



August 10th 2010, OSG Site Admin Workshop - Network Performance

Jason Zurawski, Internet2

OWAMP

Agenda

- Tutorial Agenda:
 - Network Performance Primer - Why Should We Care? **(15 Mins)**
 - Getting the Tools **(10 Mins)**
 - Use of the BWCTL Server and Client **(30 Mins)**
 - **Use of the OWAMP Server and Client (30 Mins)**
 - Use of the NDT Server and Client **(30 Mins)**
 - BREAK **(15 mins)**
 - Diagnostics vs Regular Monitoring **(30 Mins)**
 - Network Performance Exercises **(1 hr 30 Mins)**

OWAMP: What is it

- OWAMP is:
 - Command line client application
 - Policy and scheduling daemon
 - Used to determine **one way** latencies between hosts.
- Implementation of the OWAMP protocol as defined by <http://www.rfc-editor.org/rfc/rfc4656.txt>
 - Command Protocol to speak between client and server, server and server
 - Test protocol
- Different attempts to do this in the past:
 - Surveyor
 - RIPE

Why *One Way* Latency?

- Passive Measurements (e.g. SNMP)
 - Higher polling interval may mask queue depths
 - Active probing gives a better picture of *real* traffic
- Round Trip Measurements:
 - Hard to isolate the direction of a problem
 - Congestion and queuing can be masked in the final measurement
 - Can be done with a single 'beacon' (e.g. using ICMP responses)
- One Way Measurements:
 - Direction of a problem is implicit
 - Detects asymmetric behavior
 - See congestion or queuing in one direction first (normal behavior)
 - Requires '2 Ends' to measure properly

OWAMP Control Protocol

- Supports authentication and authorization of the users that will test
- Used to configure the parameters of a test
 - Endpoint controlled port numbers
 - Extremely configurable *send schedule*
 - Configurable packet sizes
- Used to start/stop tests
- Used to retrieve results
 - Provisions for dealing with partial session results in the event of a failure

OWAMP Test Protocol

- “*Lightweight*” compared to the control protocol
- Uses UDP as the transport protocol, since the protocol needs to be able to measure individual packet delivery times
- Supports varying packet sizes
- Data needed to calculate experimental errors on the final result is in every packet
- Packets can be “open”, “authenticated”, or “encrypted”

Sample Implementation

- Applications
 - owampd daemon
 - owping client
- Open source license and development
- Built upon protocol abstraction library
 - Supports *one-off* applications
 - E.g. “powstream” is a client that sends a constant stream of measurement packets.
 - Allows authentication/policy hooks to be incorporated

Functionality (owping client)

- Owping
 - client requests tests from an OWAMP server
 - Client can be *sender* or *receiver*
 - Communication can be “open”, “authenticated”, or “encrypted”
 - Supports the setup (scheduling) of many tests concurrently
 - Supports the buffering of results on the server for later retrieval

Functionality (owampd server)

- owampd
 - Accepts requests for tests
 - Responds with a message: *accepted/denied*
 - Tests are formally started with a **StartSessions** message from the client.
 - Runs tests
 - Sessions with packets received at the server are buffered for later retrieval

OWPING Example


```
boote@nms-rlat.chic.net.internet2.edu: /home/boote
boote@nms-rlat:~[360]$ owping nms-rlat.newy.net.internet2.edu
Approximately 13.0 seconds until results available

--- owping statistics from [64.57.17.34]:45355 to [nms-rlat.newy.net.internet2.edu]:44244 ---
SID:      40391162cbec228e81118c1953a5eef9
first:    2008-05-31T19:16:31.627
last:     2008-05-31T19:16:43.362
100 sent, 0 lost (0.000%), 0 duplicates
one-way delay min/median/max = 11/11/11 ms, (err=0.0442 ms)
one-way jitter = 0 ms (P95-P50)
Hops = 3 (consistently)
no reordering


--- owping statistics from [nms-rlat.newy.net.internet2.edu]:44247 to [64.57.17.34]:45356 ---
SID:      40391122cbec228ebb1bde827906fe35
first:    2008-05-31T19:16:31.608
last:     2008-05-31T19:16:41.979
100 sent, 0 lost (0.000%), 0 duplicates
one-way delay min/median/max = 10.9/11/11 ms, (err=0.0442 ms)
one-way jitter = 0 ms (P95-P50)
Hops = 3 (consistently)
no reordering

boote@nms-rlat:~[361]$
```


OWAMP GUIs – Delay/Loss Plot



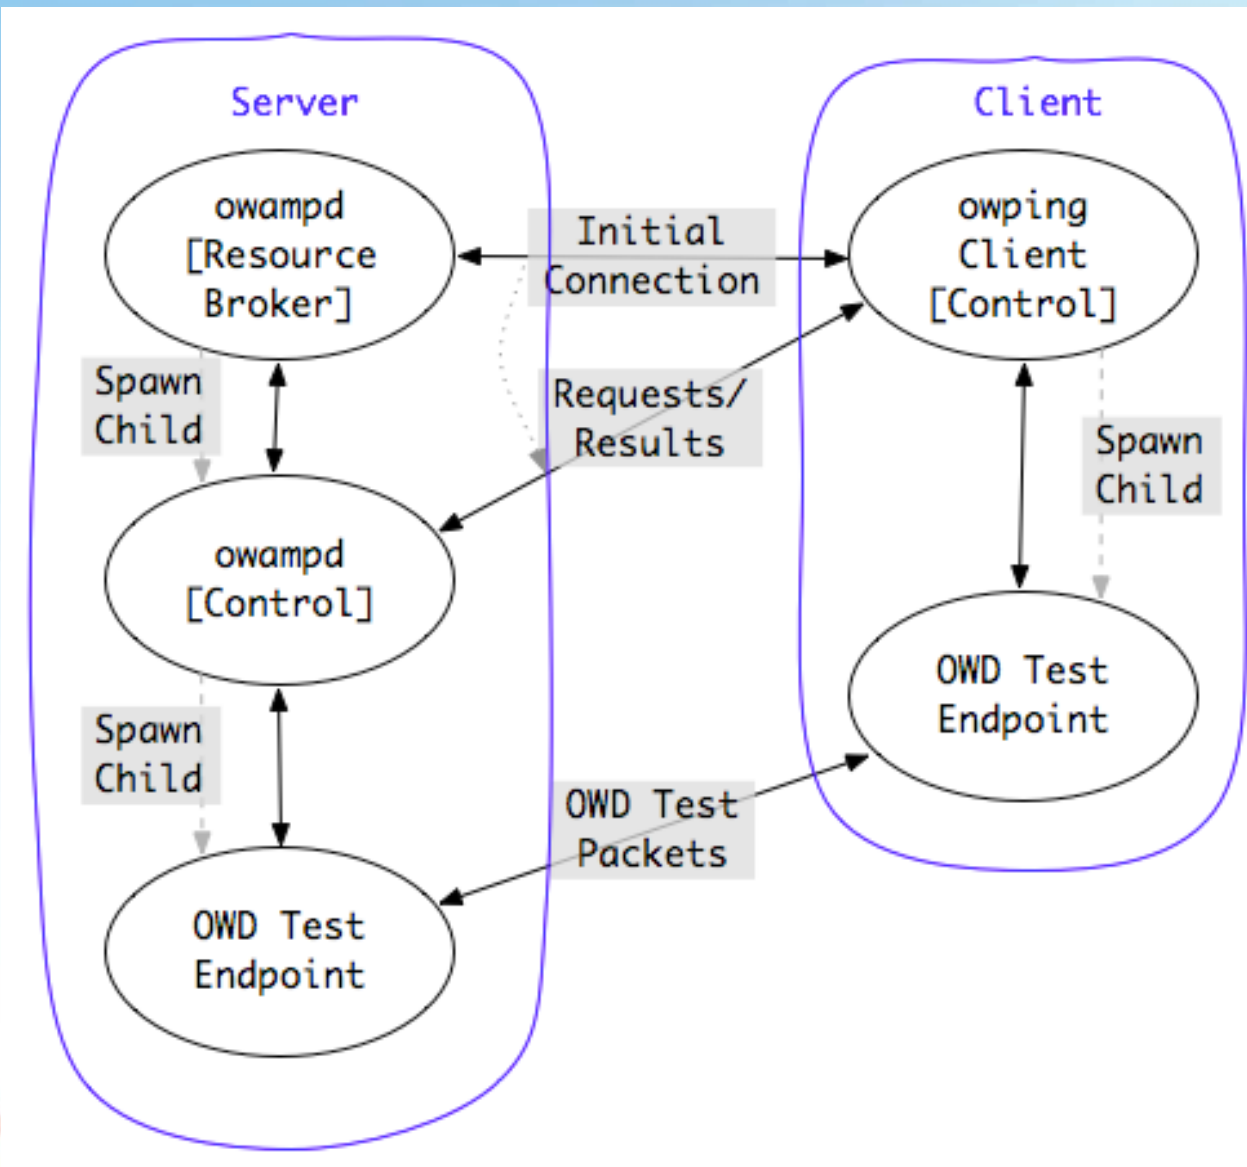
OWAMP GUIs - Jitter



Resource Allocation

- Each connection is “classified” (authentication)
- Each classification is associated with a set of hierarchical limits
 - Bandwidth (bandwidth)
 - Session buffer (disk)
 - Data retention (delete_on_fetch)
 - Connection policy (allow_open_mode)
- (no time dependent dimension to resource allocation in owampd)

Architecture



General Requirements – Time Source

- NTP (ntpd) synchronized clock on the local system
 - Configure NTP properly (don't rely on system defaults!)
 - Strictly speaking, owamp will work without NTP. However, your results will be meaningless in many cases
 - More info here:
<http://www.internet2.edu/performance/owamp/details.html#NTP>

General Requirements – Support

- Source should compile for all modern *NIX
 - *BSD, Linux, OS X
 - configure/make/make install
- RPMs compiled specifically for CentOS 5.x
 - May work with other RPM based systems (Fedora, RHEL)

Hardware

- Stable System Clock
 - Temperature controlled environment
 - No power management of CPU
 - “Smart” or “Green” CPUs/Cores may degrade performance to save energy – this may cause jitter in the measurements
- No strict requirements for CPU, Memory, Bus speed
 - More tasking schedules will require more capable hardware

Operational Concerns

- Time:
 - NTP issues predominate the problems
 - Determining an accurate timestamp “error” is in many ways more difficult than getting a “very good” timestamp
 - Working as an “open” server requires UTC time source (For predefined test peers, other options available)
- Firewalls:
 - Port filter trade-off
 - Administrators like pre-defined port numbers
 - Vendor manufactures would probably like to “prioritize” test traffic
 - Owampd allows a range of ports to be specified for the receiver
 - Read documentation for more details

Policy/Security Considerations

- Third-Party DoS source
 - Compromised server may send packets to other locations.
- DoS target
 - Excessive traffic will harm measurement results
- Resource consumption
 - Memory (primary and secondary)
 - Network bandwidth

Policy Recommendations

- Restrict overall bandwidth to something relatively small
 - Most OWAMP sessions do not require much
- Limit “open” tests to ensure they do not interfere with precision of other tests

Availability

- Main Page:
 - <http://www.internet2.edu/performance/owamp>
- Mailing lists:
 - owamp-users@internet2.edu
 - owamp-announce@internet2.edu

Hands On

- Testing OWAMP:
 - Log on to testbed
 - Test from one host to another:
 - owping HOSTNAME
 - Try different hosts. Try longer tests. What happens when we use:
 - -c (number of packets to send, try 1000)
 - -t (test to HOSTNAME only)
 - -f (test from HOSTNAME only)



OWAMP

August 10th 2010, OSG Site Admin Workshop – Network Performance
Jason Zurawski – Internet2

For more information, visit <http://www.internet2.edu/workshops/npw>