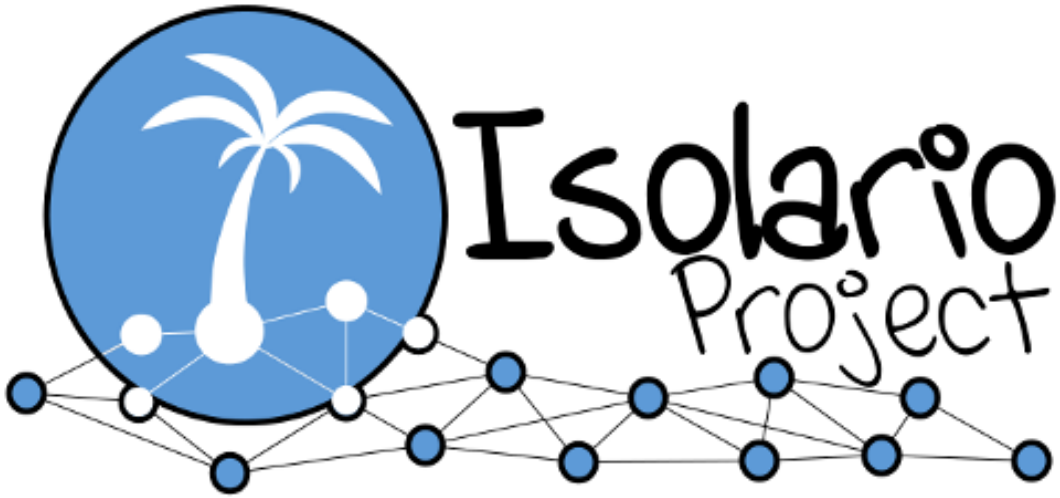

the **real-time** Internet routing observatory

Luca Sani

TOP-IX MEETING, 26 September 2017

Consiglio Nazionale delle Ricerche
Istituto di Informatica e Telematica

Our research interest: the Internet AS-level ecosystem

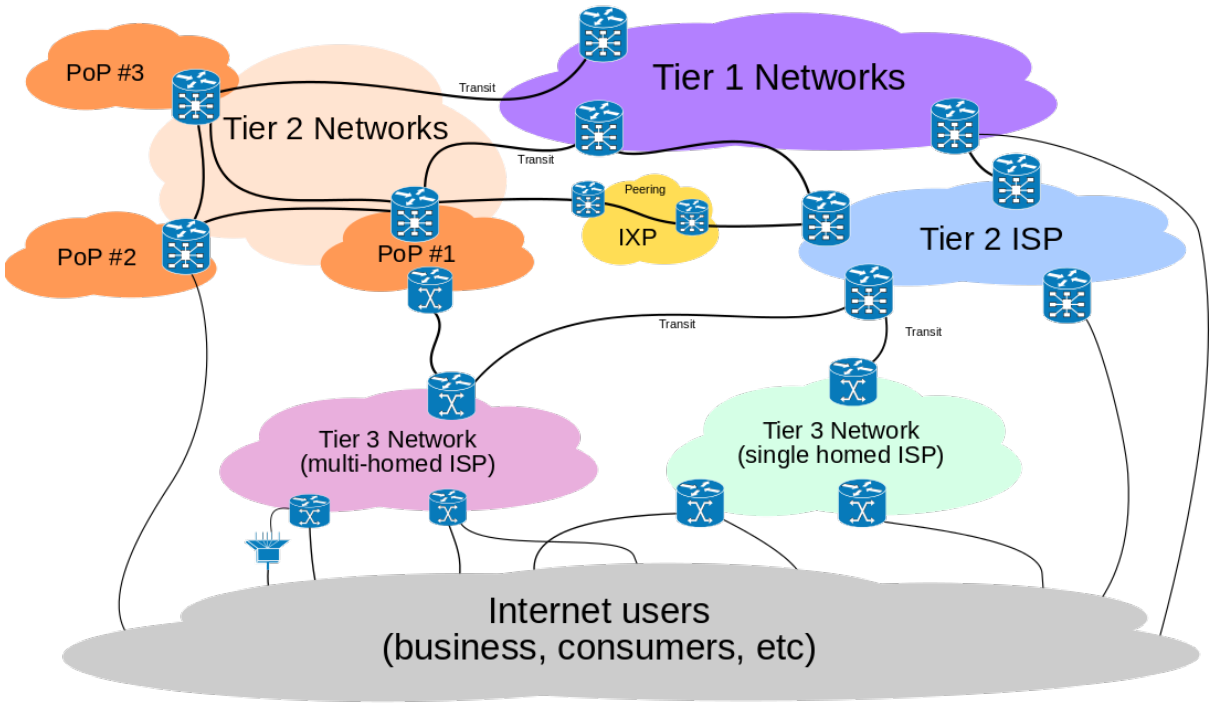

Why is it important?

- 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

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

Route collectors are devices which collects BGP routing data from co-operating ASes (feeders)

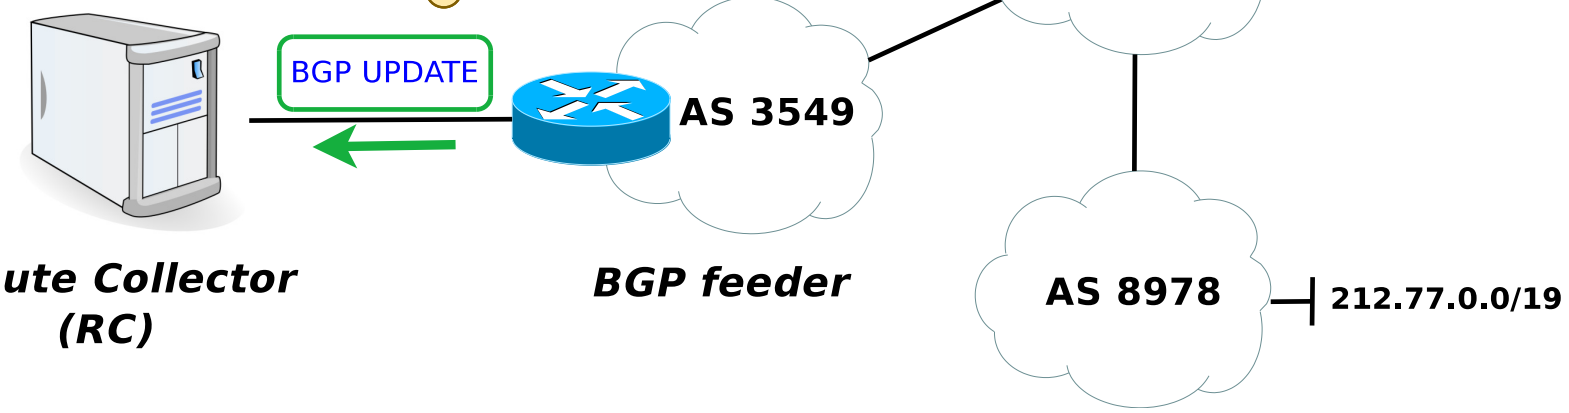

Route collectors collect routing information and **not** user traffic

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

Beware of data completeness!

Feeders connected to route collectors (September 23rd, 2017)

- 1065 ASes announcing v4 data, 644 announcing v6 data
- 263 ASes share full v4 routing table, 212 their full v6 routing table

A view from the top

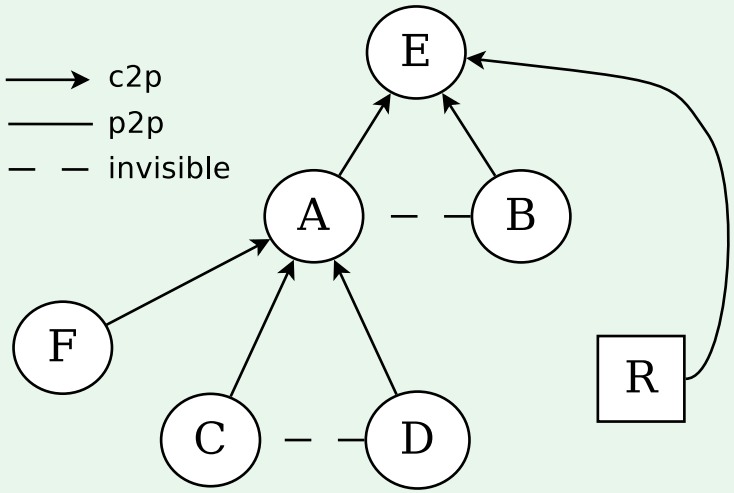

A view from the bottom

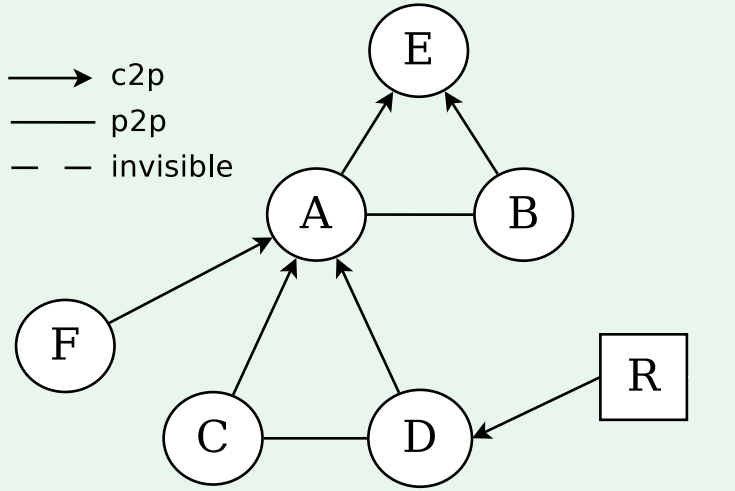

Nowadays most p2p connectivity (IXPs) is invisible to route collectors

How much incomplete?

Formal (and boring) MSC problem definition

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

subject to

$$\sum_{AS_i : n \in S^{(d)}_{AS_i}} x_{AS_i} \geq 1 \qquad \forall n \in \mathcal{N} \tag{2}$$

$$x_{AS_i} \in \{0, 1\}, \qquad \forall AS_i \in \mathcal{U} \tag{3}$$

... 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

E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani, "*A Novel Methodology to Address the Internet AS-level Data Incompleteness*", in IEEE/ACM Transactions on Networking, pp. 1314-1327, Volume 23, Issue 4, August 2015

How much incomplete?

September 2017

It was possible to discover the **full** connectivity of:

- 702 out of 9621 ASes (7.30%) which transit v4 traffic for other ASes
- 350 out of 3148 ASes (11.12%) which transit v6 traffic for other ASes

Top 10 countries (per number of transit AS covered)

	v4 ASes		v6 ASes
United Kingdom	138 (19.49%)	**Germany**	56 (23.05%)
Netherlands	104 (22.86%)	**United Kingdom**	41 (19.81%)
Germany	103 (18.39%)	**Switzerland**	33 (35.11%)
France	78 (23.56%)	**Italy**	28 (45.16%)
Italy	69 (28.40%)	**France**	22 (22.45%)
Switzerland	59 (31.38%)	**Netherlands**	21 (14.19%)
Russia	48 (4.37%)	**Austria**	21 (26.25%)
Spain	45 (20.36%)	**Spain**	12 (22.22%)
Sweden	34 (16.50%)	**Russia**	11 (5.85%)
Austria	33 (19.19%)	**Denmark**	10 (21.74%)

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

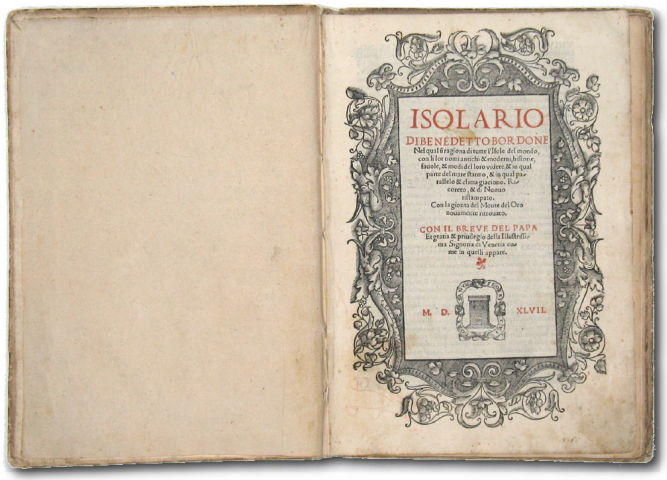

Isolario - The Book of Islands

"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

Isolario system overview

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

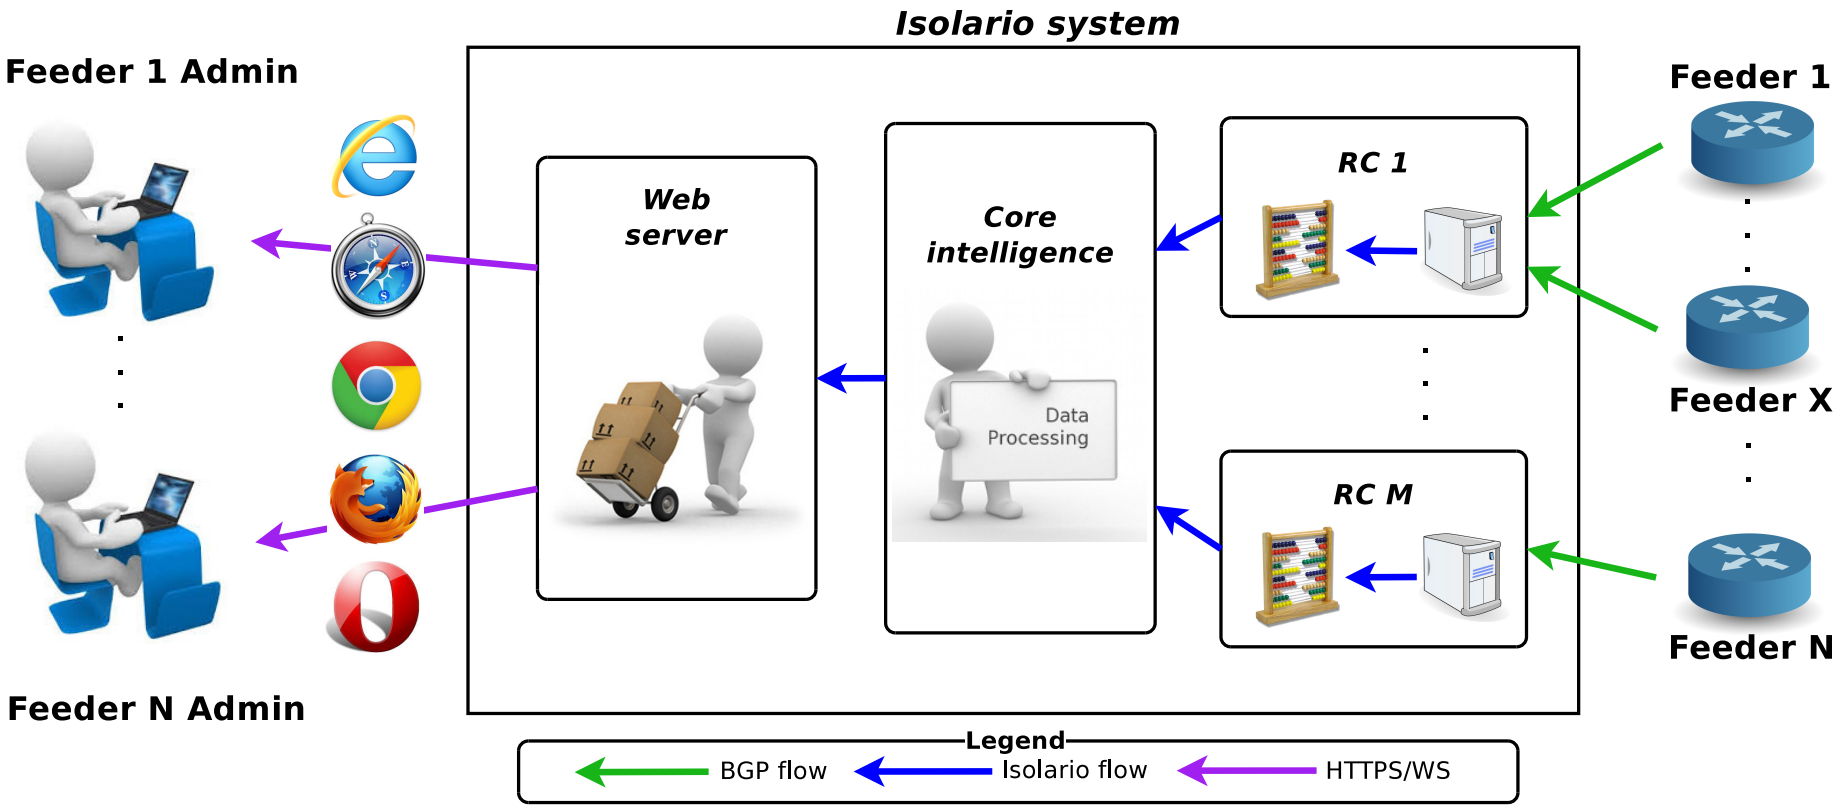

Results are provided to users via WebSockets

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

- 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*

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

- 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*

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

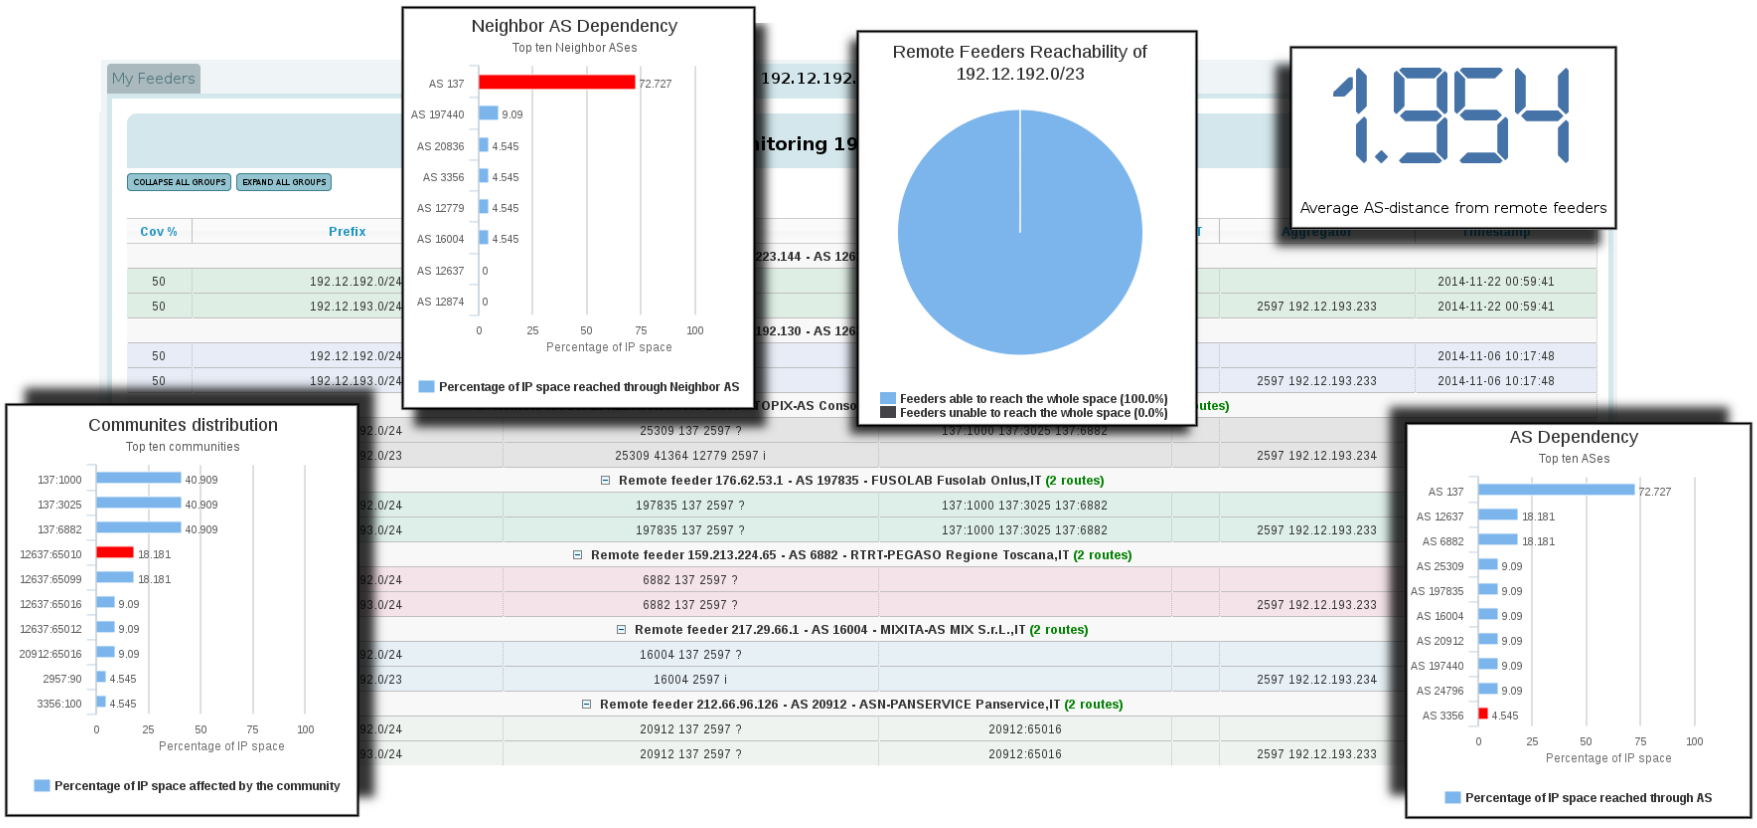

The more the feeders, the more SR is useful!

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**)

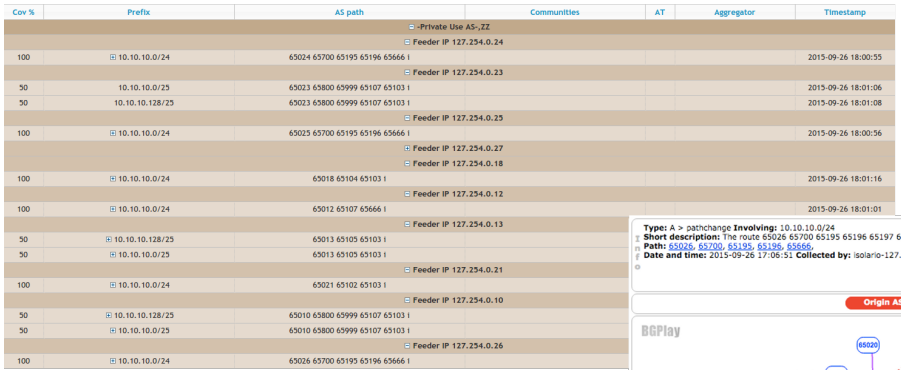

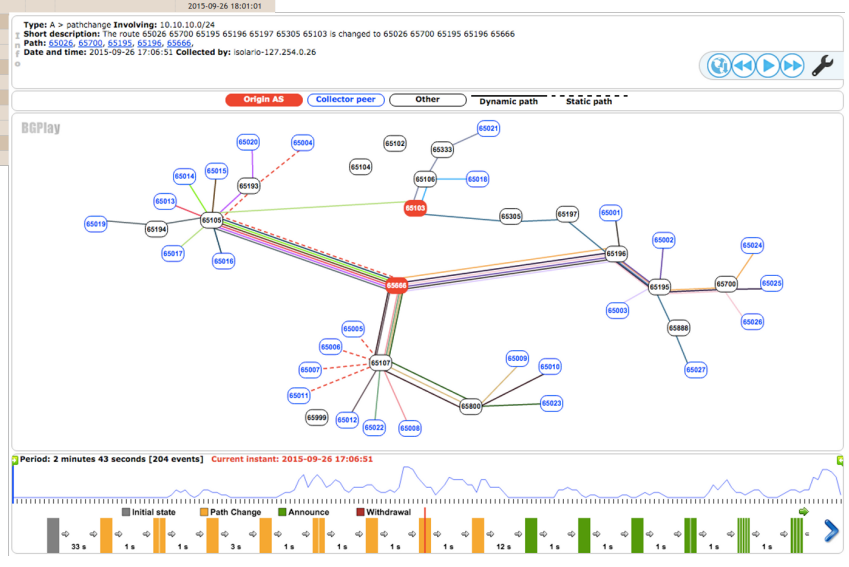

BGPlay is currently integrated in SR

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

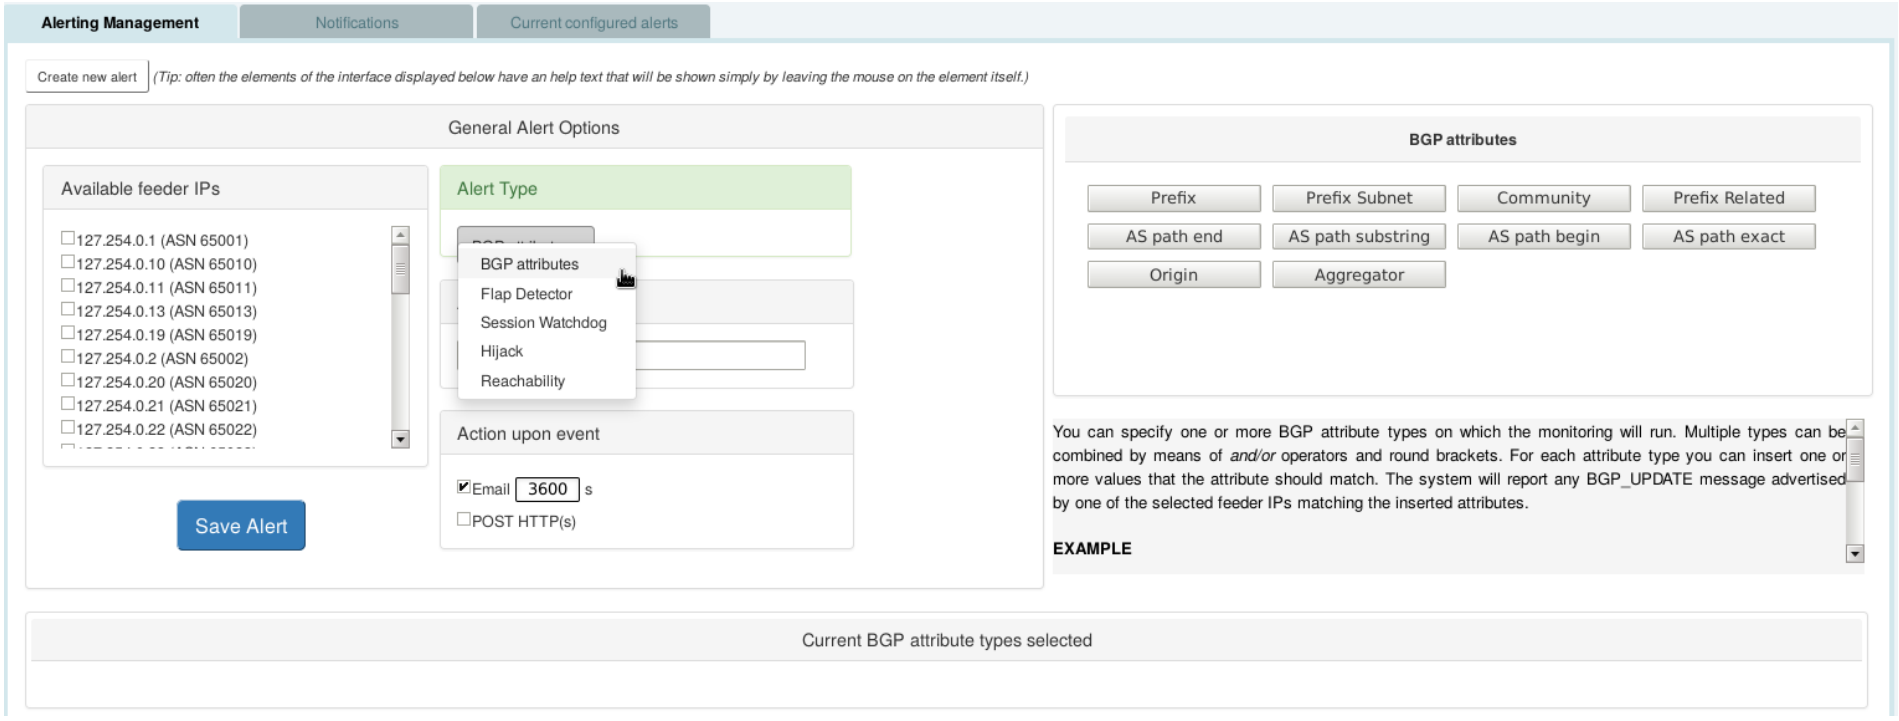

Daily report

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: 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[1]: *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[2]: *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 withdrawns related to proprietary subnets: *0*

Figure 1: Amount of packets received per hour

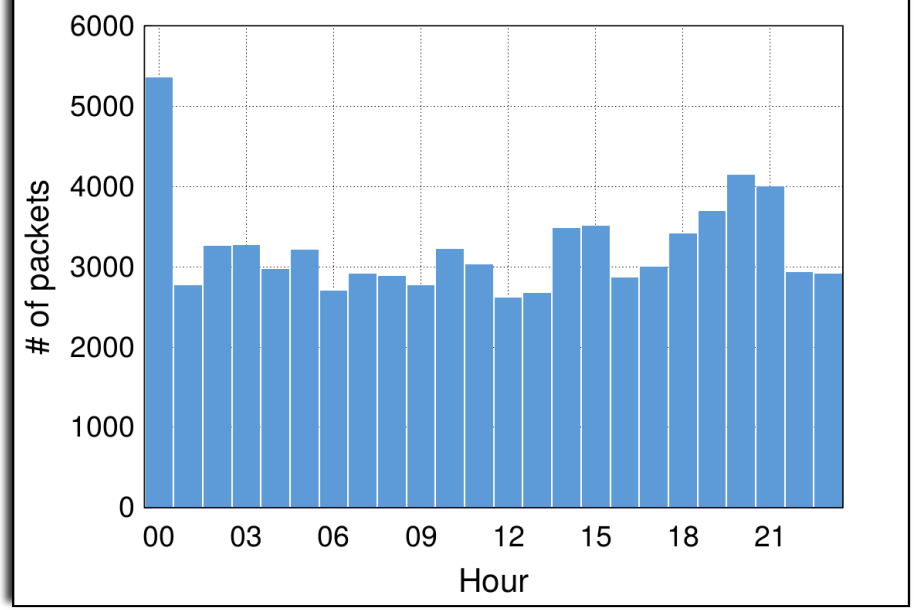

Summary: how to use Isolario?

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)

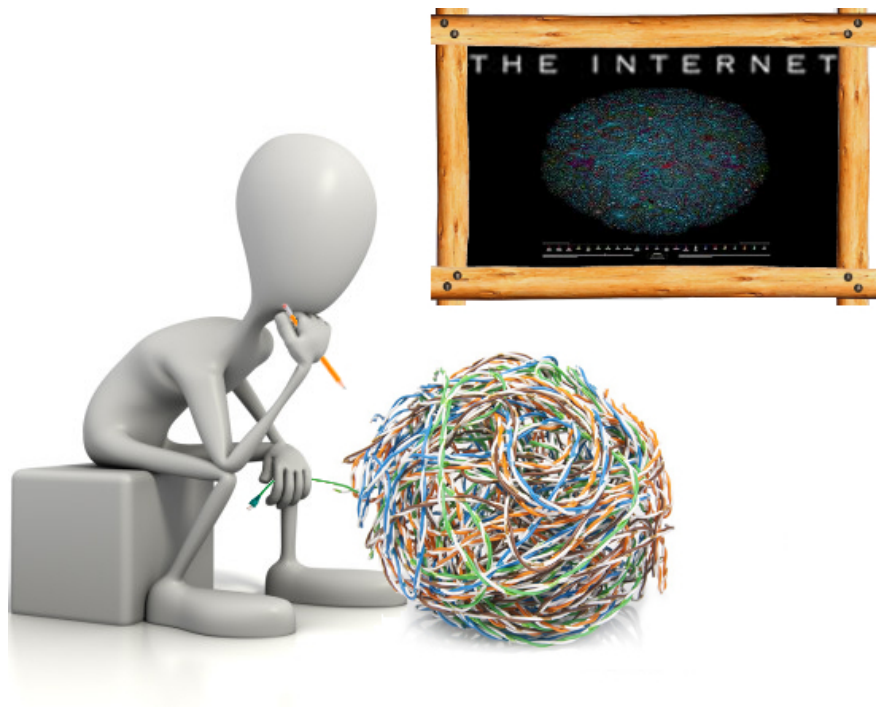

Current participants

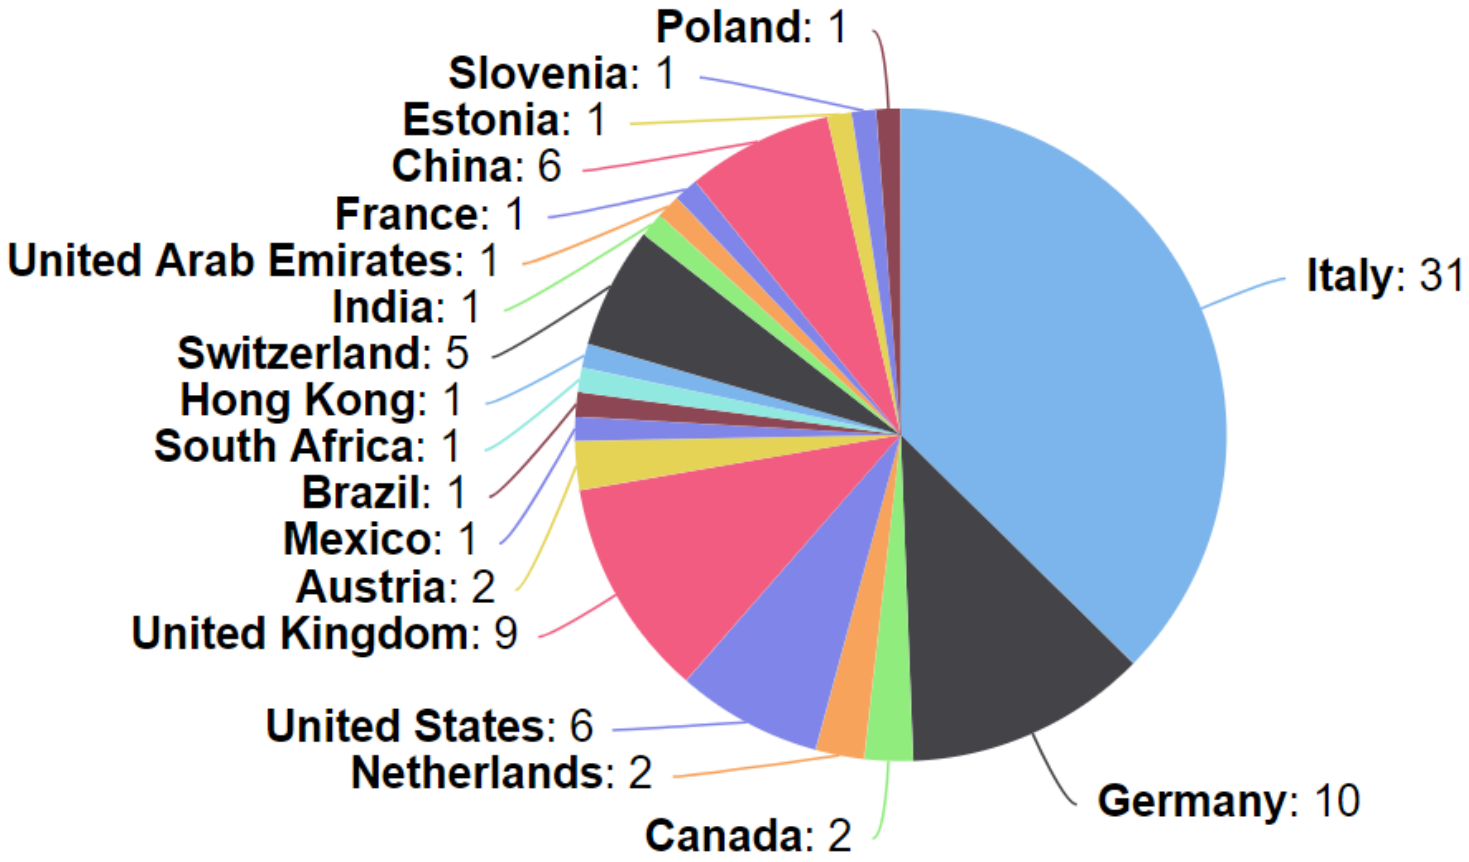

83 ASes connected

86 v4 sessions

76 v6 sessions

Only 24 out of 83 ASes are already connected to RIS and/or Route Views

What we provide to the community?

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

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

   - used in Hurricane Electric BGP Toolkit (**https://bgp.he.net**)

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
3. · · ·

Future

Other data sources

We are going to integrate external real-time sources to improve the usefulness of services

- RIPE RIS stream (under development)
- BGPmon

Deployment of Isolario collectors around the world

- Lightweight version of route collector
- Anycast deployment (latency, redundancy)
- Start with local IXPs
- Goal: deploy on global scale

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