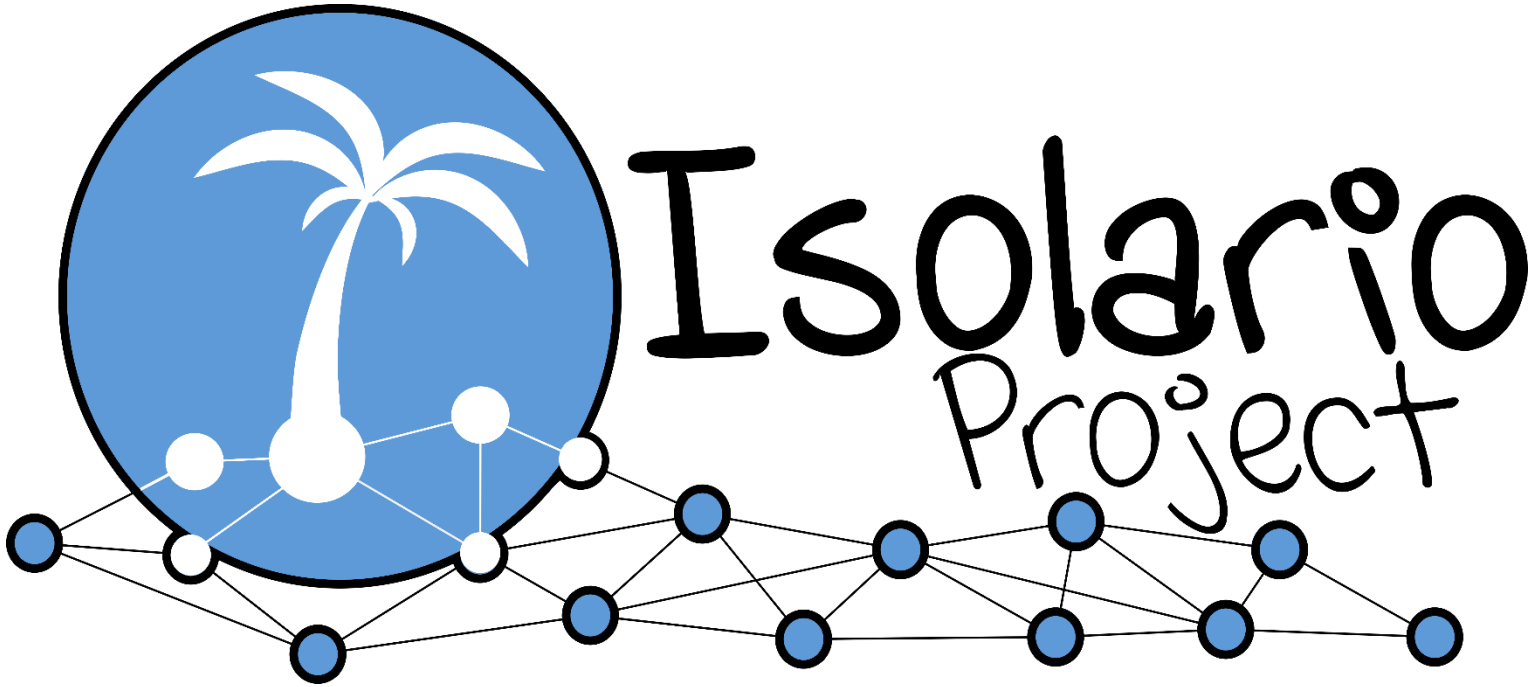

The real-time Internet routing observatory

Alessandro Improta

TREX Workshop

June 15th, 2018

Registro.it

Consiglio Nazionale delle Ricerche
Istituto di Informatica e Telematica

Our research interest: the Internet AS-level ecosystem

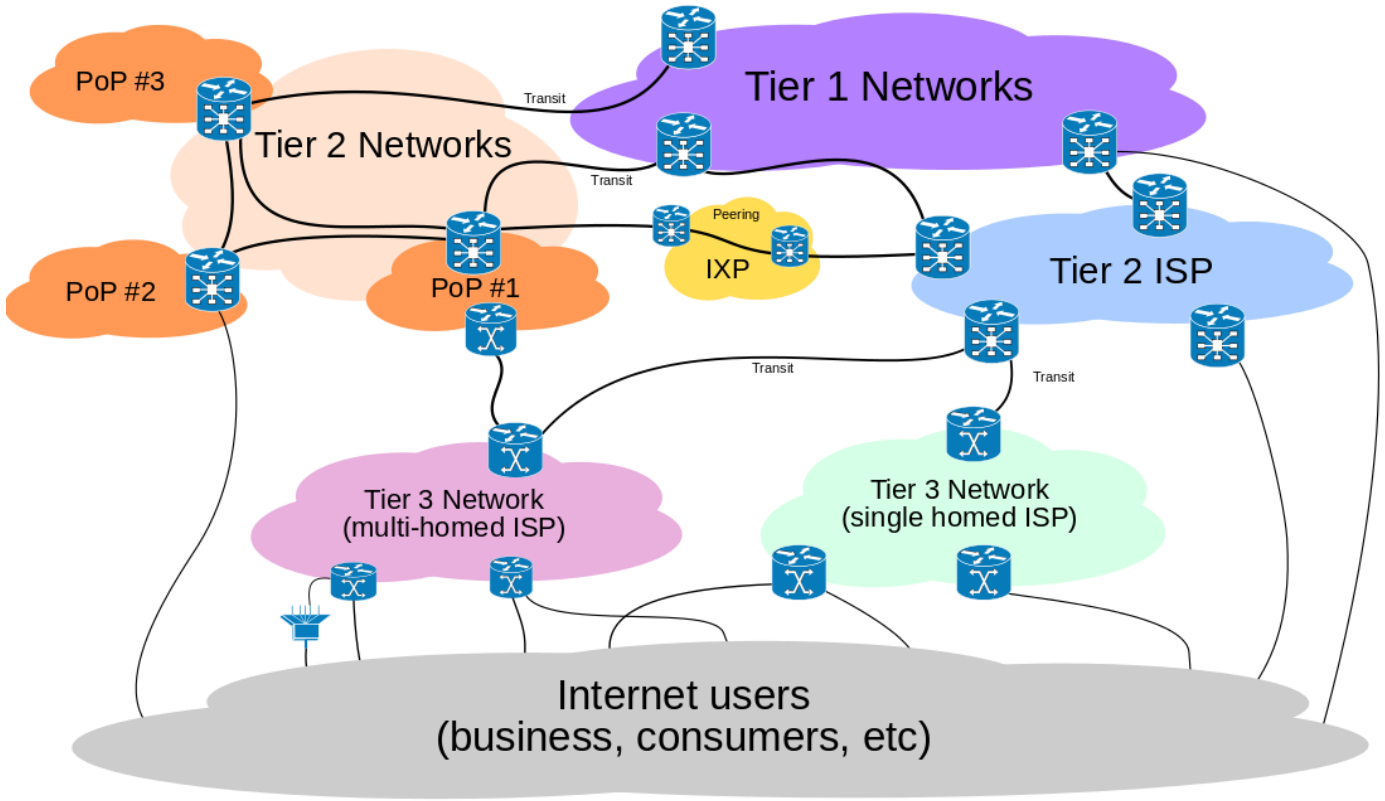

Why is it important?

- 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

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

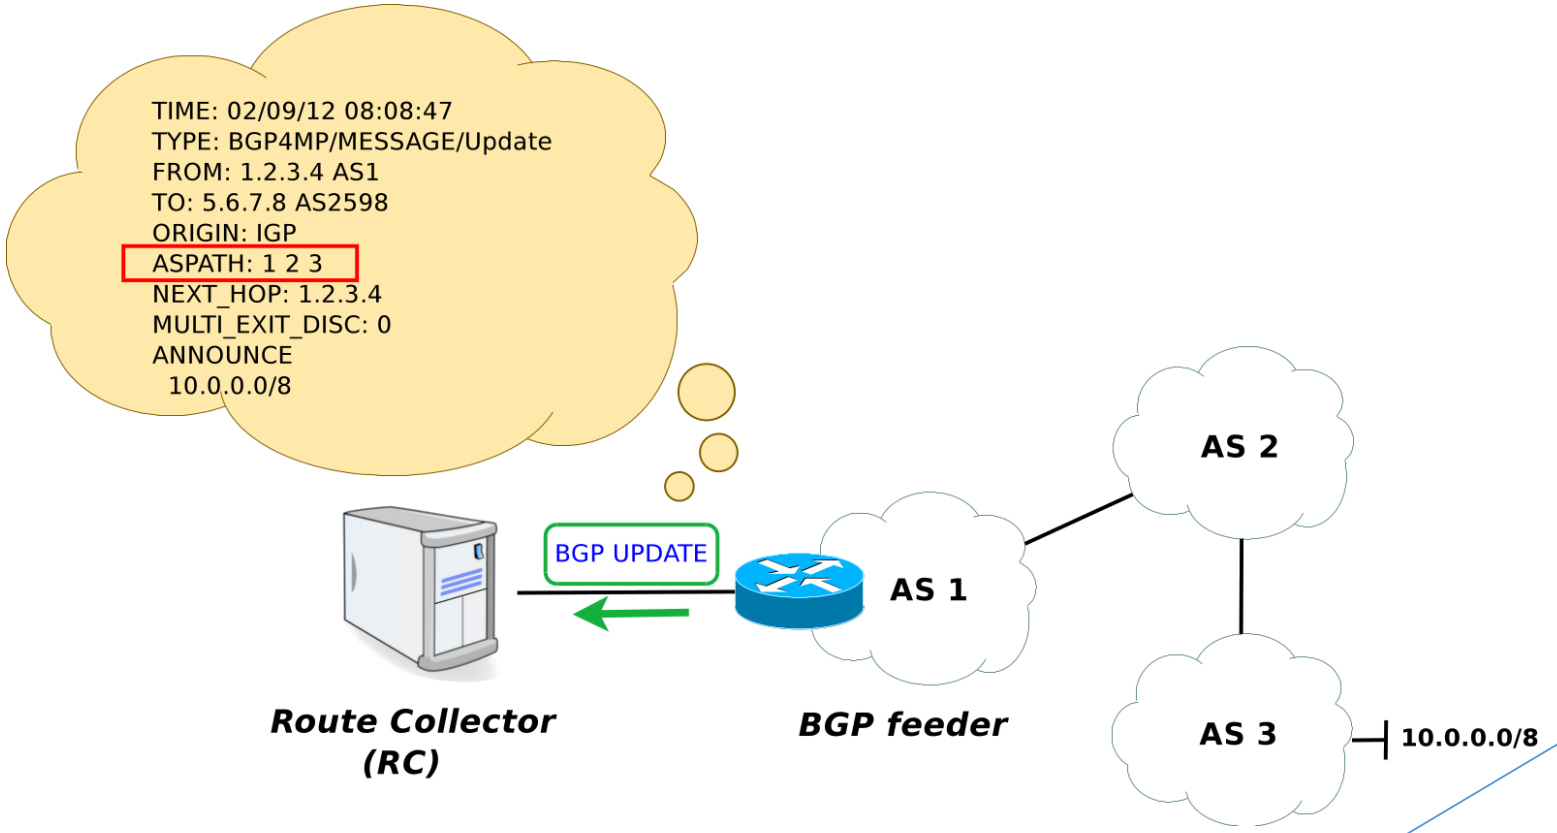

Route collectors collect routing information and **not** user traffic

BGP route collectors

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 Views, RIS and/or PCH (June 11th, 2018)

- 1092 ASes announcing v4 data, 662 announcing v6 data
- 232 ASes share full v4 routing table, 187 their full v6 routing table

A view from the top

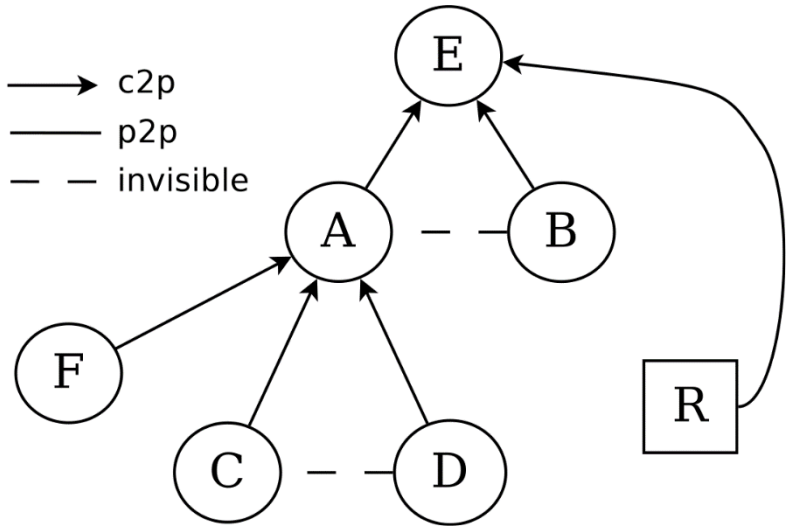

A view from the bottom

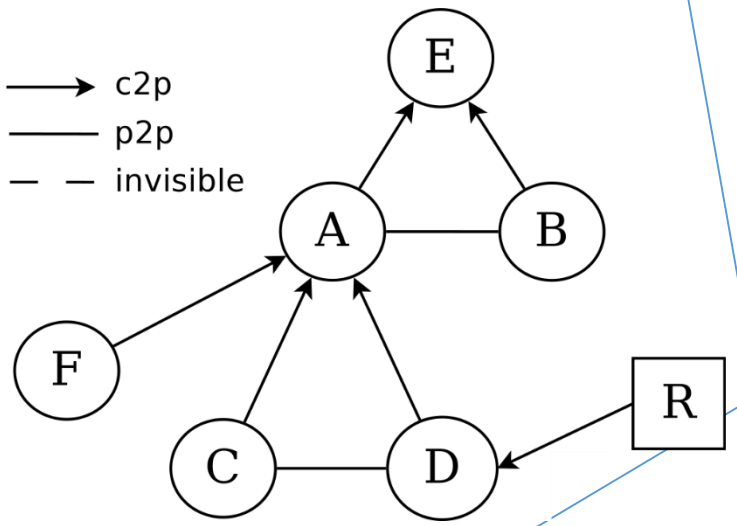

Most of p2p connectivity (IXPs) is currently invisible to route collectors

How much incomplete?

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

$$\text{Subject to} \quad \sum_{AS_i : n \in S^{(d)}_{AS_i}}^{x_{AS_i}} \geq 1 \forall n \in \mathcal{N}$$

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

… 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*

[1] Gregori E. et al. “A Novel Methodology to Address the Internet AS-level Data Incompleteness” in IEEE/ACM Transactions on Networking, pp. 1314-1327, Vol. 23(4), Aug 2015

How much incomplete?

In June 11th, 2018 it was possible to discover the full connectivity of:

- 542 out of 10,282 ASes (5.27%) which transit v4 traffic for other ASes
- 468 out of 3444 ASes (13.59%) which transit v6 traffic for other ASes

	v4 transit ASes		v6 transit Ases
Sweden	38 out of 223 (17.04%)	Sweden	15 out of 78 (9.23%)
Denmark	26 out of 117 (22.22%)	Denmark	15 out of 36 (41.67%)
Norway	24 out of 131 (18.32%)	Norway	14 out of 41 (34.15%)
Finland	**14 out of 87 (16.09%)**	**Finland**	**10 out of 41 (24.39%)**
Estonia	7 out of 41 (17.07%)	Iceland	5 out of 10 (50.00%)
Lithuania	4 out of 63 (6.35%)	Estonia	4 out of 12 (33.33%)
Latvia	4 out of 58 (6.90%)	Latvia	2 out of 16 (16.66%)
Iceland	3 out of 25 (12.00%)	Iceland	2 out of 16 (12.50%)

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

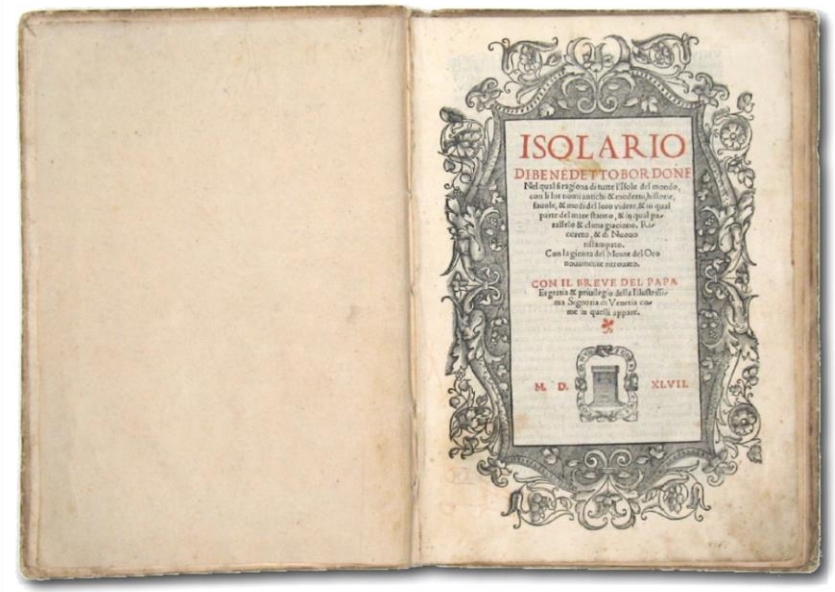

Isolario - The Book of Islands

"[…] 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 at least one v4/v6 BGP session with Isolario providing their **full** routing table
- In change, Isolario offers real-time applications based on the aggregation of every routing information collected

Isolario real-time system overview

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

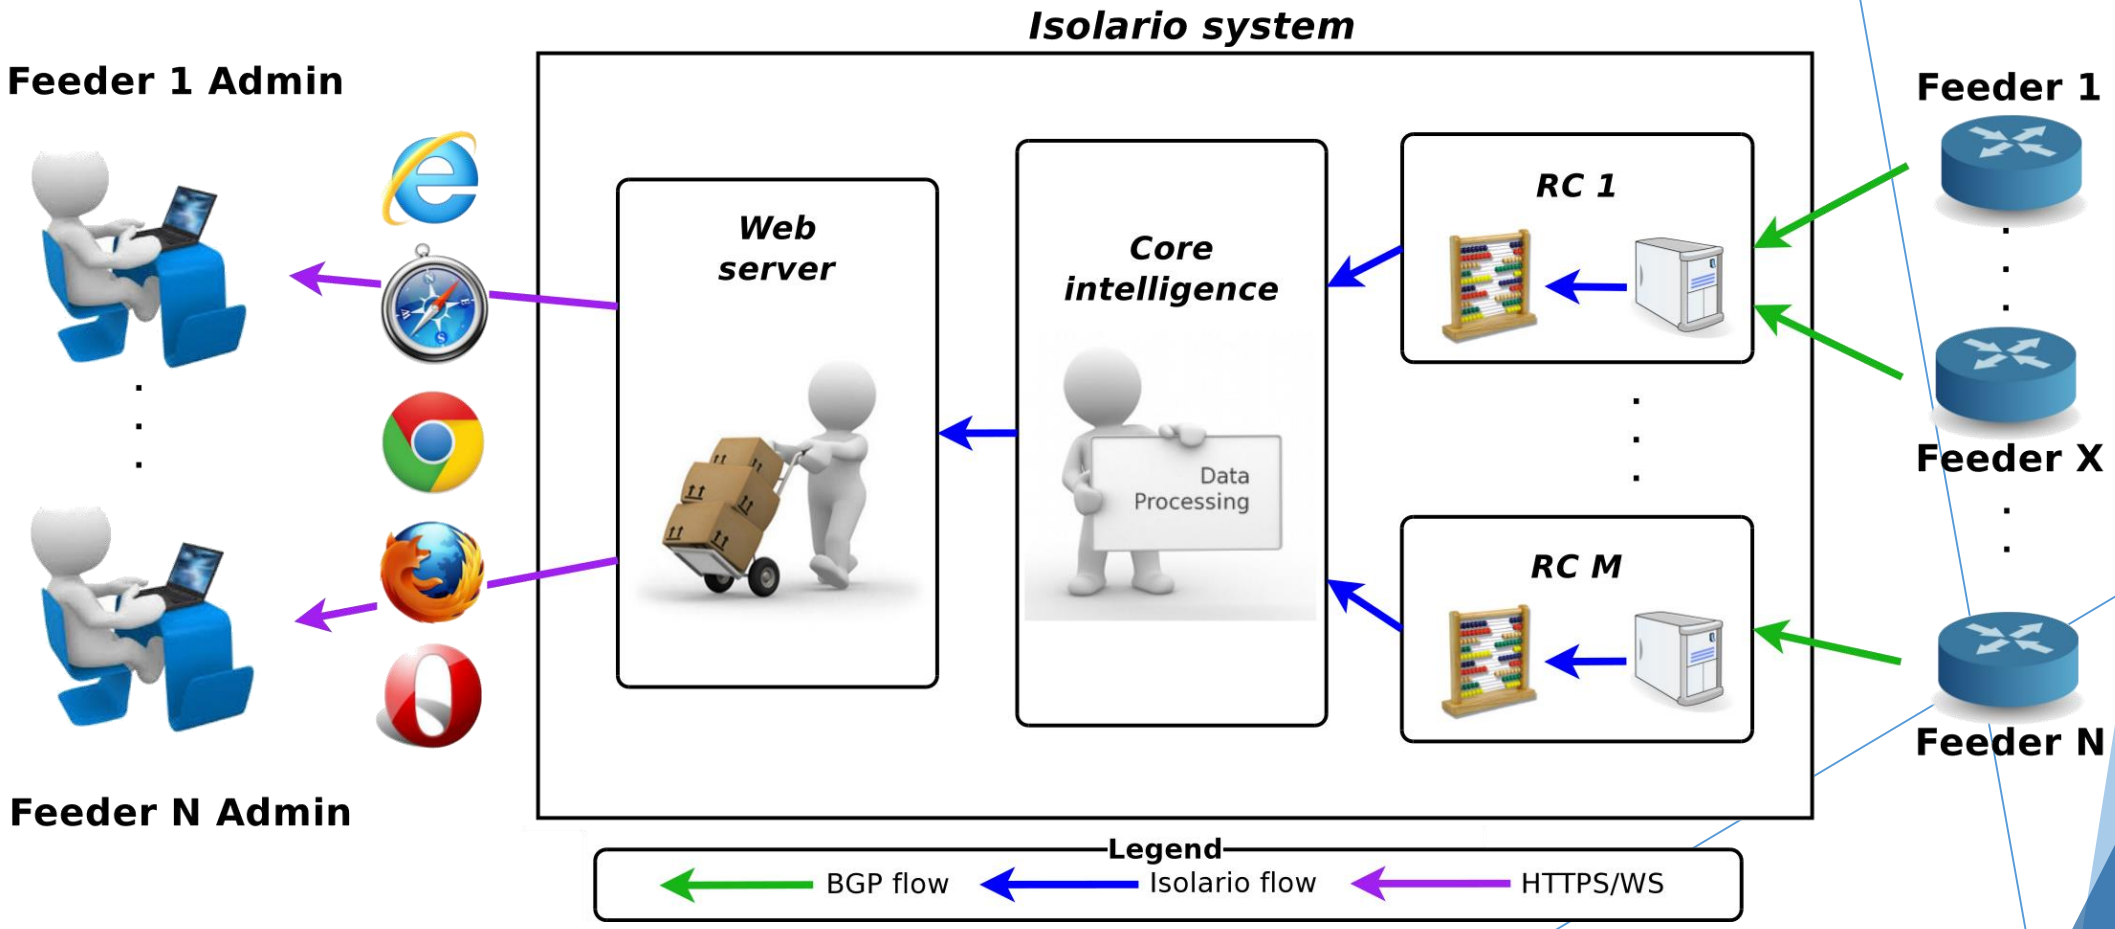

Results are provided to users via WebSockets (RFC 6455)

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 monitoring services

BGP Flow viewer

Routing table viewer

Subnet reachability

Website reachability

Diagnostic services

Alerting system

Daily reports

Historic monitoring services (work in progress)

Routing table viewer

Subnet reachability

Please, feel free to try isolario.it !

Username: *guest*

Password: *guest*

BGP Flow Viewer (BFV)

BFV allows to monitor BGP packets announced by a feeder to Isolario

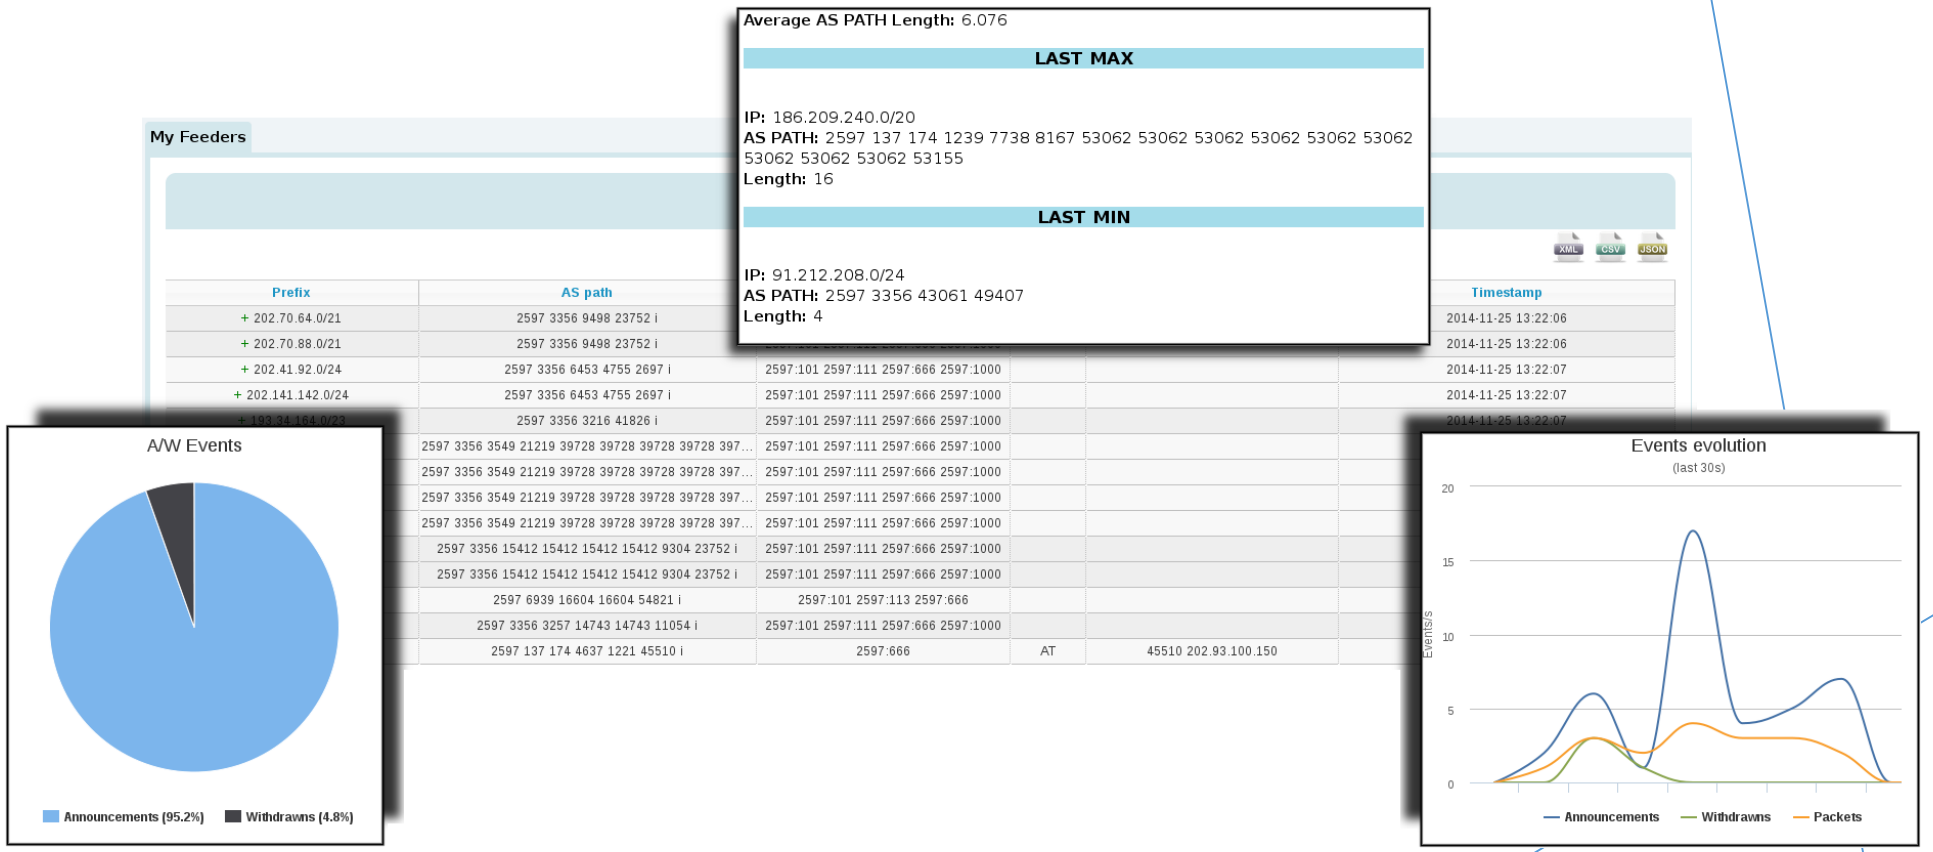

https://youtu.be/QynZqNMCyXw

Subnet Reachability (SR)

SR allows to understand how Isolario feeders are reaching subnets of interest

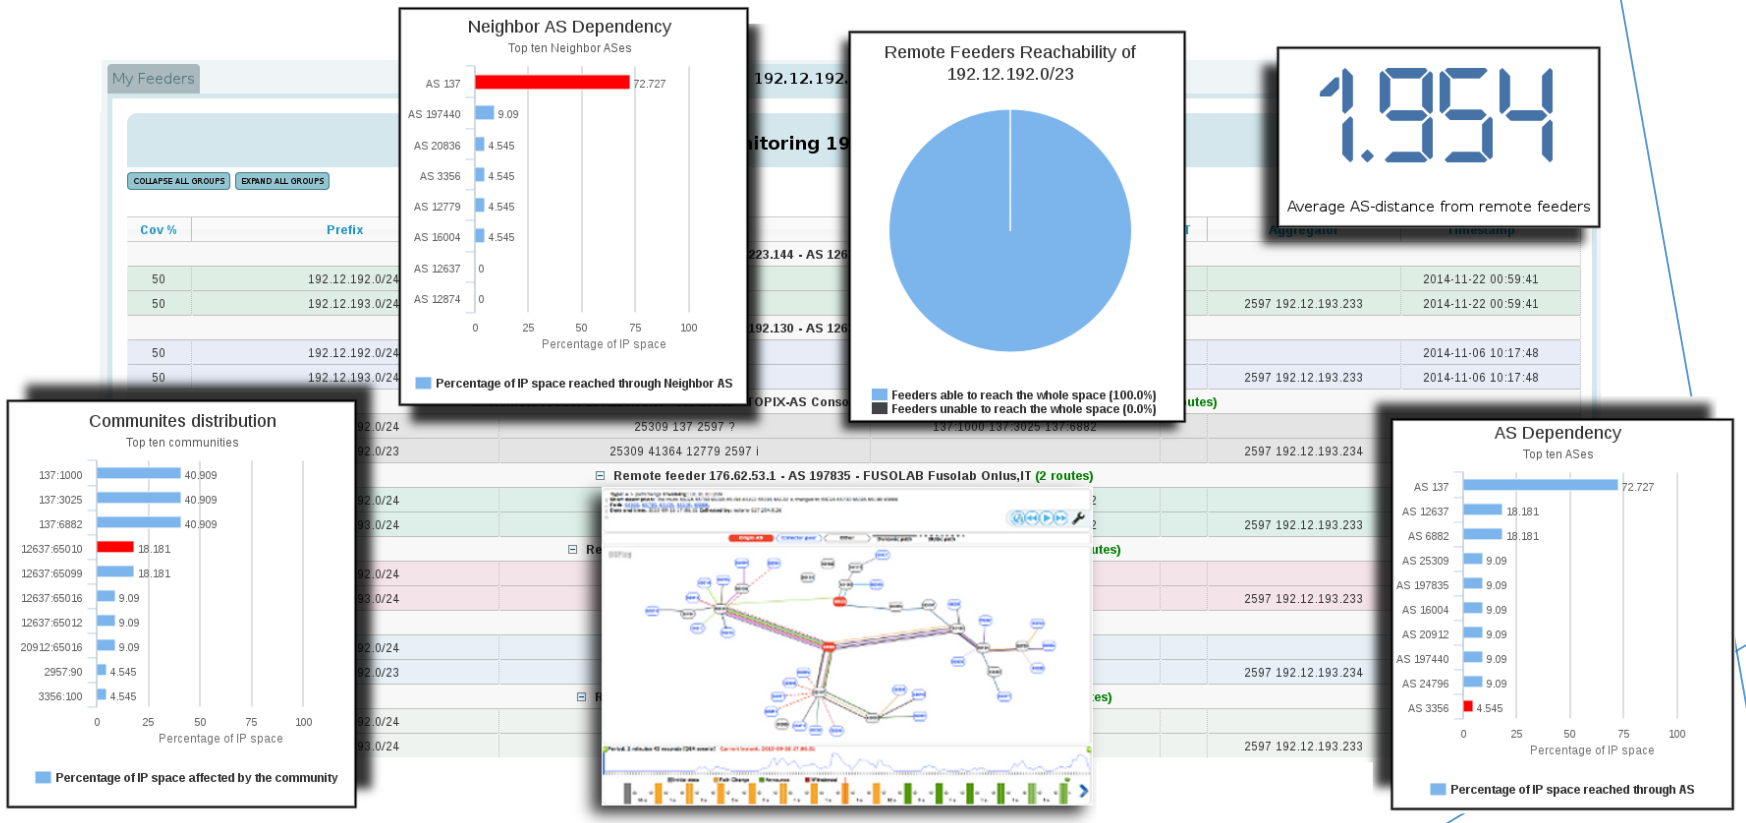

https://youtu.be/uTrLVv1PoPo

Alerting system

The alerting system allows to receive notifications as soon as any user-configured alarm is triggered

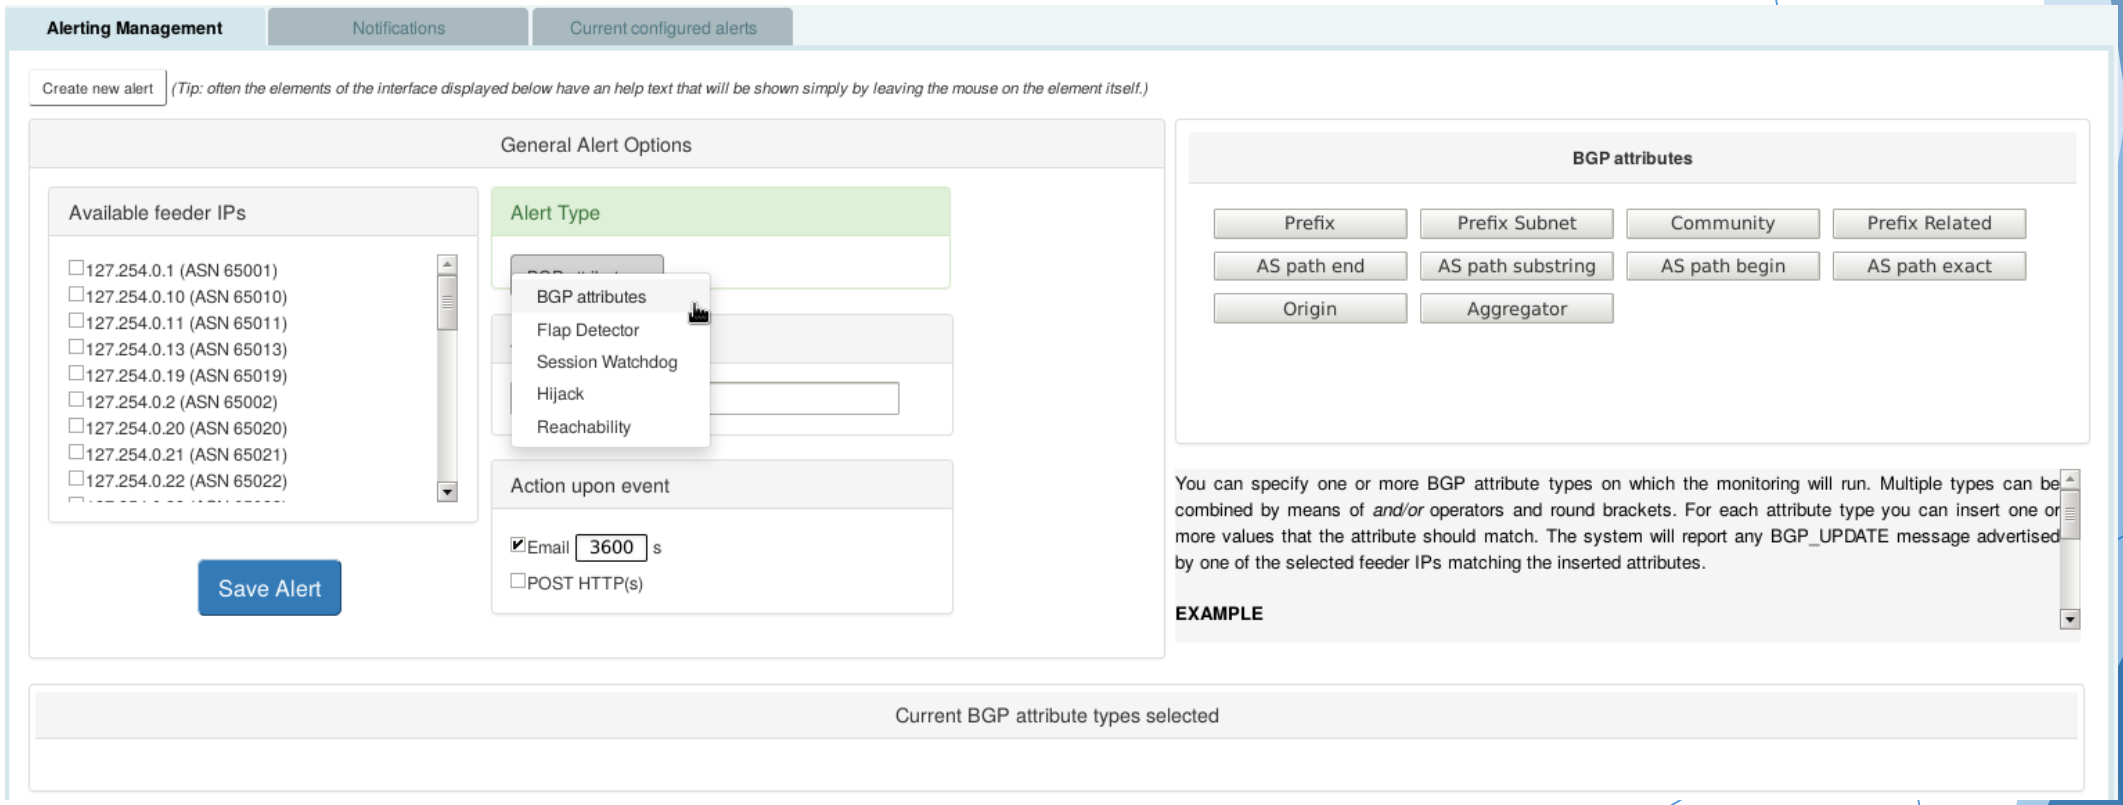

https://youtu.be/p_r2pRHK7EI

Daily reports

Summaries about inter-domain routing status as perceived by the Isolario:

- **Feeder reports** about the evolution of the feeder routing
- **AS reports** about the reachability of the network of the feeder AS

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

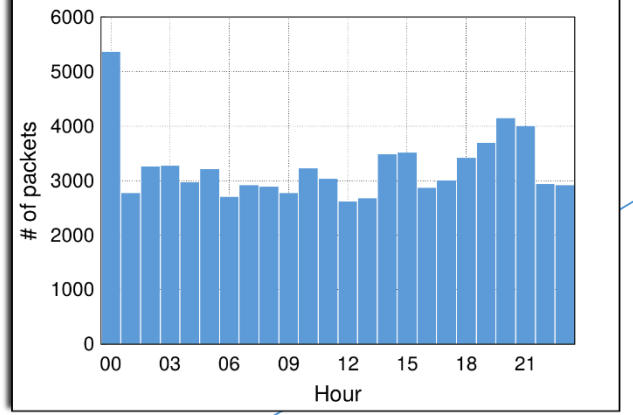

How to use Isolario?

Real-time services

Let me check what's going on…

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

Alerting system

Some particular routing event is on NOW!

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

Daily reports

Was my routing ok yesterday?

- Check historic services (soon)!
- Do something! (if needed)

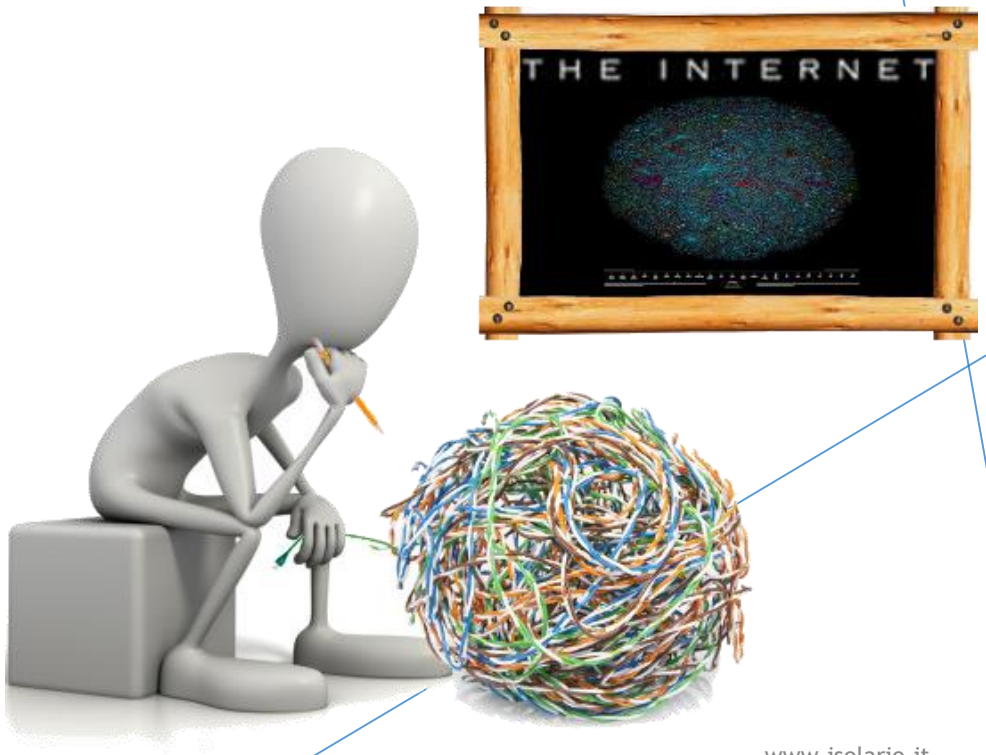

Isolario numbers (June 11th, 2018)

FEEDER ASES GEOLOCATION

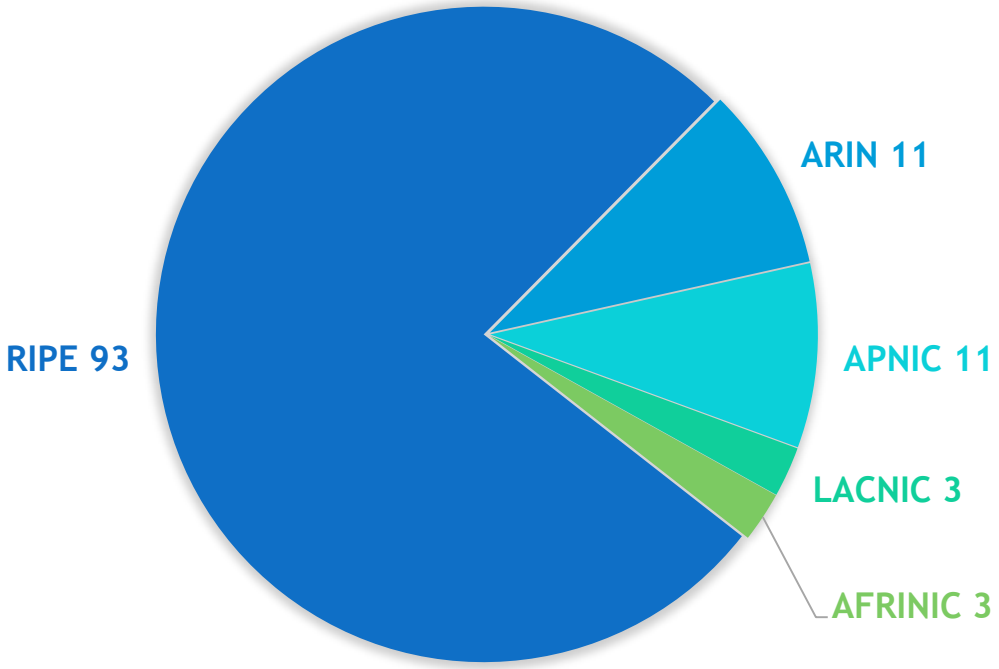

Number of ASes participating: *127*

Number of sessions configured: *368*

- **IPv4:** *192* from *120* ASes
- **IPv6:** *176* from *102* different ASes

ASES CONNECTED TO OTHER COLLECTORS

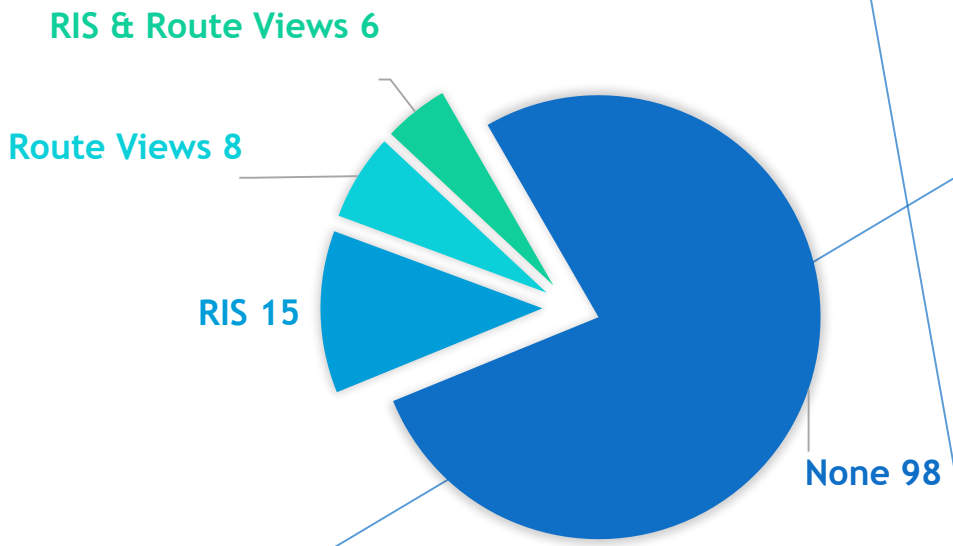

Number of full routing tables:

- **IPv4:** *122* from *78* different ASes
- **IPv6:** *133* from *82* different ASes

Isolario numbers with ADDPATH – RFC 7911 (June 11th, 2018)

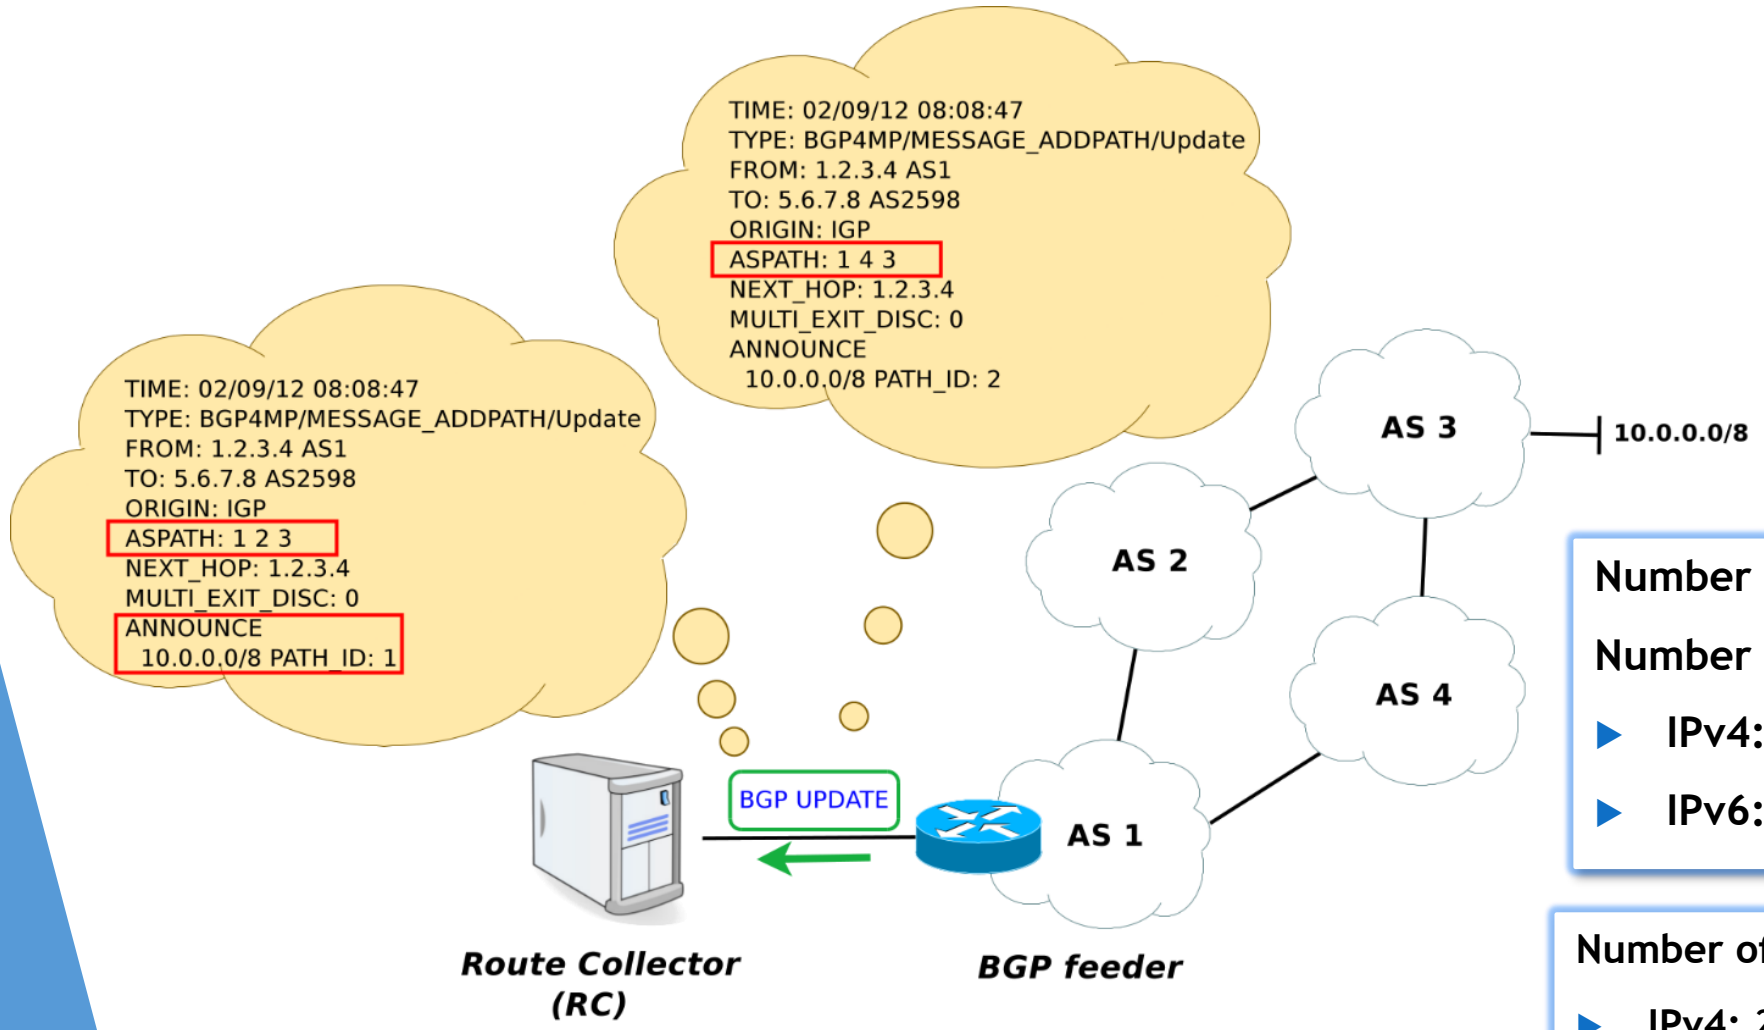

Number of ADDPATH ASes participating: *13*

Number of sessions configured: *35*

- **IPv4:** *18* from *11* different ASes
- **IPv6:** *17* from *11* different ASes

Number of full routing tables:

- **IPv4:** *221 (+99)* from *131 (+53)* different ASes
- **IPv6:** *236 (+103)* from *136 (+54)* different ASes

What do we provide to the community?

MRT data

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

  * same format as RIPE RIS and Route Views (RFC 6396, ADDPATH RFC 8050)

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

Periodic analyses

- AS characteristics
- Feeder contribution
- Total coverage of route collectors

Open source software

- Interactive Collecting Engine (ICE)
- MRT Data Reader

What's next?

New services

- Bogon real-time recognition
- Real-time looking glass
- Route collector on BMP (RFC 7854)

Our future research directions

- Real-time routing anomaly detection (e.g. prefix hijack)
- Pattern recognition in BGP attributes
- Country-focused special analyses (e.g. Internet shutdown recognition)
- Feeder data hygiene techniques

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

Questions?

alessandro.improta@iit.cnr.it

info@isolario.it