

the **real-time** Internet routing observatory


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



Consiglio Nazionale delle Ricerche  
Istituto di Informatica e Telematica



# Our research focus: the Internet AS-level ecosystem




## Why is it important?

- To identify Internet topological properties and drawbacks
- To build realistic network topology generators for simulations
- To evaluate the effectiveness of new protocols

# Classic BGP route collector concept

A Route Collector (RC) is a device which collects BGP routing data from co-operating ASes

```
TIME: 02/09/12 08:08:47
TYPE: BGP4MP/MESSAGE/Update
FROM: 67.17.82.114 AS3549
TO: 128.223.51.102 AS6447
ORIGIN: IGP
ASPATH: 3549 137 137 137 8978
NEXT_HOP: 67.17.82.114
MULTI_EXIT_DISC: 14163
ANNOUNCE
212.77.0.0/19
```



RCs only collect routing information and not user traffic

# BGP route collector projects

## University of Oregon Route Views Project

Route Views was originally conceived as a tool for Internet operators to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. It collects BGP packets since 1997, in MRT format since 1997

<http://www.routeviews.org>



## RIPE NCC Routing Information Service (RIS)

The RIPE NCC collects and stores Internet routing data from several locations around the globe, using RIS. It collects BGP packets in MRT format since 1999

<https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>

## Packet Clearing House (PCH)

PCH is the international organization responsible for providing operational support and security to critical Internet infrastructure, including Internet exchange points and the core of the domain name system. It operates route collectors at more than 100 IXPs around the world and its data is made available in MRT format since 2011

[https://www.pch.net/resources/Raw\\_Routing\\_Data](https://www.pch.net/resources/Raw_Routing_Data)



# BGP data incompleteness

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?

$$\text{Minimize} \quad \left( \sum_{AS_i \in \mathcal{U}} x_{AS_i} \right) \quad (1)$$

subject to

$$\sum_{AS_j: n \in S_{AS_j}^{(d)}} x_{AS_j} \geq 1 \quad \forall n \in \mathcal{N} \quad (2)$$

$$x_{AS_i} \in \{0, 1\}, \quad \forall AS_i \in \mathcal{U} \quad (3)$$

... or in other words

Select new BGP feeders such that each transit AS has a **finite and bounded** p2c distance from the route collector infrastructure

# How much incomplete are BGP data?

April 2017

It was possible to discover the full connectivity of:

- 935 out of 9334 ASes (10.02%) which transit v4 traffic for other ASes
- 382 out of 2978 ASes (12.83%) which transit v6 traffic for other ASes

	v4 ASes	v6 ASes		v4 ASes	v6 ASes
<b>AE</b>	9 (16.07%)	5 (15.15%)	<b>OM</b>	5 (26.32%)	4 (30.77%)
<b>BH</b>	0 (0%)	0 (0%)	<b>PS</b>	0 (0%)	0 (0%)
<b>IQ</b>	4 (9.52%)	0 (0%)	<b>QA</b>	0 (0%)	0 (0%)
<b>IR</b>	0 (0%)	2 (16.66%)	<b>SA</b>	9 (18.75%)	3 (11.11%)
<b>JO</b>	4 (21.05%)	0 (0%)	<b>SY</b>	0 (0%)	0 (0%)
<b>KW</b>	0 (0%)	0 (0%)	<b>TR</b>	16 (18.39%)	6 (14.63%)
<b>LB</b>	4 (11.76%)	0 (0%)	<b>YE</b>	0 (0%)	0 (0%)


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** applications based on the aggregation of every routing information collected

# What 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

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

- 1 AS characteristics
- 2 Feeder contribution
- 3 Total coverage of RCs

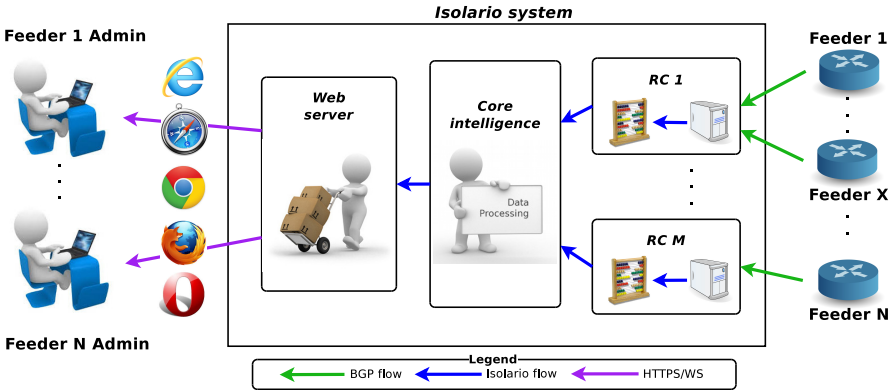
## Open source software

- 1 Interactive Collecting Engine (ICE)
- 2 MRT Data Reader



# Isolario system overview

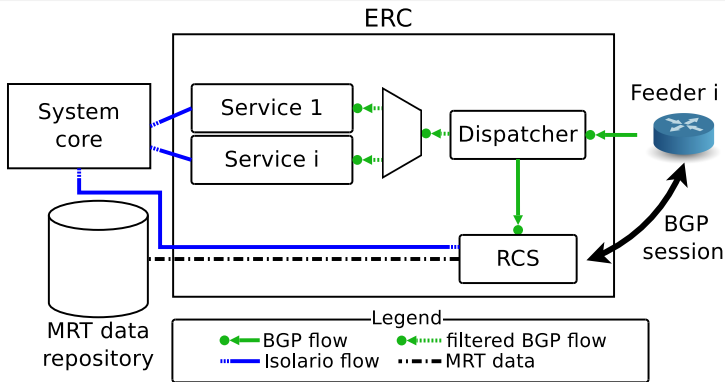
Incoming BGP flows are used as **real-time streams** for services dedicated to participants



Results are provided to users via WebSockets

# Enhanced BGP Route Collector

Incoming flows are duplicated as soon as they arrive and feed both the Route Collecting Software (RCS) and service modules



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

# 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

# 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


Please, feel free to try our real-time services!

`https://www.isolario.it`

Username: *guest*

Password: *guest*

## Real-time services



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



# Routing table viewer

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

AS 23752

NPTELECOM,NR,AS Nepal Telecommunications Corporation, Internet Services,NP

Reachability Dependency per AS 100%

Avg Distance 4.5

My Feeders

192.12.193.233 x

IP space related to 202.70.64.0/21

REMOVE PREFIX

Cev %	Prefix	AS path	Communities	AI	Aggregator	Timestamp	A	W
12.5	202.70.64.0/24	2597 3356 9498 23752	2597:101 2597:111 2597:666 2597:1000			2014-11-23 09:16:05	0	0
12.5	202.70.64.0/21	2597 3356 6453 23752	2597:101 2597:111 2597:666 2597:1000			2014-11-25 15:33:43	3	0
		2597 2396 15412 15412 15412 15412 9304 23752	2597:101 2597:111 2597:666 2597:1000					
		2597 3356 9498 23752	2597:101 2597:111 2597:666 2597:1000					

### A/W Events

Announcements: Number of events: 3 (100.0%)

- Announcements (100.0%)
- Withdrawals (0.0%)

### Reachability

- Reachable IPs (100.0%)
- Unreachable IPs (0.0%)

### Reachability Dependency per AS

Top ten ASes

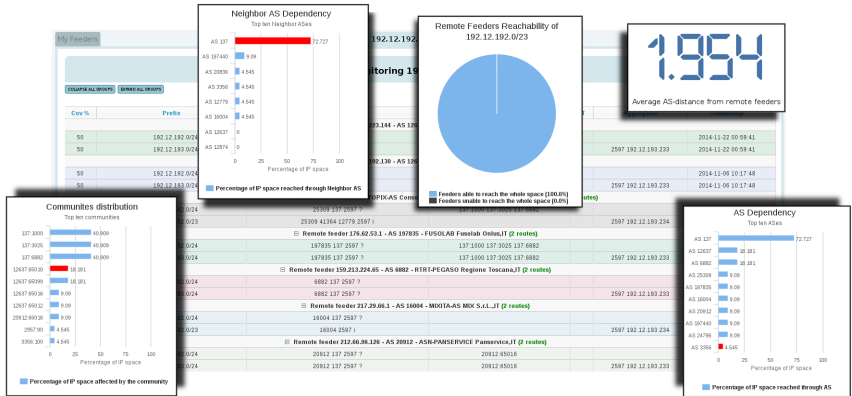
AS	Percentage of IP space reached through AS
AS 23752	100
AS 3356	100
AS 9498	62.5
AS 3549	50
AS 7473	37.5
AS 9304	25
AS 9829	12.5
AS 6453	12.5
AS 25412	12.5

Percentage of IP space reached through AS



# Subnet reachability

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

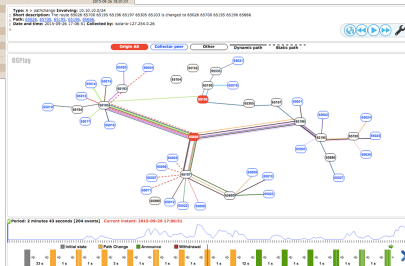


The more the feeders, the more SR is useful!

# Isolario real-time visualisation with BGPlay

- BGPlay is an **open-source** tool for the visualisation of BGP routing
- Thanks to the close collaboration with Massimo Candela (RIPE NCC) we integrated in Isolario the BGPlay **real-time version** (<http://bgplay.massimocandela.com>)

ASN	Prefix	All paths	ASes	ASes	Aggregator	Timestamp
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00
100	0.0.0.0/24	65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	ASes: 65024 63708 63105 63196 15066 1	2015-09-28 10:00:00



BGPlay is currently integrated in SR



# 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

**Alerting Management** | Notifications | Current configured alerts

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**

- 127.254.0.1 (ASN 65001)
- 127.254.0.10 (ASN 65010)
- 127.254.0.11 (ASN 65011)
- 127.254.0.13 (ASN 65013)
- 127.254.0.19 (ASN 65019)
- 127.254.0.2 (ASN 65002)
- 127.254.0.20 (ASN 65020)
- 127.254.0.21 (ASN 65021)
- 127.254.0.22 (ASN 65022)

**Alert Type**

- BGP attributes
- Flap Detector
- Session Watchdog
- Hijack
- Reachability

**Action upon event**

Email  s

POST HTTP(s)

**Save Alert.**

### BGP attributes

Prefix  Prefix Subnet  Community  Prefix Related

AS path end  AS path substrings  AS path begin  AS path exact

Origin  Aggregator

You can specify one or more BGP attribute types on which the monitoring will run. Multiple types can be combined by means of `and/or` operators and round brackets. For each attribute type you can insert one or more values that the attribute should match. The system will report any BGP\_UPDATE message advertised by one of the selected feeder IPs matching the inserted attributes.

**EXAMPLE**

Current BGP attribute types selected



# 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



# 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)



# Why Isolario?

## What's the need of *\*yet\** another routing analysis tool?

- The more (and diversified) the BGP data sources, the better
- Isolario tools are **just** an incentive to push network admins to share their BGP routing data with the research community
- Most routing analysis tools (commercial and not) either use BGP data publicly available *or* do not publish the BGP data they collect

## What's the need of *\*yet\** another route collector?

- Real-time services require a different route collecting infrastructure
- The do-ut-des paradigm **may** be appealing to some of those network admins who are not sharing data with any route collector (yet)

It is not relevant whether you decide to connect to Isolario, Route Views, RIPE NCC RIS and/or PCH, as long as you share the data!

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)**