



IBM Software Group

WebSphere Application Server I/O Performance

Michele Calcavecchia (macalcav@us.ibm.com)
WebSphere Application Server L2 Support Engineer
January 2013



WebSphere® Support Technical Exchange



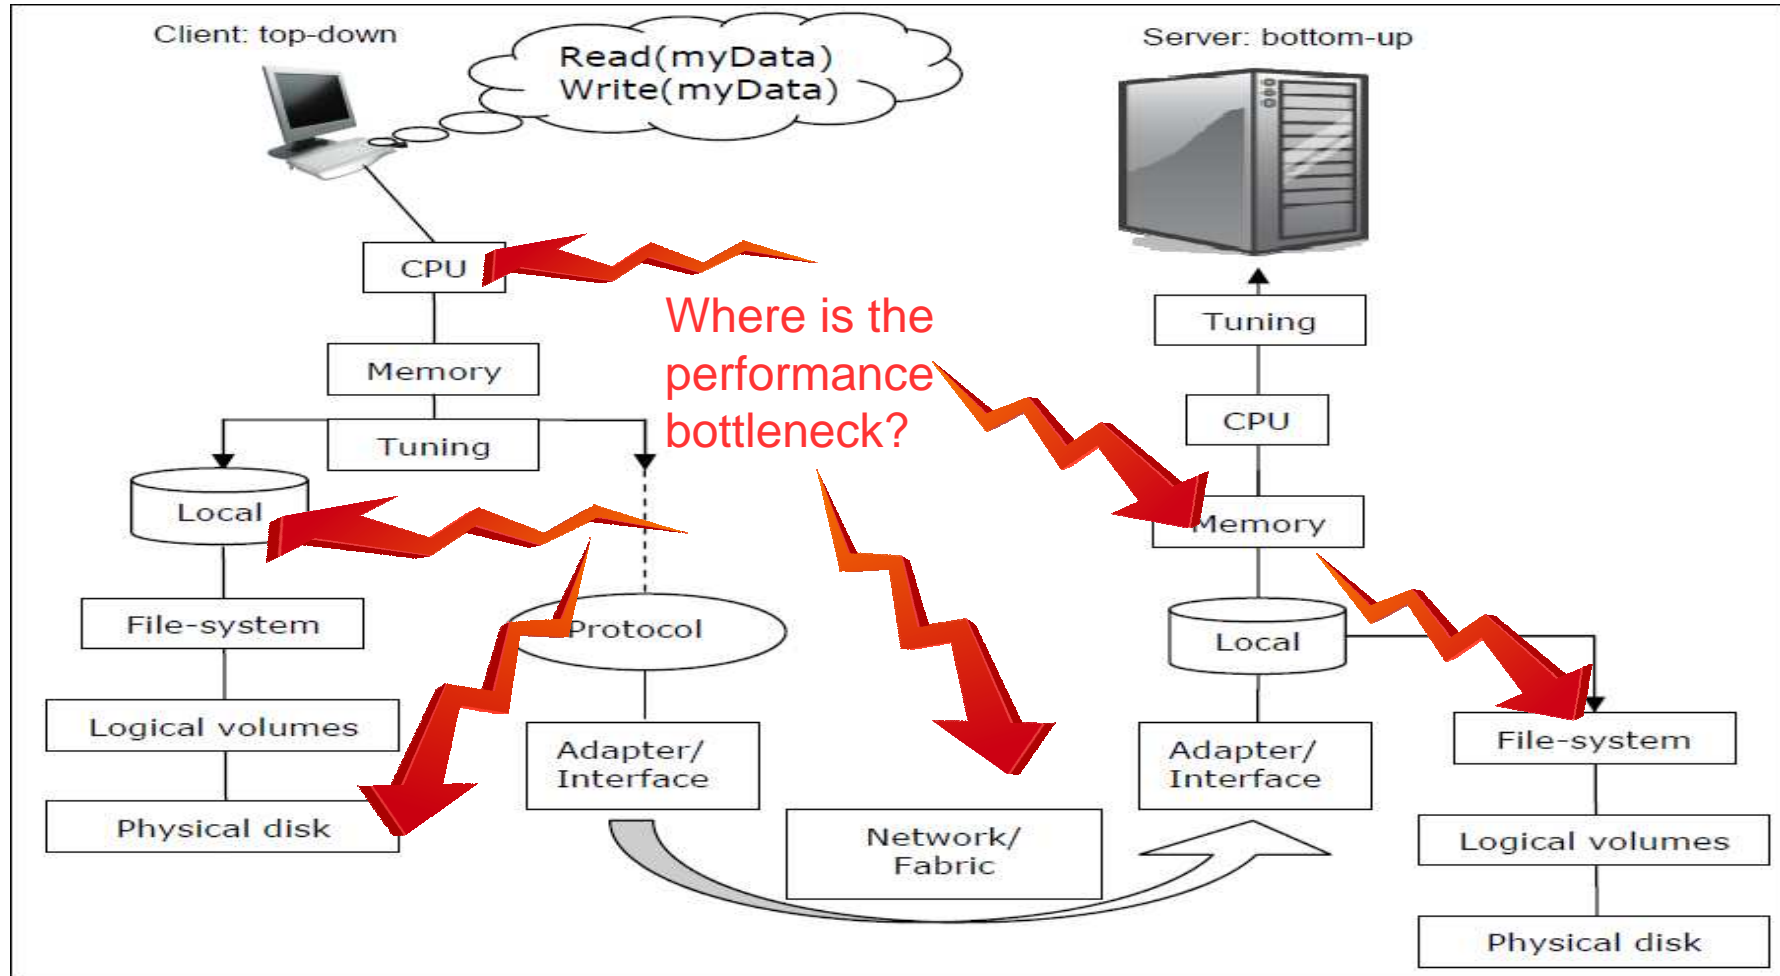
Agenda

- Definition of I/O
- Example of I/O
- Considerations for WebSphere I/O Related Problems
- “Optimal Hardware Configuration”
- WebSphere Support/Collecting the data
- Reviewing Data – vmstat
- Diagnosis/Reviewing Data – vmstat
- Underlying Cause of Performance Issues
- Application Specific Information
- Tools
- Next Step Tools
- Avoid Performance Problems
- Summary

Definition of I/O

- I/O (input/output) occurs everywhere in a computer system where input is followed by an output of information
- I/O occurs between computer systems
- I/O is not a single event, it is comprised of different parts and operations
 - ▶ Network
 - ▶ File System
 - ▶ Database
 - ▶ Internet
- There could be differing views of I/O as one can look at I/O in terms of queuing systems or CPU registers. However in this presentation we are speaking in terms of file and socket operations and the general transfer of data between computer components.

An Example of I/O Components



Considerations for WebSphere I/O Related Problems

- A request will only perform as fast as it's slowest component.
- Parts to consider for I/O:
 - ▶ Number of requests
 - ▶ Type of file system
 - Network Attached Storage
 - Virtualized
 - ▶ RAM
 - ▶ Disk latency
- The idea behind strategic I/O tuning is to avoid the slowest components (disk) and opt to use the fastest (cache).

“Optimal Hardware Configuration”

In the WebSphere Application Server Information Center section “Tuning the application serving environment” asks that you check these items for “optimal hardware configuration”:

- Optimize Disk Speed
- Increase Processor Speed
- Increase System Memory
- Run Network Cards and Network Switches at Full Duplex

WebSphere L2 Support

- WebSphere Support is responsible for identifying WebSphere Application Server performance issues reported as possible defects
- Client is asked to collect performance data during the time of the poor performance
MustGather: Performance, hang, or high CPU issues with WebSphere Application Server on AIX
<http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg21052641>
- It is very important that the data be collected during the time of poor performance and that all the data be collected during that same time frame.
- By correlating the information from the different logs, WebSphere Support can identify the areas of performance concern.

WebSphere Support – Collecting the Data

- Data collection by
 - ▶ ISA
 - ▶ “Collecting Data Manually” - (script included)
- Important logs for Support's review of I/O related issues:
 - ▶ Vmstat
 - ▶ Javacores or thread dumps
 - ▶ Also:
 - Garbage collection data (rule out Java memory problems)
 - WebSphere Application Server logs:
 - SystemOut
 - systemErr
 - native_stdout
 - native_stderr

Diagnosis

In the upcoming slides we will look at ways to identify I/O related performance problems that affect WebSphere performance.

Reviewing Data: vmstat

vmstat

- available on
 - ▶ AIX
 - ▶ Linux
 - ▶ Solaris
- Not available on Windows
 - ▶ *use other tools defined later in presentation

vmstat example

This is output from AIX vmstat

kthr	memory		page				faults				cpu					
---	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	
r	b	avm	fre	re	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa
0	0	3456	27743	0	0	0	0	0	0	131	149	28	0	1	99	0
0	0	3456	27743	0	0	0	0	0	0	131	77	30	0	1	99	0
1	0	3498	27152	0	0	0	0	0	0	153	1088	35	1	10	87	2
0	1	3499	26543	0	0	0	0	0	0	199	1530	38	1	19	0	80
0	1	3499	25406	0	0	0	0	0	0	187	2472	38	2	26	0	72
0	0	3456	24329	0	0	0	0	0	0	178	1301	37	2	12	20	66
0	0	3456	24329	0	0	0	0	0	0	124	58	19	0	0	99	0
0	0	3456	24329	0	0	0	0	0	0	123	58	23	0	0	99	0

- Columns of interest for I/O related issues:

- ▶ b – “blocked”

- Average number of kernel threads placed in the Virtual Memory Manager (VMM) wait queue (awaiting resource, awaiting input/output) over the sampling interval

- ▶ wa - “wait”

- Processor idle time during which the system had outstanding disk/NFS I/O request(s)

vmstat example (cont.)

- pi - “page in”
 - ▶ Number of pages paged in from paging space. Paging space is the part of virtual memory that resides on disk. It is used as an overflow when memory is over committed. Paging space consists of logical volumes dedicated to the storage of working set pages that have been stolen from real memory. When a stolen page is referenced by the process, a page fault occurs, and the page must be read into memory from paging space. It is likely that a page-in will force a different page to be stolen and, therefore, paged out.
- po - “page out”
 - ▶ Number (rate) of pages paged out to paging space. Whenever a page of working storage is stolen, it is written to paging space
- fr - “free memory pages”
 - ▶ Average number of free memory pages. A page is a 4 KB area of real memory. The system maintains a buffer of memory pages, called the free list, that will be readily accessible when the VMM needs space.
- sr - “scan rate”
 - ▶ Number of pages that were examined per second by the page-replacement algorithm during the interval. The higher the sr value compared to the fr value, the harder it is for the page-replacement algorithm to find eligible pages to steal.

vmstat example (cont.)


kthr		memory		page					faults			cpu				
r	b	avm	fre	re	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa
0	0	3456	27743	0	0	0	0	0	0	131	149	28	0	1	99	0
0	0	3456	27743	0	0	0	0	0	0	131	77	30	0	1	99	0
1	0	3498	27152	0	0	0	0	0	0	153	1088	35	1	10	87	2
0	1	3499	26543	0	0	0	0	0	0	199	1530	38	1	19	0	80
0	1	3499	25406	0	0	0	0	0	0	187	2472	38	2	26	0	72
0	0	3456	24329	0	0	0	0	0	0	178	1301	37	2	12	20	66
0	0	3456	24329	0	0	0	0	0	0	124	58	19	0	0	99	0
0	0	3456	24329	0	0	0	0	0	0	123	58	23	0	0	99	0

- High values in the “wa” column (in the absence of other high values like pi) is an indicator that the contention is a source on the other side of the Application Server (database? network?).
- Investigate why the system is idle while there are pending requests.
- This system may also just be showing signs of underutilization. Perhaps WebSphere has been so efficient in completing tasks that it spends more time waiting than doing work?
 - Vmstat data should not be considered unless there has already been a performance issue identified.
 - Vmstat is a quick way to identify where possible performance problems are and is not the last stop in identifying the underlying cause.

Underlying Cause of Performance Issues Identified by vmstat

- Disk Latency
- System Memory Shortage
 - ▶ not to be confused with java heap memory

Disk Latency



- Number of I/Os
- Data layout
- Type of disk
- Speed of disk
- Disk speed and configuration have a dramatic effect on the performance of application servers running applications that are heavily dependent on the database support, using extensive messaging, or processing workflow
- ★ Avoid disk I/O

vmstat suspect: disk latency

Linux Example:

```
procs  ---memory-----          swap-  --io--          -system-  --cpu---
r  b  swpd   free   buff  cache   si so   bi   bo    in   cs  us sy id wa st
1   8   5888 140136 167700 1506440  0  0   6783 17333  1738  16648  6  2  9  83  0
1   6   5888 140136 167700 1506440  0  0   6885 18353  1831  18658  8  4  9  77  0
1   6   5888 140136 167700 1506440  0  0   6482 12323  1534  17632 16 12  9  63  0
```

“wa” shows high percentage waiting on i/o while a number of threads are “b” blocked waiting on...disk i/o?

vmstat suspect: disk latency

AIX Example:

kthr		memory		page						faults			cpu			disk xfer				
r	b	avm	fre	re	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa	1	2	3	4
0	0	3456	27743	0	0	0	0	0	0	131	149	28	0	1	99	0	0	0		
0	0	3456	27743	0	0	0	0	0	0	131	77	30	0	1	99	0	0	0		
1	0	3498	27152	0	0	0	0	0	0	153	1088	35	1	10	87	2	0	11		
0	1	3499	26543	0	0	0	0	0	0	199	1530	38	1	19	0	80	0	59		
0	1	3499	25406	0	0	0	0	0	0	187	2472	38	2	26	0	72	0	53		
0	0	3456	24329	0	0	0	0	0	0	178	1301	37	2	12	20	66	0	42		
0	0	3456	24329	0	0	0	0	0	0	124	58	19	0	0	99	0	0	0		
0	0	3456	24329	0	0	0	0	0	0	123	58	23	0	0	99	0	0	0		

- “wa” shows high percentage waiting on i/o
- “disk xfer”:
 - number of transfers per second to the specified physical volumes that occurred in the sample interval. The count only represents requests to the physical device and not the amount of data requested.

vmstat suspect: disk latency

Solaris Example:

procs		memory		page							disk				faults			cpu			
r	b	w	swap	free	re	mf	pi	po	fr	de	sr	s5	s1	s3	s4	in	sy	cs	us	sy	id
0	39918	0	55724136	3595152	209	126	4	0	0	0	0	0	0	0	1	1141	5007	3253	12	2	85
0	16401004	0	55636368	3354576	639	424	0	0	0	0	0	0	0	0	0	1409	5655	3546	27	6	66
0	16401492	0	55610824	3328832	55	1205	0	0	0	0	0	0	0	0	0	1303	10084	3235	38	6	56
0	16401591	0	55595112	3313160	463	529	0	0	0	0	0	0	0	0	0	1352	2518	3583	30	6	64
0	16401615	0	55584992	3303016	771	420	0	0	0	0	0	0	0	0	11	1314	2013	3046	33	6	60

- b - the number of blocked kernel threads that are waiting for resources I/O, paging, etc.
- w - (different meaning than Linux and AIX) the number of swapped out lightweight processes (LWPs) that are waiting for processing resources to finish
- disk - number of disk operations per second. There are slots for up to four disks. The letter indicates the type of disk, the number is the logical unit number.

System Memory

■ Avoid paging

▶ Heap sizing problems

http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/understanding/mm_heapsize_problems.html

- ...The java heap should not be more than 70% occupied. The maximum heap size should, if possible, be able to be contained in physical memory to avoid paging. For the best performance, try to ensure that the heap never pages.

■ Consider memory access speed

▶ Platform 2015: Intel® Processor and Platform -Evolution for the Next Decade

http://epic.hpi.uni-potsdam.de/pub/Home/TrendsAndConceptsII2010/HW_Trends_borkar_2015.pdf

“the advantages of higher clock speeds are in part negated by memory latency, since memory access times have not been able to keep pace with increasing clock frequencies.”

- ★ One way to determine the appropriate amount of RAM for a system is to look at the largest value for avm as reported by the vmstat command. Multiply that by 4 K to get the number of bytes and then compare that to the number of bytes of RAM on the system. Ideally, avm should be smaller than total RAM

http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.prftungd/doc/prftungd/mem_usage_vmstat.htm

Vmstat AIX – system memory

kthr		memory				page				faults			cpu		
r	b	avm	fre	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa
1	0	144059	61061	0	0	0	0	0	14	154	149	1	3	97	0
1	0	144057	61063	0	0	0	0	0	12	285	145	1	3	96	0
2	0	219258	7451	0	0	21589	202426	3	20	10251	212	74	26	0	0
0	2	276288	4632	51	22219	57551	353503	6	6756	10341	2794	27	68	2	2
1	4	318982	4571	31	20527	21412	24771	0	4847	3279	1076	20	72	3	4
0	8	341294	4717	37	21518	22285	24563	0	6212	3642	1166	21	72	3	5
0	6	357028	4819	29	15234	15415	17132	0	4216	2290	784	22	70	3	5

- Every time a computational page (CP) has to be paged-in means that a running application is blocked until the CP is written back into memory from the paging space
- The ratio (sr/fr) between sr (scanned pages) and fr (freed pages) defines how many file pages were scanned in order to find one that could be freed.
- The higher the sr value compared to the fr value, the harder it is for the page-replacement algorithm to find eligible pages to steal.

Vmstat Linux: system memory

```
procs-----memory----- --swap-- --io--- --system-- ----cpu----
r b swpd free buff cache si so bi bo in cs us sy id wa
1 0 1921 113 6 2157 0 0 296 136 1598 3865 82 8 9 1
2 0 1921 108 6 2176 0 0 16 0 1453 3130 74 5 21 0
0 2 1931 119 6 2178 0 9668 568 9868 6855 3748 69 12 0 19
5 2 1931 98 6 2202 0 0 2604 8 4958 3125 66 8 3 24
1 4 1932 96 6 2204 0 1096 3112 1100 4035 2679 48 8 1 43
2 4 1932 119 4 2188 0 0 2364 36 15909 2754 57 12 9 22
0 4 1937 97 4 2206 0 4968 3464 4968 3889 3085 55 10 1 35
7 3 1937 101 0 2212 0 0 3168 160 12122 2931 63 11 1 26
1 4 1942 104 0 2210 0 4932 3348 5232 16944 3087 66 12 3 20
```

- swpd – virtual memory used in KB
- free – idle memory in KB
- If the virtual memory used (memory-swpd) remains high and the free memory (memory-free) remains low, then the system may be memory constrained and can benefit from additional RAM.

vmstat Solaris: system memory

procs			memory				page					disk			faults		cpu			
r	b	w	swap	free	re	mf	pi	po	fr	de	sr	s1	s2	s3	in	sy	cs	us	sy	id
0	0	0	644688	28240	0	1263	608	5	5	0	0	103	2	5	762	4440	714	26	40	33
1	0	0	532216	18940	7	2281	226	8	8	0	0	9	2	53	511	2892	600	14	62	24
6	0	0	528912	7652	15	2159	44	232	818	3072	302	7	3	8	611	2488	613	32	61	6
5	0	0	521648	7112	10	2369	94	222	330	1640	44	35	8	2	781	2569	742	30	70	0
6	0	0	525804	7136	21	2381	1310	672	1840	1500	599	24	5	21	637	2741	634	31	69	0
2	0	0	527992	7880	9	1349	3405	557	2066	1100	684	7	4	133	675	1457	643	16	45	39
1	0	0	530208	6780	13	1261	3553	1170	3580	1220	1079	32	4	165	725	1375	661	12	43	46
1	0	0	526996	7028	7	855	36	181	530	1576	149	10	2	4	569	4596	539	66	23	11
0	0	0	527548	8440	16	541	65	114	250	1408	55	14	1	1	505	2002	559	44	20	36
0	0	0	530664	11008	2	499	25	2	0	1032	0	16	6	1	383	1209	511	6	15	79
1	0	0	531460	35916	0	370	130	0	0	756	0	4	6	0	367	866	519	3	24	73

- de - anticipated short-term memory shortfall in KB, which gives the page scanner a target number of pages to free up. Any activity in this column should be considered a problem as this implies that the system feels need to prepare to page out active pages.
- As is true with AIX, the scanned page (sr) count should not exceed the freed page (fr) count.

Application Specific Information

- If an I/O problem is suspected look for further clues in the WebSphere logs:
 - ▶ Javacores/Thread Dumps
 - ▶ System Logs

Javacores and Thread Dumps

```
"WebContainer : 48" (TID:0x000000013D7B6D00,  
  sys_thread_t:0x000000013D7867B0, state:CW, native ID:0x000000002570099)  
prio=5  
  at java/io/FileOutputStream.open(Native Method)  
  at java/io/FileOutputStream.<init>(FileOutputStream.java:205(Compiled  
Code))  
  at java/io/FileOutputStream.<init>(FileOutputStream.java:157(Compiled  
Code))...
```

- If the javacore shows many threads in `java/io/` or a thread that does not move beyond this state over time, then I/O could be the source of the performance issue.

Javacores and Thread Dumps (cont.)

```
at java.net.SocketInputStream.socketRead0(Native Method)
at java.net.SocketInputStream.read(SocketInputStream.java:155)
at java.io.BufferedInputStream.fill(BufferedInputStream.java:229)
at java.io.BufferedInputStream.read(BufferedInputStream.java:246)
...
```

- Threads in “socketRead” for an extended period of time, mean that the work has left WebSphere and the thread will wait indefinitely for the work to return in order to continue.
- Check the database (slow? Huge query? Not available?)
- Check the network (Saturated? Blocked? Not available?)

- Other “I/O” related stack data may include similar:

```
“java.nio.channels.SocketChannel.write”
```

WebSphere Logs: “Too Many Open Files” message

- “IOException: Too many open files”
- Check the OS ulimits to see what the value is for “NOFILE”

```

User Limits (in bytes except for NOFILE and NPROC)
-----
type                soft limit      hard limit
RLIMIT_AS           unlimited      unlimited
RLIMIT_CORE         0              unlimited
RLIMIT_CPU          unlimited      unlimited
RLIMIT_DATA         unlimited      unlimited
RLIMIT_FSIZE        unlimited      unlimited
RLIMIT_LOCKS        unlimited      unlimited
RLIMIT_MEMLOCK      32768          32768
RLIMIT_NOFILE      8192          65536
RLIMIT_NPROC        1064960        1064960
RLIMIT_RSS          unlimited      unlimited
RLIMIT_STACK        10485760       unlimited

```

- General suggested value is 10000 for WebSphere Application Server although many stack products require this number to be higher.

Too Many Open Files error message

<http://www-01.ibm.com/support/docview.wss?uid=swg21067352>

Guidelines for setting ulimits


<http://www-01.ibm.com/support/docview.wss?uid=swg21469413>

Identify I/O Problems: Analysis Tools

- **WAIT- IBM Whole-system Analysis of Idle Time**
http://researcher.watson.ibm.com/researcher/view_project.php?id=1332
 - **TMDA – Thread and Monitor Dump Analyzer**
<https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=2245aa39-fa5c-4475-b891-14c205f7333c>
 - **Java Health Center**
<http://www.ibm.com/developerworks/java/jdk/tools/healthcenter/>
- ★ All above tools can be used with Windows as well as AIX, and Linux. Java Health Center is not available for Solaris.

Tools: WAIT

- IBM Whole-system Analysis of Idle Time (WAIT)
<https://wait.researchlabs.ibm.com/>
- Input multiple javacores from same process taken over the period of time where poor performance has occurred
- Output shows where threads are spending their time based on the tops of the thread stacks



Tools: WAIT (cont.)

- Looking at the “Delayed by Disk” threads, you can identify the code that may be delayed by I/O

Stack Viewer				
5 stack clusters, across all thread samples				
%	Wait State	Thread Pool	Work Unit	Thread Stack
33%	Delayed by Disk	HTTP Worker Threads	WebServicesService	<ul style="list-style-type: none"> File System Manipulation + java/io/ FileOutputStream.open() native method new java/io/ FileOutputStream() line 157 WebSphere Internals + com/ibm/wsspi/configarchive/ DefaultFileAccessor.getOutputStream() com/ibm/ws/sca/internal/admin/ SCACCommandHelper\$2.run() lin Security java/security/ AccessController.doPrivileged() line 192 WebSphere Internals + com/ibm/ws/sca/internal/admin/ SCACCommandHelper.discardSe com/ibm/ws/sca/internal/proxy/impl/ ProxyInvocationHandlerImpl.i Reflective Proxies \$Proxy313.processMessage() bytecode 18 WebSphere Internals + com/ibm/wsspi/sca/webservice/bean/impl/ WebServiceBeanImpl.pr com/ibm/ws/webservices/engine/ WebServicesEngine.invoke() lin

Tools:TMDA(Thread Monitor Data Analyzer)

<https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=2245aa39-fa5c-4475-b891-14c205f7333c>

WebContainer : 48

snapshot:

- (1)java/io/FileOutputStream.open(Native Method)
- (2)java/io/FileOutputStream.open(Native Method)
- (3)java/io/FileOutputStream.open(Native Method)

WebContainer : 52

snapshot:

- (1)com/ibm/ws/sca/internal/admin/SCACommandHelper.findChildComponent(SCACommandHelper.java:169)
- (2)java/io/UnixFileSystem.createDirectory(Native Method)
- (3)java/io/UnixFileSystem.createDirectory(Native Method)

Compare Threads : javacore.20120921.095256.4063688.0001.txt javacore.20120921.095302.4063688.0002.txt javacore.20120921.095307.4063688.0003.txt			
Thread ▲	javacore.20120921.095256.4063688.0001.txt [Sep 21 ...	javacore.20120921.095302.4063688.0002.txt [Sep 21 ...	javacore.20120921.095307.4063688.0003.txt [Sep 21 ...
WebContainer : 48	java/io/FileOutputStream.open(Native Method)	java/io/FileOutputStream.open(Native Method)	java/io/FileOutputStream.open(Native Method)
WebContainer : 49	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
WebContainer : 5	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
WebContainer : 50	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
WebContainer : 51	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
WebContainer : 52	com/ibm/ws/sca/internal/admin/SCACommandHelp...	java/io/UnixFileSystem.createDirectory(Native Method)	java/io/UnixFileSystem.createDirectory(Native Method)
WebContainer : 53	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...

TMDA output

Compare Threads : javacore.20120921.095256.4063688.0001.txt javacore.20120921.095302.4063688.0002.txt javacore.20120921.095307.4063688.0003.txt			
	javacore.20120921.095256.4063688.0001.txt [Sep 21 ...	javacore.20120921.095302.4063688.0002.txt [Sep 21 ...	javacore.20120921.095307.4063688.0003.txt [Sep 21 ...
1er: 48	⌘ java/io/FileOutputStream.open(Native Method)	⌘ java/io/FileOutputStream.open(Native Method)	⌘ java/io/FileOutputStream.open(Native Method)
1er: 49	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
1er: 5	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
1er: 50	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
1er: 51	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...
1er: 52	com/ibm/ws/sca/internal/admin/SCACommandHelp...	⌘ java/io/UnixFileSystem.createDirectory(Native Method)	⌘ java/io/UnixFileSystem.createDirectory(Native Method)
1er: 53	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...	com/ibm/ws/sca/internal/admin/SCACommandHelp...

- This tool is also available in ISA

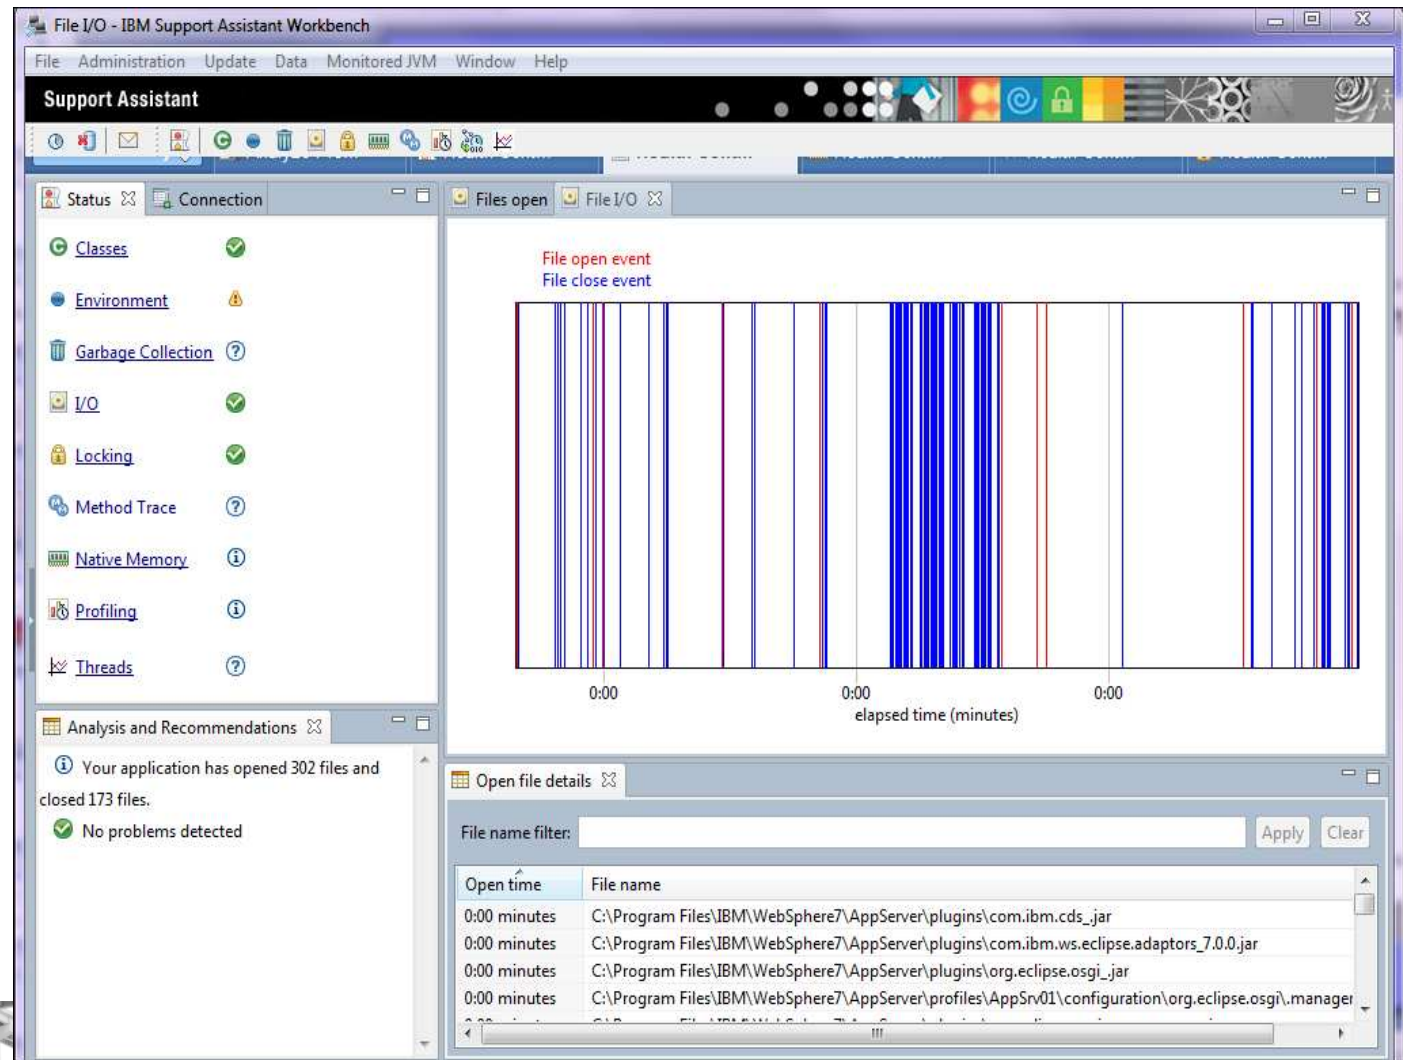
Download IBM Support Assistant (ISA) Lite for WebSphere Application Server

<http://www-01.ibm.com/support/docview.wss?uid=swg24020502>



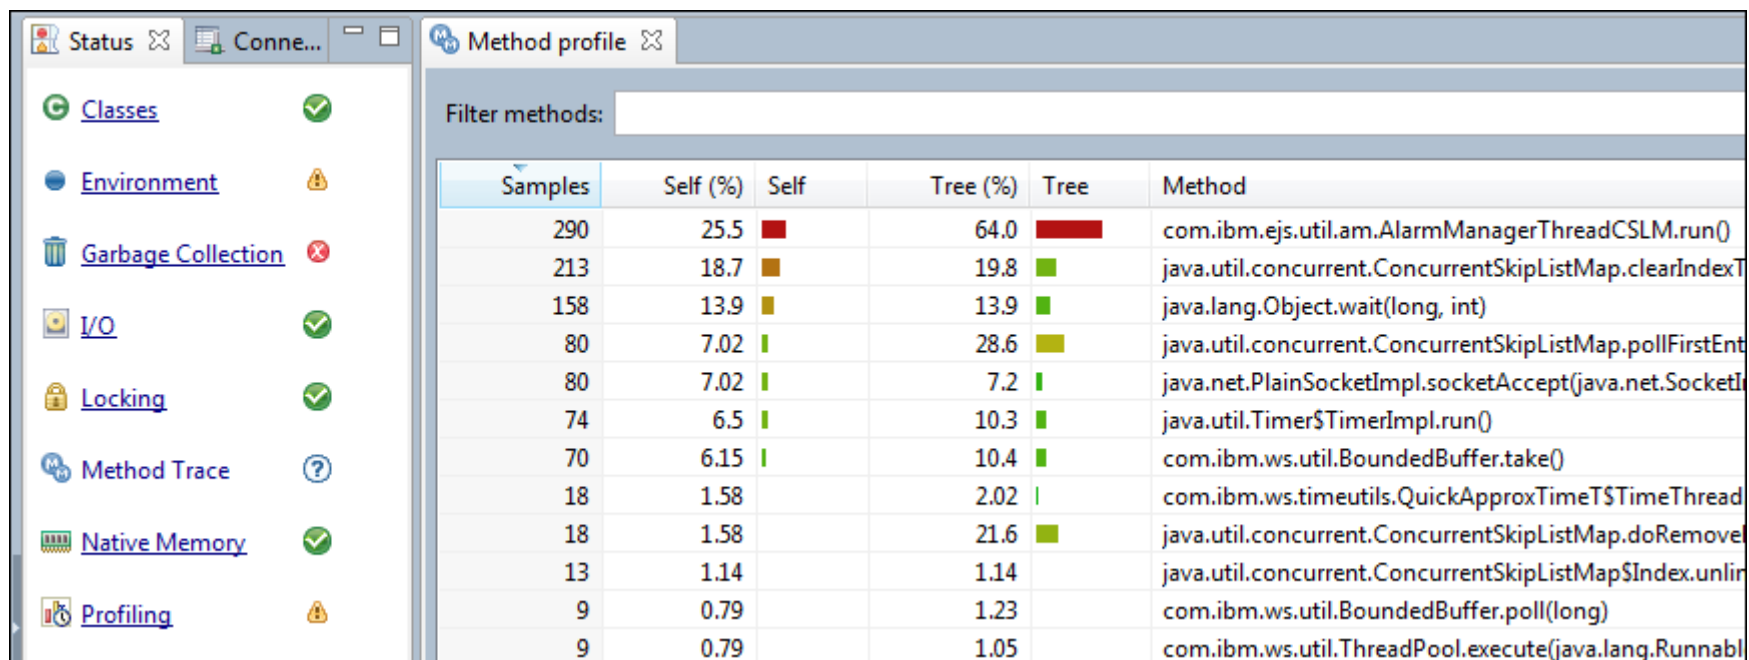
Tools: Health Center - I/O

- Java Health Center monitors I/O for a Java application
- Only for opening and closing of files
- May not be particularly useful for performance monitoring
- Health Center as a whole is useful for monitoring application health



Tools: Health Center - Profiling

- Also available in Java Health Center is Method Profiling.
- Although this is based on CPU usage, the pictorial output may easily highlight where I/O activity is contentious.



Samples	Self (%)	Self	Tree (%)	Tree	Method
290	25.5	■	64.0	■	com.ibm.ejs.util.am.AlarmManagerThreadCSLM.run()
213	18.7	■	19.8	■	java.util.concurrent.ConcurrentSkipListMap.clearIndexT
158	13.9	■	13.9	■	java.lang.Object.wait(long, int)
80	7.02	■	28.6	■	java.util.concurrent.ConcurrentSkipListMap.pollFirstEnt
80	7.02	■	7.2	■	java.net.PlainSocketImpl.socketAccept(java.net.SocketI
74	6.5	■	10.3	■	java.util.Timer\$TimerImpl.run()
70	6.15	■	10.4	■	com.ibm.ws.util.BoundedBuffer.take()
18	1.58	■	2.02	■	com.ibm.ws.timeutils.QuickApproxTimeT\$TimeThread
18	1.58	■	21.6	■	java.util.concurrent.ConcurrentSkipListMap.doRemove
13	1.14	■	1.14	■	java.util.concurrent.ConcurrentSkipListMap\$Index.unlin
9	0.79	■	1.23	■	com.ibm.ws.util.BoundedBuffer.poll(long)
9	0.79	■	1.05	■	com.ibm.ws.util.ThreadPool.execute(java.lang.Runnabl

**This is not an example of high CPU caused by I/O. Example above is to show what the method profile output looks like.

Next Step Tools

- Once the performance issue has been identified as related to I/O, the system administrator can use the following tools to inspect closer:

- ▶ **iostat**: Reports CPU statistics, AIO and input/output statistics for the entire system, adapters, TTY devices, disks CD-ROMs, tapes and file systems.

lostat Command:

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.commands/doc/aixcmds3/iostat.htm>

- ▶ **nmon**: Displays local system statistics in interactive mode and records system statistics in recording mode.

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.commands/doc/aixcmds4/nmon.htm>

- ▶ **mpstat**: Collects and displays performance statistics for all logical processors in the system

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.commands/doc/aixcmds3/mpstat.htm>

Avoid Performance Problems

WebSphere Application Server provides several tunable parameters and options to match the application server environment to the requirements of your application:

- Review the hardware and software prerequisites for WebSphere Application Server.
- Install the most current refresh pack, Fix Pack, and recommended interim fixes.
- Check the hardware configuration and settings.
- Review your application design.
- Tune the Solaris Operating System.
- Tune the Java Virtual Machine settings.
- Use a Type-4 (or pure Java) JDBC driver.
- Tune WebSphere Application Server JDBC resources and associated connection pools.
- Enable the Pass by Reference option, if applicable.
- Ensure that the transaction log is on a fast disk.
- Tune related components, such as data bases, messaging providers, and so on.
- Disable functions that are not needed.

Summary

- WebSphere Application Server, like any other application can only be as robust as the environment in which it resides
- Using the available tools, we can identify areas of contention and seek streamlined methods of resolution
- Sometimes we need to look beyond the application server and application to resolve performance bottlenecks

Resources

Assessing disk performance with the vmstat command

http://pic.dhe.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.prfungd/doc/prfungd/assess_disk_perf_vmstat.htm

Basic Monitoring of I/O on AIX

[http://www-](http://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/380d85576e3b900a8625755900805026/$FILE/Basic_Monitoring_of_IO_on_AIX.pdf)

[03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/380d85576e3b900a8625755900805026/\\$FILE/Basic_Monitoring_of_IO_on_AIX.pdf](http://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/380d85576e3b900a8625755900805026/$FILE/Basic_Monitoring_of_IO_on_AIX.pdf)

Checking hardware configuration and settings

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=compass&product=was-base-dist&topic=trpf_tunehdwcap

Disk IO Tuning in AIX 6.1

http://www.ibm.com/developerworks/wikis/download/attachments/104533522/AIX_Disk_IO_Tuning_093011.pdf

File system caching configurations

<http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/topic/com.ibm.db2.luw.admin.dboj.doc/doc/c0051304.html>

Guidelines for setting ulimits

<http://www-01.ibm.com/support/docview.wss?uid=swg21469413>

Memory usage determination with the vmstat command

http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.prfungd%2Fdoc%2Fprfungd%2Fmem_usage_vmstat.htm

Resources (cont.)

MustGather: Performance, hang, or high CPU issues with WebSphere Application Server on AIX

<http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg21052641>

[Optimizing](#) AIX 7 performance

<http://www.ibm.com/developerworks/aix/library/au-aix7optimize1/index.html>

Optimizing AIX 7 performance: Part 2, Monitoring logical volumes and analyzing the results

<http://www.ibm.com/developerworks/aix/library/au-aix7optimize2/index.html>

Platform 2015: Intel® Processor and Platform -Evolution for the Next Decade

http://epic.hpi.uni-potsdam.de/pub/Home/TrendsAndConceptsII2010/HW_Trends_borkar_2015.pdf

Reading vmstat in linux – Part 2

<http://www.helpmehost.com/tag/linux-performance/>

Too Many Open Files error message

<http://www-01.ibm.com/support/docview.wss?uid=swg21067352>

Under the Hood: Of POWER7 Processor Caches

http://www-03.ibm.com/systems/resources/systems_power_software_i_perfmgmt_underthehood.pdf

User space memory access from the Linux kernel

<http://www.ibm.com/developerworks/linux/library/l-kernel-memory-access/index.html>

vmstat Command

<http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.cmds%2Fdoc%2Faixcmds6%2Fvmstat.htm>

Tools

IBM Whole-system Analysis of Idle Time (WAIT)

http://researcher.watson.ibm.com/researcher/view_project.php?id=1332

IBM Thread and Monitor Dump Analyzer for Java

<https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=2245aa39-fa5c-4475-b891-14c205f7333c>

IBM Monitoring and Diagnostic Tools for Java - Health Center

<http://www.ibm.com/developerworks/java/jdk/tools/healthcenter/>

Contact Me

- Michele Calcavecchia
- Email:
macalcav@us.ibm.com
- On Connections:
https://w3-connections.ibm.com/profiles/html/profileView.do?userid=6f63d140-ea64-102b-87fa-a1760d15112d&lang=en_us
- WebSphere Application Server L2 Support on Connections:
<https://w3-connections.ibm.com/communities/service/html/communityview?communityUid=2e920240-407d-471c-8c0a-358f89ab3448>

Additional WebSphere Product Resources

- Learn about upcoming WebSphere Support Technical Exchange webcasts, and access previously recorded presentations at:
http://www.ibm.com/software/websphere/support/supp_tech.html
- Discover the latest trends in WebSphere Technology and implementation, participate in technically-focused briefings, webcasts and podcasts at:
<http://www.ibm.com/developerworks/websphere/community/>
- Join the Global WebSphere Community:
<http://www.websphereusergroup.org>
- Access key product show-me demos and tutorials by visiting IBM® Education Assistant:
<http://www.ibm.com/software/info/education/assistant>
- View a webcast replay with step-by-step instructions for using the Service Request (SR) tool for submitting problems electronically:
<http://www.ibm.com/software/websphere/support/d2w.html>
- Sign up to receive weekly technical My Notifications emails:
<http://www.ibm.com/software/support/einfo.html>

Connect with us!

1. Get notified on upcoming webcasts

Send an e-mail to wsehelp@us.ibm.com with subject line “wste subscribe” to get a list of mailing lists and to subscribe

2. Tell us what you want to learn

Send us suggestions for future topics or improvements about our webcasts to wsehelp@us.ibm.com

3. Be connected!

Connect with us on [Facebook](#)

Connect with us on [Twitter](#)

Questions and Answers

