



IBM Software Group

Troubleshooting DataPower Appliances

Matthias Siebler (msiebler@us.ibm.com)
Alfred Williamson (alfredq@us.ibm.com)
8/15/2012



WebSphere® Support Technical Exchange



Agenda

- Error Reports
- Status Providers
- Packet Captures
- Out-of-Memory (OoM)
- Large Debug Logs
- Q&A
- Summary

Summary

The WebSphere DataPower family of appliances has a wealth of tools available for troubleshooting problems in the field. However putting all these tools together is a difficult task for even experienced developers of the platform. Customers are often overwhelmed by the volume of data and how to interpret the data properly. This session will describe how to systematically approach troubleshooting several common scenarios.



Must gather

- <http://www-01.ibm.com/support/docview.wss?uid=swg21515489>
- Error reports contain most status providers
 - ▶ Some cannot fit into the error report due to size or time constraints
 - ▶ Error report content is continually being updated & improved
 - ▶ Reports can be useful even some time after the event
 - ▶ Status snapshot after the fact should be augmented by historical trends before the event
- Best practice is to have some minimal archives & trend graphs
 - ▶ But beware: Do not monitor the boxes to death!
 - ▶ All data is orthogonal to the method; via SNMP, CLI, webGUI, etc.

Error report analysis

- Audit log will have history of restarts
- Grep for errors; have a history of 'expected' errors & unexpected
 - ▶ i.e. loadbalancer health checks

Basic data collections

- Minimal: (every 5 minutes)
 - ▶ Memory
 - ▶ Load
 - ▶ Established TCP connections
- Throttler status log target is an easy way to collect all this
- Create a dedicated log target file or syslog to keep from having it rotate away
- Logs persist after a crash but status provider data is lost
- Know what is typical for your system

Advanced data

- All log files:
 - ▶ Additional logs not put into the error report will be in 'logtemp'
 - ▶ Top level & for the specific domain(s); unless too many domains
 - ▶ Automated scripts to get files via CLI or SOMA are helpful to build in advance
 - ▶ Log files can be under 'logstore' if using the log to RAID option

Status Providers

- These provide information about the system
 - E.g. filesystem, environment sensor, domain status....
- Information can be accessed thru
 - Webgui
 - CLI
 - SOMA
 - SNMP



Memory & Load Statistics

- Show load

```
xi50(config)# show load
```

Task Name	Load	Worklist	CPU	Memory	Filecount
-----	---	-----	---	-----	-----
main	2	0	0	33	109
wtx	0	0	0	0	20
ssh	0	0	0	0	29

- Show memory

```
xi50(config)# show mem
```

```
Memory Usage: 9 %  
Total Memory: 4148536 kilobytes  
Used Memory: 385722 kilobytes  
Free Memory: 3762814 kilobytes  
Requested Memory: 1197680 kilobytes  
Hold Memory: 811958 kilobytes
```

Throttler log

- Usage log:

```
20120314T102536Z [slm][debug] throttle(Throttler): tid(943):  
Memory(3570524/4098982kB 87.107579 free) Pool(1041874)  
Ports(31756/31850) Temporary-FS(224/242MB 92.561983 free) File(OK)
```

- ▶ Memory: same as 'Free Memory' from 'show memory'
- ▶ Pool: same as 'Hold Memory' from 'show memory'
- ▶ Ports: number of free ports (internal structure; 'show connections')
- ▶ File: generic test of all filesystems access (4 possible answers)
 - Cannot access filesystem due to low memory
 - Router has too many open files (may need to reload)
 - System has too many open files (may need to reboot)
 - other?

Discrepancies

- What to look for in the status providers:
 - ▶ 'show tcp'; 'show connections' & 'show handles'
 - ▶ All give slightly different results but roughly map one-to-one
 - ▶ If one is out-of-range by an order of magnitude could indicate an issue

- 'show load' vs. 'show cpu'
 - ▶ Load is an instantaneous measure
 - ▶ CPU is averaged
 - ▶ Load can jump around a lot; however ...
 - Mismatch can indicate 'spinning' ports

Packet Captures

- Definitive answer to protocol interoperability.
- Not necessarily the same as the Probe!
- Now can capture on loopback, VLANs or all interfaces at once.
- TCPdump format; viewable by Wireshark etc.

Packet Capture

Start Packet Capture [Help](#)
No Packet Capture Available for Downloading
Interface Type ***Ethernet Interface** *
Mode *
Maximum Duration seconds *
Maximum Size KB *
Maximum Packet Size bytes *
Filter Expression

Stop Packet Capture [Help](#)
Interface Type *
Ethernet Interface *

Packet Capture Filtering

- Expression format should follow 'pcap-filter(7)'
 - ▶ <http://www.unix.com/man-page/FreeBSD/7/pcap-filter>
 - ▶ e.g.

```
xi50(config-if[eth0])# pack temporary:zero.cap -1 1000 'tcp port 80'  
Trace begun.  
xi50(config-if[eth0])# █
```
- Supports basic and advanced filtering capabilities
- Provides ability to filter on
 - ▶ IP address
 - ▶ Port
 - ▶ MAC address
 - ▶ and many other qualifiers

FFDC background packet capture

- Always-on background packet capture has low overhead
- Captures packets on all interfaces simultaneously
- Capture automatically generated when
 - ▶ the system experiences an outage, such as a crash
 - ▶ user requested - Must-Gather operation
- When FFDC triggers report generation – information is current
- Enables a packet capture to be compressed and stored automatically in an Error Report and optionally sent off-box
- To enable:
<http://www-01.ibm.com/support/docview.wss?uid=swg21445832>

Service Probe

- Multistep Probe shows the payload as it moves through the processing policy – not meant to be on-the-wire

view	trans#	inbound-url	outbound-url	type	rule	client-ip
	528197	http://10.10.36.11:2053/	http://10.10.36.11:2053/	request	Transformer_request	10.10.100.177

Content of context 'INPUT':

```

<?xml version='1.0' encoding='UTF-8'>
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP-ENV:Body>
    <cbc:Order xmlns:cbc="urn:oasis:names:specification:ubl:schema:xsd:CommonBasicComponents-1.0" xmlns:cbc="urn:oasis:names:specification:ubl:schema:xsd:CommonBasicComponents-1.0" xmlns:cbc="urn:oasis:names:specification:ubl:schema:xsd:CommonAggregateComponents-1.0">
      <cbc:BuyerID>20031234-1</cbc:BuyerID>
      <cbc:IssueDate>2003-01-23</cbc:IssueDate>
      <cbc:LineExtensionTotalAmount amountCurrencyCodeListVersionID="0.3" amountCurrencyID="USD">438.50</cbc:LineExtensionTotalAmount>
    </cbc:Order>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
  
```

List of service variables:

name	type	value
var://service/client-service-address	string	'10.10.100.177:4269'
var://service/current-call-depth	string	'0'
var://service/domain-name	string	'default'
var://service/error-code	string	'0x00000000'
var://service/error-headers	string	'(empty string)'
var://service/error-message	string	'(empty string)'
var://service/error-protocol-response	string	'0'
var://service/error-subcode	string	'0x00000000'
var://service/formatted-error-message	string	'(empty string)'
var://service/header-manifest	node-set	'(show nodeset)'
var://service/input-size	string	'6593'
var://service/local-service-address	string	'10.10.36.11:2053'
var://service/max-call-depth	string	'128'
var://service/persistent-connection-counter	string	'1'
var://service/processor-name	string	'Transformer'
var://service/processor-type	string	'XML Firewall'
var://service/system/forward	string	'local://Transformer_generated.wsd'
var://service/time-elapsed	string	'37'
var://service/time-forwarded	string	'0'
var://service/time-response-complete	string	'0'
var://service/time-started	string	'3'

Probe Triggers

- Can be used to limit capture in case of high traffic

Edit Probe Triggers

Client-IP Match	<input type="text"/>
Inbound URL Match	<input type="text"/>
Outbound URL Match	<input type="text"/>
Processing Type	Any <input type="button" value="v"/>
Processing Rule Match	<input type="text"/>
XPath Expression Match	<input type="text"/> <input type="button" value="XPath Tool"/>

Memory or other resource leaks

- Generally must always have a baseline
- What can be leaked?
 - ▶ Memory
 - ▶ File handles/sockets/file descriptors
 - ▶ Ports (slightly different from sockets)
 - ▶ Inodes (very rare)
- Tracing must be turned on before the resource is leaked
- Currently leak detection requires a reboot; development is planning for always-on resource leak tracing



Memory logs

- Each log message captures a snapshot at that time
 - ▶ Not cumulative; can go up & down
 - ▶ Not exhaustive; some actions or protocols can allocate memory outside
- Added in 3.8.2: units are in bytes

```
20110224T110806Z [memory-report][debug] mpgw(sender): tid(8464):  
Response Finished: memory used 21968644
```

```
20101108T232236Z [memory-report][debug] mpgw(sftp-ftp-mpgw):  
tid(6000)[response][9.42.102.172]: Processing [Rule (sftp-ftp_policy_rule_1),  
Action ('sftp-ftp_policy_rule_1_results_output_0', results()), Input(INPUT),  
Output(NULL)] finished: memory used 595732
```

Services Memory Status Provider



Services Memory Usage

[Refresh Status](#)

Show Current Domain

Domain "33558":

Service	Name	current (kilobytes)	1 min (kilobytes)	1 min - 5 min (kilobytes)	5 min - 10 min (kilobytes)	10 min - 1 hour (kilobytes)	1 hour - 12 hours (kilobytes)	12 hours - 1 day (kilobytes)	Lifetime (kilobytes)
xMLFirewallService	ContentValidator	0	0	0	0	0	0	0	0
xMLFirewallService	loop	0	0	0	0	0	0	0	4536
xMLFirewallService	loopback	0	0	0	0	0	0	0	0
xMLFirewallService	tester	0	0	0	0	0	0	0	5632

Domain "default":

Service	Name	current (kilobytes)	1 min (kilobytes)	1 min - 5 min (kilobytes)	5 min - 10 min (kilobytes)	10 min - 1 hour (kilobytes)	1 hour - 12 hours (kilobytes)	12 hours - 1 day (kilobytes)	Lifetime (kilobytes)
xMLFirewallService	gzip	0	0	0	0	0	0	0	0
multiProtocolGateway	test	0	6059	0	0	0	0	0	6059



Scalability concerns

- SLM vs. Monitors
- Determining scalability requires a methodical and sensible approach
- Debugging can be surprisingly tricky
- How much traffic is the box actually taking?
 - ▶ How many requests?
 - ▶ What kind?
 - ▶ How big?
 - ▶ What actions?

Message Count Monitors

- Most accurate method for determining exactly how much traffic a service is processing
- More lightweight than SLM
- Does not have as many options&features

The image displays two overlapping configuration windows from the IBM WebSphere administrative console. The background window is titled "Configure Message Count Monitor" and shows the "Main" tab. It includes fields for Name, Administrative State (set to "enabled"), Comments, Message Type (set to "(none)"), Measure (set to "Requests"), Source (set to "All"), Header, and Maximum Distinct Sources (set to "10000"). The foreground window is titled "Configure Message Matching" and shows the "Main" tab. It includes fields for Name, Administrative State (set to "enabled"), Comments, IP Addresses, Excluded IP Addresses, HTTP Method (set to "POST"), Request URL (set to "*"), and a "Done" button. A third window is partially visible behind the others, showing a browser address bar with "DataPower XI50 | C" and "9.70.154.27 https".

SLM

- Shaping can be used to smooth traffic
- Should not be used to hide a broken backend
 - ▶ Plan on shaping for a few seconds; not minutes
- Reliability should be end2end; not hop-by-hop

Audit log

- Polling for uptime is best practice for monitoring restarts
 - ▶ Except when the 32 bit counter wraps
- Monitoring the audit log is also useful
 - ▶ If the uptime goes down then the box has rebooted; otherwise it reloaded

- ▶ But note this message:

```
20120217T083717Z [eventlog][failure] (SYSTEM:default:*:*):  
Booting build 205760 on 2011/11/15 11:02:50 count 32. Uptime  
777830
```

- ▶ Booting message w/ type failure is just an indication that the audit log has rotated; not that the box has restarted



Log event trigger

- Should have the ability to execute any CLI command or CLI script
- Needs to match on Log Message ID and message text using an optional regular expression
- Examples of ability to execute any CLI command or CLI scripts:
 - start a packet capture on a specific event
 - stop a packet on the next occurrence of the same event
 - perform a must-gather Error Report generation

Log Target: default-log [up]

Apply Cancel Undo

[Export](#) | [View Log](#) | [View Status](#) | [Help](#)

Event Triggers

Message ID	Regular Expression	Only Once	Only this Trigger	CLI Command	
0x80400036		on	on	interface eth0; packet-capture temporary:///my pcap	↑ ↓ ✎ ✕
0x80400036		on	on	interface eth0; no packet-capture temporary:///my pcap	↑ ↓ ✎ ✕
					Add



Large debug logs

- The goal is that we want to capture all possible data.
- Debug logging will do that; with some few exceptions:
 - ▶ RBM (optional)
 - ▶ webGUI (optional)
 - ▶ logging about logging (not possible.)
- In the default domain; create a new log target
 - ▶ type file
 - ▶ format text
 - ▶ timestamp numeric
 - ▶ archive rotate
 - ▶ event all debug
- maximum total is 50MB times 100 = 5 GB
 - ▶ Do we have that much?
 - ▶ Make sure the space is available!



Large debug logs – on device

- Best practice: log to RAID
- Rotate; do not archive files
- Better to pull via HTTP rather than push
- If push must be used FTP is the preferred approach
- The file log target cannot rotate more than once per second
 - ▶ minimum size of the log file should be able to contain more than 1 seconds worth of data; otherwise you will certainly be losing messages.
- Dropped messages are also in the log file: “Buffer Overflow: X event(s) lost”

Large Debug logs

- Always check to make sure they work
 - ▶ Check on the log target status
 - ▶ Should have zero dropped events

The running configuration of the domain contains unsaved changes. [Review changes.](#)

Debug-Level Logging and Debug Logging (in default domain) are enabled, which impacts performance. [Manage debug settings.](#)



Logging Targets

[Refresh Status](#)

name	status	memory (kbytes)	processed (events)	dropped (events)	pending (events)	error info
default-log	Active	6	358	0	0	none
test	Active	2	2	0	0	none

Large Debug logs – off device

- Best practice – syslog over UDP
- Using syslog-tcp may cause bottleneck (if using firmware 4.0.2 or before)
 - ▶ DataPower opens many simultaneous connections
 - ▶ Can bring down some servers
- Always set a static route to the syslog servers to force outbound traffic over the correct interface
- Adding a syslog log target is a lightweight addition to a busy box
- Note: UDP syslog may truncate some longer messages



Support resources

- IBM WebSphere DataPower SOA Appliance Handbook
- IBM Support Portal for DataPower
 - http://www.ibm.com/support/entry/portal/Overview/Software/WebSphere/WebSphere_DataPower_SOA_Appliances
- developerWorks articles
- WebCasts
- Forum: <https://www.ibm.com/developerworks/forums/forum.jspa?forumID=1198>
- User Groups: <http://www.websphere.org/websphere/Site?page=ugdetail&groupId=165>

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