

In this issue:

Research & Development

InterRoute tool & border router stability metrics

Testbeds

Experimental INTERMON testbed Europe – Latam
America

Research & Development

InterRoute tool - Evaluation of border router stability

Contribution from Pedro Gutiérrez
(Telefonica)

1. Introduction to inter-domain topology analysis tool of INTERMON – the InterRoute

The INTERMON InterRoute tool provides the interface between the policy based INTERMON database (topology repository part) and public or private BGP-4 route repositories. Public route repositories have been established by the Oregon Routeviews project, the RIPE and others at selected points of the Internet, which normally serve as Internet traffic exchanges. Private route repositories replicate the setup proposed by the Oregon Routeviews project in an INTERMON user's autonomous system and feed the data to the INTERMON database. All public BGP-4 route repositories are based on the standard format defined in Merit's Multithreaded Routing Tooling MRT. This format has also been chosen for the INTERMON toolkit, thus enabling the project to use the public route repositories for its own research.

The InterRoute information processing chain

The INTERMON InterRoute tool implements a three stage processing chain for the interdomain topology information generated by the different domains along the path for a time window. The front-end process derives the domains along the end-to-end path using traceroute data. Autonomous systems and Internet exchanges are identified. The domains, which provide BGP-4 data, are selected. Basing the route repository selection on traceroute data narrows the data sets to those, which will be indeed used by the Internet to transfer the packets between the source and the destination. This keeps the amount of information, which has to be processed at later stages at a reasonable minimum. The use of this strategy is confirmed by other tools, which combine different sources of topology data to provide accurate topology

maps of the Internet like the Rocketfuel tool [Rocketfuel].

The second stage in the information processing chain collects all routing information available for the time window of the observation. A snapshot of the routing table immediately before the observation time window is obtained and forms the basis for the rest of the computations. The evolution of the routing table is obtained from the routing update collections, which are available for the significant route repositories. This results in a potentially large database. Routing table snapshots are around 400k routes long and 600k-1M updates a day are normal.

The last stage filters the updates, which can be significant for an end-to-end scenario, simulating the behaviour of a router as closely as possible. Routers work on a longest prefix match basis when selecting the output interface for a given IP address. Since the probability of a route withdraw affecting the longest prefix match is not nil, all possible prefix matches are stored in the database. Despite this relaxed filtering strategy, this second stage process reduces the amount of storage required for the topology databases by several orders of magnitude. The following figure shows all components involved in the InterRoute tool architecture.

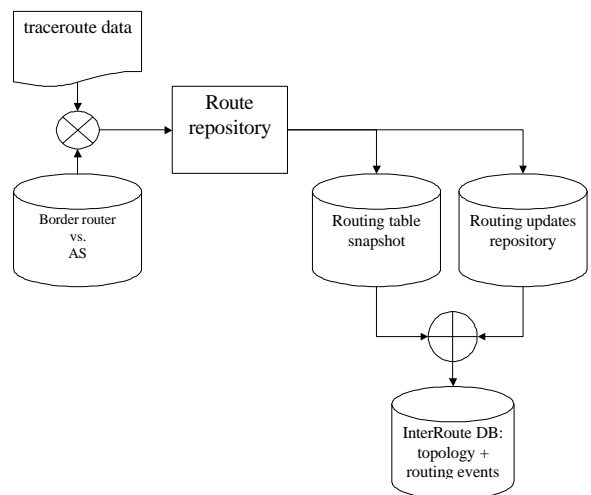


Figure 1: Information flow in the InterRoute architecture

2. Routing instability in the Internet

Behaviours different to the expected router behaviour should be regarded as routing instability. The reasons for routing instability vary from case to case and are beyond the scope of this paper. The importance of routing instabilities for us is that BGP-4 routing policies introduce metastability in the routing protocol. This, in turn, may lead to resonances in the network, as has occurred in the past, with different routing storms, which have rendered the Internet useless for significant periods of time. A network's life cycle from the point of view of the BGP-4 routing protocol consists of two events.

When the network is commissioned, an advertisement for it is generated. Routers, which receive this advertisement, will install the route in their routing tables, provided the routing policies established by the network administrators allow it. When the network is decommissioned, the router responsible for providing connectivity for it will generate a BGP-4 withdrawal, which will propagate according to routing policies, causing the routing table entry for this network to be deleted.

3. A simple router stability metric

Router stability metrics taking into account the router failures are extremely interesting from a network administrator's point of view. Adequately mapped, router stability metrics allow a network administrator to locate problems in a network promptly. The router stability metric supported by InterRoute are derived from the intrinsic behaviour of the BGP-4 routing protocol, show the potential impact of routing events on an end-to-end QoS/SLA and guide the network administrator when navigating through the inter-domain network map.

The 'cold-warm-hot' router stability metric

We propose a simple, intuitive metric for router stability, which can be represented by any network administration system easily and will draw the attention of the network administrator to the different potential hot spots in his network in a prioritised way, focusing on the points which pose a bigger potential threat to the end to end SLA, based on deviations from the expected BGP-4 protocol behaviour. This metric is graded in three different levels, ranging from cold to warm to hot.

Cold routers will be those, which had a stable behaviour during the complete observation period; a route was present in the routing table at the beginning of the observation time window and they received no advertisements related to the end-to-end path during it.

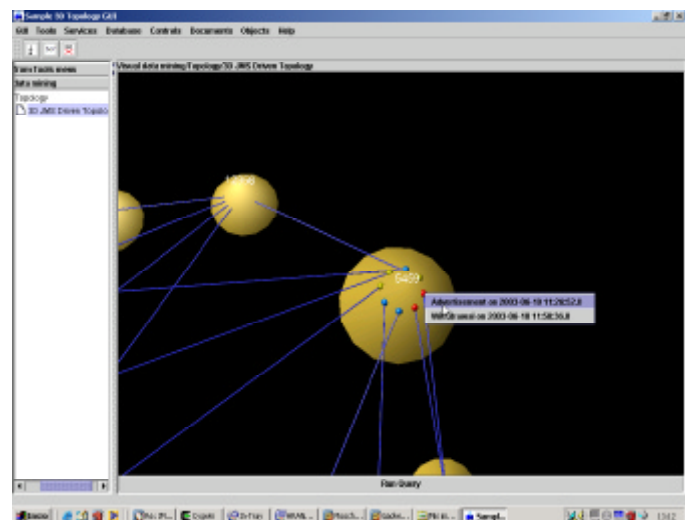
A router will be considered as **warm**, if a route was present at the beginning of the observation period but advertisements of the route were received during the

observation period. The warm state indicates the situation where route stability is not compromised but the deviations from the expected BGP-4 behaviour hint potential problems in the preceding autonomous systems along the path. A router will be **hot** if it received a withdrawal for the route during the observation period. Withdrawals indicate three potentially critical situations with regard to an end-to-end SLA:

- The provider withdrew the route (either accidentally or intentionally) and the end to end connectivity was lost
- A transit provider withdrew the route. This was caused by changes in routing policies or operational errors and might have impacted the end-to-end SLA.
- A provider accidentally injected the route violating the routing policies agreed with his peers and had to withdraw it when the situation was detected. This might have created unwanted alternate paths, which will have impacted the end-to-end SLA.

Use of stability metric in the InterRoute user interface

The following figure shows an autonomous system with its border routers represented using the 'cold-warm-hot' metric translated as a 'blue-yellow-red' colour scale:



This snapshot was directly taken from the INTERMON GUI. The autonomous system shows cold (blue), warm (yellow) and hot (red) border routers. The user has just selected one of the hot routers and its events are presented as a dropdown menu. A route advertisement, followed by a route withdrawal can be observed. This dropdown menu launches other datamining applications, which show the behaviour of

the end-to-end QoS parameters during a time window, which is centred at the time of the routing event.

4. Conclusions

Routing stability is a key for providing a predictable end-to-end service through the Internet. A metric for border router stability based on this principle has been proposed for integration in InterRoute tool of the INTERMON toolkit architecture. This simple stability metric has high expressive potential, but might require fine-tuning to enhance its accuracy. Although first trials indicate that the metric helps detect potential problems, there are several directions for further study.

Future work:

Fine-graining of the metric to accommodate more router states: This metric directs the network administrator towards potential problem areas, but further studies need to clarify if a finer graded metric with more intermediate values could enhance the tool's expressive power.

Restructuring event categories: A route advertisement, which will cause a new route to be installed in the router, might indicate routing problems. In any case, it is a routing event, which should prompt observation of the end-to-end SLA in a time window around its occurrence. It therefore might be necessary to reclassify this router.

Other similar cases need to be studied:

Basing the metric on the actual impact on the routing tables of the border routers: an implementation of the portions of the BGP-4 routing tables which manage the routing tables should tighten the relationship between the metric and the routing packet forwarding paths, thus the QoS experienced by packets in a particular router.

5. References

- INTERMON IST-2001-34123 <http://www.ist-intermon.org>
- Oregon Routeviews Project - <http://antc.uoregon.edu/route-views/>
- RIPE Route Repositories - <http://http://www.ripe.net/>
- MRT - <http://www.mrtd.org>
- Rocketfuel - <http://www.acm.org/sigcomm/sigcomm2002/papers/rocketfuel.pdf>

Testbeds

Experimental INTERMON testbed Europe –Latin America

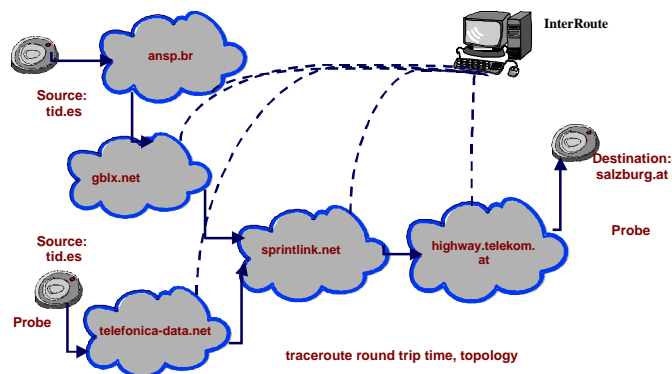
An experimental testbed is established to test and validate the INTERMON toolkit architecture in large-scale international inter-domain environment involving countries from Europe (Spain, Austria, Germany) and Latin America (Brazil).

Currently, two tools are integrated in the INTERMON toolkit, which is installed in the testbed:

- CM Toolset agents are installed at the end systems for QoS measurement of emulated traffic in inter-domain environment.
- InterRoute tool is integrated in the testbed to analyse the inter-domain connectivity and border router stability based on BGP-4 protocol monitoring.

The testbed is established for scenarios aimed at topology and performance analysis on the selected inter-domain paths. Additional scenarios will be aimed to study relationships of end-to-end topology and QoS behaviour with the spatio-temporal QoS analyser considering QoS patterns (outliers).

The following picture shows the topology of the end-to-end inter-domain paths involved in the testbed on autonomous systems.



In particular, the goals of the topology analysis based on the InterRoute are aimed to investigate:

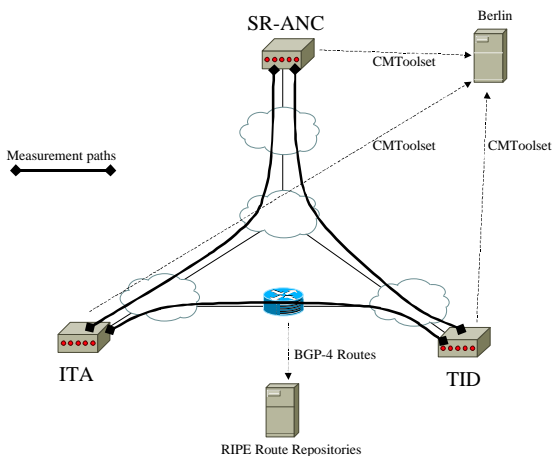
- Stability of the border router of the inter-domain connection.
- Changes in router of the inter-domain path.

- Multihoming issues, and usage of different transit autonomous systems.

The first topology measurements using CM Toolset have shown that the inter-domain connection Brazil-Austria is much more unstable than Spain-Austria. Looking on the inter-domain connectivity picture, there are two transit autonomous systems – Sprintlink and Highway Telekom, which are part of the two inter-domain connections. Therefore, the cause for the instability could be found more probably for the rest of the Brazil-Austria connection – the Global Crossing backbone (gblx.net) and ansp.br network.

Considering earlier experiments between Spain and Austria, special attention should be given to the core router, which could cause poor behaviour due performance degradation when core routers are changed.

The measurement scenario is aimed to study the emulated traffic behaviour on inter-domain connections with the same destination from different sources using different transit autonomous systems.



The emulated traffic is defined for each of the inter-domain connection.

This testbed will be extended in the next time with INTERMON measurement components installed at end systems in Canada in order to study more complex inter-domain communication scenarios.