Colin Henderson – Executive IT Architect, IBM Software Services for WebSphere
27 August 2009

AIX Tuning for IBM HTTP Server and WebSphere Application Server

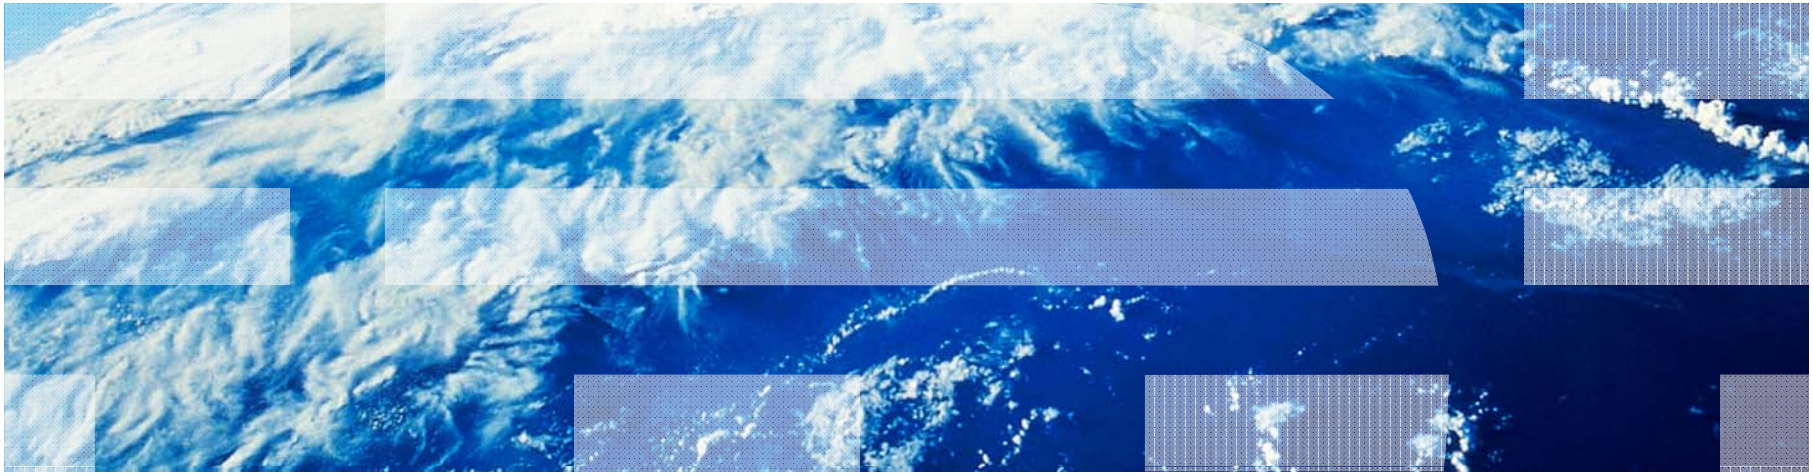

Agenda

Performance tuning considerations

General AIX tuning for a WebSphere infrastructure

AIX tuning for the IBM HTTP Server

AIX tuning for WebSphere Application Server

Scope of this presentation

What this presentation includes

This presentation will discuss AIX tuning as it relates to the best performance of the primary components of a WebSphere infrastructure, these include the IBM HTTP Server and the WebSphere Application Server versions 6.0, 6.1, and 7.0

This presentation will discuss only enough of the architecture and internal workings of the IBM HTTP Server and WebSphere Application Server to put the AIX tuning in context.

What this presentation does not include

This presentation does not address:

- Tuning of either the IBM HTTP Server or WebSphere Application Server
- Tuning of external network components
- Tuning of AIX supporting backend components
- Tuning of backend components or stacked products (Portal, Process Server, … etc.)

General Considerations for Tuning a WebSphere Infrastructure

- Tuning a WebSphere infrastructure should be a holistic process. Individual elements or components of the infrastructure should not be tuned in isolation.
- Rarely is a performance problem the result of a single component “There is always a next bottleneck”.
- Be wary of performance tuning advice from the Internet … it may be dated.
- The good news: most of the AIX defaults work well with IHS and WebSphere App Server.
- One size never fits all. Although the standard defaults may provide satisfactory service levels in 95% of cases there will always be 5% of installations where tuning is required to achieve satisfactory service levels for users.
- Consider the worst case scenario. Anticipate events that will stress the environment. These may include peak times of day, week, month or year, and special events such as marketing campaigns or sporting events

The WebSphere "Gold Standard" Infrastructure

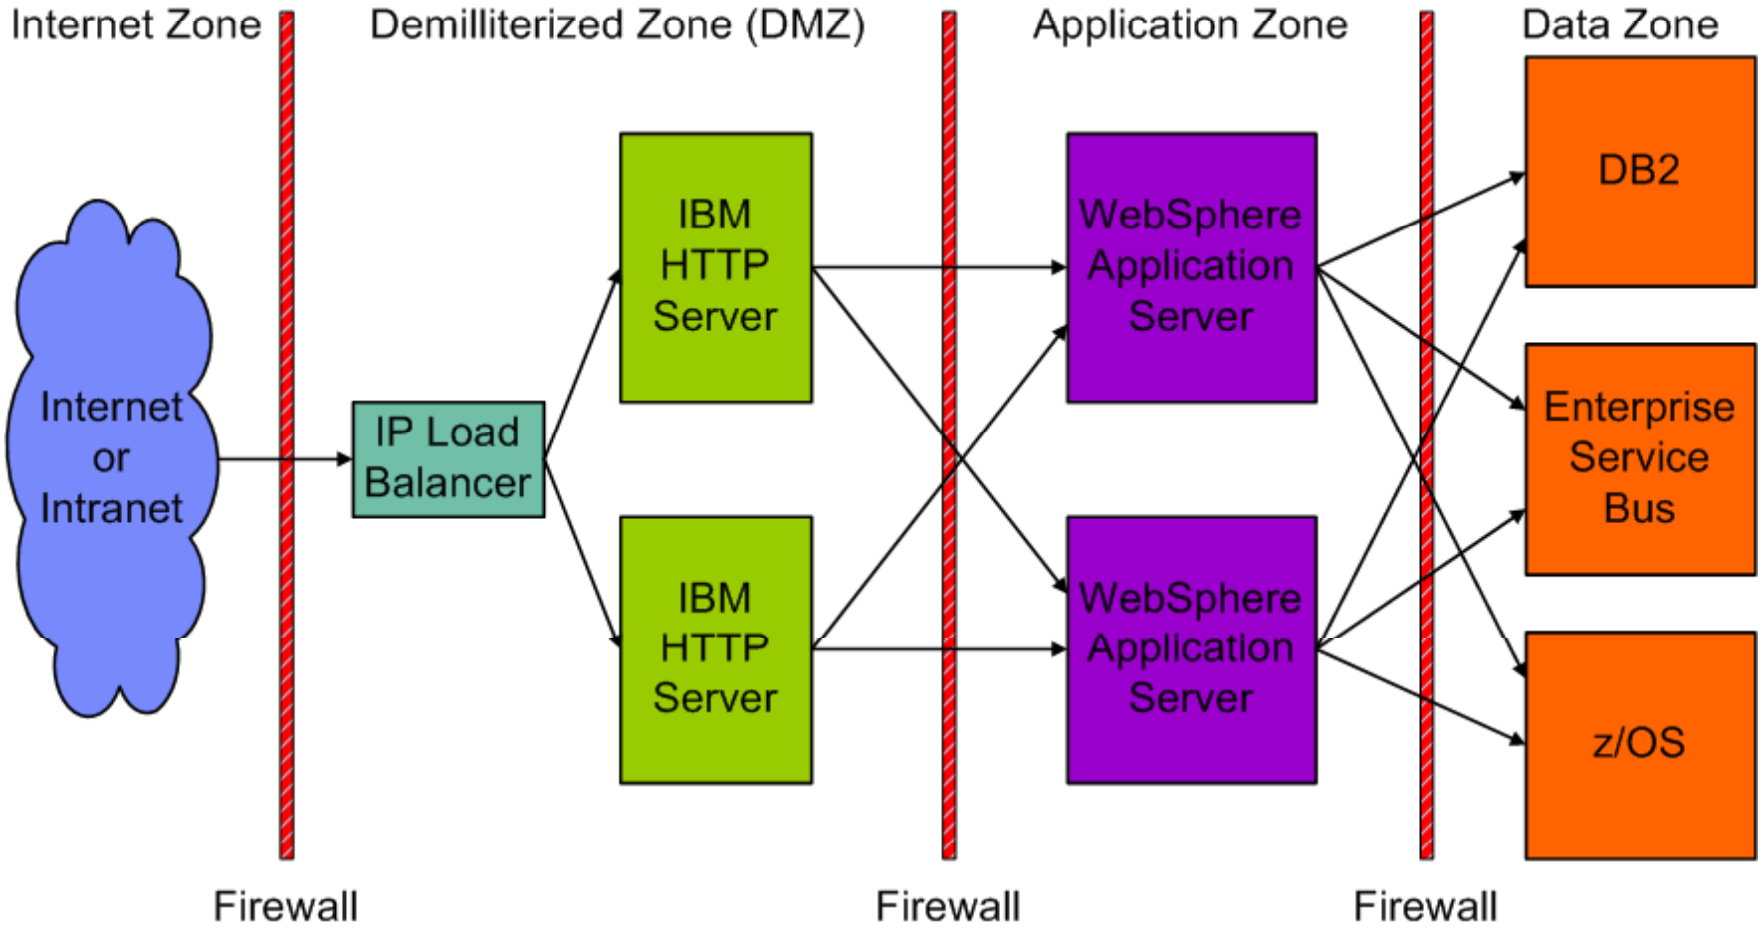

General AIX tuning for a WebSphere infrastructure - SMT

Enable Simultaneous Multi-Threading (SMT)

Both IBM HTTP Server and WebSphere Application Server are highly multi-threaded and can benefit significantly from Simultaneous Multi-Threading, up to 40% increased throughput on Power5 systems and 50% increased throughput on Power6.

Note: Early implementations of Java – prior to Java V1.2 – used “green threads”. These threads were emulated threads that all ran on a single OS thread. Some individuals still have the misconception that Java runs on a single OS thread and therefore disable SMT. This has not been true since Java V1.2 was introduced in 1999.

General AIX tuning for a WebSphere infrastructure – TCP/IP

tcp_timewait – The time between connection close and release of resources

Default: 5 (75 seconds – 5 x 15 seconds)
Recommendation: 1 (15 seconds)

tcp_finwait2 – The time to wait before closing a connection in the FIN_WAIT_2 state

Default: 1200 (10 minutes – 1200 half seconds)
Recommendation: 120 (1 minute)

Example from a live system:

CLOSE_WAIT	4
CLOSING	56
ESTABLISHED	224
FIN_WAIT_1	13
FIN_WAIT_2	46
LAST_ACK	4
TIME_WAIT	73

TCP/IP State Transitions

2 Page Pocket Guide available from W. Richard Stevens Web site, refer to the references at the end of the presentation for the link

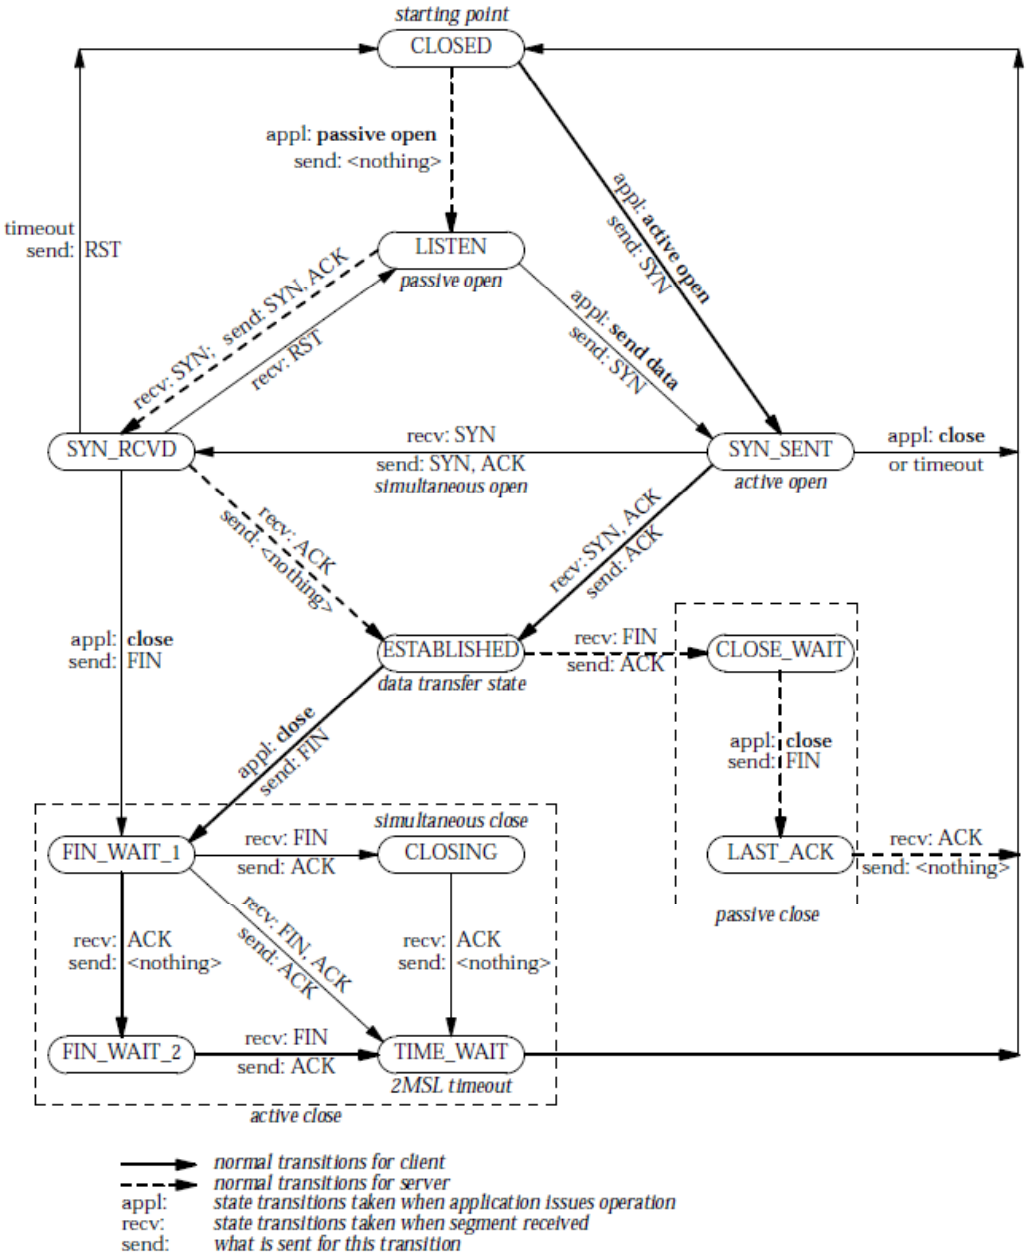

General AIX tuning for a WebSphere infrastructure – TCP/IP (con't)

tcp_keepidle – The time to keep an idle connection in the ESTABLISHED state

Default: 14400 (2 hours – 14400 half seconds)
Recommendation: 600 (5 minutes)

tcp_keepintvl – The time between keepalive packets (validates a connection)

Default: 150 (75 seconds – 150 half seconds)
Recommendation: 10 (5 seconds)

tcp_keepinit – The TCP initial connection timeout (allows failover to take place quickly)

Default: 150 (75 seconds – 150 half seconds)
Recommendation: 10 (5 seconds)

Web page network timing for www.time.gov

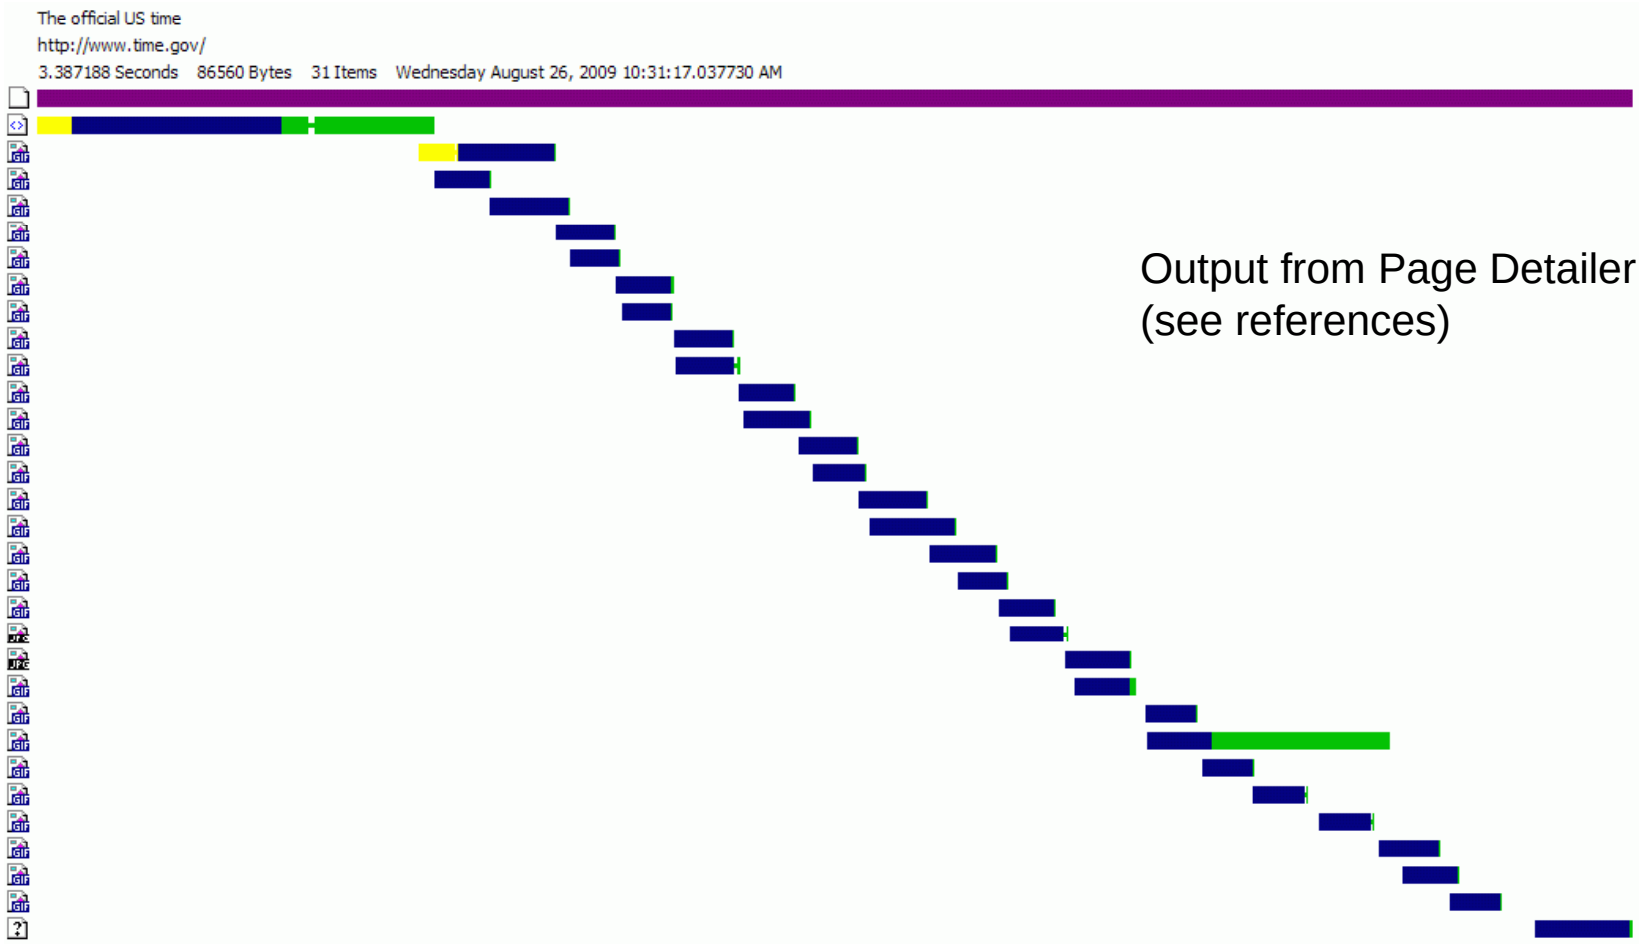

Output from Page Detailer
(see references)

General AIX tuning for a WebSphere infrastructure – ulimits

Set the Hard limits to unlimited

In /etc/security/limits:

```
core_hard = -1
cpu_hard = -1
data_hard = -1
fsize_hard = -1
stack_hard = -1
nofiles_hard = -1
```

Set the default Soft limits to unlimited for AT LEAST the WebSphere userid(s)

In /etc/security/limits:

```
core = -1
cpu = -1
data = -1
fsize = -1
stack = -1
nofiles = -1
```

General AIX tuning for a WebSphere infrastructure – Full Duplex

Network interfaces should always be set in Full Duplex mode.

I highly recommend that network interfaces be manually configured for the correct speed to match their peer network components and configured for Full Duplex. Network interfaces running in Half Duplex mode will severely impact server throughput.

General AIX tuning for a WebSphere infrastructure – cipherspecs

Cipherspecs ARE NOT an AIX tunable, however, if SSL connections are heavily used the cipherspecs used can SEVERELY impact CPU utilization. The following information is provided for your awareness in the event that you experience high CPU utilization with moderate to high numbers of connections

When SSL connections are established the client and server negotiate the cipherspec to use for traffic encryption. The server will select the strongest cipher that is supported by both the client and the server. By default triple-DES (3DES) is the strongest cipher supported by the IBM HTTP Server. Since 3DES is almost universally supported by client browsers if it is supported on the server then it will be the cipher chosen.

Unless 3DES cipher strength is required and CPU utilization is not a concern I recommend removing 3DES from the server list of available ciphers.

See: “IBM HTTP Server Performance Tuning” in the references for a more complete discussion of ciphers.

The request queuing funnel

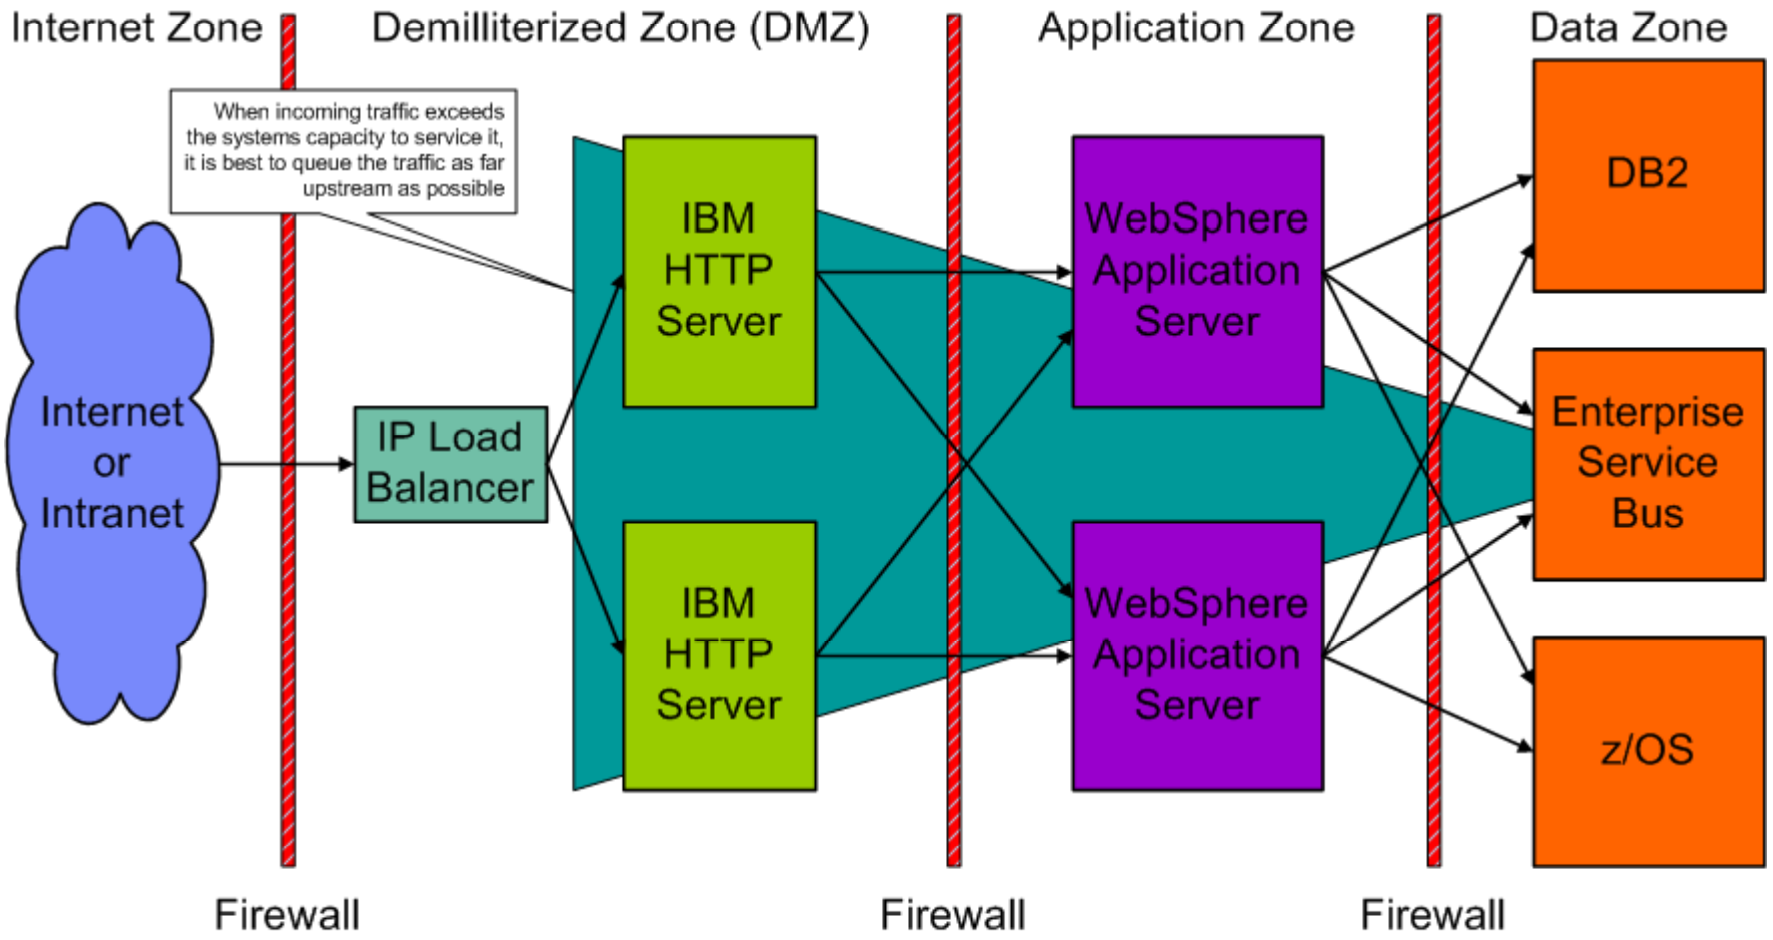

Architecture of the IBM HTTP Server

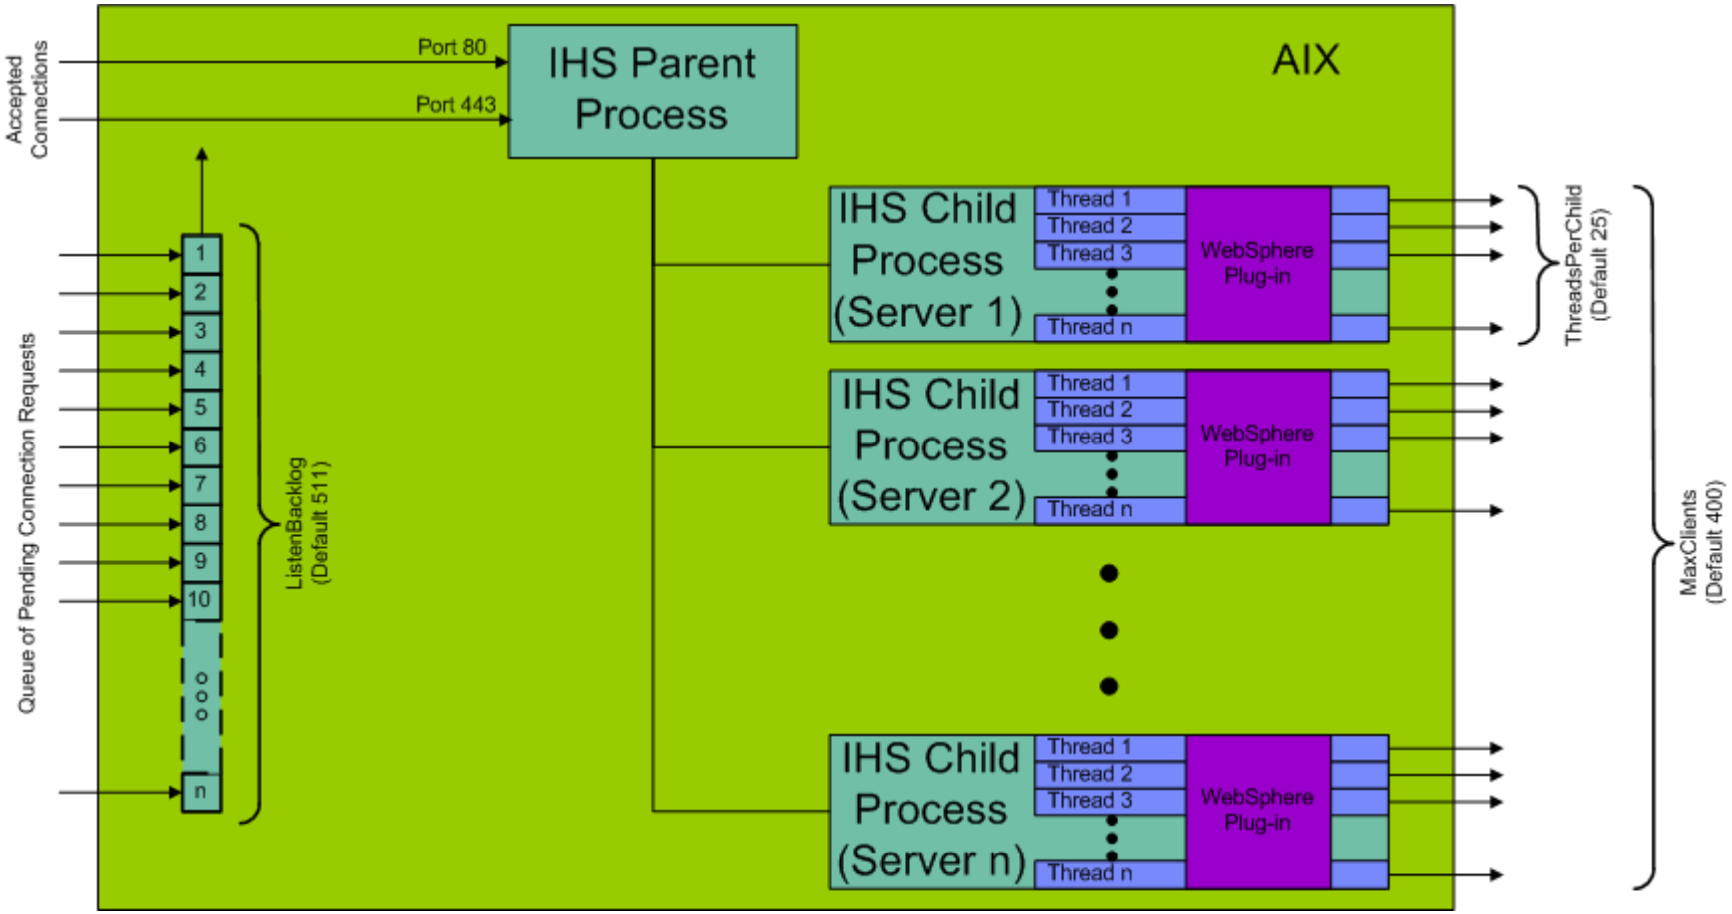

AIX tuning for the IBM HTTP Server – connection backlog

somaxconn – the maximum number of connections in a connection request queue

Default: 1024
Recommend: 4096 or larger

The actual size of the IHS connection backlog is determined by the IHS default or the ListenBacklog directive in the IHS configuration file. During peaks or bursts of inbound Web traffic we want to queue as much of the traffic at the Web server as possible rather than flood the downstream servers with too much traffic. If somaxconn is set to a reasonably large number then IHS can easily be reconfigured to use a larger connection request (ListenBacklog) queue. Otherwise, somaxconn can only be altered temporarily and a reboot will be required to increase the maximum queue size permanently.

AIX tuning for the IBM HTTP Server – nodelayack

tcp_nodelayack – disable delay of acknowledgements

Default: 0
Recommend: 1

TCP deliberately delays sending acknowledgements by up to 200ms to try an piggyback packet acknowledgements with response packet. For large POST requests (multi-megabyte) this can cause low data transfer rates due to 150-200ms delays between packets.

AIX tuning for WebSphere Application Server
Typical JVM Architecture

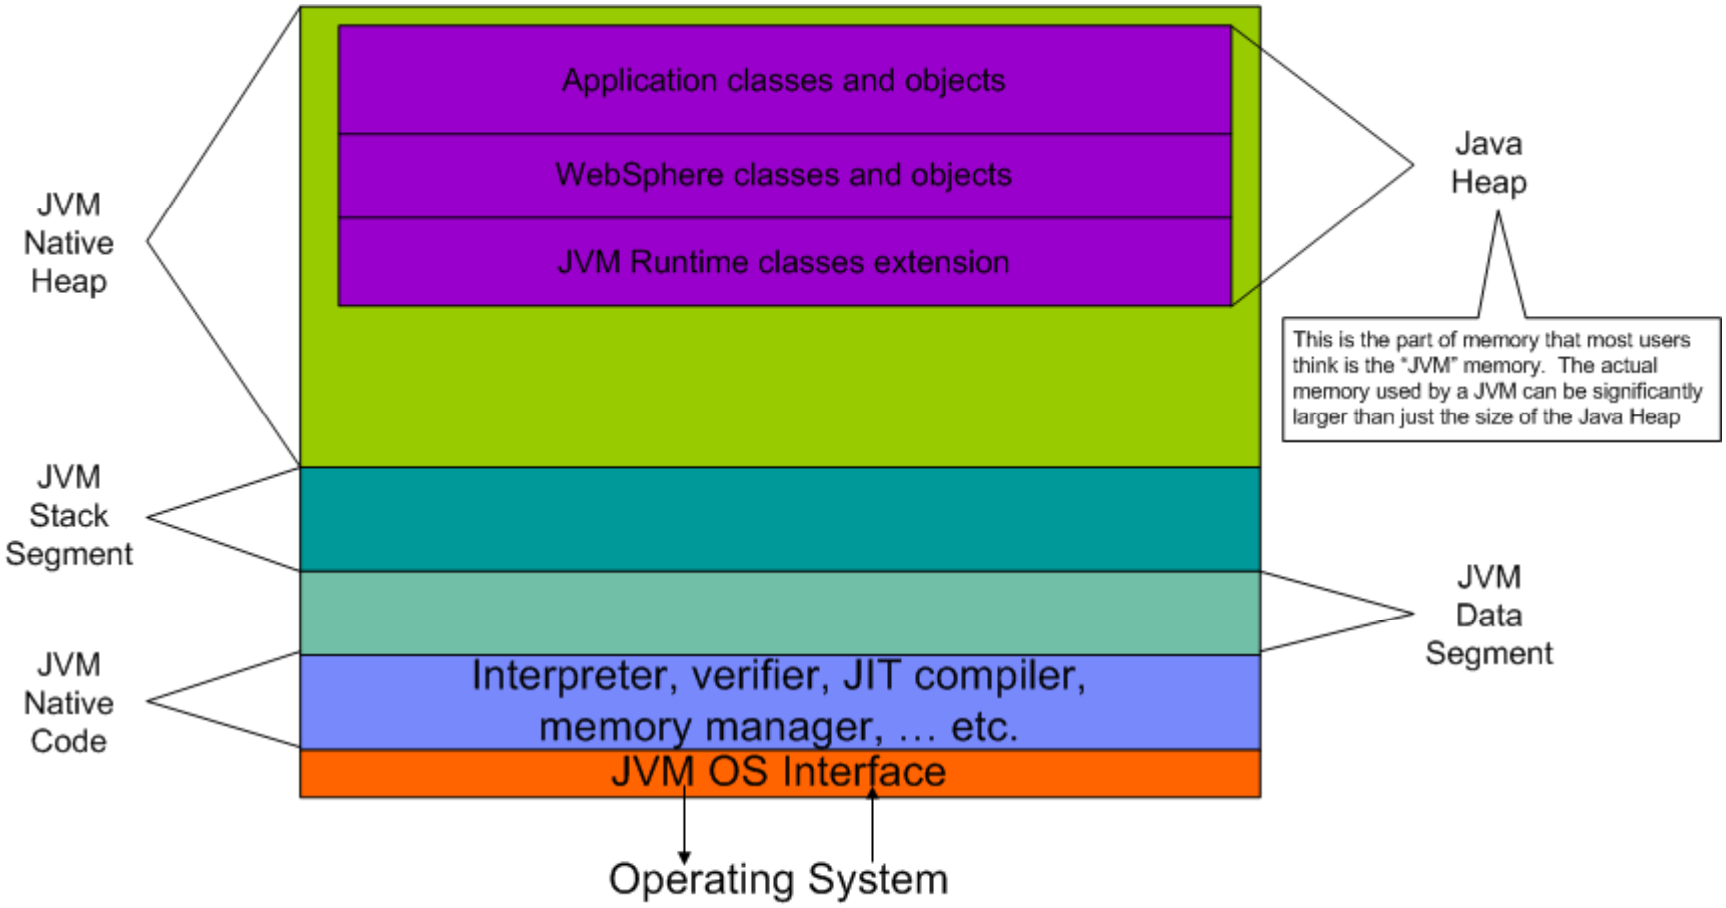

AIX tuning for WebSphere Application Server - memory

Apart from the tuning which has been discussed thus far which is common to both IHS and WebSphere Application Server, the most important consideration for WebSphere performance is memory. Java performance, and therefore WebSphere performance is critically dependant on JVM memory being resident and not paged.

Java users tend to think that the Java Heap equates to JVM memory. As the diagram on the previous page showed there is significant memory used by the JVM outside of the Java Heap. When planning a WebSphere infrastructure it is important to account for all of the memory that is required by all JVMs on a server plus other system support processes to ensure that memory is not overcommitted thus forcing JVM paging to occur.

The following data is taken from a live production WebSphere system on AIX:

Java Heap : 256MB (determined using Tivoli Performance Viewer)

Actual Memory Resident Set Size (RSS): 550MB (determined using ps aux)

AIX tuning for WebSphere Application Server – connection backlog

somaxconn – the maximum number of connections in a connection request queue

Default: 1024
Recommend: 128

The actual size of the WebSphere connection request queue is determined by the TCP Transport Channel default (511) or the setting of the TCP Transport Channel custom property listenBacklog. Reducing the size of the connection request queue improves failover in the event of an application server failure. This of course assumes that there are more than one application servers that inbound requests can failover to.

Since listenBacklog is a custom property it is often overlooked (usually overlooked) and the connection request queue therefore defaults to a size of 511. By setting somaxconn at the OS level this limits the maximum connection request queue size regardless of the default or explicit setting of listenBacklog.

References

- IBM alphaWorks: http://alphaworks.ibm.com
- IBM DeveloperWorks: http://www.ibm.com/developerworks/
- WebSphere V6.0.2 InfoCenter: Tuning AIX Systems
  http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/topic/com.ibm.websphere.nd.doc/info/ae/ae/tprf_tuneaix.html
- WebSphere V6.1 InfoCenter: Tuning AIX Systems
  http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/topic/com.ibm.websphere.nd.doc/info/ae/ae/tprf_tuneaix.html
- WebSphere V7.0 InfoCenter: Tuning AIX Systems
  http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.nd.doc/info/ae/ae/tprf_tuneaix.html
- *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, 1994, ISBN 0-201-63346-9
  Pocket Guide: http://www.kohala.com/start/pocketguide1.ps
- IBM HTTP Server Performance Tuning:
  http://publib.boulder.ibm.com/httpserv/ihsdiag/
- IHSDiag (Diagnosis tools and documentation for IBM HTTP Server)
  http://www-01.ibm.com/support/docview.wss?uid=swg24008409

References (cont'd)

- Apache HTTP Server Online Documentation: http://httpd.apache.org/
- Page Detailer (browser traffic monitoring tool): http://www.alphaworks.ibm.com/tech/pagedetailer
- Crash on AIX produces no core or a truncated core: http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg21052642

Contact Information

Colin Henderson
Executive IT Architect
Software Services for WebSphere
IBM Software Group

Office Phone: (919) 654-6378
E-mail: cjhender@us.ibm.com