interfaces and related control blocks, including (but not limited to) ...
 - CICS Kernel Anchor and Domain Table separated, ...
 - Kernel Stack, Stack Segments, Larger Save Areas, ...
 - Kernel Task Entry (TAS), ...
 - Common System Area (CSA) – significantly changed ,...
 - Task Control Area (TCA) – 15% smaller, ...
 - Control block changes for the following domains ...
 - Trace, Message, Loader, Temporary Storage, Monitoring, ...

Notes

In CICS TS V4.2 we have enabled the CICS domain architecture to run in, and exploit, 64-bit addressing mode, allowing CICS domains to use stack storage, domain anchor storage, and all associated domain control blocks in virtual storage above the bar.

CICS and 64-bit Support – Overview ...

- The following CICS domains run AMODE(64)
 - Trace Domain, Message Domain, Temporary Storage Domain, ...
 - Kernel Domain, Monitoring Domain, Storage Manager Domain, ...
 - Lock Manager Domain, ...
- z/OS MEMLIMIT
 - New minimum requirement of 4G
 - CICS will not start with less – message DFHSM0602
- EDSALIM
 - Changes in CICS over the years mean that the minimum and default EDSALIM sizes are no longer adequate!
 - Current values ...
 - Minimum size is 10M, Default size is 34M
 - New values ...
 - Minimum size is 48M, Default size is 48M

Notes

The listed domains now run AMODE(64).

The z/OS MEMLIMIT parameter limits the amount of 64-bit (above-the-bar) storage that the CICS address space can use. This storage includes the CICS Dynamic Storage Areas above the bar (collectively called the GDSA) and any MVS storage in the CICS region outside the GDSA.

A CICS region requires at least 4 GB of 64-bit storage. You cannot start a CICS region with a z/OS MEMLIMIT value that is lower than 4 GB. If you attempt to do so, message DFHSM0602 is issued, a system dump with the dump code KERNDUMP is produced, and CICS terminates. Note: CICS does not try to obtain the z/OS MEMLIMIT amount of storage when initializing. 64-bit storage is obtained as required.

The EDSALIM system initialization parameter specifies the upper limit of the total amount of storage within which CICS can allocate the individual Extended Dynamic Storage Areas (ExDSAs) that reside in 31-bit (above-the-line) storage; that is, above 16 MB but below 2 GB.

The minimum and default EDSALIM values have changed to 48 MB to ensure that there is sufficient storage for CICS initialization. These changes have also been reflected in the supplied System Initialization Tables (SITs) and the Installation Verification Procedures (IVPs).

CICS and 64-bit Support – Storage Manager

- Improved 64-bit Storage Manager Statistics ...
 - Additional 64-bit Storage Usage Metrics from z/OS ...
 - Number of memory objects, number of shared memory objects, ...
 - Allocated private storage memory objects, high-water-mark, ...
 - Allocated shared storage memory objects, high-water-mark, ...
 - Current and Peak GDSA allocated, ...
- SOS infrastructure for above the bar storage implemented
 - SMNT Storage_Notify infrastructure implemented
 - RS Domain now Notified when SOS occurs – rather than noticing!
 - For future exploitation ...
 - No current CICS 64-bit exploiter can do much about it!
- Metrics useful to measure the impact of 64-bit JVMs
 - Storage for 64-bit JVMs is outside the CICS DSAs

Notes

The CICS Storage Manager statistics now provide additional information about 64-bit storage. The storage manager global statistics, mapped by the DFHSMDS copy member, and the storage manager dynamic storage areas statistics, also mapped by the DFHSMDS copy member, have been enhanced with new performance metrics.

The Storage Manager reports produced by the DFHSTUP and DFH0STAT statistics programs show these new statistics.

CICS and 64-bit Support – Exploiters

- CICS Java
 - Move to 64-bit JVM (Java 6) for pooled JVM and JVMServer
 - 31-bit Java not supported
- CICS Trace
 - Internal Trace Table above the bar
 - Internal trace table in 64-bit storage only if Transaction isolation inactive (TRANISO=NO) *or* APAR OA34311 applied on z/OS 1.12
 - Transaction Dump Trace Table in 64-bit storage
 - Many trace domain control blocks above the bar
 - 64-bit GTF tracing (OA32611 for z/OS 1.11 and 1.12)
- Message tables are above the bar
 - Subject to TRANISO restriction

Notes

All JVMs now run in AMODE(64) instead of AMODE(31), increasing the capacity for running more JVMs in a CICS region. JVM servers and pooled JVMs use 64-bit storage, significantly reducing the storage constraints in a CICS region for running Java applications. You can therefore reduce the number of CICS regions that run Java to simplify system management and reduce infrastructure costs. You can also use System z Application Assist Processors (zAAPs) to run eligible Java workloads. CICS uses the IBM 64-bit SDK for z/OS, Java Technology Edition, Version 6.0.1. You must download this version of the SDK to run Java applications in CICS. You can continue to build Java programs using a different version of the SDK. If you have any programs that use the Java Native Interface (JNI), including other products, you must ensure that these programs can run in a 64-bit environment.

CICS can obtain 64-bit (above-the-bar) storage, rather than 31-bit (above-the-line) storage for the internal trace table, depending on the version of the z/OS operating system, and whether the CICS region operates with transaction isolation.

CICS and 64-bit Support – Exploiters ...

- Dump Domain
 - 64-bit storage SDUMP support
 - Requires z/OS APAR OA32271 (z/OS 1.11 and 1.12)
 - Use is now made of DUMPPRIORITY for key control blocks in 64-bit storage that should have priority when taking an SDUMP
 - Improves chances of storage being in the SDUMP in case the dump data sets are not big enough!
 - Recommend always use SMS managed dump data sets
- User Exits – eXit Programming Interface (XPI)
 - XPI is based on CICS domain architected interfaces
 - Domain Linkage conventions have changed!
 - **All Global User Exit programs using the XPI MUST be reassembled**

Notes

CICS uses the IARV64 DUMPPRIORITY keyword when allocating 64-bit z/OS storage to ensure that key control blocks, like the trace table, are captured in a dump. You are recommended to use SMS managed dump data sets so that the dump datasets are big enough.

The eXit Programming Interface (XPI) used in Global User Exit programs (GLUEs) is based on the CICS domain architecture. Because significant changes have been made to the domain architecture and the linkage conventions used, all user exit programs that contain XPI calls must be reassembled.

CICS and 64-bit Support – Temporary Storage

- Many Temporary Storage control blocks above the bar
- TS Main above the bar (subject to TRANISO restriction)
- Limiting Main Temporary Storage use
 - New SIT parameter TSMAINLIMIT={64M|amount}
 - Maximum is 32G – But must not be >25% of the z/OS MEMLIMIT for the CICS region else CICS will not start – message DFHTS1608
 - Available on INQUIRE and SET TEMPSTORAGE commands
 - SET TEMPSTORAGE TSMAINLIMIT ...
 - *cannot decrease by too much*
 - or ...
 - *increase to >25% of the z/OS MEMLIMIT for the CICS region*
 - *INVREQ and message DFHTS1606 or DFHTS1607*
 - TSMAININUSE option on INQUIRE TEMPSTORAGE provides current amount of storage occupied

Notes

Main temporary storage queues can now use 64-bit (above-the-bar) storage. CICS provides new facilities so that you can check the storage use of main temporary storage queues and limit that storage use.

Main temporary storage is in 64-bit storage rather than 31-bit (above-the-line) storage, depending on the version of the z/OS operating system and whether the CICS region operates with transaction isolation.

If your CICS applications use large amounts of main temporary storage, the move to 64-bit storage can increase the available storage elsewhere in your CICS regions.

Auxiliary temporary storage queues and shared temporary storage queues continue to use 31-bit storage.

If your CICS applications currently specify the location of temporary storage, you can review this. If an application specifies that auxiliary temporary storage is used, and you do not require recoverable temporary storage, you can change it to specify main temporary storage. The advantages of this action are that space becomes available in 31-bit storage, and input/output activity to write data to disk is reduced.

You can control how much storage the CICS region makes available to main temporary storage queues by using the TSMMAINLIMIT system initialization parameter. The default is 64 MB. This limit on storage use does not include auxiliary temporary storage queues and shared temporary storage queues.

Temporary Storage – Automatic TS Cleanup

- EXPIRYINT can be specified on a TSMODEL definition
 - Applies to ...
 - MAIN and non-recoverable AUXiliary TSQUEUEs
 - Does NOT apply to ...
 - Remote queues – use TSMODEL in QOR instead
 - Recoverable AUX
 - Shared TS queues
 - CICS internal queues
 - Queues that don't match a TSMODEL
 - Default is zero (no expiry) – interval can be specified in hours
 - 0|1-15000
 - Available via CEDA, DFHCSDUP, CPSM BAS, and CICS Explorer
 - SPI to INQUIRE and SET via CPSM and CICS Explorer

Notes

You can specify that CICS deletes temporary storage queues automatically when they are no longer required.

Automatic deletion of eligible temporary storage queues reduces the unnecessary use of virtual storage. To use this feature, you set suitable expiry intervals in the temporary storage models (TSMODEL resource definitions). The expiry interval is available for main temporary storage queues and non-recoverable auxiliary temporary storage queues that match TSMODEL resource definitions in the local CICS region.

Temporary Storage – Automatic TS Cleanup ...

- CICS system task attached every 30 minutes
 - Algorithm to minimise how often full scan of tsqueues is performed
 - Quick scan of TSMODELS to pick up additions/deletions
 - If no TSMODELS with an expiry interval task terminates immediately
 - Not an exact match on expiry interval
 - Tsqueues will be deleted by interval+10% subject to 30 minute minimum
 - When task performs a scan, it issues message DFHTS1605
 - Shows the number of temporary storage queues that were scanned and the number that were deleted
 - If the cleanup task ends abnormally, it issues message DFHTS0001 and does not run again until CICS is restarted
- If you change the expiry interval in a TSMODEL
 - Existing tsqueues that match the model are not affected
 - Use the expiry interval that applied when they were first created!

Notes

EXPIRYINT specifies the expiry interval, in hours, for a temporary storage queue that matches this model. The interval count begins after each use of the temporary storage queue. If the temporary storage queue is not used again before the expiry interval is reached, the temporary storage queue becomes eligible for CICS to delete it automatically.

When the CICS cleanup task performs a scan, it issues message DFHTS1605. This message shows the number of temporary storage queues that were scanned and the number that were deleted. If the cleanup task ends abnormally, it issues message DFHTS0001, and does not run again until CICS is restarted.

If you change the expiry interval in a TSMODEL resource definition, existing temporary storage queues that match the model are not affected. Those temporary storage queues continue to use the expiry interval that applied when they were created. If all the TSMODEL resource definitions with a nonzero expiry interval are deleted from a CICS region, CICS stops scanning for expired temporary storage queues.

The existing behaviour of the TSAGE parameter specified on a macro TST is unchanged.

More VSAM LSR Pools

- CICS historically has allowed use of 8 LSR pools
 - In the days of local DLI, DLI used the other 8
 - When local DLI was removed, CICS File Control was not changed
- Since z/OS V1.4 ...
 - DFSMS allowed use of 256 LSR pools per address space
- In CICS TS V4.2 we will allow use of 255 LSR pools (1 - 255)
 - We will not use LSR pool 0
 - Today LSRPOOL number is held in one byte on the FCTE
 - With 0 meaning no LSR (i.e. NSR)
- Potential performance optimization where greater subdivision of VSAM files across LSRPOOLS is required
 - i.e. Place highly-used files in their own LSRPOOL

Notes

You can now define up to 255 LSR pools using the LSRPOOLNUM attribute in the FILE and LSRPOOL resource definitions.

The existing attribute LSRPOOLID in the FILE and LSRPOOL resource definitions has been replaced with LSRPOOLNUM to increase the maximum number of LSR pools available in a CICS region from 8 to 255.

The LSRPOOL resource defines the size and characteristics of the local shared resource (LSR) pool. Increasing the number of LSR pools available in a CICS region can improve transaction response times for CICS workloads that use VSAM files.

- For example an application that processes an entire data set every time the transaction is used can cause excessive I/O operations. Placing the data set into its own LSR pool can help eliminate I/Os and prevent buffer stealing in a shared LSR pool.

- You can also slow down transactions that are causing buffer steals in a shared VSAM LSR pool by placing the VSAM data sets in a detuned LSR pool to limit the number of data component buffers. In many cases a dedicated LSR pool performs better than a shared LSR pool. Dedicated LSR pools are also an excellent replacement for CICS data tables, when the capacity of the CICS data tables has been exceeded.

More VSAM LSR Pools ...

- CEDA DEFINE FILE and DEFINE LSRPOOL
 - Today CEDA uses one character field for LSRPOOLID
 - New LSRPOOLNUM keyword introduced allowing values
 - NONE|1-255 – Default is 1
 - LSRPOOLID now only set via compatibility mode for back level systems with a range 1-8
 - For existing definitions, the value of LSRPOOLID is transferred to LSRPOOLNUM
- DFHCSDUP
 - Allows specification of LSRPOOLNUM with range 1-255
 - Allows setting of LSRPOOLID with range 1-8 in compatibility mode
 - Existing job using LSRPOOLID will run ...
 - Value transferred to LSRPOOLNUM

Notes

You can use the new LSRPOOLNUM attribute on the FILE and LSRPOOL resource definitions to specify LSR pool numbers in the range 1 – 255. The existing LSRPOOLID value which takes values in the range 1 – 8 is maintained for compatibility.

More VSAM LSR Pools ...

- EXEC CICS INQUIRE and SET FILE
 - LSRPOOLNUM defined as alias of LSRPOOLID
 - Existing programs using LSRPOOLID will still work
- EXEC CICS CREATE FILE and CREATE LSRPOOL
 - Supports use of LSRPOOLNUM
 - Existing use of LSRPOOLID will still work
- CICS Explorer
 - Existing parameter CSDLSRNO changed to support range of 1-255
- CPSM
 - Supports LSRPOOLNUM for Operations and Admin
- Statistics
 - Reports on up to 255 LSR pools

Notes

All Externals that today use LSRPOOLID in the range 1-8 have been changed to support LSRPOOLNUM in the range 1-255.

Summary

- CICS TS V4.2 delivers new and enhanced capabilities in the area of scalability ...
 - New **concurrency** option for greater OTE exploitation
 - Threadsafe Mirrors for Function Shipping over IPIC connections
 - Threadsafe CICS-DBCTL interface when connected to IMS V12
 - More threadsafe EXEC CICS API and SPI commands
 - New 64-bit infrastructure and exploitation
 - More VSAM LSR Pools
- Foundation for even greater capacity and throughput

Notes

The scalability enhancements provided by CICS Transaction Server for z/OS Version 4.2 provide, via the OTE enhancements, the ability for more workload to exploit the power of the mainframe. The 64-bit enhancements provide the ability to scale vertically, and therefore do more work in a single CICS region, as well as providing a foundation for even greater capacity in the future.