Optical Multi-Domain Routing

Xavi Masip and Marcelo Yannuzzi
Technical University of Catalonia, Avda. Victor Balaguer, s/n, 08800 Vilanova i la Geltrú, Spain
{xmasip, yannuzzi} @ac.upc.edu

Abstract: Optical networks provide a clear opportunity to redesign the multi-domain routing paradigm. This paper reviews the current limitations in multi-domain routing as well as some of the research lines in the optical area.

OCIS codes: (060.4250) Networks; (060.4251) Networks, assignment and routing algorithms

1. Introduction

The Internet is a decentralized set of networks known as domains or Autonomous Systems (ASs), each one managed by a single authority and under a common routing policy. Today's Internet is built by the interconnection of more than 29.000 ASs on a well-known network hierarchy. This hierarchical structure is rooted in the two different relationships that could exist between transit ASs, namely, peer-to-peer or customer-provider, hence motivating the definition of different transit ISPs according to their established relationship (Tier-1, Tier-2 or Tier-3).

The problem of selecting routes to forward traffic on an end-to-end scenario is known as multi-domain (or inter-domain) routing. Traditionally, multi-domain routing has been a well-known problem largely analysed in pure IP (packet) networks. The Border Gateway Protocol (BGP) is currently the de facto standard inter-domain routing protocol in the Internet (BGP-4 is the current official release revised in 2006). BGP was basically designed as a protocol to distribute and exchange reachability information. Unfortunately, despite the large efforts the scientific and industrial communities have devoted to multi-domain routing, many issues remain unsolved. These unsolved problems are rooted in: *i)* current network and users' requirements have substantially changed from those motivating the initial protocol design; *ii)* the intricate interactions and dependencies between domains;, and *iii)* ISPs are generally reluctant to introduce changes and test them if there is no clear source of revenue, what definitely hinders real progress.

As the optical technology is getting improved, new features and capacities are to be allocated on the optical side. This means that several of the network issues traditionally addressed only on the IP layer, will need to be faced again in the optical layer. Multi-domain routing is clearly one of these issues, so it is of utter importance to take advantage of the lessons learned in the IP layer and avoid the pitfalls of the past. Yet, the main issue here is that the multi-domain routing problem has nor even been solved completely in the IP layer, so finding the appropriate solution in the optical layer is not expected to be easy.

2. Research Challenges in Optical Multi-Domain Routing

The main objective of whatever solution to be proposed for multi-domain routing is scalability. Scalability is in fact a key aim in multi-domain routing as dealing with the whole network structure. Unfortunately, this is not a single problem to manage, as there are many issues crossly affecting scalability that cannot be addressed on an isolated way, but rather it is fundamental to understand their relationships and dependencies to actively contribute with feasible solutions.

At present, the main set of problems inherent to multi-domain routing in IP networks is perfectly known [1], and is represented on the left-hand side hexagon in Fig.1. This set of problems has never been faced as a whole (this indeed would mean to substitute BGP) but rather only incremental or partial solutions have been adopted. An interesting point here is that this set of problems can be mapped almost directly into a new hexagon (cf., right-hand side of Fig.1), hence defining a set of research challenges [2] for optical multi-domain routing. Next a short description of each of these challenges is presented:

Security: This is a critical issue nowadays. The main problems related to security are: *i)* TCP connections supporting peering ASs' sessions are vulnerable by attackers (TCP end-points can be easily determined) and *ii)* lack of authentication in BGP.

Convergence properties: Current BGP convergence is quite slow (hundreds of seconds), and is mainly rooted in the path exploration process, which is definitely not desirable in optical networking. At the packet layer, most of the suggested solutions aim at reducing the BGP convergence time, but this strategy will not suit an optical scenario. This is mainly due to the fact that during a convergence, the expected traffic loss would be huge as optical connections are supposed to transport large volumes of data. A recently proposed line of work pushes for avoiding

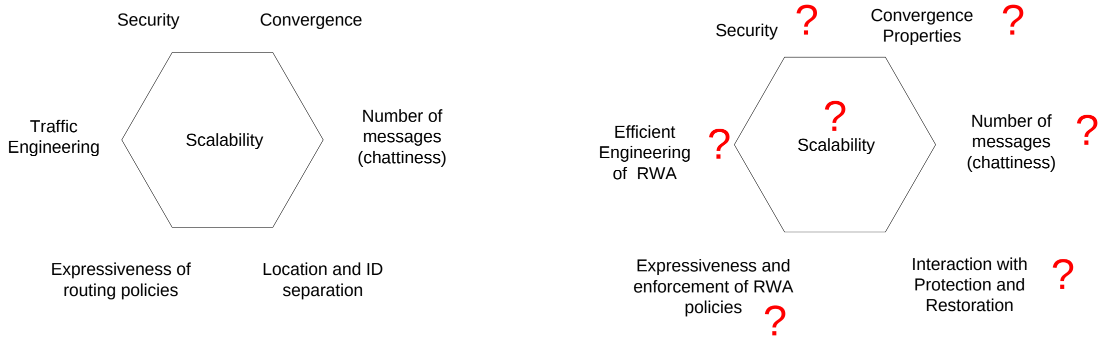

Fig. 1 Research challenges in multi-domain routing (IP layer on the left-hand side, and optical layer on the right-hand side)

elaborating on the convergence time, but rather working on local resilience [3]. This strategy might perfectly fit optical networks constraints.

Chattiness: Keeping more than 280.000 entries on the routing table, means that a large number of messages are expected to be exchanged between domains whenever a change occurs in the network. Optimally, the strategy to be defined must avoid exchanging storms of advertisements, what indeed, is not easy to achieve. Tentative solutions for BGP rely on: *i)* adding a timer (MRAI, minimum route advertisement time) to distance two advertisements for the same prefix, or; *ii)* implementing BGP route flap damping to ignore routes changing too often. None of these solutions are effective, as it is very difficult to define the most suitable value of the MRAI timer, whereas flap damping increases the convergence time. Thus, new solutions should be sought. A promising work in this area, directly affecting scalability, focuses on aggregation. Topology aggregation indeed limits the amount of information to be distributed as well as the bandwidth required for signalling [4].

Interaction with protection and restoration: Protection and restoration clearly impact on convergence, and hence in scalability. Most of the existing proposals target protection between border routes, what despite being robust lacks of domain diversity, so negatively impacting on resilience. Combining protection and restoration is a key research challenge in a multi-domain scenario. The research target aims at defining the right combined strategy fitting best both the autonomic routing policies as well as the connection/traffic requirements of different domains. Besides, another challenging topic with important implications to protection/restoration is grooming.

Expressiveness of routing policies: Current routing performance on the Internet is based on the interaction of different and independent policies implemented on different ASs. This autonomic behaviour can potentially lead to well-known routing problems, such as route oscillations or routing inconsistencies after failures.

Efficient engineering of routing and admission capabilities: Routing and Wavelength Assignment (RWA) on a multi-domain scenario must be designed chasing two main skills, traffic engineering capabilities and QoS support. Both issues are traditional weaknesses in current packet networks, despite being appointed as mandatory requirements for current and future transport networks. Also critical is to define a more effective framework between intra- and inter-domain routing policies. Moreover, secure reservation and the admission scheme should also be linked to guarantee the network and users' agreements as well as to suitably match the defined billing policy.

The main point beyond the conceptual issues defined so far is that all of them have been studied separately in the past. Facing the crossed interactions among them all on a global, comprehensive, integrated, and large-scale solution is what makes the design definitely challenging.

3. Tentative Approaches

There are basically three conceptual paths to tackle the multi-domain routing problem in optical networks. The first line of work pushes for extending BGP to support the advertisement and signalling of optical information between routing domains. The resulting protocol is named OBGP [5]. The main advantage of OBGP is that it is expected to scale as much as BGP. On the other side, its main drawback is that all BGP weaknesses will be present in the optical layer, specially in terms of: *i)* routing and traffic engineering control, since only network reachability information (NRI) is exchanged between domains; *ii)* lack of path diversity since only the best path information is disseminated, and *iii)* slow convergence and chattiness rooted in the path exploration problem.

The second line of work boils down to incrementally improve BGP/OBGP by adding new features to a new protocol. An elemental protocol enhancement could focus on tackling the fact that the BGP/OBGP routing process will generally select the path minimizing the number of traversed ASs (as any path-vector routing protocol), since only NRI is exchanged between network domains. It is widely accepted that some basic functionalities for current networks, such as load balancing, network recovery, or flexible traffic engineering, cannot coexist in this scenario.

This issue may be solved by allowing neighbouring domains to exchange NRI along with aggregated path state information (PSI). Three open concerns should be deeply analyzed here. Firstly, the metrics with the PSI must be defined, for example, an aggregated value capturing the availability of wavelengths or their load along a path [6]. Secondly, a metric dissemination procedure must also be defined. And lastly, the update of the PSI metrics might lead to an increase in both the frequency and number of routing messages exchanged between domains. To avoid this, the PSI updates could be piggy-backed either on Keepalive or signalling messages [6]. However, the fundamental drawbacks of this approach are that it still suffers from limited TE capabilities, slow convergence and chattiness, due to the path exploration problem.

The third and last scenario is the simplest one in terms of wording; replace any reference to BGP, that is, to generate a new multi-domain route control model not based on legacy BGP limitations. A first step forward drives to apply and extend the Path Computation Element (PCE) basics [7], that's to say having a decoupled control plane (splitting control and data planes). Aligned to this concept, the Inter-Domain Routing Agent (IDRA)-based architecture was recently proposed in [8]. In this architecture an agent (the IDRA) is allocated on each AS with the basic functions of: *i)* distributing routing and signalling information between network domains, and *ii)* computing and establishing the inter-domain lightpaths. The main weaknesses of whatever solution presented in this area are that, the solution must be widely tested and analysed (what, unluckily, is almost impossible in a real large-scale scenario), and that standardization efforts must be initiated to converge on a commonly accepted solution (which would be a challenging task given the need to deal with a fully new architecture).

In any case, whatever the proposed solution is going to be, there are three vital issues that must be undeniably taken into account. First, different ASs must not only coexist but also collaborate (for example, through pre-established SLAs) on the traffic forwarding process. This means that the current selfish behaviour must be replaced by a more regulated one, but providing, at the same time, the ability for domains to supply different views of their connectivity status. Upon doing so, the network information distributed to a peering domain would depend on the agreement defined between them. Second, the interaction between the SLAs and the future export policies between optical domains are the instruments to provide these different views. Third and last, it must be noticed that the packet and optical layers will work as one in the coming years. This means that providing similar solutions to both layers (i.e. duplicating functions) would neither be optimal nor efficient. Instead, research efforts must be devoted to explore and design new cross-layer solutions for future multi-domain optical networks on a multi-layer control architecture.

4. Conclusions

Many tentative approaches have recently risen to address the multi-domain routing problem in optical networks. Most of them are based on either extending or adapting the current solutions for the packet layer. The issue here is that the key limitations in multi-domain routing have not been solved in the packet layer; hence many issues remain still open. In this sense, devoting efforts to develop protocol extensions will only drive to solutions partially solving the whole problem. Instead, multi-domain optical networks represent a clear break-point offering an opportunity to make a change, and explore alternative routing paradigms to avoid inheriting the well-known issues in BGP.

Furthermore, including additional capabilities to the optical layer should not involve having duplicate functions on the packet and optical layers. Hence, new cross-layer features must be added toward the integration of a multi-layer control plane, with the aim of optimizing both the network performance as well as its management.

Finally, it is not a given that the proposed solution must be based on current protocols. In other words, if optical networks offer the chance to address from scratch the multi-domain routing problem, why should not we push for a fully new cross-layer multi-domain architecture?

Acknowledgments: This work was partially funded by the MCyT (Spanish Ministry of Science and Technology) under contract TEC2005-08051-C03-01 and the Catalan Government under contract 2005 SGR00481

References

[1] M.Yannuzzi, O.Bonaventure, X.Masip, "Open Issues in Interdomain Routing: A Survey", IEEE Network, Vol.19, nº.6, November 2005.
[2] N.Ghani at al, "Control Plane Design in Multidomian/Multilayer Opical Networks", IEEE Communications Mag., Vol.46, nº.6, June 2008.
[3] O.Bonaventure, "Reconsidering the Internet Routing Architecture", IRTF draft, <draft-bonaventure-irtf-rrg-rira-00.txt>, March 2007.
[4] G.Maier, C.Busca, A.Pattavina, "Multi-Domain Routing Techniques with Topology Aggregation in ASON Networks", ONDM2008, Vilanova i la Geltrú, Spain, March 2008.
[5] M.Blanchet, F.Parent, B.St-Arnaud, "Optical BGP (OBGP): InterAS Lightpath Provisioning". <ietf-draft-parent-obgp-01>, March 2001.
[6] M.Yannuzzi, X.Masip, G.Fabregó, S.Sánchez, J.Solé, "OBGP+: A Simple Approach to Drastically Improve OBGP", ONDM 2008, Vilanova i la Geltrú, Spain, March 2008.
[7] Path Computation Element (PCE) WG, IETF, http://www.ietf.org/html.charters/pce-charter.html.
[8] M.Yannuzzi, X.Masip, G.Fabregó, S.Sánchez, A.Sprintson, A.Orda, "Toward a New Route Control Model for Multidomain Optical Networks", IEEE Communications Mag., Vol.46, nº.6, June 2008.