Isolario: a Do-ut-des Approach to Improve the Appeal of BGP Route Collecting

Pietro Giardina,Enrico Gregori,Alessandro Improta,Alessandro Pischedda,Luca Sani
Institute of Informatics and Telematics, Italian National Research Council
Pisa, Italy
{name.surname}@iit.cnr.it

ABSTRACT

The incompleteness of data collected from BGP route collecting projects is a well-known issue. As pointed out in recent works, one of the possible solutions is to increase the amount of ASes connected to a route collecting project, starting from those ASes located in the Internet periphery. In this paper we present Isolario, a project based on the do-ut-des principle to increase the appeal of BGP route collecting. Different from other projects, Isolario aims at persuading network administrators in small-medium organizations to share their routing information by offering services in return, ranging from real-time analyses of the incoming BGP session(s) to historic analyses of routing reachability. These services are implemented on the Isolario infrastructure and are accessible by the network administrators of the ASes participating the project via commonly used web browsers.

1. INTRODUCTION

The route collectors deployed by the Route Views project of the University of Oregon and the Routing Information Service (RIS) project of the Réseaux IP Européens Network Coordination Center (RIPE NCC) have been an invaluable source of information for researcher all over the world for the past twenty years. However it is not all a bed of roses. Data collected by Route Views and RIS is indeed known to be largely incomplete [1, 2]. Moreover, most of the organisations which decided to participate in these projects are very large Internet Service Providers (ISPs) [1] and, as a consequence, data collected by these projects could be biased, depending on the analysis carried out. For example, they largely fail to reveal the largest part of the peering ecosystem [3, 4] mostly due to the nature of their participants. We believe that the main cause of these drawbacks is the voluntary basis on which these projects rely on. Most of the largest ISPs are indeed attracted by the opportunity to exhibit their interconnections to potential customers, while – in our opinion – the administrators of the smallest organisations may not find any strong motivation to join. Given the economic relationships ruling the Internet [5], these latter companies are however the most important sources of information to reveal characteristics of the Internet previously unknown via their hypothetical BGP feed to a route collector [1]. For example, most of these ASes would be able to discover part of the peering connections established at IXPs by them directly (or by their providers), thus potentially contributing to turn upside-down the hierarchical view of the Internet.

In this paper we introduce Isolario, a route collecting infrastructure which aims to improve the knowledge about the Internet ecosystem by enhancing the appeal of the classic concept of BGP route collector, trying to attract smaller ISPs and CDNs into route collecting. To do that, in addition to simple data collection, Isolario provides the network administrators of each AS participating to the project a set of *real-time* and *historic* services focused on their inter-domain routing, in exchange for (at least) one BGP session from one of their border routers. To achieve that, Isolario parses, stores, filters and redirects each packet from the incoming BGP flows towards dedicated modules which implement the services. Then, the result is provided to Isolario users through an HTML5 website which exploits the *WebSocket* protocol (RFC 6455) to update web pages only when new data related to the user becomes available, without additional polling traffic being generated. Since Isolario is lead by the *do-ut-des*[1] paradigm its services are dedicated to its participants only, while MRT data is made publicly available as every other BGP route collecting project described above with the aim to improve the state of the art about the Internet structure knowledge.

2. ISOLARIO ARCHITECTURE

Isolario is a distributed system devised to collect, parse and elaborate BGP data sent from the routers of its participants – hereafter *feeders* – and to provide results to network administrators of the participating ASes by introducing the minimum amount of delay as possible. To achieve this target, Isolario makes use of three main components, as depicted in Figure 1: *i*) Web core, *ii*) System core, and *iii*) Enhanced Route Collec-

[1] Latin expression for "*I give so that you will give*"

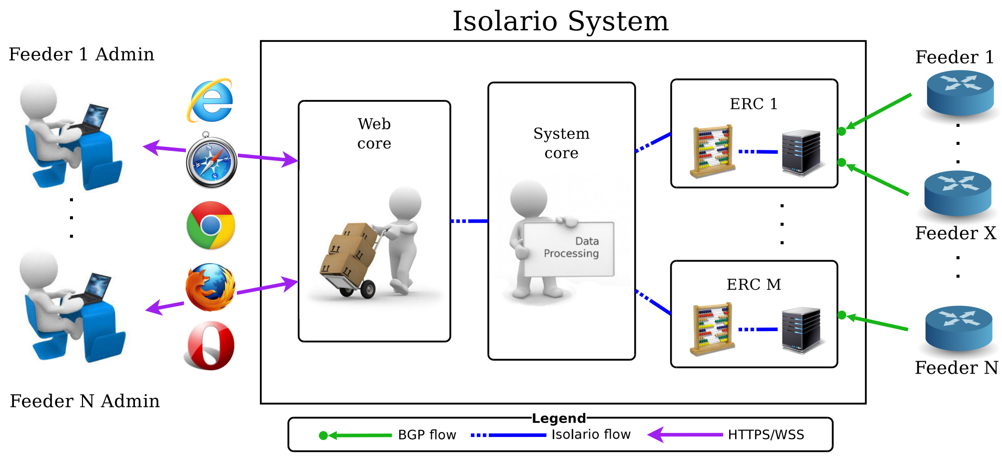

Figure 1: Isolario system overview

tors (ERCs). Each component is designed to be modular and scalable, to allow the introduction of new pieces of equipment in a plug-and-play fashion without affecting the user experience.

Isolario system has been designed to fulfill two main purposes. On one side, it must establish and maintain BGP sessions with each feeder, and has to collect routing information in the MRT repository. On the other side, it must react to user service inputs by computing and elaborating BGP incoming packets or historic data and providing the requested service output(s). An Isolario user interacts with the *web core* component from a private area of the Isolario website, where every service available can be found. As soon as a service has been chosen, it is created a dedicated WebSocket session between the Web core and the user web browser on which user requests and the results will flow in real-time. WebSockets are chosen to allow the server to send data directly to the user browser as soon as new data is available, avoiding costly and unuseful polling traffic from the web browser or explicit refresh requests from the user. On the Web core side, HTTPS and WebSocket flows are dispatched by a proxy module to either the Webserver (HTTPS) or to the related service module (WebSocket), which will propagate the traffic to the system core. To minimize the amount of unuseful routing data in the system, each service also identifies which portion of IPv4/IPv6 space is interested by the user query, and generates dedicated messages to configure filters on the proper ERC(s). The *system core* component (Figure **??**) then handles user requests propagated by the web core as well as filter configuration messages. This component is composed by a set of real-time service modules and a set of historic service modules. The former set of service modules dispatches user requests and filter configuration messages to the related service modules located on ERCs, and aggregates results based on the partial results received from each ERC. On the opposite, historic services fetch data from a dedicated software which is designed to allow fast access to stored routing data. Lastly, the *enhanced route collectors* handle user requests and filter configuration messages. Each ERC has the role to establish and maintain active BGP sessions towards a set of feeders, dispatching every single BGP packet which matches filtering criteria towards service modules. Every BGP session on ERCs is maintained by a custom route collecting software, which also shoots every 2 hours a snapshot of its RIB and dumps every 5 minutes the full sequence of BGP packets received in the last 5 minutes in MRT format. To reduce the delays in elaboration, we introduced a dedicated module between each feeder and our system, which forwards the BGP packets to the RCE and applies the filter configuration messages received from the system core. Every single incoming BGP packet matching the filter criteria is then forwarded also to the proper service module, which will perform all the required elaborations. Results obtained from service modules are finally propagated back to the user through the Isolario system and the WebSocket session.

3. SERVICE OVERVIEW

Services are the *do* part of the *do-ut-des* principle, and thus the main reason of the participation of the user to the project and the key to attract new potential users. So far, we developed four main categories of services, depending on the data source exploited: *flow-based* services, *subnet-based* services, *historic* services and *alerting* services. Here in this section we provide a brief description of each of these categories, together with some examples of services already available to use.

3.1 Flow-based services

Some of the possible routing problems that a common network administrator could face during a normal workday are caused by the behaviour of its BGP neighbours. In theory a network administrator should monitor the routing tables of ASBRs to understand at glance whether the AS is experiencing these problems, but this is not feasible. Isolario provides a valid solution by offering to each Isolario user a set of *real-time* monitoring services based on the BGP flow(s) that the user organisation established with Isolario, without increasing the CPU load on the peering ASBR with multiple CLI accesses or requiring the installation of additional software on the user-side.

The schema followed by this class of services is depicted in Figure 2. As soon as the user triggers the requests from the web browser, the related service generates a proper set of filter configuration messages to allow BGP data of interest to flow in the system. The set of filter configuration message and their effect on the system strictly depend on the needs of each service. The filtered UPDATE messages are then parsed and elaborated from the given service module on ERCs and the requested results are provided back to the user, together with some related statistics.

Up to date, every Isolario user can use **BGP Flow Viewer (BFV)**, which shows the `UPDATE` messages that the Isolario route collector is receiving from the ASBR(s) managed by the user, and which allows to quantify the amount of routing traffic currently generated by their ASBR, **Routing Table Viewer (RTV)**, which allows users to analyse portions of the routing table announced by their ASBR(s) – which coincide with the best routes identified by the BGP decision process of the ASBR – and to monitor their evolution, allowing users to investigate possible reachability problems in real-time, and **Route Flap Detector (RFD)**, which allows users to detect route flaps in real-time, allowing them to take the opportune countermeasures.

3.2 Subnet-based services

One of the most (work related) desires of a common network administrator is that the networks he/she manages are reachable from everyone everywhere at anytime. This desire cannot be achieved simply with correct management of the networks of the AS managed,

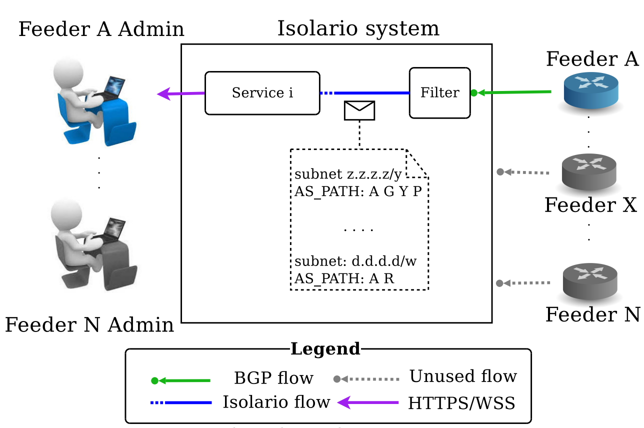

Figure 2: Flow-based services overview

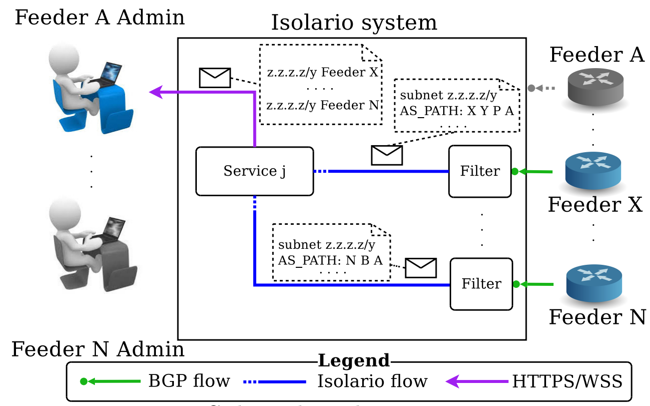

Figure 3: Subnet-based services overview

and it requires also excellent network planning skills to choose the best BGP neighbours and their role (provider or peer). The main problem is that it is extremely hard for any network administrator to understand and prove the efficacy of the choice made. This because every single AS has its own perspective of the Internet, and there is no collaboration in creating a big picture. The only public tools that network administrator could use so far to obtain a different perspective about their networks are looking glasses (LGs) or MRT data collected by Route Views and RIS. On one side, LGs are easy to access but can offer just a snapshot of the reachability of the network under analysis from the single perspective of the AS which provide the LG access. On the other side, MRT data can provide more points of view than LGs, but their analysis require know-how and the development of dedicated software. Moreover, they can just show the past behaviour of network reachability, and not the present. Given its innovative structure, Isolario is able to merge the pros of the two approaches *and* to provide results in *real-time*, thus allowing network administrators to check the reachability of their networks with just a couple of clicks.

The schema followed by this class of services is depicted in Figure 3. Each service generates a set of filter configuration messages to allow BGP data to flow in the system but, differently from flow-based services, these messages are sent to filters regulating *every* incoming BGP flow. The filtered UPDATE messages are then parsed and elaborated from the given service module locally on ERCs and then each partial piece of information is merged in the service module in the system core. Finally, the requested results are provided to the user, again together with some related statistics. So far, the only service currently available in this class is the self-descriptive **My Subnet Reachability (MSR)**, but we are currently working on the development of the **Outage detector (OD)** service, which would allow each Isolario user to understand which part of the Internet are experiencing problems by aggregating routing information from every feeder.

3.3 Historic data services

Historia magistra vitae[2] is an old proverb which means that the study of the past should serve as a lesson to the future. This also applies to network administrators, which may learn from the analysis of `UPDATE` messages announced in the past something about the resilience and robustness of their routing system. A window in the past could help in troubleshooting current routing problems and could help in choosing the correct peers and providers to diversify the connectivity of their AS. Historic information are thus extremely useful, but to date network administrators were required to analyse MRT data by themselves, or to create their own historic repository privately inside their own AS, creating their private route collecting systems. Isolario exploits a dedicated module able to retrieve as fast as possible historic routing information, to provide each of its users something similar to a time machine. So far, we developed the historic aliases of some of the flow-based and subnet-based services described above, **Historic Routing Table Viewer (HRTV)** and **Historic My Subnet Reachability**.

3.4 Alerting services

Services described so far assume the physical presence of users, which need to actively interact with Isolario via their web browser. According to Murphy's law, outages and peculiar network events are more likely to happen when the network administrator is not checking real-time services for whatever reason than when they are logged in Isolario. Needless to say, in several cases it is extremely important that network administrators act as fast as possible. To help them in that, Isolario devised a special class of services aimed at monitoring BGP Flows provided by feeders in order to catch and notify interesting events for users, like route flaps, temporary network unreachability or prefix hijack attempts.

Every service in this class runs 24/7 following the rules that each user set up in the service web page. Depending on the service chosen, filter configuration messages will flow following the schema described in flow-based or subnet-based services, with the service configuration set by the user. As soon as BGP message of interest reach the system core, the service will identify whether an alarm has to be triggered or not. If an alarm is triggered, every data related will be delivered to the interested user through the communication channels indicated by the user in the alarm configuration phase (e.g. e-mail, HTTP/HTTPS POST messages).

Besides the alerting features described above, a further interesting example of service in this class is the **Prefix Hijack Detector (PHD)**. PHD is currently set up to identify from every single perspective available in Isolario any attempt of hijack of any subnet belonging to one of the Isolario feeders and communicate that to the related user.

[2]Latin expression for *History is life's teacher.*

4. CONCLUSIONS AND FUTURE WORK

The lack of interest in route collecting project is the major cause of the incompleteness of BGP data collected by Route Views and RIS. To increase the interest of network administrator in sharing their routing information, we proposed Isolario, an enhanced route collecting project based on the do-ut-des principle. To achieve that, Isolario offers a set of useful services to help network administrators in troubleshooting network reachability problems in exchange for their IPv4/IPv6 BGP full routing tables. These services and the whole Isolario architecture have been developed with a particular care to user experience in order to attract as many network administrators as possible.

Isolario is currently about to complete the private test phase and start the deployment phase, accepting new feeders. In the new phase MRT data collected from the feeders will be made available and a free trial of Isolario services will be offered to the public. To join Isolario, we require feeders to establish (at least) one BGP session with one of our route collectors and to announce us their routes towards all the Internet destinations. Interested organisations can request to participate from the "Join Us" page on Isolario website (https://www.isolario.it). In the very near future we plan to increase the number of services available to broaden the usefulness of Isolario and to ease the user access to services by developing a mobile application and a set of dedicated APIs.

5. REFERENCES

[1] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani, "On the Incompleteness of the AS-level Graph: a Novel Methodology for BGP Route Collector Placement," in *IMC'12*, pp. 253–264.

[2] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (In)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM ToN*, vol. 18, no. 1, pp. 109–122, 2010.

[3] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy, "Lord of the Links: A Framework for Discovering Missing Links in the Internet Topology," *IEEE/ACM ToN*, vol. 17, no. 2, pp. 391–404, 2009.

[4] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is More to IXPs Than Meets the Eye," *ACM CCR*, vol. 43, pp. 19–28, 2013.

[5] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM ToN*, vol. 9, no. 6, pp. 733–745, 2001.