

XORP BGP Test Harness

Version 0.1

XORP Project
International Computer Science Institute
Berkeley, CA 94704, USA
feedback@xorp.org

December 11, 2002

1 Introduction

This document describes a test harness that was built primarily to test the XORP BGP implementation. It may be possible to augment the harness to use it for testing other protocols.

A single BGP process is placed in the harness and tests are performed on it. The tests can range from testing the decision process to verifying that a session is dropped when the hold timer expires.

Not all of the features mentioned below have been implemented yet. Section 6 has a list of these features.

2 Requirements

A major requirement was to allow the testing of any BGP process, not just our own. In the case of our own BGP process it was essential that regression tests could be run and results verified totally from within scripts. Without the need for any manual configuration. The XORP BGP regression tests can be run directly from the “Makefiles”.

The test harness supports testing at various levels:

- Decision process.

The BGP decision process can be tested by sending update packets to BGP. Then verifying that the correct update packets are sent by the BGP process.

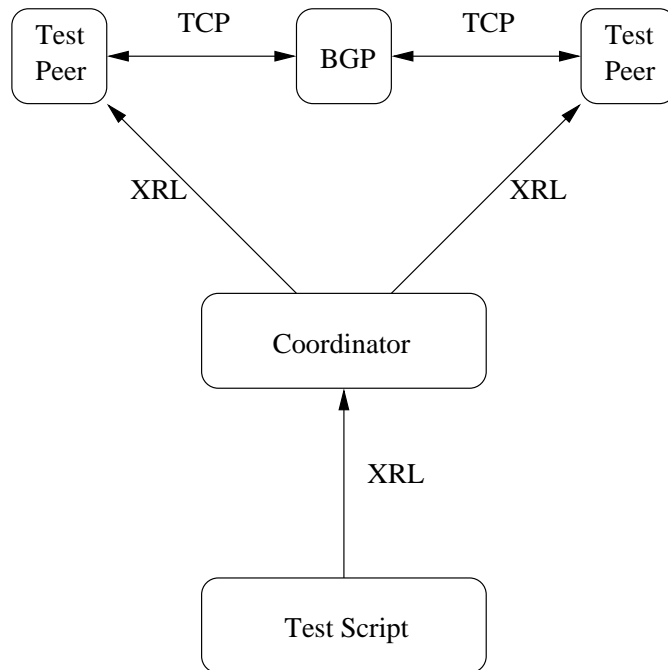


Figure 1: Test harness processes.

The actual packets sent by the BGP process can be compared to expected packets. Another way that the BGP decision process can be tested is by testing the routing tables held at the peers.

- Low level testing.
Testing responses to deliberately corrupted packets.
- Load testing.
Testing reaction to introducing a large number of routes back to back.
- Timer testing.
Correct operation of timers such as the HOLD TIME timer.

3 Usage

Figure 1 shows a BGP process under test that is connected to two test peers. The test harness has been split into a number of separate processes. The main reason

```
reset
target xorp 179
initialise attach peer1
initialise attach peer2
peer1 establish active true AS 1 keepalive true holdtime 0 id 10.10.10.10
peer1 assert established
peer2 establish active true AS 2 keepalive true holdtime 0 id 20.20.20.20
peer2 assert established
```

Figure 2: Establish BGP Sessions

for using multiple processes is that a third party BGP processes may not be able to accept multiple connections from the same host (IP address). The harness was simpler to implement by splitting functionality into separate processes.

The test harness consists of a *coordinator* process through which all interactions with the harness are mediated. There are also one or more *test peer*'s. Each *test peer* is capable of forming one BGP session with the BGP process under test. The *coordinator* process communicates with the *test peer*'s using XRL's [1]. The *coordinator* process accepts commands via XRL's from test scripts. Currently our test scripts are written in the shell programming language. The scripts could be written in any language. The full set of commands accepted by the *coordinator* can be found in 4.1.

Figure 2 shows an example of a simple program that might be sent from a test script to a *coordinator*. It is assumed that before the script is sent that all the processes are already running. One important point to note is that due to the asynchronous nature of XRL's a command returning does not necessarily mean that it has completed successfully. In the XORP test scripts there are delays between certain operations. This example show the *coordinator* being *reset*, then its given the hostname and port number of the BGP process under test. Then the *coordinator* attaches to the two *test peer*'s peer1 and peer2, these are the XRL names by which the *test peer*'s are known. The majority of the commands are sent to the *test peer*'s themselves. In order to send a command to a *test peer* the command is preceded by the *test peer* name. In our example each *test peer* forms a session with the BGP process and then the *assert* command is used to assert that a session is still established. As noted above the *establish* command completing does not mean that a session was established. Of course attempting to send a *packet* on a session that has not yet been established will generate an error.

The *expect packet . . .* command is used to create a queue of expected packets. Whenever a packet arrives on a peer it is checked against the queue of expected packets. If there are no expected packets on the queue then no action is taken. If

```

...
# Update error, missing well known attribute.
peer1 expect packet notify 3 3
# An update packet without an origin
peer1 send packet update aspath 1 nexthop 20.20.20.20 nlri 10.10.10.0/24
# Delay
peer1 assert queue 0
...

```

Figure 3: Wait for a notify packet

```

...
peer1 trie rcv lookup 212.174.196.0/24
peer1 trie rcv lookup 212.174.196.0/24 aspath 1
...

```

Figure 4: Performing lookup in trie

there is a packet on the expect queue it is compared against the incoming packet, if they match all is fine and the packet is removed from the queue. If the incoming packet does not match the packet at the head of the queue then an error is flagged and the non matching packet is saved. At the end of a set of tests the *assert queue* command can be used to verify that queue is at the expected length. If an error has occurred this is the point at which the non matching packet and the expected packet are returned, along with an error status. Figure 3 is a example of a code fragment that is waiting for a notify packet on peer1. A notify packet is added to the queue of expected packets, then an update packet without an origin is sent to the BGP process. An update packet without an origin is an error and should generate the notify packet that we are waiting for. Note that it is the responsibility of test script to add a delay between sending the *update packet* and the *assert*.

Each *peer* has two tries associated with it, a sent and a received trie. Each update packet that is sent to, or received from, the BGP process is passed to the appropriate trie. The update packet is processed and saved if it contains any NLRI's. Update packets containing only WITHDRAW's are not saved at this time. At this time two types of *lookups* can be performed on a trie (Figure 4). If a *lookup* fails then an error is returned. A test might therefore involve sending many update packets from different peers and then verifying that the routing tables at the various peers are correct. Saving the update packets in the tries makes it possible to dump the routing tables.

4 Commands

4.1 COORD XRL's

- Command("command string")
Accept commands via XRL's.
- Pending()
Returns true while there are any uncompleted commands. Can be used to poll the coordinator to verify that the previous command has completed.

4.1.1 Commands currently accepted by coord

- reset
Reset all the state in the coordinating process.
- target <hostname> <port>
The BGP process under test.
- initialise attach/create peername
Form an association with a test_peer. If the second argument is attach then it is assumed that the test_peer is already running. If the second argument is create then the test_peer is started (not currently supported).

4.1.2 Peer specific commands

- connect
Connect to the target
- listen
Listen for a connection from the target.
- establish active <true/false> AS <value> keepalive <true/false> holdtime <value> id <ipv4>
The active, AS, keepalive, holdtime and id arguments are optional. Active defaults to being true and actively makes a connection, setting active to false sets up a listener. The AS value is recommended if a connection is wanted.
- send packet update origin <num> aspath <path> nexthop <ip> localpref <num> nlri <net> withdraw <net>

- send dump mrt d update filename
Given a file in mrt d dump format send the update packets in this file.
- trie <rcv/sent> lookup <net>
Test to see if this net is in the trie.
- trie <rcv/sent> lookup <net> not
Test to see if this net is not in the trie.
- trie <rcv/sent> lookup <net> aspath <path>
Test to see if this net is in the trie and associated with the provided AS path.
- expect packet notify <error code> <sub error code>
Place a notification packet on the expect queue. The <error code> is mandatory. The <sub error code> is optional.
- expect packet update origin <num> aspath <path> nexthop <ip> localpref <num> nlri <net> withdraw <net>
Place an update packet on the expect queue.
- expect packet open asnum <value> bgpid <ipv4> holdtime <value>
Place an open packet on the expect queue. All fields shown are mandatory.
- expect packet notify
Place a notify packet on the expect queue.
- assert queue <queue length>
Check the queue length of the expect queue. Every message that matches removes an entry from the queue. If an error has previously occurred then this call will return the error. The length of the queue check is optional.
- assert established
Verify that a session has actually been established. Some tests can pass without a BGP process being present. These tests require this interface.
- dump <rcv/sent> <mrt d/text> <traffic/routeview/current/debug> <filename>
A mechanism for saving conversations or dumping routing tables. The received and sent cases can be dealt with independently. Four types of dumps are supported:

1. Traffic.
This is basically all the traffic which is sent and received. The dumping can be disabled by making a call with the <filename> argument removed.
2. Routeview.
The current state of the routing table.
3. Current.
Travels through the routing table and dumps all the update packets that have caused entries in the routing table. The packets are dumped in the order in which they arrived. This is only an approximation as update packets containing only withdraws will not be saved.
4. Debug.
Visit all nodes in the trie and dump the update packet that was responsible for this entry. Update packets can have multiple NLRI's associated with them so a packet can be in the dump many times.

The save file can be either in mtrd dump format or in xorp text format.

4.2 TEST PEER XRL's

Commands that are accepted by the test peer.

- Register("coord")
This is an external registration to the test peer. All packets received by the test peer are sent to the "coord".
- Packetisation("bgp")
Tell the test peer to treat incoming packets as BGP packets packetise them accordingly. Otherwise just packetise the the packets the way they appear from the connection.
- Connect("host", "port")
Connect to the named host and port.
- Listen("address", "port")
Listen for connections on this address and port.
- Send("Data")
Send data on the TCP connection.

- Disconnect()
Drop the current TCP connection.
- Terminate()
Terminate the process.

4.3 TEST PEER CLIENT XRL's

- Packet("peer", "status", "time", "data")
 - "peer"
The peer that the packet came from.
 - "status"
If the remote peer had been asked to perform packetisation. Then if a bad message is received signify this. Also after a bad is received packetisation is disabled.
 - "time"
The time when the packet was received in micro seconds since 1970-1-1.
 - "data"
The raw data that was read on the connection.

5 Outstanding Issues

- At the time of writing the harness has only been used against the XORP BGP process. There is no reason to believe that it could not be used against implementations.

6 Not yet implemented or TODO list

- Constructing corrupted packets.
- Dumping current and routeviews.

References

- [1] XORP Inter-Process Communication Library. XORP technical document.
<http://www.xorp.org/>.