

WebSpeed vs REST Adapter Benchmarks, Statistics, and other BS

Dustin Grau, Software Architect



BRAVEPOINT

5000 Peachtree Ind. Blvd.

Suite 100

Norcross, GA 30071

- Senior developer and consultant at BravePoint, Inc.
 - Founded in 1987 with currently ~125 employees
 - Consulting, training, and placement services
- WebSpeed application developer since 1999
- Implementing JS/AJAX frameworks since 2010
- Lead architect for modernization framework “Application Evolution”

Is it worth switching to REST for performance gains?

This is the question we always get asked by anyone looking at OE 11.
Poll: How many people are using the new REST adapter?

- Technology introduction: the OE REST Adapter
- A primer on JMeter: configuration and running
- Baseline configs: WebSpeed (Apache) and OE Rest (Tomcat)
- Establishing thresholds and parameters for testing
- Benchmarking and results
- Crunching the numbers

- Built-in feature, similar to the Web Services Adapter
- Included since OpenEdge version 11.1 (or thereabout)
- Requires a license for the AppServer component
- When you need an open web interface other than SOAP/XML
- Uses Tomcat as a bridge between HTTP and an AppServer
- Should be serving many small requests via client API requests
- We want to see how this compares to home-grown WS/CGI



- Open-source, Java-based load testing suite
- Can be obtained at <http://jmeter.apache.org>
- Similar to ApacheBench (ab), Siege, and others
- Includes both benchmarking and reporting tools
- Useful for building complex/looping test cases
- Let us explore how this works...



The screenshot displays the Apache JMeter GUI. On the left, a tree view shows a test plan named 'OE REST vs WebSpeed'. It contains two main user groups: 'OE REST Users' and 'WebSpeed Users'. 'OE REST Users' includes a 'Simple Controller' with sub-elements: 'Tomcat 10x10', 'REST 5x5', 'REST 10x10', 'REST 20x20', and 'REST 40x40'. 'WebSpeed Users' includes a 'Simple Controller' with sub-elements: 'Apache 10x10', 'WSCGI 5x5', 'WSCGI 10x10', 'WSCGI 20x20', and 'WSCGI 40x40'. Both user groups also have 'View Results in Table', 'View Results Tree', and 'Response Time' elements. At the bottom of the tree are 'Summary Report', 'Aggregate Report', and 'WorkBench'.

The right pane shows the configuration for the selected 'Thread Group' 'OE REST Users'. The 'Name' field is 'OE REST Users'. The 'Comments' field is empty. Under 'Action to be taken after a Sampler error', the 'Continue' radio button is selected. The 'Thread Properties' section includes: 'Number of Threads (users):' set to 1000, 'Ramp-Up Period (in seconds):' set to 20, 'Loop Count:' with 'Forever' unchecked and '1' entered, and checkboxes for 'Delay Thread creation until needed' and 'Scheduler', both of which are unchecked.

7

All testing begins with "threads" which defines the load you'll be placing on your server.

The screenshot shows the Apache JMeter interface with the 'HTTP Request Defaults' configuration window open. The left sidebar shows a test plan named 'OE REST vs WebSpeed' with two user groups: 'OE REST Users' and 'WebSpeed Users'. The 'HTTP Request Defaults' window is configured with the following settings:

- Name:** HTTP Request Defaults
- Comments:** (empty)
- Web Server:**
 - Server Name or IP: localhost
 - Port Number: 8980
 - Timeouts (milliseconds): Connect: (empty), Response: (empty)
- HTTP Request:**
 - Implementation: (dropdown menu)
 - Protocol [http]: (empty)
 - Content encoding: (empty)
 - Path: (empty)
- Parameters:**

Send Parameters With the Request:			
Name:	Value	Encode?	Include Equals?
msdelay	0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
- Proxy Server:**
 - Server Name or IP: (empty)
 - Port Number: (empty)
 - Username: (empty)
 - Password: (empty)
- Embedded Resources from HTML Files:**
 - Retrieve All Embedded Resources
 - Use concurrent pool. Size: 4
 - URLs must match: (empty)

By assigning HTTP defaults, you make your test scenarios easier to manage.

HTTP Header Manager

Name: HTTP Header Manager

Comments:

Name	Value
Accept-Encoding	gzip, deflate

Buttons: Add, Add from Clipboard, Delete, Load, Save

Headers can include things like REALM authentication, or as seen here compression settings.

The screenshot displays the Apache JMeter GUI. On the left is a tree view of the test plan, with 'REST 10x10' selected under 'WebSpeed Users'. The main panel shows the configuration for an 'HTTP Request'.

HTTP Request Configuration:

- Name: REST 10x10
- Comments: 4kb Sample Size (Dynamic)
- Web Server: Server Name or IP: [], Port Number: [], Timeouts (milliseconds): Connect: [], Response: []
- HTTP Request: Implementation: [], Protocol [http]: [], Method: GET, Content encoding: []
- Path: /DemoPublic/rest/public/benchmark
- Options: Redirect Automatically, Follow Redirects, Use KeepAlive, Use multipartform-data for POST, Browser-compatible headers
- Parameters (Body Data):

Name	Value	Encode?	Include Equals?
numFields	10	<input type="checkbox"/>	<input checked="" type="checkbox"/>
numRecords	10	<input type="checkbox"/>	<input checked="" type="checkbox"/>
- Send Files With the Request: File Path: [], Parameter Name: [], MIME Type: []
- Proxy Server: Server Name or IP: [], Port Number: [], Username: [], Password: []
- Embedded Resources from HTML Files: Retrieve All Embedded Resources, Use concurrent pool. Size: 4, URLs must match: []
- Source address: IP/Hostname: []
- Optional Tasks: Use as Monitor, Save response as MD5 hash?

Test scenarios can include URL parameters and a specific URL.

- Mid-2012 MacBook Pro (Quad-Core i7, 16GB)
- VMWare running Windows 7 32-bit
 - 2-CPU with 2 GB Memory
 - All requests go to localhost
 - Rebooted VM between tests
- Static test is “10x10” JSON object
 - 10 rows of data (temp-table)
 - 10 fields of consistent size
 - 4.4 kilobytes (4,481 bytes)



Your mileage will always vary, and is about as ridiculous as assigning MPG on an electric car...

- 2000 threads (users) ramped up over 10 seconds, 4kb static file
- Default Configurations
 - Apache: 48 requests/sec @ 220 KB/sec, 0% errors
 - Tomcat: 52 requests/sec @ 240 KB/sec, 0% errors
- Compression Setup
 - Apache: 4kb -> 377 bytes (gzip, mod_deflate, level 1)
 - Tomcat: 4kb -> 353 bytes (gzip, as-is)
- Optimized Configurations
 - Apache: 160 requests/sec @ 117 KB/sec, 0% errors
 - Tomcat: 191 requests/sec @ 123 KB/sec, 0% errors



You should always be using compression. Period.

- What are we testing?
 - Dynamic data packets, created by WebSpeed or AppServer
- CPU and memory concerns?
 - I'm using a VM; CPU and memory could be increased
 - Expect different results with different hardware (surprise!)
- I/O, HDD vs SSD?
 - Not a factor, all data being generated never touches disk
 - No attached databases, but we are logging (minimal info)

- Using 10 agents (min/max/initial) for WebSpeed and AppServer
- WebSpeed is not using Nameserver; directly to port (cgiip)
- Same for AppServer, using AppserverDC in runtime.props
- No special performance tuning (wsbroker1, restbroker1)
- No changes to default Tomcat memory
- No further tweaks to Apache workers
- Only making GET requests

- Same configuration (2000 users, 10 seconds), 4kb dynamic packet
 - WebSpeed: 85 requests/sec @ 153 KB/sec, 82% errors
 - AppServer: 87 requests/sec @ 161 KB/sec, 84% errors
- Looking at broker statistics (maximums)
 - WS: Queue Depth 0; Req. Wait 2ms; Req Duration 2,726ms
 - AS: Queue Depth 30; Req Wait 1,843ms; Req Duration 2,230ms
- Errors indicate refused connections, HTTP-500, or other problems
- We want to test near the limits of our system, not exceed them...

- New configuration (1000 users, 10 seconds), 4kb dynamic packet
 - WebSpeed: 67 requests/sec @ 67 KB/sec, 36% errors
 - AppServer: 80 requests/sec @ 86 KB/sec, 30% errors
- Looking at broker statistics (maximums)
 - WS: Queue Depth 8; Req Wait 618ms; Req Duration 722ms
 - AS: Queue Depth 8; Req Wait 178ms; Req Duration 1,774ms
- Errors have been reduced but the load on the brokers is still high
- So it's better, but let's try this again...


- Final configuration (1000 users, 20 seconds), 4kb dynamic packet
 - WebSpeed: 49 requests/sec @ 31 KB/sec, 0% errors
 - AppServer: 51 requests/sec @ 32 KB/sec, 0% errors
- Looking at broker statistics (maximums)
 - WS: Queue Depth 3; Req Wait 307ms; Req Duration 347ms
 - AS: Queue Depth 2; Req Wait 46ms; Req Duration 1,405ms
- Lesson: the HTTP servers can keep up better than the brokers
- We have a threshold, so that's good enough to continue!

- Ramping up data packets (file size)
 - 5x5 = 1kb (ideal)
 - 10x10 = 4kb (common)
 - 20x20 = 18kb (large)
 - 40x40 = 70kb (object)

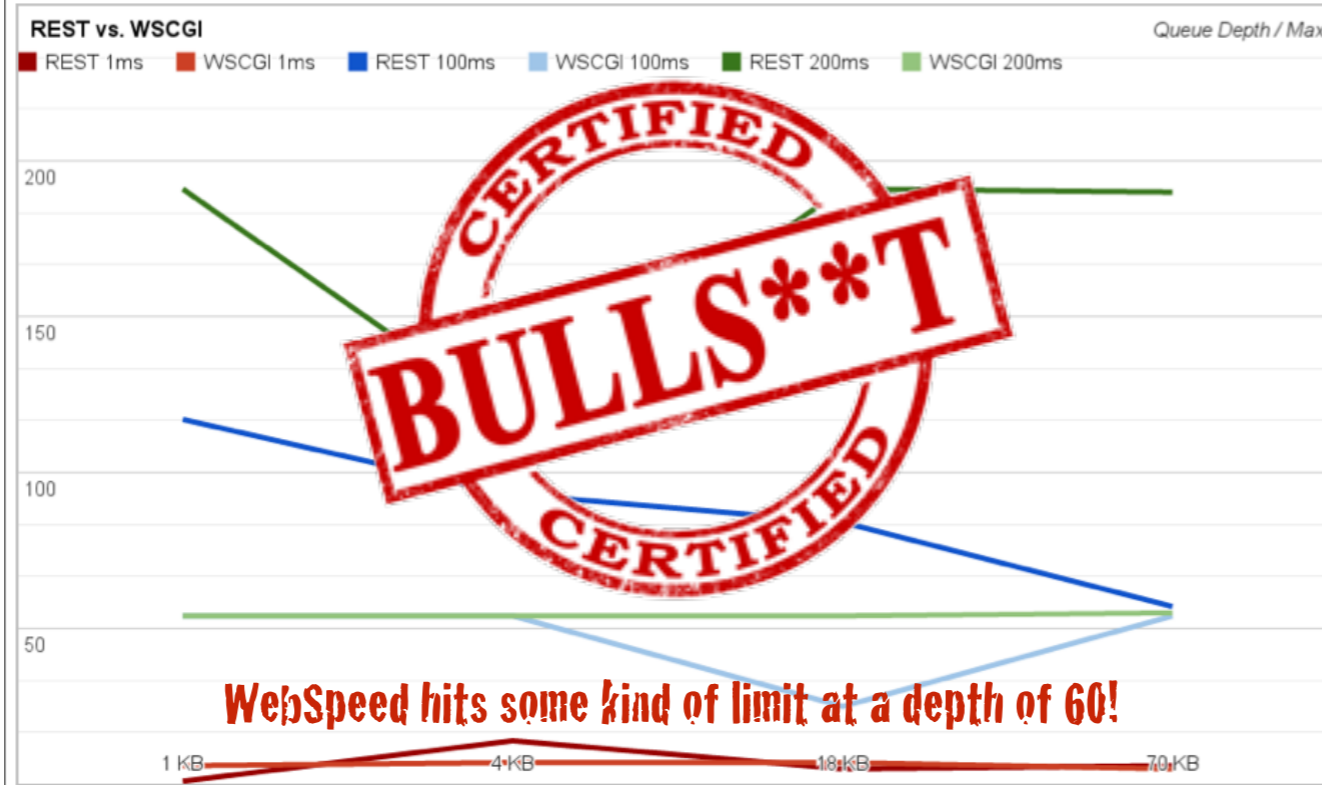
- Ramping up processing time (latency)
 - 0-5ms (base tests)
 - 100ms (typical case)
 - 200ms (heavy query)
 - Using sub-second pause




We can't live on averages alone, because the actual test results show more information.




We're still working with averages, just from the Progress perspective.




We may be limited by some kind of process limit in Apache, so this can't be taken as truth.




Theory: Tomcat needs to ramp up memory to handle the initial load, but afterwards can process and clear the backlog.
WebSpeed manages to keep up with the load relatively well for small packet sizes.




Nearly the same view as before, just overall longer response times.




Now we're seeing some results. As packet size increases, the playing field is more level.




Large packet sizes are no good for anyone over extended periods of time.




Starting over with our small packet, but taking longer to process on the back-end.
This is the same pattern we saw earlier, just with slightly longer response times.




Again, once Tomcat has ramped up memory and threads, it can handle the load better. WebSpeed begins to show it's limits, probably related to that queue depth limit we saw earlier.



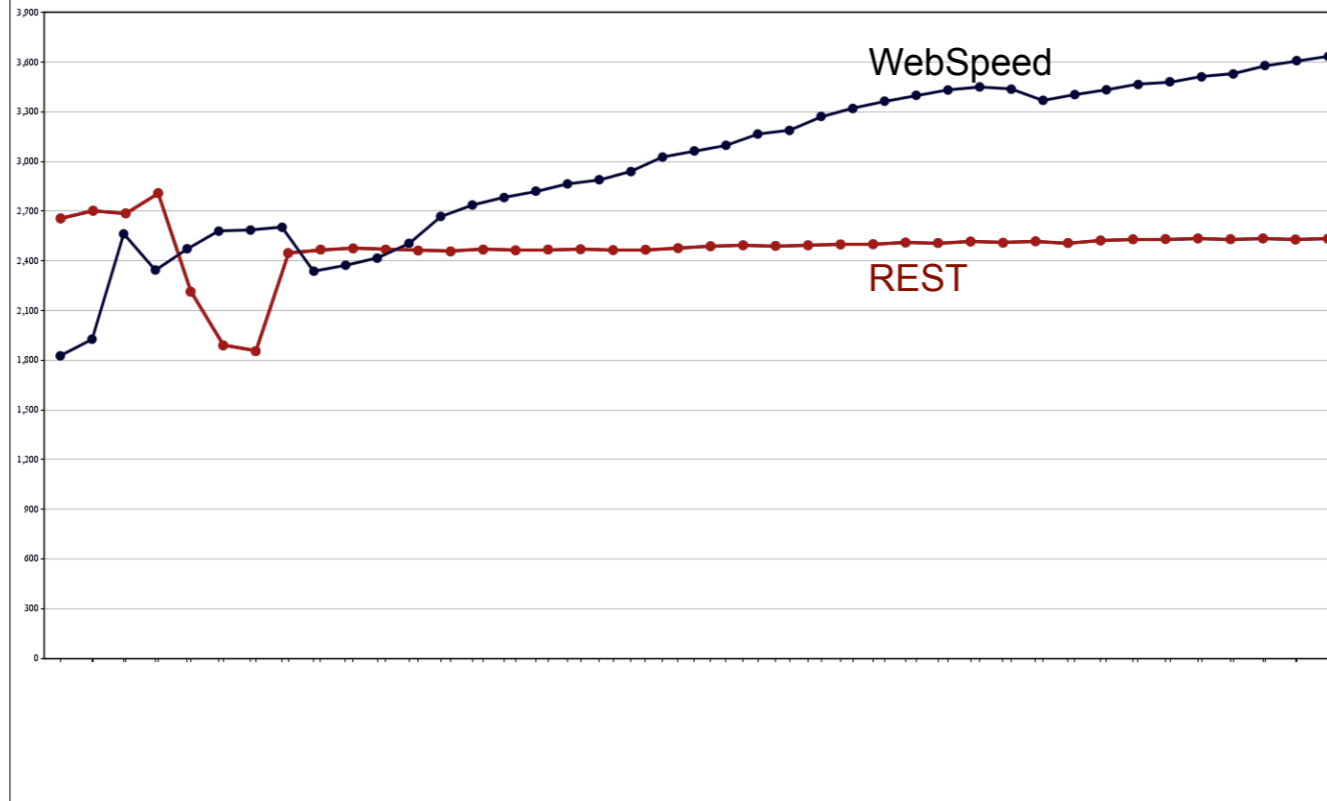
The response times for WebSpeed are stil consistent with the last test, further proving a possible limit. Meanwhile Tomcat is just taking slightly longer to respond with data.



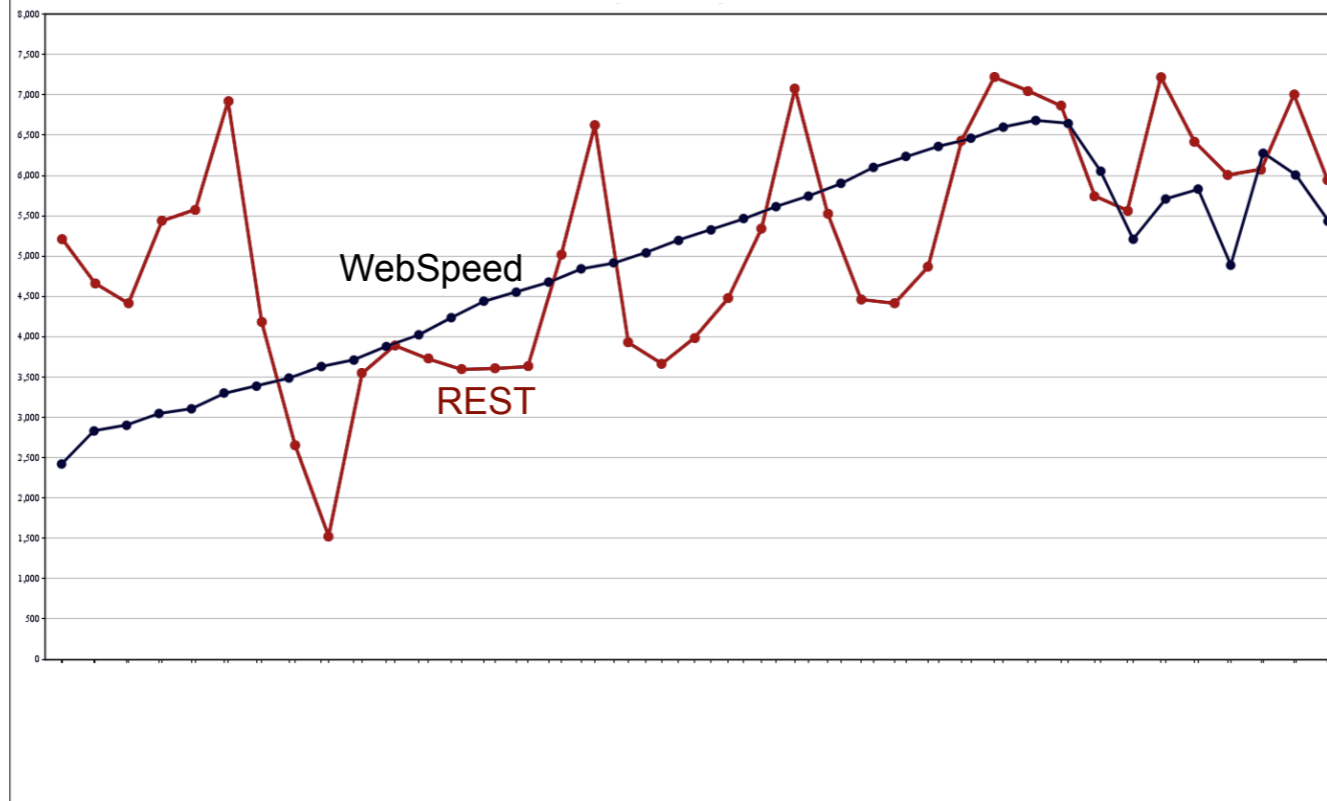
And we're back to our largest packet size with a moderate delay on the server. Nobody likes this much.



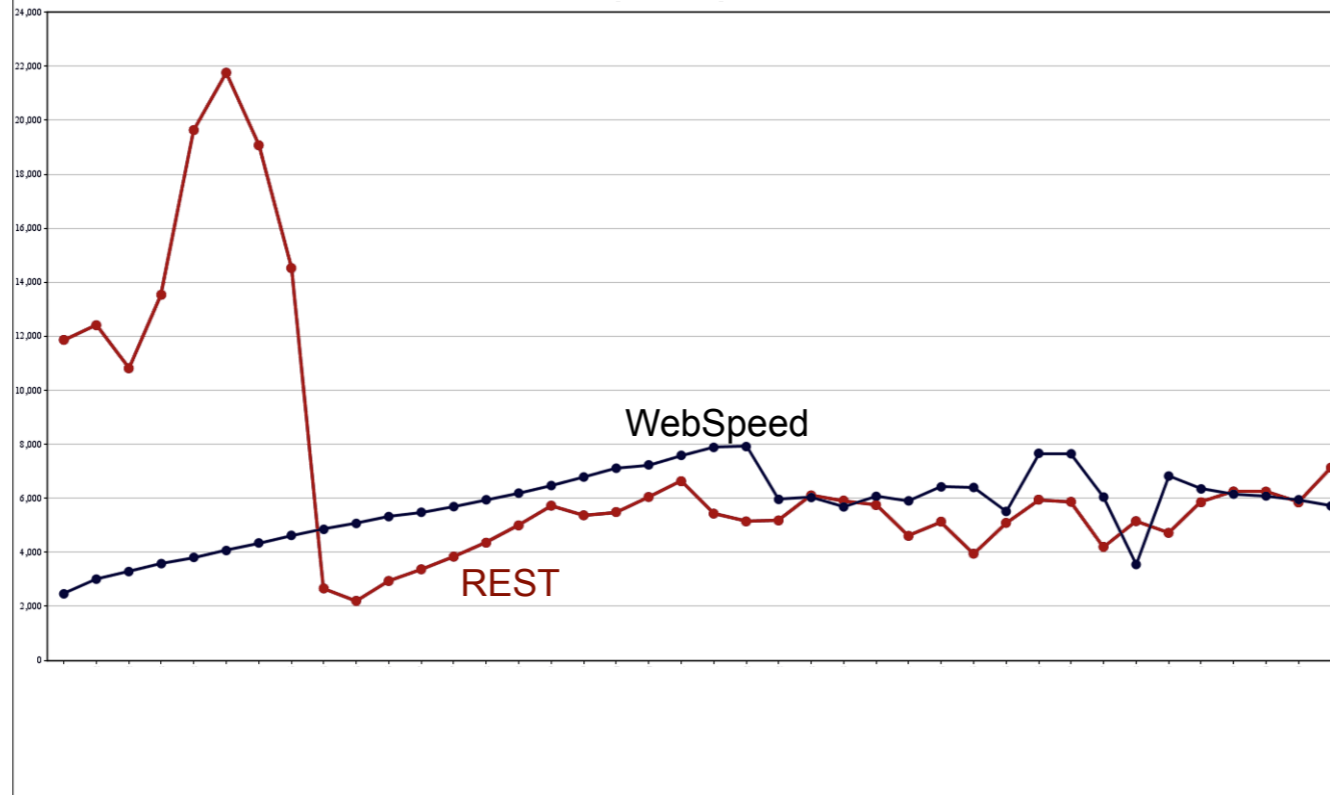
Interesting. Increasing the time to return data from the brokers means little to your HTTP daemon.
At this point the bottleneck is at the agents, which are busy. Requests are waiting for the next available.
I ran this benchmark numerous times, all resulting in the same jagged appearance.



The longer your query, the worse the response overall. Packet size means little at this point.




Ouch. This should be the worst case you should experience...



Ok, maybe this is the worst. That's a minimum 2 second wait before we even start to see data.
At this point everyone is waiting on the next available agent.

- How many requests can we handle without error?
 - Depends on your hardware/VM configuration
 - Possibly dependent on number of agents
- What makes the biggest difference, if any?
 - Concurrency: agents can only handle so many requests
 - Payload: puts a bottleneck at the network, HTTP daemon
 - Processing: puts the bottleneck at the OE broker's agent

- There is a tradeoff in processing of requests
 - Tomcat: few java.exe processes, many threads
 - WebSpeed: many cgiip.exe processes, few threads
- Tomcat seems to lag with initial requests but levels off
 - Java memory increases over time, subject to GC
- WebSpeed remains relatively consistent under most loads
 - Takes time to fork and exec each CGIIP process
- Not shown: Error % increases with payload size and latency
- Requests/sec and KB/sec were very similar between servers
- Fewer requests show even less drastic results...



For lower traffic, difference is only 10-20ms between technologies!

The previous tests are still relevant. They show that from a cold start there are irregularities with the results.

- Averages slightly favor WebSpeed, but the trends favor REST
- The larger the response packet, the more level the playing field
- Longer processing time affects all results relatively, to a point
- There is a tight connection between Tomcat and the AppServer
- If we could avoid the CGIIP process, WebSpeed would improve
- There is still room for performance tuning (Apache, Tomcat, OE)
- Going forward, just use the best tool for your situation!

Configuration time for REST is significantly greater than WebSpeed.
So consider that if you want things up and running quickly.

Thank You!

Please direct any angry emails or complaints about omissions via `/dev/null` :)

BRAVEPOINT

Dustin Grau

dgrau@bravepoint.com