Measuring routing (in)security

Disclaimer – this is not a scientific research

Andrei Robachevsky

manrs@isoc.org

Why to measure?

Provide a factual state of routing security as it relates to MANRS

- Support the problem statement with data
- Demonstrate the impact and progress
- Network, country, region, over time

Inform MANRS members about their degree of commitment

- Improve reputation and transparency of the effort

Automate the process

- Make it more comprehensive and consistent
- Reduce the load
- Allow preparation (self-assessment)

How to measure?

Transparent.

- The measurements should use publicly available data sources and the code should be made open source.

Passive

- No cooperation is required from a network.

MANRS Actions

Filtering

Prevent propagation of incorrect routing information

Ensure the correctness of your own announcements and announcements from your customers to adjacent networks with prefix and AS-path granularity

Anti-spoofing

Prevent traffic with spoofed source IP addresses

Enable source address validation for at least single-homed stub customer networks, their own end-users, and infrastructure

Coordination

Facilitate global operational communication and coordination between network operators

Maintain globally accessible up-to-date contact information in common routing databases

Global Validation

Facilitate validation of routing information on a global scale

Publish your data, so others can validate

What can we measure?

Action 1: Filtering

Metric	Description
M1	route leak by the AS
M2	route hijack by the AS
M1C	route leak by a customer and not filtered by the AS
M2C	route hijack by a customer and not filtered by the AS
M3	announcement of bogon prefixes
M4	announcement of bogon ASNs (unallocated/reserved)

Action 2: Anti-spoofing

Metric	Description
M5	spoofable IP blocks
M5C	spoofable IP blocks of client AS’es

Action 3: Coordination

Metric	Description
M8	contact registration (RIR, IRR, PeeringDB)

Action 4: Facilitate global validation

Metric	Description
M6	policy documented in an IRR (aut-num w/import/export, as-set)
M7IRR	registered routes (% of routes registered)
M7RPKI	valid ROAs (% of routes registered)
M7CIRR	registered customer routes (% of routes registered)
M7CRPKI	valid ROAs for customer routes (% of routes registered)

How to calculate? E.g. M2 - route hijack by an AS?

Impact

- M2 = $f(\#prefixes, address\ span, duration)$
- Not all prefixes are equal
- Does size matter?
- Hard to normalize/define thresholds

Conformity

- M2 = $f(\#distinct\ incidents, resolution\ time)$
- # incidents and resolution time show the degree of negligence
- What is an incident?
- Finite number – easy to define thresholds

Events and incidents. E.g. M2C

Weight

- Events are weighted depending on the distance from the culprit
- M1C (ASPATH-1), 0.5*M1C(ASPATH-2), 0.25*M1C(ASPATH-3)… min 0.01

Incident

- Events with the **same weight** that share the **same time span** are merged into an **incident**.

Duration

- Non-action is penalized
- < 30mins -> 0.5 * weight
- < 24hours -> 1.0 * weight
- < 48hours -> 2.0 * weight

Example: direct customer hijacks prefixes

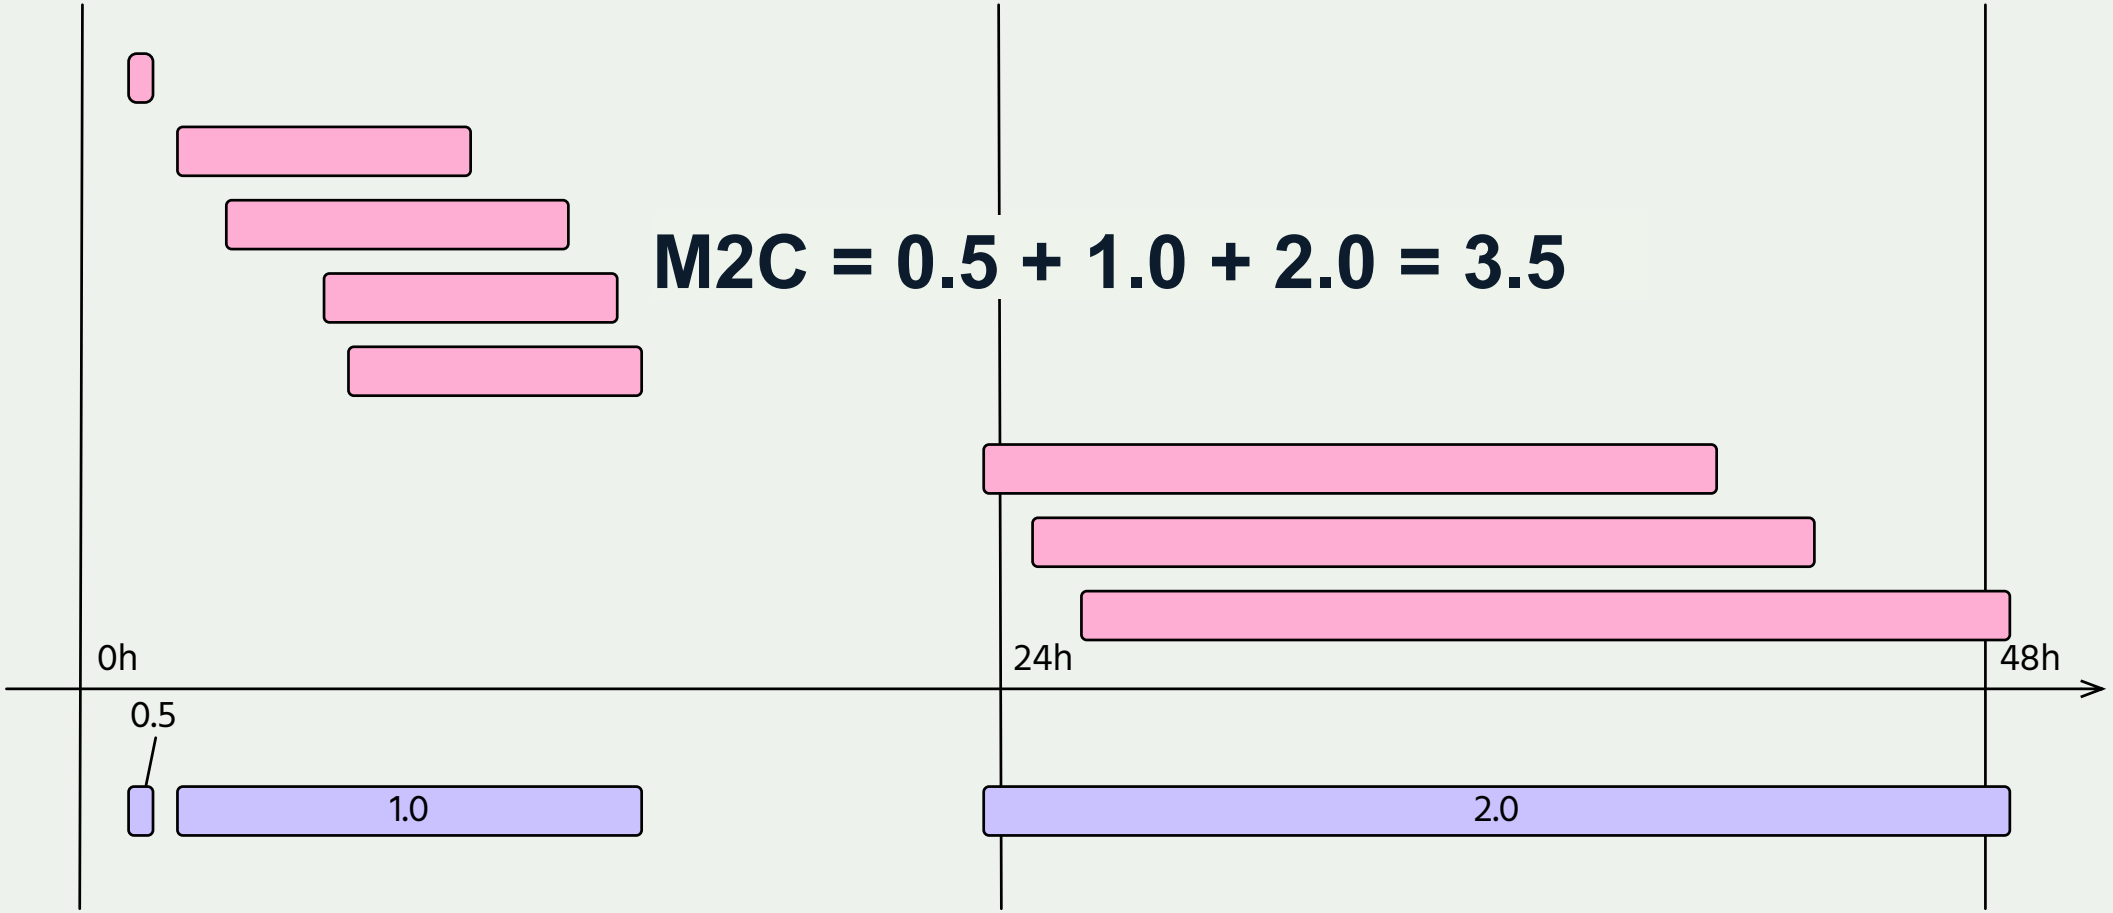

Feedback and ideas are welcome!

robachevsky@isoc.org

Backup slides

How does it all fit together?

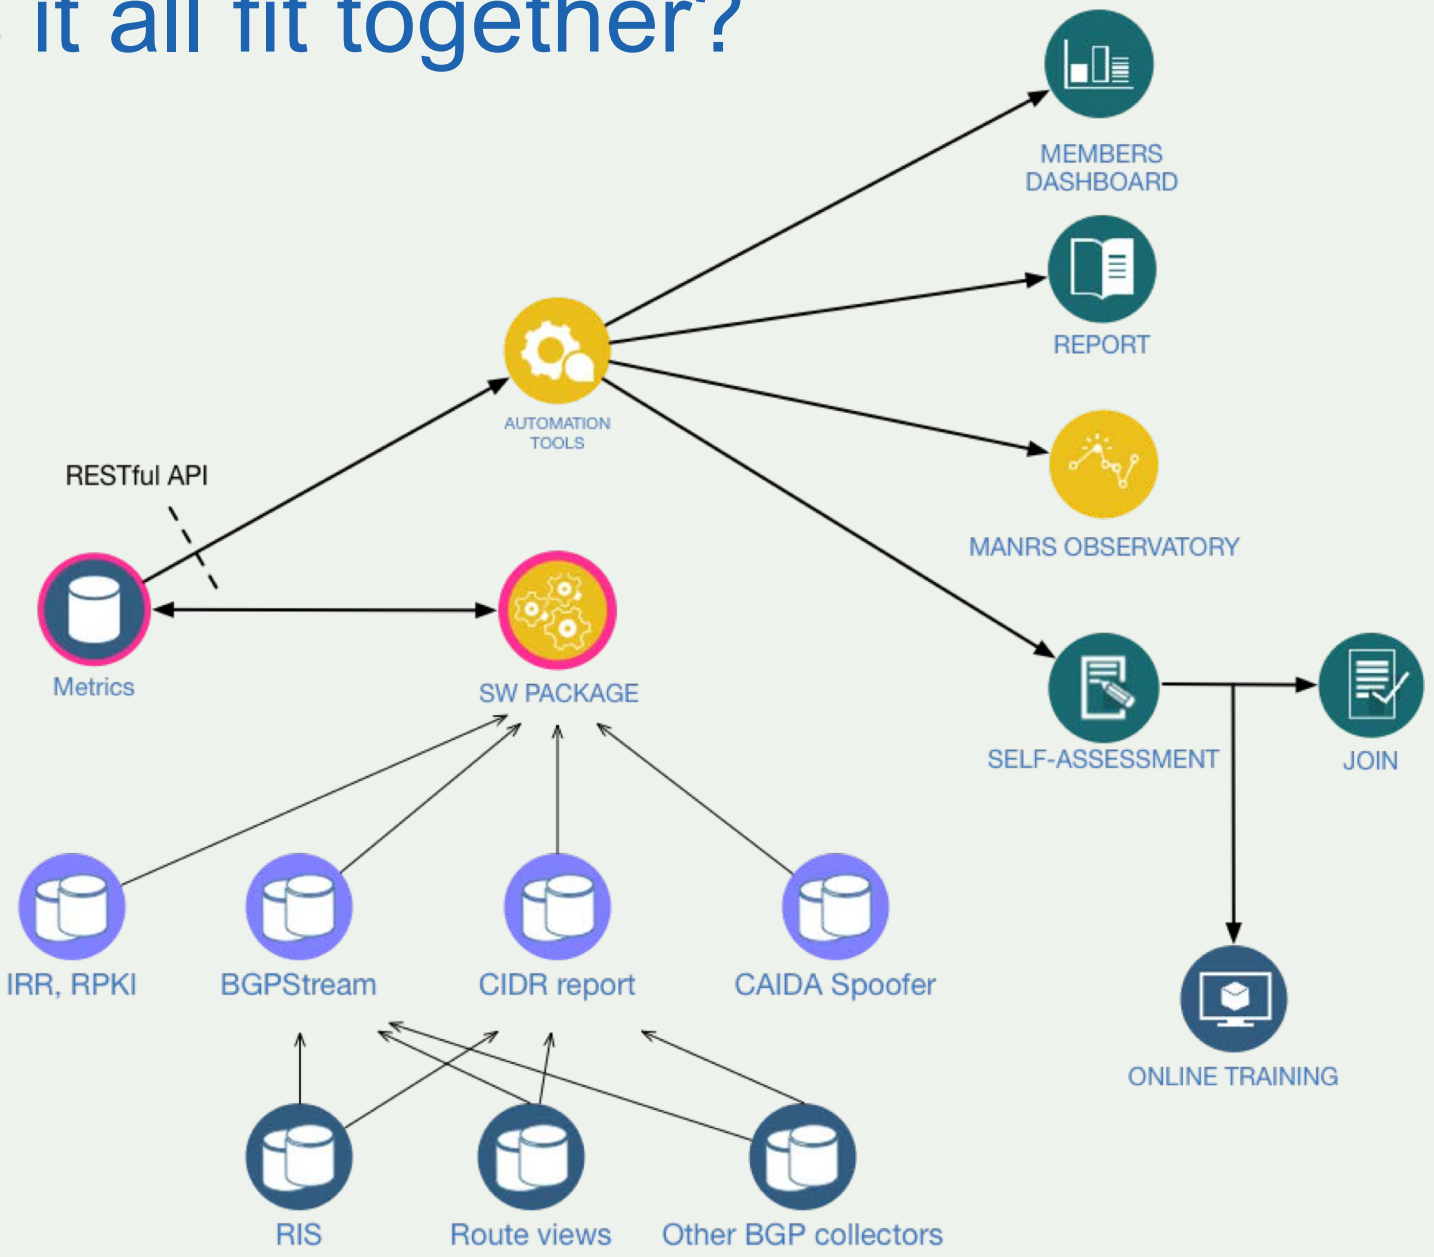

Thank you.

MANRS@isoc.org

manrs.org