

Isolario: the real-time Internet routing  
observatory

**Alessandro Improta**  
alessandro.improta@iit.cnr.it

**Luca Sani**  
luca.sani@iit.cnr.it

ITNOG-2, 3 November 2016, Bologna



Consiglio Nazionale delle Ricerche  
Istituto di Informatica e Telematica



# Unveiling the Internet structure with BGP data

## BGP route collectors

BGP data collected up to date has been unvaluable to reveal the Internet inter-domain characteristics, but it is known to be largely *incomplete*



## How much incomplete? (October 2016)

It was possible to discover the full connectivity of:

- 961 out of 9799 ASes (9.81%) which transit traffic for other ASes
- 70 out of 277 ASes (25.27%) of those operating in Italy


Main cause: small number of small ASes connected

Do AS administrators see any direct outcome in sharing their routing information?

# Isolario project

Objective: push more ASes to join

The more the ASes, the more the completeness of public BGP data



Isolario - The Book of Islands

*"where we discuss about all islands of the world, with their ancient and modern names, histories, tales and way of living..."*

Benedetto Bordone  
(Italian cartographer)

Approach: Do-ut-des

- Participants open a BGP session with Isolario providing the BGP full routing table and its evolution over time
- In change, Isolario offers **real-time** and **historic** analysis applications based on the aggregation of every routing information collected

# Data we plan to provide to research community

## MRT data (same format as RIPE RIS, Route Views, ...)

- 1 RIB feeder snapshots every 2 hours
- 2 UPDATE collections every 5 minutes

[https://isolario.it/Isolario\\_MRT\\_data/](https://isolario.it/Isolario_MRT_data/)

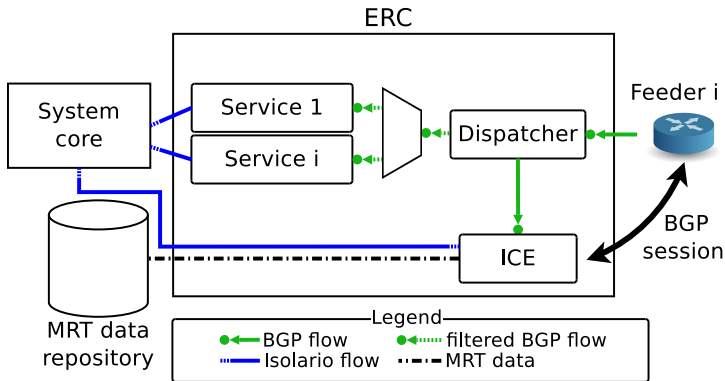
## Periodic analyses (daily, weekly, monthly, ...)

- 1 AS-level Topologies (Global and Geographic)
- 2 AS characteristics
- 3 Feeder contribution
- 4 Total coverage of RCs



# Enhanced BGP Route Collector

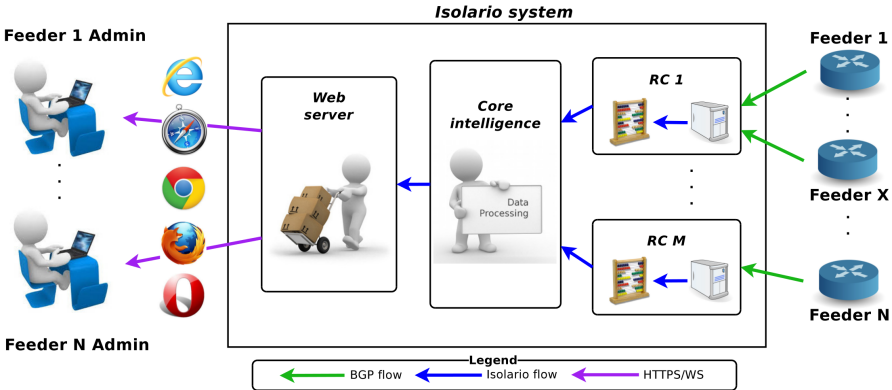
Incoming flows are duplicated as soon as they arrive and feed both the Interactive Collecting Engine (ICE) and service modules



As usual, RCs only collect routing information and **not** user traffic

# Isolario system overview

Incoming BGP flows are used as **real-time streams** for services dedicated to participants



Results are provided to users via WebSockets

# Isolario free services for feeders

Every feeder has free access to a set of services tailored to monitor and analyse BGP data coming into Isolario system

## Real-time services

- BGP flow viewer
- Routing table viewer
- Website reachability
- Subnet reachability

## Historic services




work in progress

- Routing table viewer
- Subnet reachability

## Diagnostic services

- Alerting system
- Daily report

# Real-time services



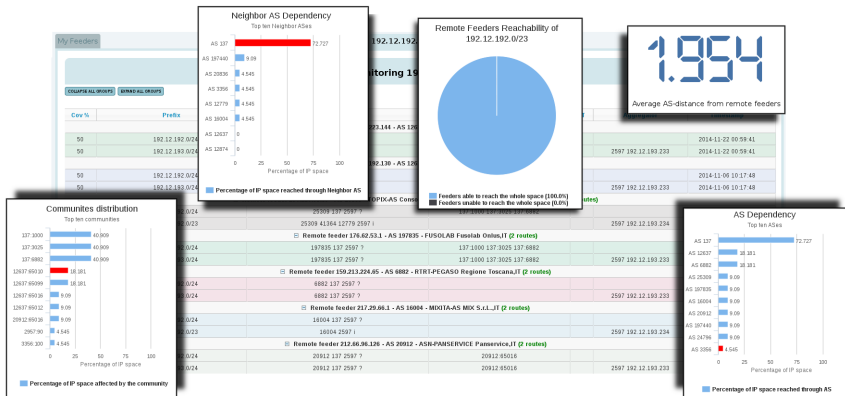
Real-time services allow to monitor BGP data flowing into Isolario system





# Subnet reachability

Allows to analyse in real-time the routes that every Isolario feeder is announcing to Isolario to reach a portion of the IP space



The more the feeders, the more SR is useful!



## Diagnostic services



Diagnostic services exploit incoming BGP flows and/or historic data to report anomalies of the inter-domain routing status

# Alerting system

## Alerting system

- **BGP attributes:** BGP UPDATEs matching attributes of interest
- **Flap events:** a prefix UPDATE rate is larger than a threshold
- **Hijack attempts:** BGP UPDATEs hijacking a feeder subnet
- **Prefix reachability:** (un)reachability of prefixes of interest

The screenshot shows the 'Alerting Management' interface with the following components:

- Navigation tabs:** Alerting Management (active), Notifications, Current configured alerts.
- Header:** Create new alert (Tip: often the elements of the interface displayed below have a help text that will be shown simply by leaving the mouse on the element itself.)
- General Alert Options:**
  - Available feeder IPs:** A list of IP addresses with their corresponding ASNs, each with a checkbox. A 'Save Alert' button is located below this list.
  - Alert Type:** A dropdown menu with 'BGP attributes' selected. Other options include Flap Detector, Session Watchdog, Hijack, and Reachability.
  - Action upon event:** Includes a checked checkbox for 'Email' with a text input field containing '3600' and a 's' unit, and an unchecked checkbox for 'POST HTTP(s)'.
- BGP attributes:** A section containing several buttons: Prefix, Prefix Subnet, Community, Prefix Related, AS path end, AS path substrings, AS path begin, AS path exact, Origin, and Aggregator.
- Text:** A paragraph explaining that multiple BGP attribute types can be combined using .and/or operators and round brackets. It states that the system will report any BGP\_UPDATE message advertised by one of the selected feeder IPs matching the inserted attributes.
- EXAMPLE:** A section for providing an example configuration.
- Footer:** A label 'Current BGP attribute types selected' above a list area.

# Daily report

Summary about the feeder inter-domain routing status as perceived by the Isolario system

For example...

## Routing statistics

- #Announce, #Withdrawn
- Most (un)stable prefixes

## Reachability statistics

- Inbound reachability

## BGP attributes statistics

- AS path anomalies



Daily report

Feeder 192.65.131.235 (AS 2598)

Thursday 21<sup>st</sup> May, 2015



Consiglio Nazionale delle Ricerche  
Istituto di Informatica e Telematica



# Daily report: Summary of statistics

## 1 General statistics

Analysis start date: *Thursday 21 May 2015 at 00:00:00*

Analysis end date: *Thursday 21 May 2015 at 23:59:59*

Number of non overlapping IPv4 space covered<sup>1</sup>: *2739704260 (98.581001 %)*

The remaining 1.418999 % is covered by a default route

Packets received: *227490*

Feeder status at end date: *up*

Downs experienced since start date: *0*

## 2 Route statistics

Subnets: *532099*

Unstable subnets: *57727 (10.848 %)*

Stable subnets: *474372 (89.151001 %)*

Number of reserved subnets: *1* – see Sect. 2.4 for further details

Geolocated subnets<sup>2</sup>: *475610 (89.383003 %)*

## 5 AS statistics

ASes seen: *50241*

Private ASes: *34 (0.067 %)*

Public ASes: *50207 (99.931999 %)*

Public ASes on 16 bits: *42864 (85.316002 %)*

Public ASes on 32 bits: *7343 (14.615 %)*

Number of public ASes at start date: *50089*

Number of public ASes at end date: *50142*

Difference: *+53 ASes (+0.105 %)*

## 7 My subnet statistics

Total number of subnets perceived as proprietary: *1*


Subnet  
192.65.131.0/24

Number of events related to proprietary subnets: *0*


Number of announcements related to proprietary subnets: *0*

Number of withdrawals related to proprietary subnets: *0*

Figure 1: Amount of packets received per hour



# Historic services



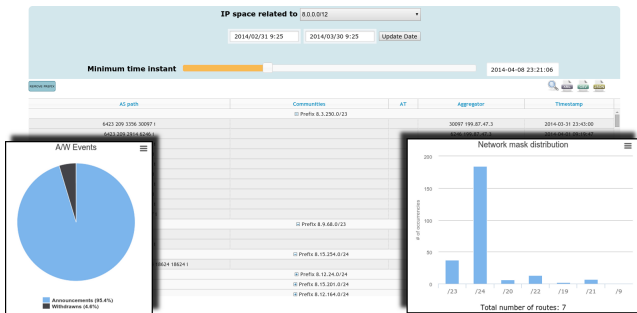
Historic services exploit every BGP data available (Route Views, RIPE NCC RIS, Isolario) to show how routes evolved in the past



## Applications



- **Routing table viewer:** Allows to analyse portion(s) of the routing table that each feeder announced to Isolario
- **Subnet reachability:** Allows to analyse the reachability of the IP space portions from every feeder available in the past



# Summary: how to use Isolario?

## Real-time services

### **Something is happening**

How is my RIB(s) evolving?  
How is my reachability affected?

## Alerting System

### **Something is happening NOW!**

Check real-time services!  
Do something! (if needed)

## Daily report

### **Did something happen yesterday?**


Check historic services!  
Do something! (if needed)

## Historic services



### **Something happened**

How was my RIB(s) evolving?  
How was my reachability affected?



# Summary: how to use Isolario?

## Real-time services

### Something is happening

How is my RIB(s) evolving?  
How is my reachability affected?

## Historic services

### Something happened

How was my RIB(s) evolving?  
How was my reachability affected?



## Alerting System

### Something is happening


Check real-time  
Do something!

Please, try Isolario real-time services!

<https://www.isolario.it>

Username: *guest*

Password: *guest*



## Daily reports

### Did something happen yesterday?

Check historic services!  
Do something! (if needed)



# Current status

## Feeders

- 38 ASes
  - 1 AE, 1 BR, 1 CH, 4 DE, 1 EE, 24 IT, 1 MX, 1 NL, 2 UK, 3 US
- 50 IPv4 sessions
- 36 IPv6 sessions

## Hardware (everything located in Pisa, IT)

- 6 route collectors (Dell PowerEdge R420/R430)
- 1 real-time core (Dell PowerEdge R620)
- 1 non real-time core (Dell PowerEdge R810)
- 4 storages (Dell PowerEdge R420/R430)
- 1 webserver (Dell PowerEdge R420)

# Open-source software (C++)

## ICE - Interactive Collecting Engine

- *Interactive* BGP route collecting software
- Establishes and maintains BGP sessions, dumps MRT files
- *Multithread* and – thus – very responsive to human/automatic queries!
- Possibility to activate LZW-like compression to reduce memory consumption

## MDR - MRT Data Reader

- Tool to parse MRT files (RIB snapshots and updates)
- Easy to integrate in custom software

# Open-source software (C++)

## SG - Subnet Geolocator



- Tool to map subnets and/or ASes to continents/countries
- Takes as input a mapping list e.g. the *GeoLite City DB* provided by MaxMind [1]
- e.g. 223.64.0.0/11 → CN|HK
- e.g. 37514 → KE

---

[1] [http://geolite.maxmind.com/download/geoip/database/GeoLiteCity\\_CSV/GeoLiteCity-latest.zip](http://geolite.maxmind.com/download/geoip/database/GeoLiteCity_CSV/GeoLiteCity-latest.zip) (v4)  
<http://geolite.maxmind.com/download/geoip/database/GeoLiteCityv6-beta/GeoLiteCityv6.csv.gz> (v6)

Future directions

# IXPs services

We started a discussion with IXP people about possible services that could be useful for IXP participants

## Real-time looking glass

- An enhanced version of the classic looking-glass software
- Real-time visualization of routing events
- A BGP session is established between the router and ICE
- Queries are handled by ICE and not by the router
- e.g. Real-time monitoring of route-server BGP tables

## (Almost) Zero-configuration alerting service

Notify IXP participants whenever a routing event (i.e. a BGP\_UPDATE) involving his/her networks is received by Isolario route collectors



# Global deployment of route collectors

## Distribute route collectors at several locations around the world

- Route collector anycast
- Multiple web servers?

## Collaborations

- Packet Clearing House (PCH)
- UniLaSalle (Brazil)
- We are open to any kind of collaboration

## Main objective

- To improve the knowledge of Internet structures of developing/third-world countries
- To improve the effectiveness of monitoring services

Thank you for your attention



**Join us and help us to unveil the Internet AS-level structure!**

To participate, contact us at:  
**[info@isolario.it](mailto:info@isolario.it)**



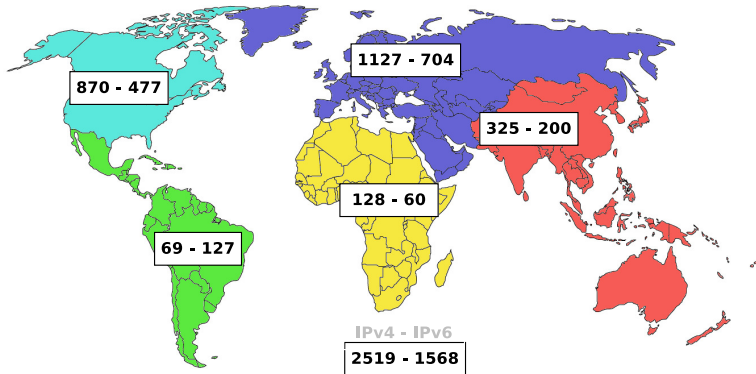
# BGP Route Collector Status (Oct 2016)



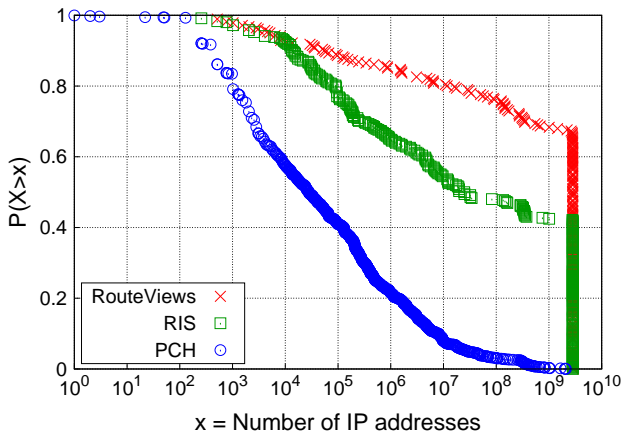
Total

**N. of RC**  
**N. of v4 feeders**  
**N. of v6 feeders**

19	17	123	159
281	358	1887	2526
197	228	1148	1573

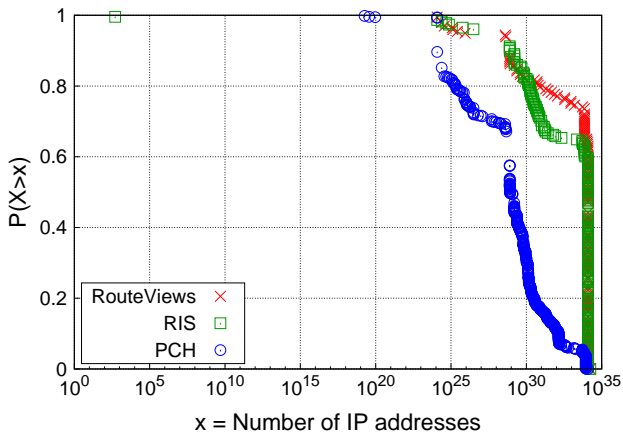


# Feeder Contribution (v4)



Only 343 IPv4 feeders announce to the RCs their full routing table

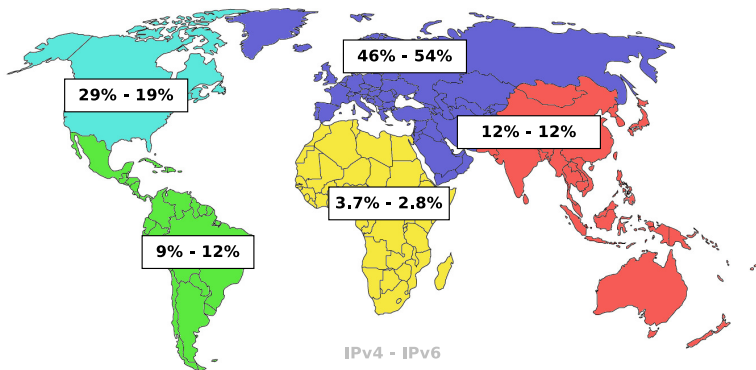
# Feeder Contribution (v6)



Only 267 IPv6 feeders announce to the RCs their full routing table

# Full feeder geographical distribution

Data collected represent mostly the Internet as viewed from Europe and North America than the real Internet



# Feeder characterization

About 80% of full feeders have a degree higher than 100

 SAVVIS



 Level 3  
COMMUNICATIONS

 TATA  
COMMUNICATIONS

 verizon

 SPARKLE  
INTERNET SERVICES

 Sprint

 TeliaSonera

 NTT  
Communications

 cogent  
COMMUNICATIONS  
Optical Internet

 HURRICANE ELECTRIC  
INTERNET SERVICES



# Conclusions on data analysis

## Conclusions

- Several p2p-connectivity is hidden from RC sight
- Several Internet regions are basically uncovered
- The typical profile of an ideal feeder is a multi-homed stub AS

## Questions

- Why there is a scarcity of participation to classical route collector projects?
- How to attract new participants?
- Is it just a case poor “marketing”?