



www.internet2.edu

OWAMP was developed by members of the Internet2 Engineering team as a reference implementation of a proposed standard going through the Internet Engineering Task Force (IETF). Internet2's E2Epi team is deploying OWAMP, traceroute, and Iperf as the first components of the E2E piPEs framework.

What is E2E piPEs? It is the End-to-End Performance Initiative (E2Epi) Performance Environment System, a project under development by Internet2's E2Epi; more information can be found at: <http://e2epi.internet2.edu>.

Even without accurate time sources, users can measure such characteristics as loss with quality acceptable for many practical purposes (including network operations).

OWAMP:

One-Way Ping

With roundtrip-based measurements, it is hard to isolate the direction in which congestion is experienced. One-way measurements solve this problem and make the direction of congestion immediately apparent. Since traffic can be asymmetric at many sites that are primarily producers or consumers of data, this allows for more informative measurements.

One-way measurements allow the user to better isolate the effects of specific parts of a network on the treatment of traffic.

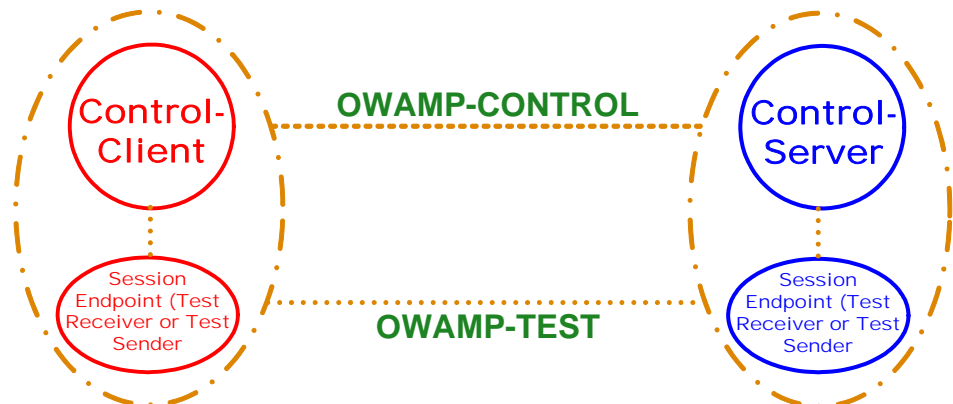
With better measurement tools and techniques like **One-Way Ping (OWAMP)**, an implementation of the IETF's One-Way Active Measurement Protocol, available, network providers will be able to better know the exact behavior of their networks and apply resources where improvement is most likely. (Note: Passive observation of average link use misses the transient queues – active measurement could see them.) Users would be more informed about network performance. This would prompt a better allocation of resources by network providers, decreasing areas of congestion where possible.

The increasing availability of precise time sources allows network hosts to timestamp packets with typical errors that are substantially smaller than the delays seen on

the Internet. This makes it possible for one-way measurements to be collected across a broad mesh of Internet paths. In addition, the open-source nature of OWAMP makes it possible for one-way metrics to become as common as roundtrip metrics have become (from tools like ping).

Using OWAMP also simplifies the analysis of measurement results – explicit send and receive timestamps for every measurement packet make analysis more straightforward because one does not need to assume return path reliability, preservation of inter-packet spacing by the roundtrip measurement reflector, etc. For example, packet reordering, which can have implications for TCP performance, can be measured under a variety of input scenarios, with separation of reordering on the forward and return paths.

The figure below illustrates the OWAMP architecture. The protocol consists of two parts: OWAMP-Control and OWAMP-Test. OWAMP-Control is used for session initiation, setup, tear-down, confirmation, and for the retrieval of the results. OWAMP-Test is the protocol for the actual measurement, which largely amounts to a packet format convention.



OWAMP session control uses traditional client-server communication between a control-client and a server, using the OWAMP-Control protocol. The server is implemented using the conventional accept-fork model. The two sides negotiate one-way test parameters using the OWAMP-Control protocol. The implementation then forks additional processes that speak to OWAMP-Test and implement the session endpoints for the Sender and Receiver roles on both sides. The OWAMP-Test protocol is used to conduct the actual test.

Using OWAMP, it is possible to collect active measurement data sufficient to determine a broad class of singleton characteristics (e.g., loss probability, median delay, jitter, 90th percentile of delay). Non-singleton characteristics, such as the expected inter-arrival gap of packets that were sent back-to-back, can be measured as well. Note: All measurements are done with synthetic traffic; application simulation is outside of the scope of OWAMP. The protocol is not designed to be able to send a packet as soon as a response to the previous packet arrives, but can send on any predetermined schedule.

OWAMP has been designed to be deployable on as many systems as possible. Just as it is possible to ping most hosts on the network today, widespread deployment of OWAMP would make it possible to conduct more accurate measurement sessions. To further this goal, an implementation of OWAMP is publicly available at: <http://owamp.internet2.edu/>.

The OWAMP development team recognized that network measurement systems become more unwieldy as their size grows. When a full-mesh measurement architecture is used, the **amount of disk space and network capacity** used by the system will grow as the square of the number of measurement nodes. There is nothing a measurement protocol can do to alleviate this problem. OWAMP, however, was designed not to introduce any new scalability problems. It allows the user to conduct only those measurement sessions desired and to retain as much (or as little) data as desired. OWAMP also does not dictate a choice of site(s) where measurement results are stored: it is possible to have all data stored

at a central site or to store data at each receiver and fetch it as needed.

OWAMP has been deployed on Internet2's Abilene network. Each of the 11 core router nodes of that network has a rack of measurement hardware deployed with it. OWAMP has been deployed on a FreeBSD system with an attached CDMA clock in each one of these measurement racks. Each of these systems is configured to run ongoing OWAMP tests to each of the other systems. This provides a full mesh of one-way delays for the entire Abilene network. This data can be seen at <http://owamp.internet2.edu/abilene/>.

For more information, contact Stanislav Shalunov (shalunov@internet2.edu) or Jeff Boote (boote@internet2.edu).

How many resources are necessary?

The daemon can be configured to consume a limited amount of resources (network bandwidth for tests and disk space for intermediate storage of results). A flexible access control mechanism is provided with controls based on both established prior relationship (i.e., a user name and password) and IP address of the requester; it is, however, hoped that those who deploy OWAMP will leave a small allocation for open use.