

EULER Technical Advisory Board Meeting
UCL, Louvain-la-Neuve, Belgium
11th December, 2012

Participants:

Dimitri Papadimitriou, EULER
Nicolai Leymann, Chief Architect at Deutsche Telekom, Germany
Thomas Beckhaus, Senior Manager Fixed Mobile Engineering, Deutsche Telekom, Germany.
Jerome Galtier, Researcher at Orange Labs, France.
Davide Careglio, EULER

The TAB meeting is structured following a presentation of the ongoing research and results of the EULER project. The main idea is to confirm with the industry the directions of the research within the EULER project.

Initially the discussion focusses on the project objectives, i.e. to find a routing scheme with a tradeoff between the memory space required for the routing table and the stretch (length of the paths). Additionally it is desired that the scheme dynamically adapts to network changes (due to changes in the topology, policy, etc) with reasonable communication and computation costs. EULER members classify the efforts in three directions: i) to improve current BGP routing scheme, ii) to change BGP to another Path-Vector routing scheme and iii) to change BGP with a new completely different routing scheme. Therefore EULER works in all three directions.

To better understand BGP, its limits and its problems, EULER has carried out a number of experiments. The results of these experiments help to see how BGP can be improved. With this goal in mind the routing stability has been studied and results are presented for discussion within the Board. Industry members express their interest in the results of these experiments. They agree that sometimes there are several routes with stability problems however to experiment with the core network is a very difficult thing. They are also reluctant to give access to their networks or data traces from their companies in order to carry out experiments even if they agree that there is a need to better understand the behavior of BGP.

EULER proposed a path stability metric that could even be used instead of or along with the local preference parameter in order to select more stable routes. However, the network operators say that they need a strong incentive / motivation in order to use this (or other) modifications of BGP and make public the stability (or instability) of their own networks. Also it is clear that there could be a gain in stability from EULER results but the operators express their doubts about what is actually the gaining regarding to short and long term instabilities, which are difficult to analyze. Operators recognize that when they detect changes in BGP routes they are afraid that some neighboring AS may be changing policies or other things that results in much more traffic through their networks, affecting their costs.

The operators consider that it is also worth to study the relationship of the BGP instabilities with the load of the network and where this instabilities are originated (or in what zone of the topology there are more instabilities) the edge or the core? They are also interested in knowing the cause of the instabilities. EULER members said that it is easier to study inter-AS instabilities but that there also are intra-AS instabilities. The last ones are very difficult to study due to they may be caused by many conditions such as the existence of MPLS tunnels, load balancing mechanisms, congestion scenarios, etc. All participants agree that it is interesting that EULER project continues analyzing the BGP stability problems although they are aware of the problem of lack of real operator information.

The discussion then moves to the issue of having a real interdomain multicast protocol. EULER representative shows the results of the analysis of a huge amount of traffic from an Internet exchange point in Spain encompassing 50,000 users. As it is known, it does not exist interdomain multicast traffic nowadays. Thus, the unicast traffic has been classified by its application/type (e.g. multimedia, web, VoIP, P2P, SSH, SSL, etc.). Particular interest shows the web traffic that includes for instance the YouTube streaming, and the P2P traffic that also includes the P2P video streaming service. The study focus on how many times within a specific time interval (of seconds) the users access to the same contents. I.e. it could have been used a multicast access to those contents instead of multiple unicast accesses. The conclusion is that it is not worthy to use multicast with general contents. However, another experiment, using the retransmission of a football match in streaming show that in this case the use of multicast would have saved an important amount of bandwidth.

Initial discussion focus on what would be the “killer” commercial application that will need to be multicast and that will reduce the cost for the network operators. They agree that it is not a technical issue, it is a political issue. There is no need for interdomain multicast nowadays. One of the key areas where it is important to follow its growth and where it could be “the” place where multicast would be needed in the future is the Mobile Phone Network. Every day the Mobile Phone Network uses more and more the Internet infrastructure.

Industry people said that multicast is in fact used intra-domain. Special cases on video-on-demand service that have for instance 5 replication sites, data centers, special customers that works in parallel in different branches, etc but they said that all these cases have in common that they are almost manually configured and static (i.e. no need of a multicast protocol). The lack of interest from the industry also affects the possible collaborations in giving access to video-on-demand services for data analysis, for instance, or having an implication in performing experiments in their networks. First they need to know what will be the service using multicast. Finally the industry says that some P2P streaming protocols are very bad designed and waste a lot of bandwidth. May be it would be worth to study the inclusion of multicast in some P2P streaming services.

The session concludes with the presentation of the Multicast Compact Routing scheme, designed by the EULER team. This scheme is distributed, dynamic and have many advantages from the compact routing ideas (memory space savings, reduced communication cost, etc). It has been tested using simulation (up to 30,000 nodes and emulation). People from industry show their interest in this protocol but it should still be tested extensively focusing on the adaptability of the method to different services and what would happen in real scenarios, which may have a very different behavior from the simulations. They insist that the Mobile (Phone Network) Multicast scenario, with more and more traffic every day and a highly dynamic situation could be a target to be explored.