Tuning z/OS for WebSphere 4.0.1 (A View from 100,000 Feet)

Clark L. Kidd
Watson & Walker, Inc.
Session 2928; SHARE 101 in Washington, D.C.
August 12, 2003

Agenda for This Presentation

- Overview
- z/OS Base Tuning
- TCP/IP Tuning
- z/OS UNIX Tuning
- Language Environment (LE) Tuning
- Workload Manager Tuning
- Hardware Upgrades
- References
- Questions?

Overview: The Big WebSphere Picture

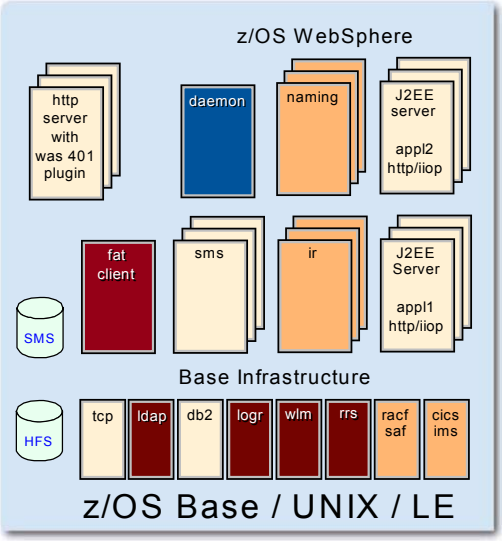

- Base Infrastructure
  - z/OS and z/OS UNIX
  - HFS and zFS
  - Language Environment (LE)
  - TCP/IP
  - WLM
  - System Logger
  - DB2
  - RACF (SAF)
  - RRS
- System Servers
  - DAEMON
  - NAMING
  - SMS (System Management Server)
  - IR (Interface Repository)
- J2EE Servers
  - Hosting J2EE applications
- HTTP
  - IBM HTTP Server (plugin)
  - J2EE Server HTTP Transport Handler
- Backend
  - CICS
  - IMS
  - DB2
  - MQSeries (JMS)

z/OS Base Tuning

- Hints for a smooth install
  - Pull PSP for WebSphere AND related products
  - Follow installation instructions to the letter
- Keep maintenance current
  - z/OS (including z/OS UNIX)
  - WebSphere Application Server (WAS)
  - Java
  - Language Environment (LE)
  - DB2 and JDBC
- Much new function being delivered via service

z/OS Base Tuning

- WebSphere storage consumption
  - Nine+ new address spaces
  - Each needs 384 MB minimum to initialize
  - Expect even more as workload increases
  - Add 200 MB if WebSphere modules not in LPA
- Virtual storage recommendation #1
  - Code "REGION=0M" in the JCL for:
    - WebSphere tasks (both run-time and server)
    - IBM HTTP Server (IHS or IMWEBSRV)
    - TCP/IP

z/OS Base Tuning

- Virtual storage recommendation #2
  - Investigate exits and thresholds that can limit region size
    - JES2 exit EXIT06
    - JES3 exit IATUX03
    - SMF exit IEFUSI
    - UNIX thresholds in BPXPRMxx
    - UNIX thresholds in RACF OMVS segments

z/OS Base Tuning

- Virtual storage recommendation #3
  - Place WAS load modules in dynamic LPALIB
  - Use dynamic option because of PDSE libraries
  - Add libraries to PROGxx PARMLIB member:
    - LPA ADD DSNAME(WAS40.WAS.SBBOLPA) MASK(*)
    - LPA ADD DSNAME(WAS40.WAS.SBBOLOAD) MASK(*)
  - Use COM='SET PROG=xx' in COMMNDxx to activate
  - Check and increase ECSA size accordingly
  - Check and increase size of CSA page data set
  - Remove any unneeded STEPLIBs from JCL

z/OS Base Tuning

- Virtual storage recommendation #4
  - If storage resources are still a problem, consider:
    - Reducing JVM heap size (256 MB is the default)
    - Reducing LE heap size (80 MB is the default)

z/OS Base Tuning

- GRS considerations
  - ENQ/DEQ used for global WebSphere transactions
  - Suggest you set GRS options based on configuration:
    - For a sysplex, use GRS=STAR
    - For a monoplex, use GRS=NONE

z/OS Base Tuning

- Component tracing
  - Check this debugging option that affects performance
  - Identify components being traced
    - D TRACE,COMP=ALL
  - Disable trace if not needed
    - TRACE CT,OFF,COMP=xxxx (temporary)
    - SYS1.PARMLIB(CTncccxx) (permanent)

z/OS Base Tuning

- System Logger
  - Common sysplex-wide logging facility
    - One logger task per image
    - Provides a single-system view of logged data
  - Supports two types of logstreams
    - DASD-only logstream
    - Coupling Facility logstream
  - Recommendations
    - Use Coupling Facility logstream
    - Else, use fast devices with DASD Fast Write (DFW)
    - Else, allocate logstreams with large CI sizes

z/OS Base Tuning

- System Logger logstream sizes
  - Small logs cause overhead during offload
  - Large logs may exhaust space in Coupling Facility
  - SMF record type 88 may be used to monitor and tune
    - Is CF size adequate?
    - What is the offload frequency?
- Use of the RRS Archive Log
  - Expect 1-3 records for each WebSphere transaction
  - Consider disabling log because of extra overhead
    - Create log during application development and tuning
    - Disable log when application stabilizes

TCP/IP Tuning

- TCPIP.PROFILE changes
  - Increase send/receive buffer sizes from 16K default

```
TCPCONFIG TCPSENDBFRSIZE 65535
          TCPRCVBUFRSIZE 65535
```

  - Add 'NODELAYACKS' to PORT statement(s)
    - Shortens delay in sending ACK response back to client
    - Should improve throughput for trivial transactions
    - Adds slightly more overhead for complex transactions
  - Specify 'SHAREPORT OFF' (default)
    - A value of 'ON' allows sharing of the same port number
    - Avoid 'ON' if possible due to excessive overhead

TCP/IP Tuning

- TCP/IP and VTAM
  - If mixing web-based and legacy applications...
  - ...Same dispatching priority for TCP/IP and VTAM
- HiperSockets for IP traffic in a Sysplex
  - Faster and less overhead than using XCF
- Advanced TCP/IP network options
  - Multiple TCP/IP stacks (availability)
  - Connection Optimization (WLM and DNS)
  - WebSphere Edge Server

z/OS UNIX Tuning

- Changing UNIX threshold values
  - Update BPXPRMxx for a permanent change
  - Use SETOMVS command for a temporary change
  - Check values when configuration changes
    - HTTP Server operational mode (Standalone/Scalable)
    - HTTP Server versus HTTP Transport Handler
    - Transaction load
  - Consider RACF overrides for some changes
  - Monitor regularly to prevent problems

z/OS UNIX Tuning

- MAXFILEPROC
  - Defines maximum files open for any one process
  - Use 10000 for up to 150 concurrent users
  - Or ... 120 X number of web requests per second
  - Consider RACF override for this threshold
- MAXTHREADTASKS
  - Defines maximum active threads for any one process
  - Use 1000-5000 based on transaction load

z/OS UNIX Tuning

- MAXTHREADS
  - Defines maximum total threads for any one process
  - Use 2 X value of MAXTHREADTASKS
- MAXSOCKETS
  - Defines maximum sockets in AF_INET socket family
  - Use (2 X value of MAXFILEPROC) + 500
- IPCSHMMPAGES
  - Defines maximum pages for shared memory segments
  - Use 256 (for IHS running in Scalable mode)

z/OS UNIX Tuning

- zSeries File System (zFS)
  - More efficient alternative to HFS
  - An enhancement to the DCE Episode file system
  - Supported starting with z/OS V1R2
  - Available to z/OS V1R1 and OS/390 V2R10 via service
  - Two types of file systems supported
    - Compatibility Mode aggregate
    - Multi-File System aggregate

z/OS UNIX Tuning

- zSeries File System (zFS) considerations
  - Transparent to applications (no changes needed)
  - Performs much better for random access
  - May perform better for sequential access
  - Other enhancements
    - File system cloning
    - Space Quotas and Access Lists
    - Better error recovery after system outages
  - But...
    - More difficult to administer
    - Cannot replace HFS in some cases

Language Environment (LE) Tuning

- General tuning hints
  - Put reentrant LE routines in common storage
    - Update SYS1.PARMLIB(LPALSTxx)
    - Add library name CEE.SCEELPA
  - Place non-reentrant LE routines in the link list
    - Update SYS1.PARMLIB(LNKLSTxx)
    - Add library name CEE.SCEERUN
  - But … when running multiple LE versions
    - Do not add modules to LPALIB/link list
    - Use Run-Time Library Services (RTLS) instead

Language Environment (LE) Tuning

- Application tuning hints
  - Update JCL for the server regions
    - Set PARM='RPTOPTS(ON),RPTSTG(ON)'
    - This will display LE options and storage statistics
  - Results will appear in joblog after a clean shutdown
    - VARY WLM,APPLENV=xxxx,QUIESCE
    - P server_name
  - Evaluate results and change options as needed
  - Remove options from JCL when finished

Language Environment (LE) Tuning

- Heap size tuning hints
  - Default of 80MB may be too large
  - Add 'RPTOPTS(ON),RPTSTG(ON)' to JCL as shown earlier
  - Evaluate output and reduce defaults if appropriate
    - HEAP(iiM,ssM)
    - ii = Initial heap allocation size in MB
    - ss = Secondary heap allocation size in MB
  - Test results and adjust values as needed
  - As before, remove diagnostic options when finished

Workload Manager (WLM) Tuning

- Control Region Service Classes
  - Assigned with the STC subsystem classification rules
  - Assign a high-velocity Service Class (or SYSSTC)
  - Consider a separate Reporting Class
  - Assign Service Class to the following tasks:
    - Daemon Server Control Region
    - System Management Server (SMS) Control Region
    - Naming Server Control Region
    - Interface Repository Control Region
    - Application Server Control Regions (J2EE and MOFW)

Workload Manager (WLM) Tuning

- Application Environments
  - A feature of Workload Manager used by:
    - WebSphere Server Regions
    - DB2 Stored Procedures
    - MQSeries Workflow
    - SOM
  - Implemented using Enclaves
  - Composed of a Queuing Manager and Server Regions

Workload Manager (WLM) Tuning

- Application Environments
  - Queuing Manager
    - Usually only one
    - Receives work and passes it to WLM
    - WLM passes work on to a Server Region
  - Server Regions
    - Usually multiple; running in separate address spaces
    - Accepts work from Queuing Manager and WLM
    - WLM will start the first one when work arrives (optional)
    - WLM will start/stop others to meet Service Class goals
    - Server Regions may run across a Sysplex

Workload Manager (WLM) Tuning

- Defining WebSphere Application Environments
  - Required for run-time and application server regions
    - System Management Server (SMS) Server Region
    - Naming Server Server Region
    - Interface Repository Server Region
    - Application Server Server Regions (J2EE and MOFW)
  - Define with subsystem type CB
  - Specify the PROC name to be started by WLM
  - Set number of servers to 'No limit'
    - Consider using '1' just during new application testing
    - Directives MIN_SRS and MAX_SRS also control this

Workload Manager (WLM) Tuning

- Server Region Service Classes
  - Started with the STC subsystem classification rules
    - Assign a high-velocity Service Class (or SYSSTC)
  - Arriving transactions put into new Service Classes
    - Service Class assignment is driven by Control Region
    - Transaction is then passed to WLM
    - WLM assigns to appropriate Server Region
  - Only one active Service Class per Address Space
    - WLM starts MIN_SRS=x Regions for each Service Class
    - MAX_SRS=x limits upper bound and may cause waiting
    - Set MAX_SRS = (MIN_SRS * maximum Service Classes)

Workload Manager (WLM) Tuning

- WebSphere transactions and WLM subsystems

Transaction Source	Subsystem
HTTP Server (Standalone mode)	CB
HTTP Server (Scalable mode)	IWEB
HTTP Server (Enclave mode)	IWEB
HTTP Transport Handler	CB
IIOP Client	CB

Workload Manager (WLM) Tuning

- Transaction Service Class guidelines
  - Suggest using a response-time based goal
  - Set priorities based on business goals
  - Use default CB/IWEB Service Class with high priority
    - Allows you to detect classification errors without penalty
    - Otherwise, default is SYSOTHER (low priority)
  - Subsystem CB classification can be based upon
    - Collection Name (CN); Logical server name
    - Subsystem Instance (SI); Logical server instance
    - Userid (UI)
    - Transaction Class (TC)

Hardware Upgrades

- Processor upgrade
  - Model G5+ provides native IEEE floating point support
  - Java workloads are floating-point intensive
- Memory upgrade
  - Expect to use 512 MB for a lightly-loaded system
  - Expect to use up to 2 GB for heavy workloads
- Other upgrades
  - Hardware encryption for SSL-heavy workloads
  - Network bandwidth
  - Sysplex configuration

References

- WebSphere tuning
  - *WebSphere for Dummies;* **Hilon Potter; IBM; SHARE 99 Session 2913; August 2002**
  - *z/OS for e-business - An introduction to System Tuning*; **SG24-6542-00**; *www.redbooks.ibm.com*
  - *Performance Tuning for WebSphere on z/OS*; **Bob St. John; IBM; SHARE 100 Session 2565; February, 2003**
  - *Using OS/390 WLM to Manage WebSphere Performance;* **Glenn Anderson; IBM; SHARE 99 Session 2547; August, 2002**

References

- WebSphere tuning
  - *WLM Application Environments;* **Kathy Walsh; IBM; SHARE 99 Session 2513; August 2002**
  - *Configuring WebSphere V4 for z/OS for Availability and Performance*; **John Hutchinson; IBM; SHARE 100 Session 2962; February, 2003**
  - *WebSphere Transaction Classification*; **Cheryl Watson's TUNING Letter 2003, No. 1;** *www.watsonwalker.com*
  - *Maximizing WebSphere Performance on z/OS*; 4.5 days; **IBM Learning Services; Course Code OZ850**
  - Questions?