# RPKI Deployment: 2020 SEACOM Update

Mark Tinka

Head of Engineering

SEACOM

# Why RPKI

- 99% of all mis-originations are accidental:
    - Pakistan/Youtube.
    - MainOne/Google.
    - Init7 at DE-CIX.
    - AS7007 incident of April, 1997.
    - e.t.c.

- Won't solve AS_PATH attacks.
    - Demo'd by Anton Kapela & Alex Pilosov at DEFCON 2008.
    - Solution is Path Validation.
    - Still years away.

# RPKI Validation

- There are a number of free validation tools:
    - RIPE NCC RPKI Validator.
    - Dragon Research RPKI Toolkit.
    - Relying Party Security Technology for Internet Routing (RPSTIR).
    - Routinator (NLnet Labs).
    - RTRlib.

# RPKI Validation

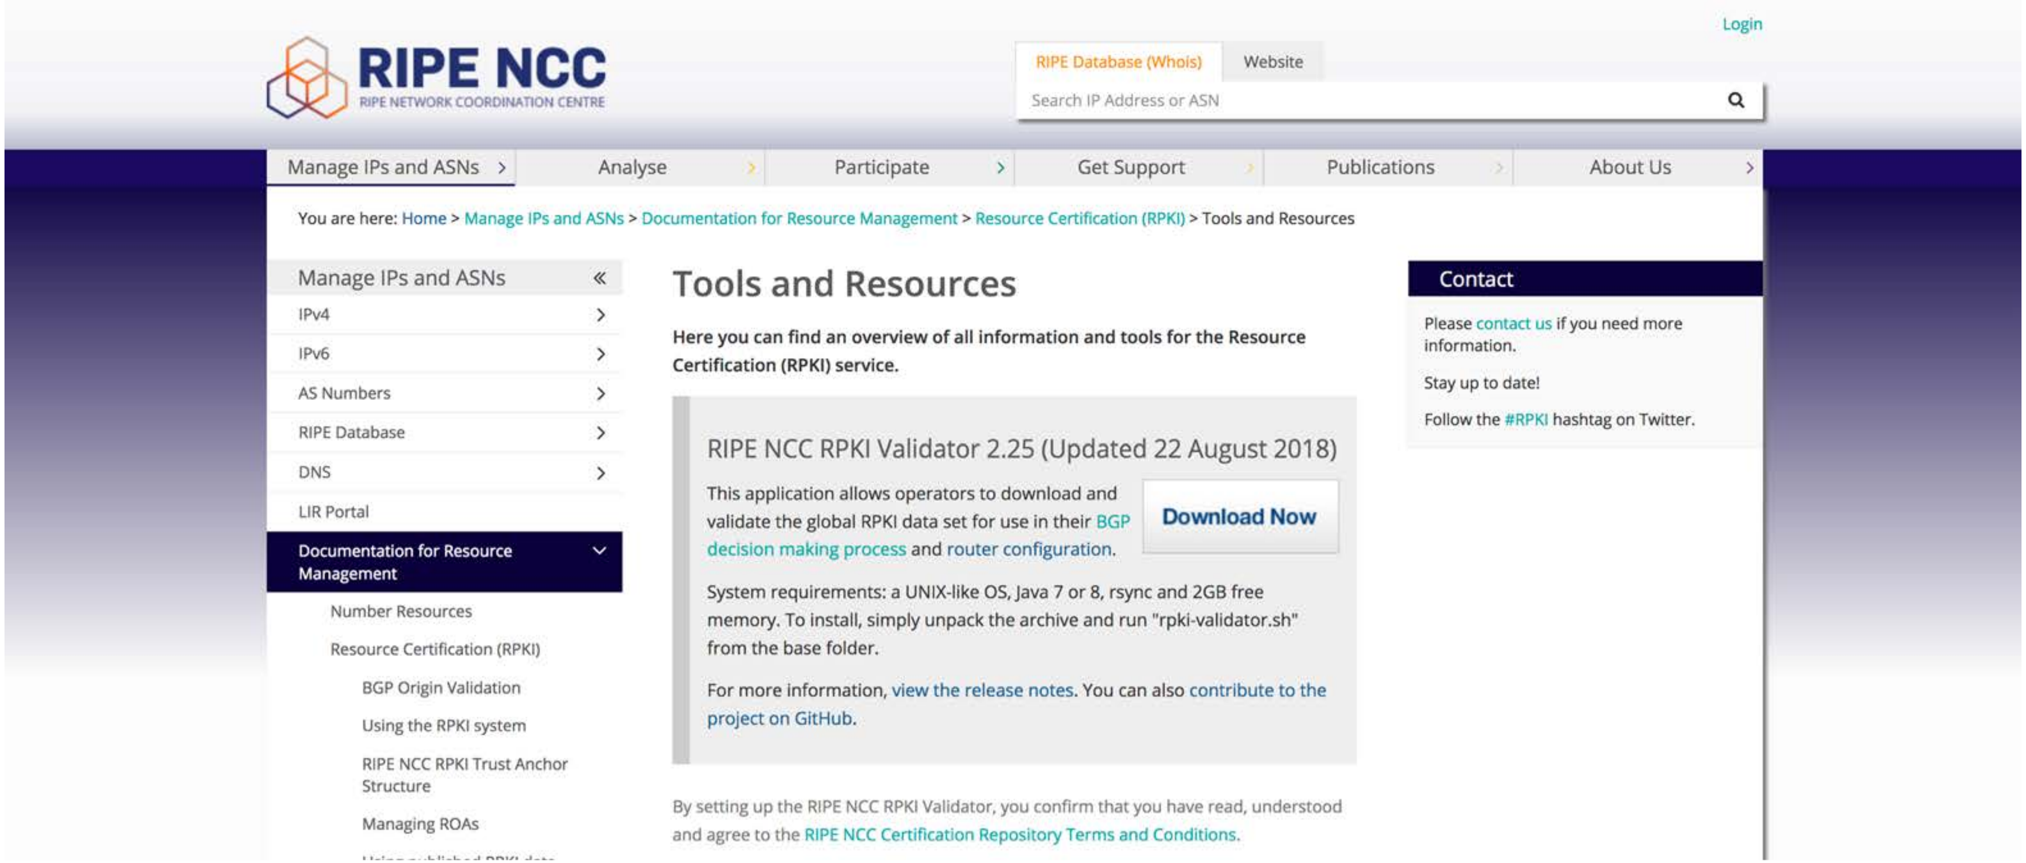

You are here: Home > Manage IPs and ASNs > Documentation for Resource Management > Resource Certification (RPKI) > Tools and Resources

## Tools and Resources

**Here you can find an overview of all information and tools for the Resource Certification (RPKI) service.**

### RIPE NCC RPKI Validator 2.25 (Updated 22 August 2018)

This application allows operators to download and validate the global RPKI data set for use in their BGP decision making process and router configuration.

System requirements: a UNIX-like OS, Java 7 or 8, rsync and 2GB free memory. To install, simply unpack the archive and run "rpki-validator.sh" from the base folder.

For more information, view the release notes. You can also contribute to the project on GitHub.

By setting up the RIPE NCC RPKI Validator, you confirm that you have read, understood and agree to the RIPE NCC Certification Repository Terms and Conditions.

### Contact

Please contact us if you need more information.

Stay up to date!

Follow the #RPKI hashtag on Twitter.

# Live Network

```
tinka@er-01-jnb.za-re0# run show validation session
Session                          State   Flaps     Uptime #IPv4/IPv6 records
105.16.aaa.b                       Up         1 2w1d 21:17:43 103943/17348
105.16.ccc.d                      Up         0 2w1d 23:02:53 103943/17348
2c0f:feb0:X:Y::Z                    Up         1 2w1d 21:17:42 103943/17348
2c0f:feb0:U:V::W                     Up         0 2w1d 23:02:53 103943/17348

{master}[edit]
tinka@er-01-jnb.za-re0#
```

# Live Network

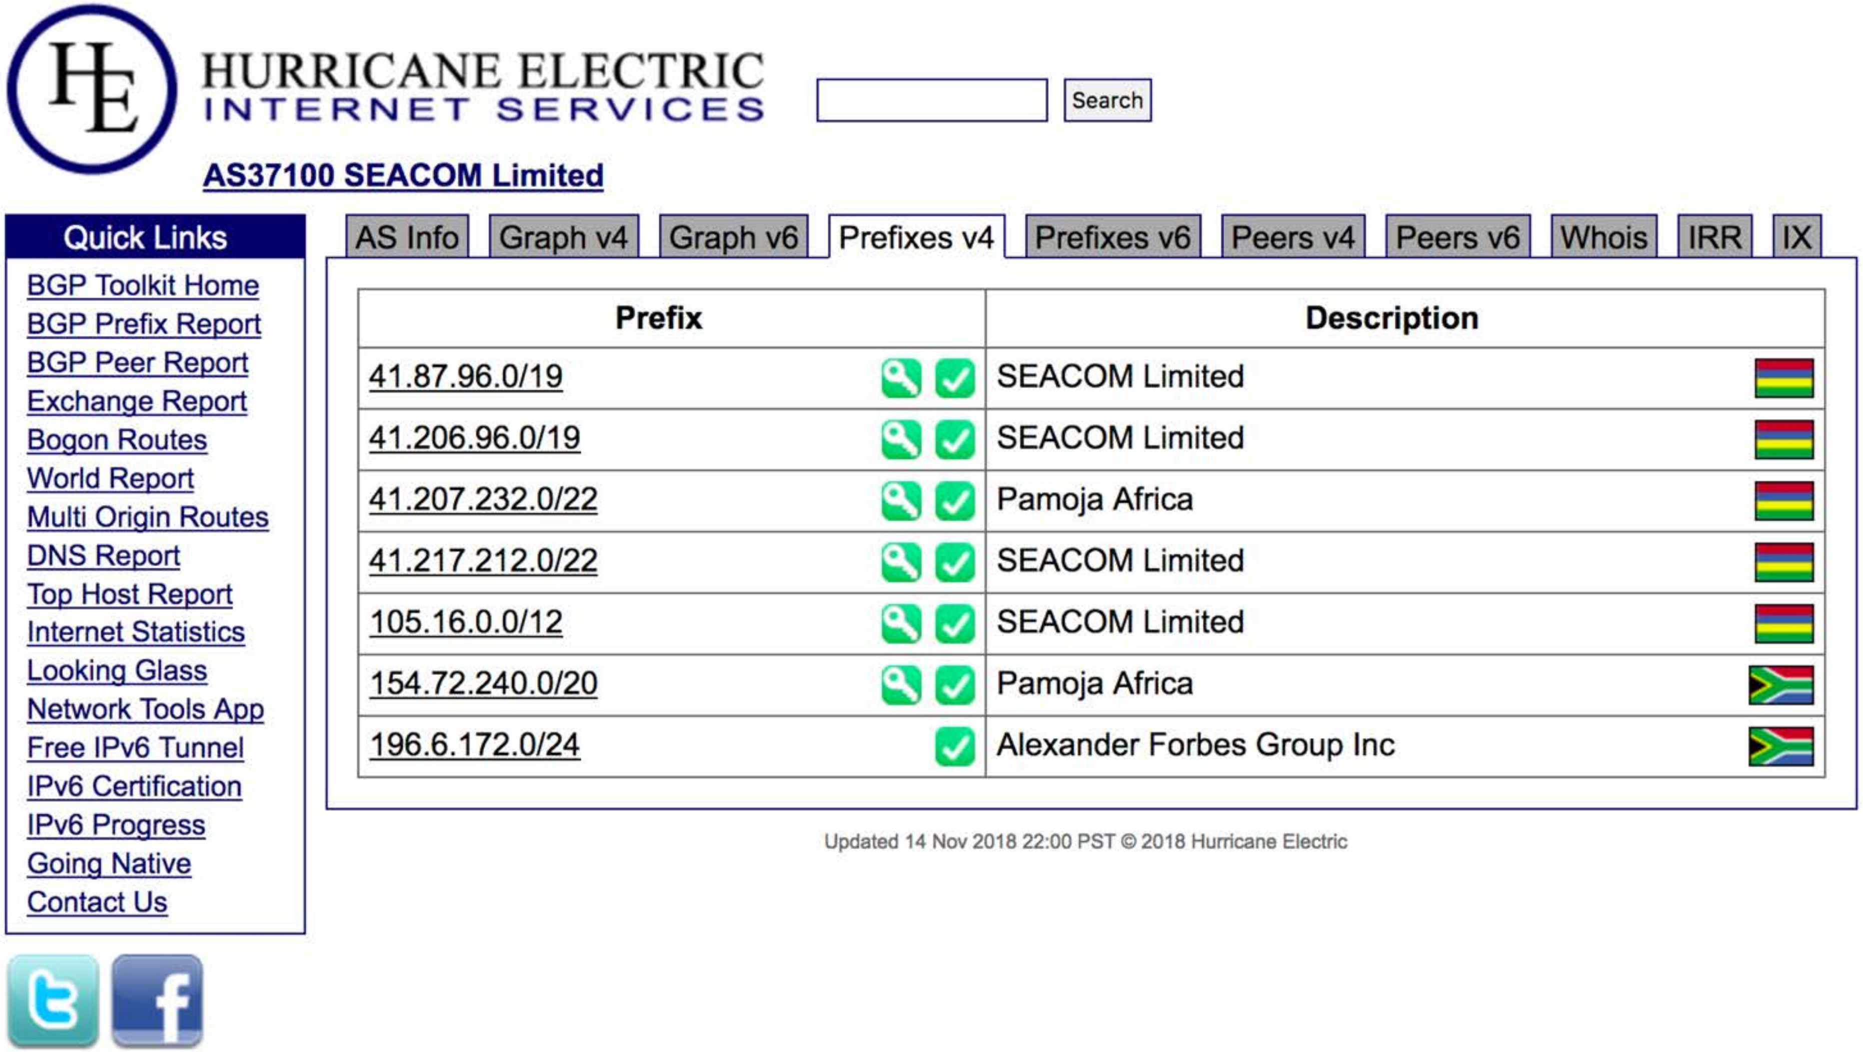

HURRICANE ELECTRIC
INTERNET SERVICES

Search

**AS37100 SEACOM Limited**

**Quick Links**

- BGP Toolkit Home
- BGP Prefix Report
- BGP Peer Report
- Exchange Report
- Bogon Routes
- World Report
- Multi Origin Routes
- DNS Report
- Top Host Report
- Internet Statistics
- Looking Glass
- Network Tools App
- Free IPv6 Tunnel
- IPv6 Certification
- IPv6 Progress
- Going Native
- Contact Us

AS Info | Graph v4 | Graph v6 | Prefixes v4 | Prefixes v6 | Peers v4 | Peers v6 | Whois | IRR | IX

| Prefix | | Description | |
|---|---|---|---|
| 41.87.96.0/19 | 🔍 ✓ | SEACOM Limited | 🇲🇺 |
| 41.206.96.0/19 | 🔍 ✓ | SEACOM Limited | 🇲🇺 |
| 41.207.232.0/22 | 🔍 ✓ | Pamoja Africa | 🇲🇺 |
| 41.217.212.0/22 | 🔍 ✓ | SEACOM Limited | 🇲🇺 |
| 105.16.0.0/12 | 🔍 ✓ | SEACOM Limited | 🇲🇺 |
| 154.72.240.0/20 | 🔍 ✓ | Pamoja Africa | 🇿🇦 |
| 196.6.172.0/24 | ✓ | Alexander Forbes Group Inc | 🇿🇦 |

Updated 14 Nov 2018 22:00 PST © 2018 Hurricane Electric

# Live Network

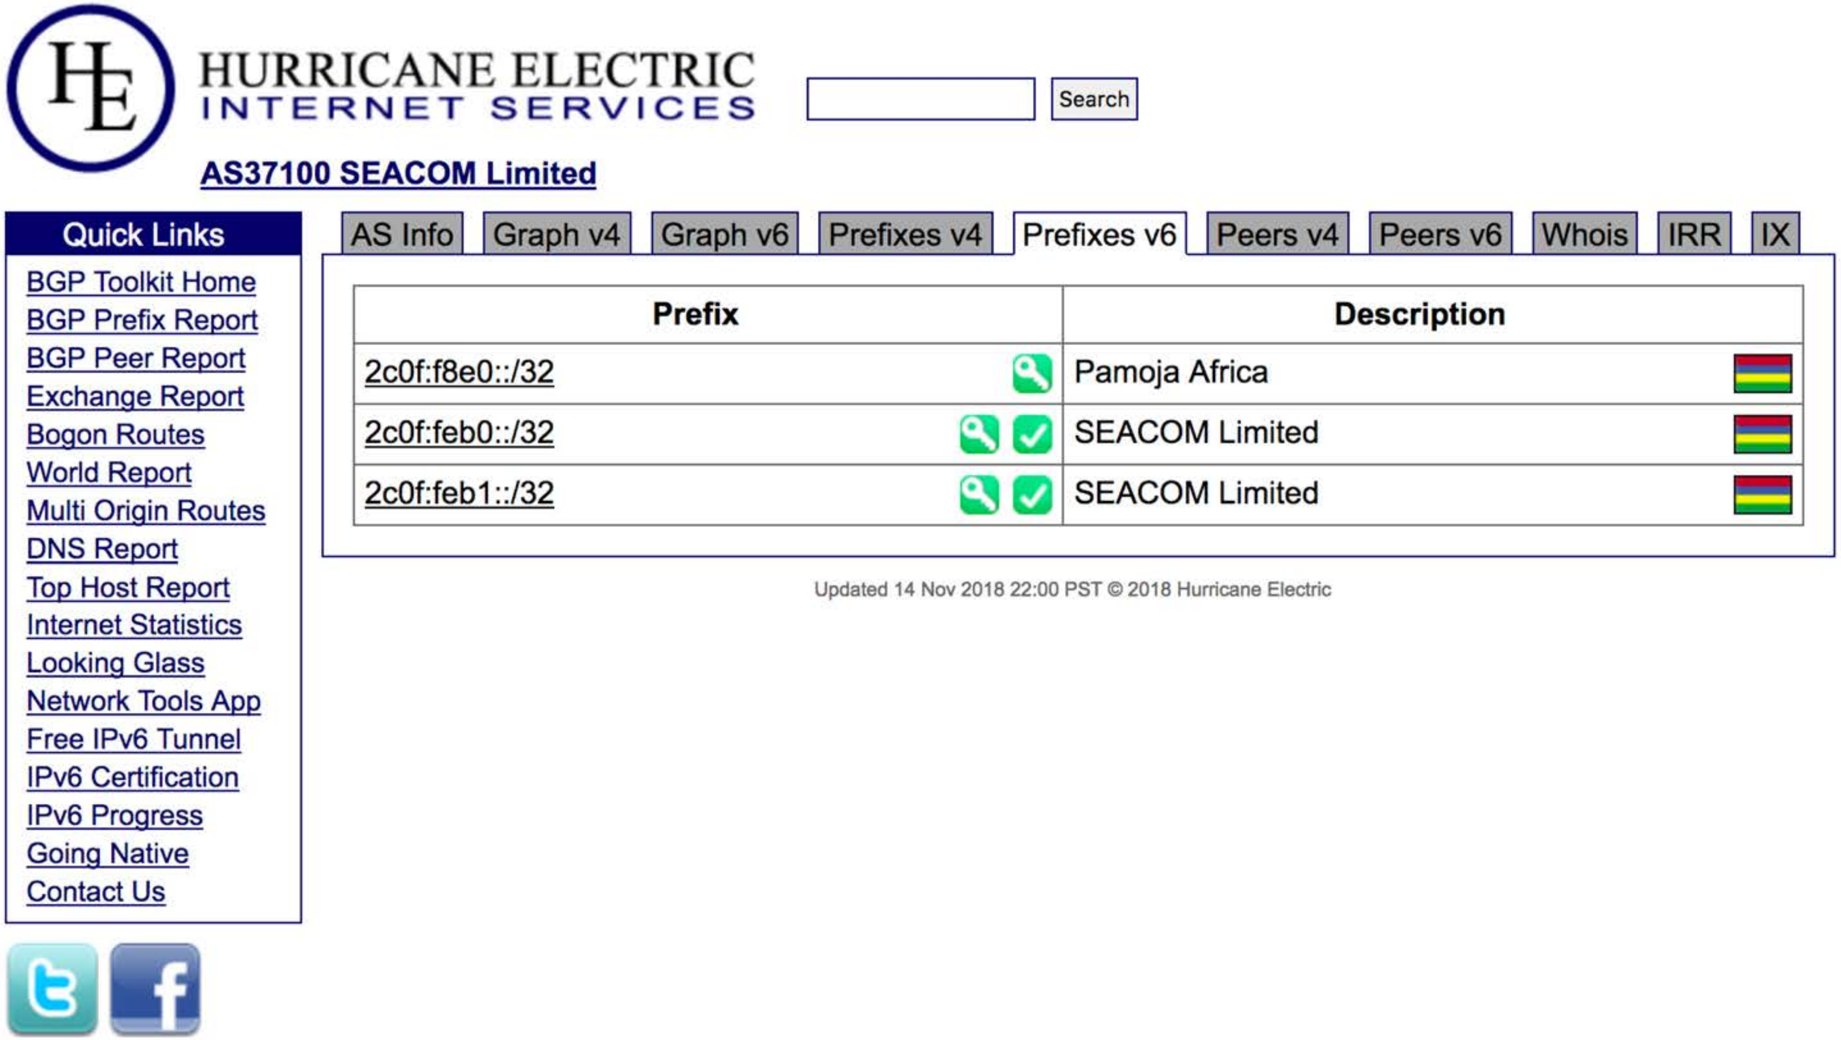

| Prefix | Description |
|---|---|
| 2c0f:f8e0::/32 🔍 | Pamoja Africa |
| 2c0f:feb0::/32 🔍 ✓ | SEACOM Limited |
| 2c0f:feb1::/32 🔍 ✓ | SEACOM Limited |

Updated 14 Nov 2018 22:00 PST © 2018 Hurricane Electric

# Things To Look Out For

- ARIN TAL (Trust Anchor Locator):
  - The ARIN TAL does not ship with RP tools.
  - Users must first specifically agree with ARIN’s RPA (RP Agreement).

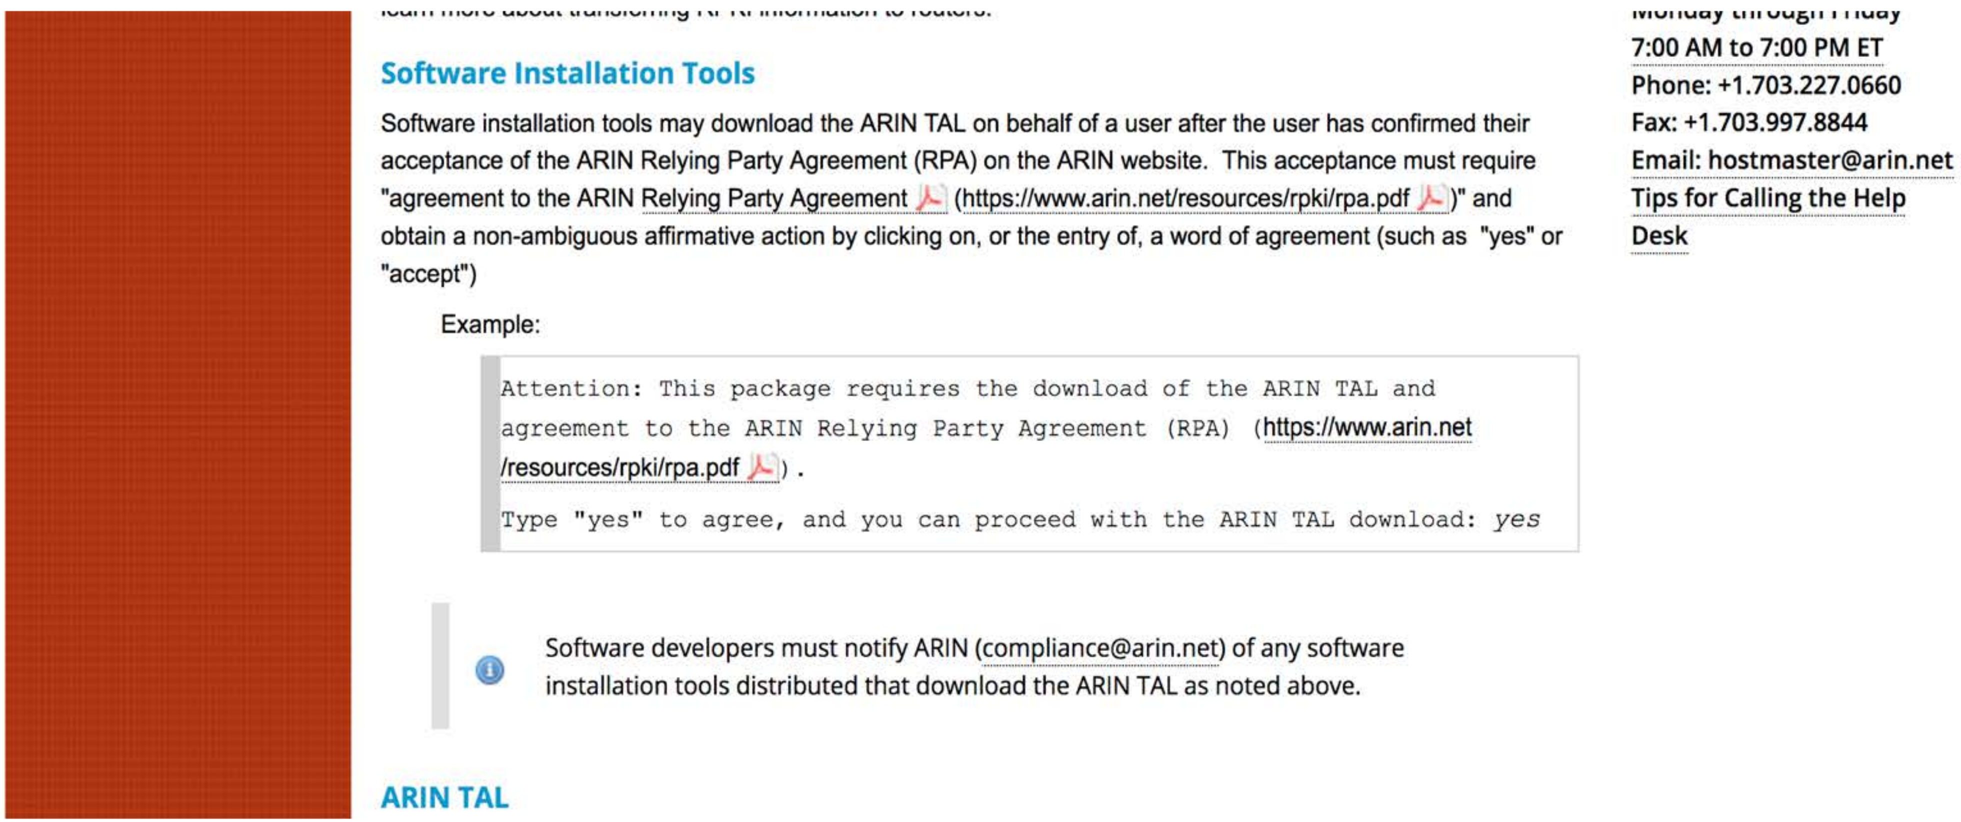

learn more about transforming RPKI information to routers.

## Software Installation Tools

Software installation tools may download the ARIN TAL on behalf of a user after the user has confirmed their acceptance of the ARIN Relying Party Agreement (RPA) on the ARIN website. This acceptance must require "agreement to the ARIN Relying Party Agreement (https://www.arin.net/resources/rpki/rpa.pdf)" and obtain a non-ambiguous affirmative action by clicking on, or the entry of, a word of agreement (such as "yes" or "accept")

Example:

```
Attention: This package requires the download of the ARIN TAL and
agreement to the ARIN Relying Party Agreement (RPA) (https://www.arin.net
/resources/rpki/rpa.pdf) .

Type "yes" to agree, and you can proceed with the ARIN TAL download: yes
```

Software developers must notify ARIN (compliance@arin.net) of any software installation tools distributed that download the ARIN TAL as noted above.

## ARIN TAL

Monday through Friday
7:00 AM to 7:00 PM ET
Phone: +1.703.227.0660
Fax: +1.703.997.8844
Email: hostmaster@arin.net
Tips for Calling the Help Desk

# Things To Look Out For

- RIPE NCC Validator Memory Requirements:
    - Default memory settings may lead to a crash.
    - Increase ”Maximum” memory to at least 4GB.

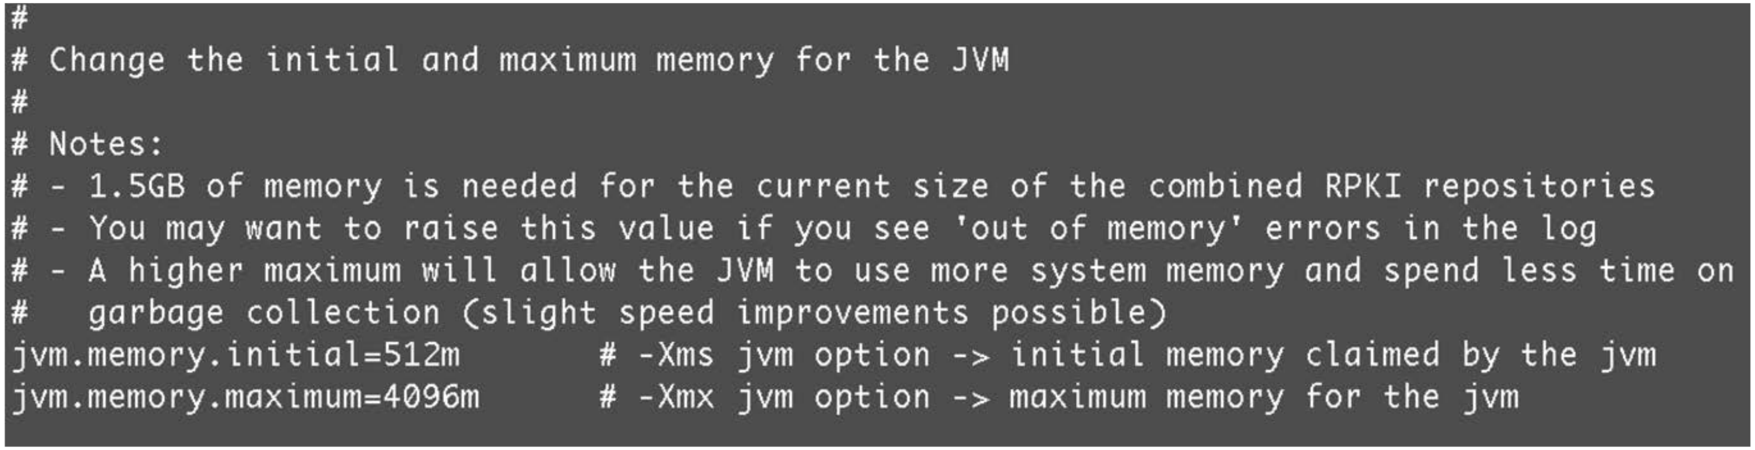

```
#
# Change the initial and maximum memory for the JVM
#
# Notes:
# - 1.5GB of memory is needed for the current size of the combined RPKI repositories
# - You may want to raise this value if you see 'out of memory' errors in the log
# - A higher maximum will allow the JVM to use more system memory and spend less time on
#   garbage collection (slight speed improvements possible)
jvm.memory.initial=512m       # -Xms jvm option -> initial memory claimed by the jvm
jvm.memory.maximum=4096m      # -Xmx jvm option -> maximum memory for the jvm
```

# Things To Look Out For – IOS/IOS XE Doing Its Own Things

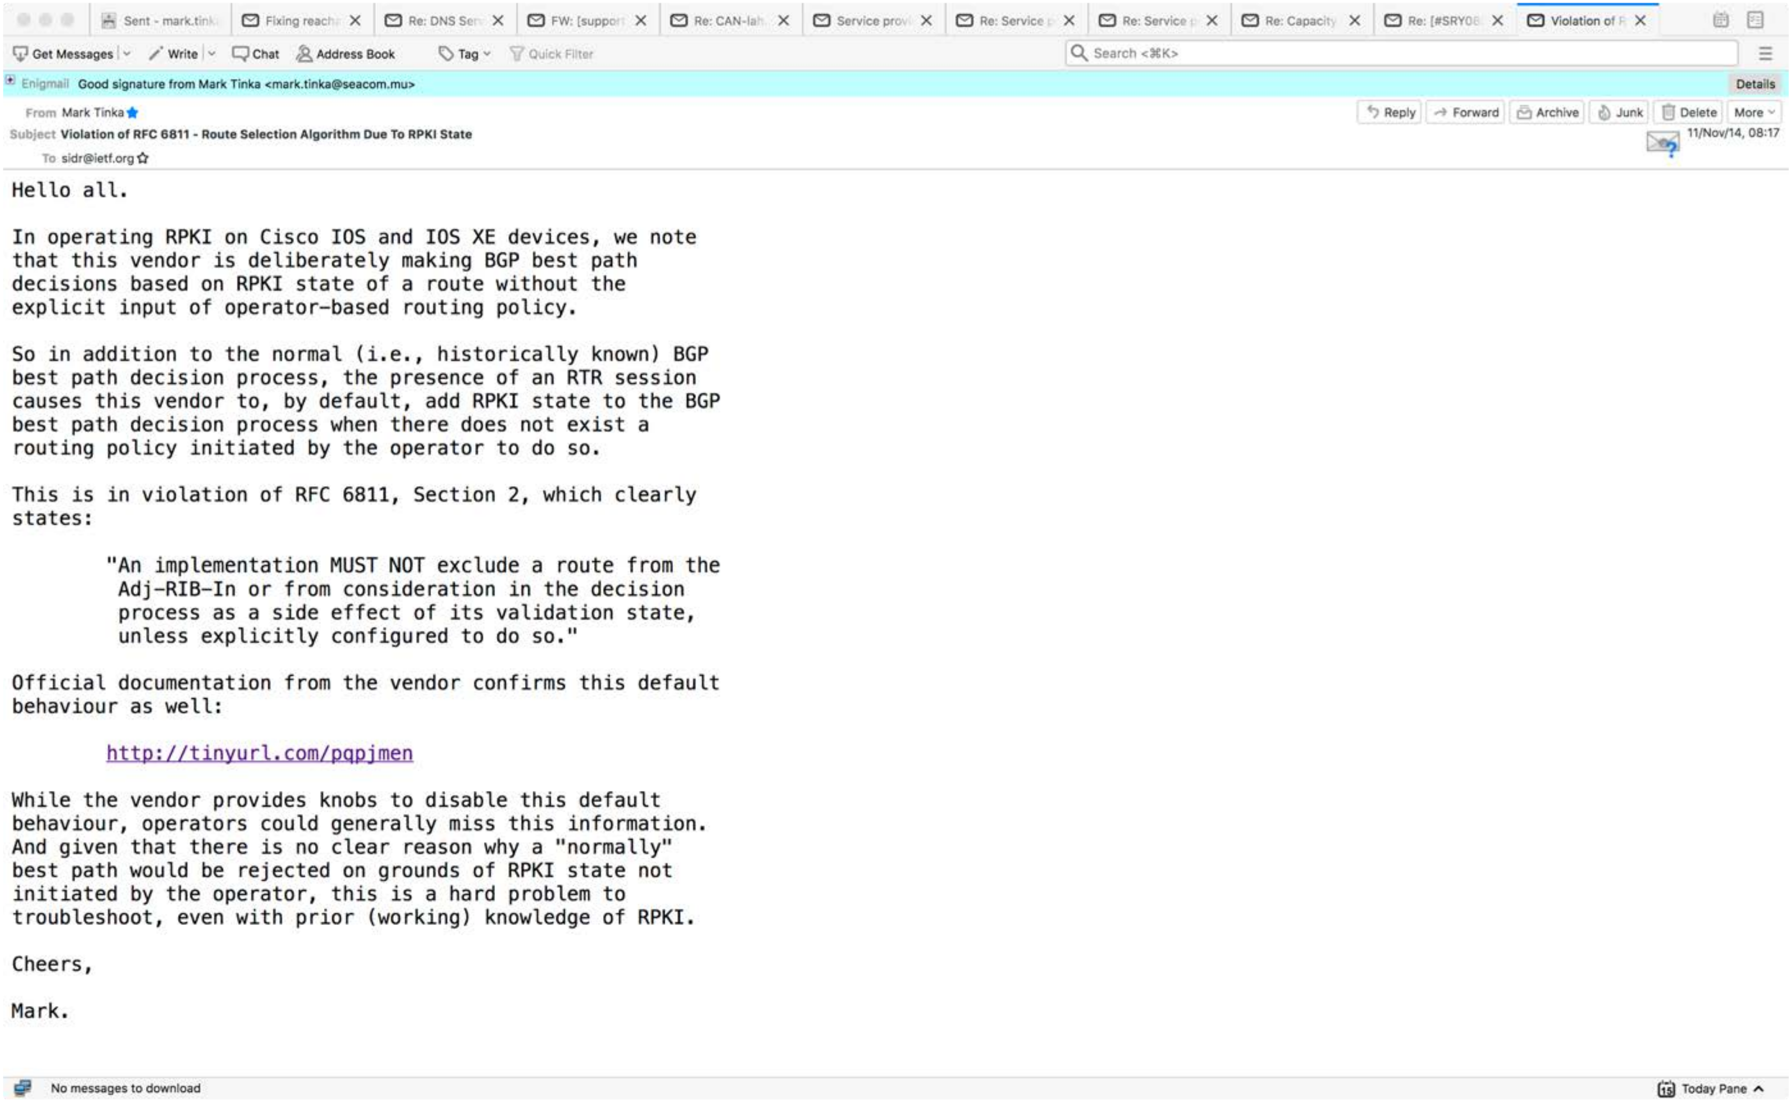

Enigmail Good signature from Mark Tinka <mark.tinka@seacom.mu>

From Mark Tinka

Subject **Violation of RFC 6811 - Route Selection Algorithm Due To RPKI State**

To sidr@ietf.org

11/Nov/14, 08:17

Hello all.

In operating RPKI on Cisco IOS and IOS XE devices, we note that this vendor is deliberately making BGP best path decisions based on RPKI state of a route without the explicit input of operator-based routing policy.

So in addition to the normal (i.e., historically known) BGP best path decision process, the presence of an RTR session causes this vendor to, by default, add RPKI state to the BGP best path decision process when there does not exist a routing policy initiated by the operator to do so.

This is in violation of RFC 6811, Section 2, which clearly states:

> "An implementation MUST NOT exclude a route from the Adj-RIB-In or from consideration in the decision process as a side effect of its validation state, unless explicitly configured to do so."

Official documentation from the vendor confirms this default behaviour as well:

> http://tinyurl.com/pqpjmen

While the vendor provides knobs to disable this default behaviour, operators could generally miss this information. And given that there is no clear reason why a "normally" best path would be rejected on grounds of RPKI state not initiated by the operator, this is a hard problem to troubleshoot, even with prior (working) knowledge of RPKI.

Cheers,

Mark.

# Things To Look Out For – IOS/IOS XE Doing Its Own Things

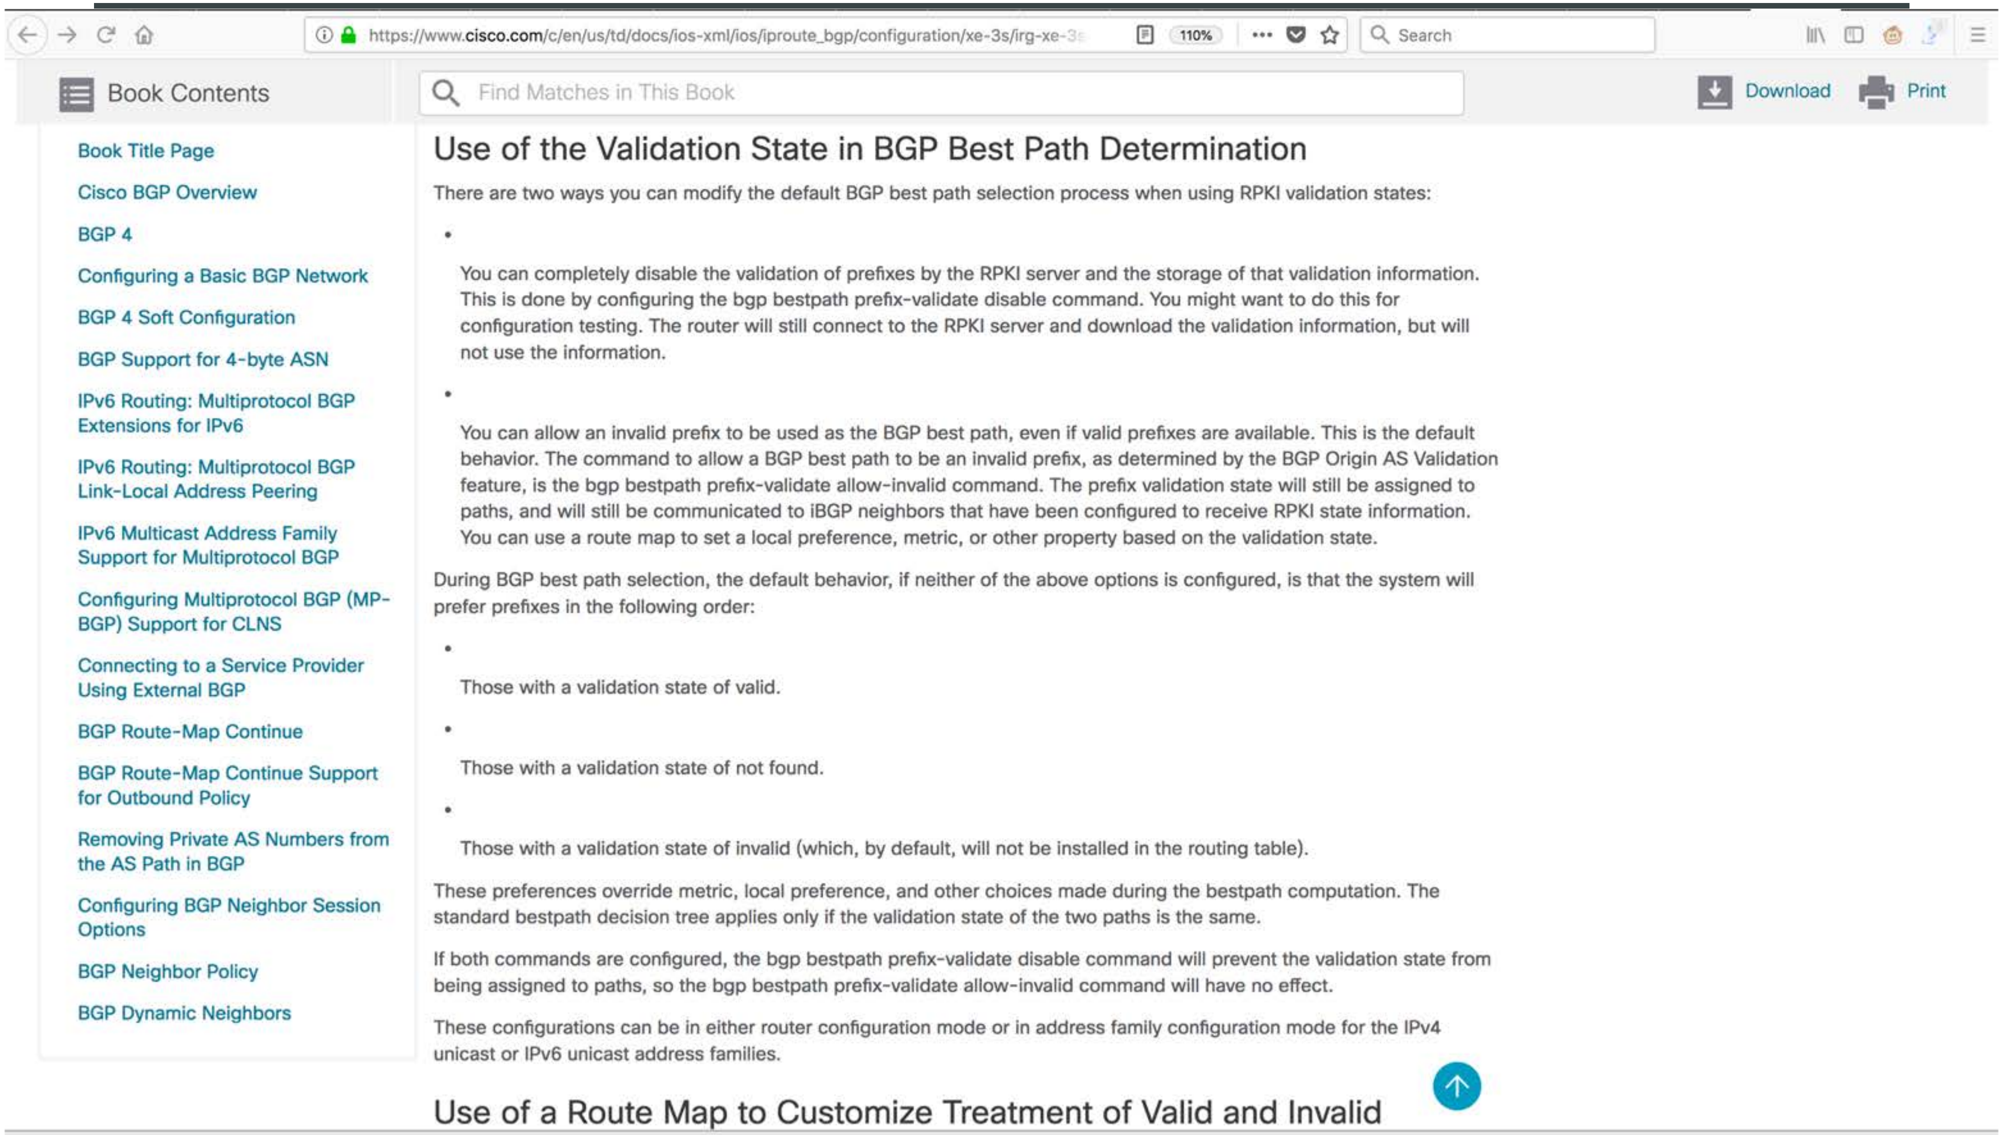
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s
110%
Search
Book Contents
Find Matches in This Book
Download
Print
Book Title Page
Cisco BGP Overview
BGP 4
Configuring a Basic BGP Network
BGP 4 Soft Configuration
BGP Support for 4-byte ASN
IPv6 Routing: Multiprotocol BGP Extensions for IPv6
IPv6 Routing: Multiprotocol BGP Link-Local Address Peering
IPv6 Multicast Address Family Support for Multiprotocol BGP
Configuring Multiprotocol BGP (MP-BGP) Support for CLNS
Connecting to a Service Provider Using External BGP
BGP Route-Map Continue
BGP Route-Map Continue Support for Outbound Policy
Removing Private AS Numbers from the AS Path in BGP
Configuring BGP Neighbor Session Options
BGP Neighbor Policy
BGP Dynamic Neighbors

## Use of the Validation State in BGP Best Path Determination

There are two ways you can modify the default BGP best path selection process when using RPKI validation states:

- You can completely disable the validation of prefixes by the RPKI server and the storage of that validation information. This is done by configuring the bgp bestpath prefix-validate disable command. You might want to do this for configuration testing. The router will still connect to the RPKI server and download the validation information, but will not use the information.

- You can allow an invalid prefix to be used as the BGP best path, even if valid prefixes are available. This is the default behavior. The command to allow a BGP best path to be an invalid prefix, as determined by the BGP Origin AS Validation feature, is the bgp bestpath prefix-validate allow-invalid command. The prefix validation state will still be assigned to paths, and will still be communicated to iBGP neighbors that have been configured to receive RPKI state information. You can use a route map to set a local preference, metric, or other property based on the validation state.

During BGP best path selection, the default behavior, if neither of the above options is configured, is that the system will prefer prefixes in the following order:

- Those with a validation state of valid.

- Those with a validation state of not found.

- Those with a validation state of invalid (which, by default, will not be installed in the routing table).

These preferences override metric, local preference, and other choices made during the bestpath computation. The standard bestpath decision tree applies only if the validation state of the two paths is the same.

If both commands are configured, the bgp bestpath prefix-validate disable command will prevent the validation state from being assigned to paths, so the bgp bestpath prefix-validate allow-invalid command will have no effect.

These configurations can be in either router configuration mode or in address family configuration mode for the IPv4 unicast or IPv6 unicast address families.

## Use of a Route Map to Customize Treatment of Valid and Invalid

# Things To Look Out For – IOS/IOS XE Doing Its Own Things

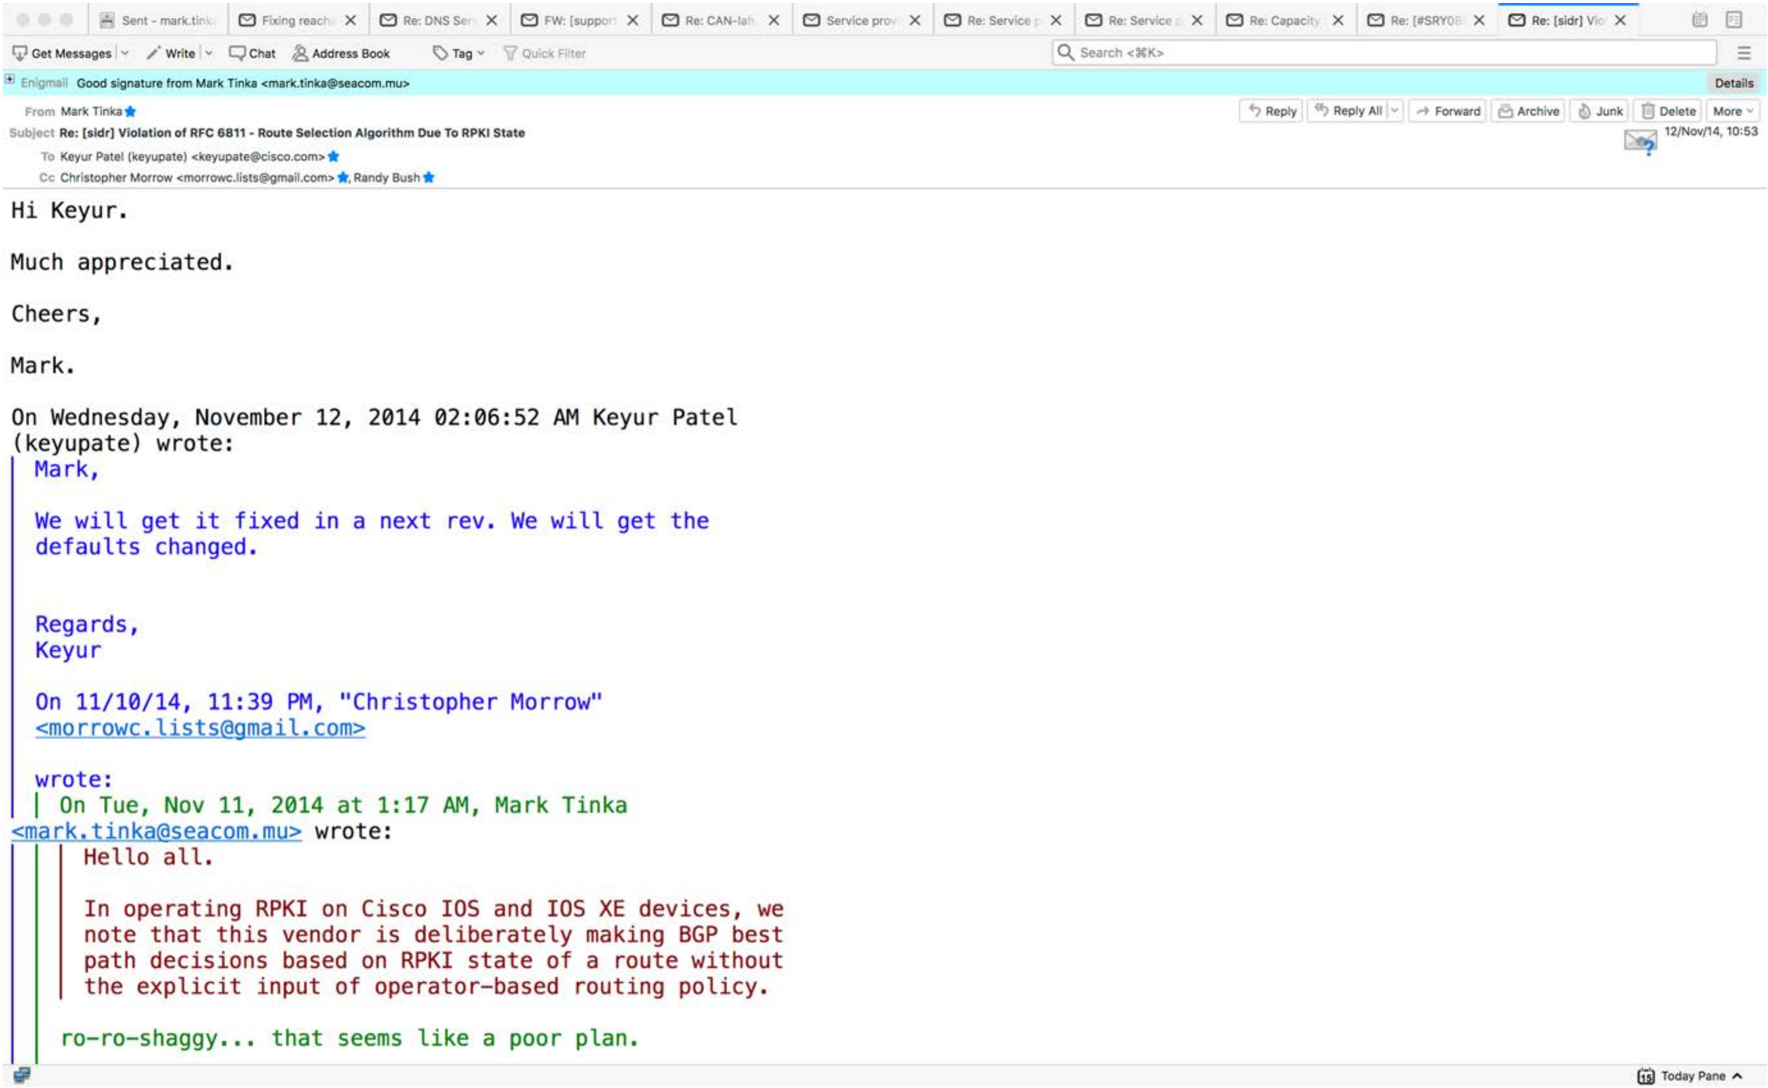

Enigmail Good signature from Mark Tinka <mark.tinka@seacom.mu>

From Mark Tinka

Subject **Re: [sidr] Violation of RFC 6811 - Route Selection Algorithm Due To RPKI State**

To Keyur Patel (keyupate) <keyupate@cisco.com>

Cc Christopher Morrow <morrowc.lists@gmail.com>, Randy Bush

12/Nov/14, 10:53

Hi Keyur.

Much appreciated.

Cheers,

Mark.

On Wednesday, November 12, 2014 02:06:52 AM Keyur Patel (keyupate) wrote:

> Mark,
>
> We will get it fixed in a next rev. We will get the defaults changed.
>
> Regards,
> Keyur
>
> On 11/10/14, 11:39 PM, "Christopher Morrow" <morrowc.lists@gmail.com>
>
> wrote:
>
>> On Tue, Nov 11, 2014 at 1:17 AM, Mark Tinka <mark.tinka@seacom.mu> wrote:
>>
>>> Hello all.
>>>
>>> In operating RPKI on Cisco IOS and IOS XE devices, we note that this vendor is deliberately making BGP best path decisions based on RPKI state of a route without the explicit input of operator-based routing policy.
>>
>> ro-ro-shaggy... that seems like a poor plan.

# Things To Look Out For – IOS/IOS XE Doing Its Own Things

- Over iBGP sessions on IOS and IOS XE, all routes are ”Valid”.
- Violates Section 4 of RFC 8481 and Section 2 of RFC 6811:

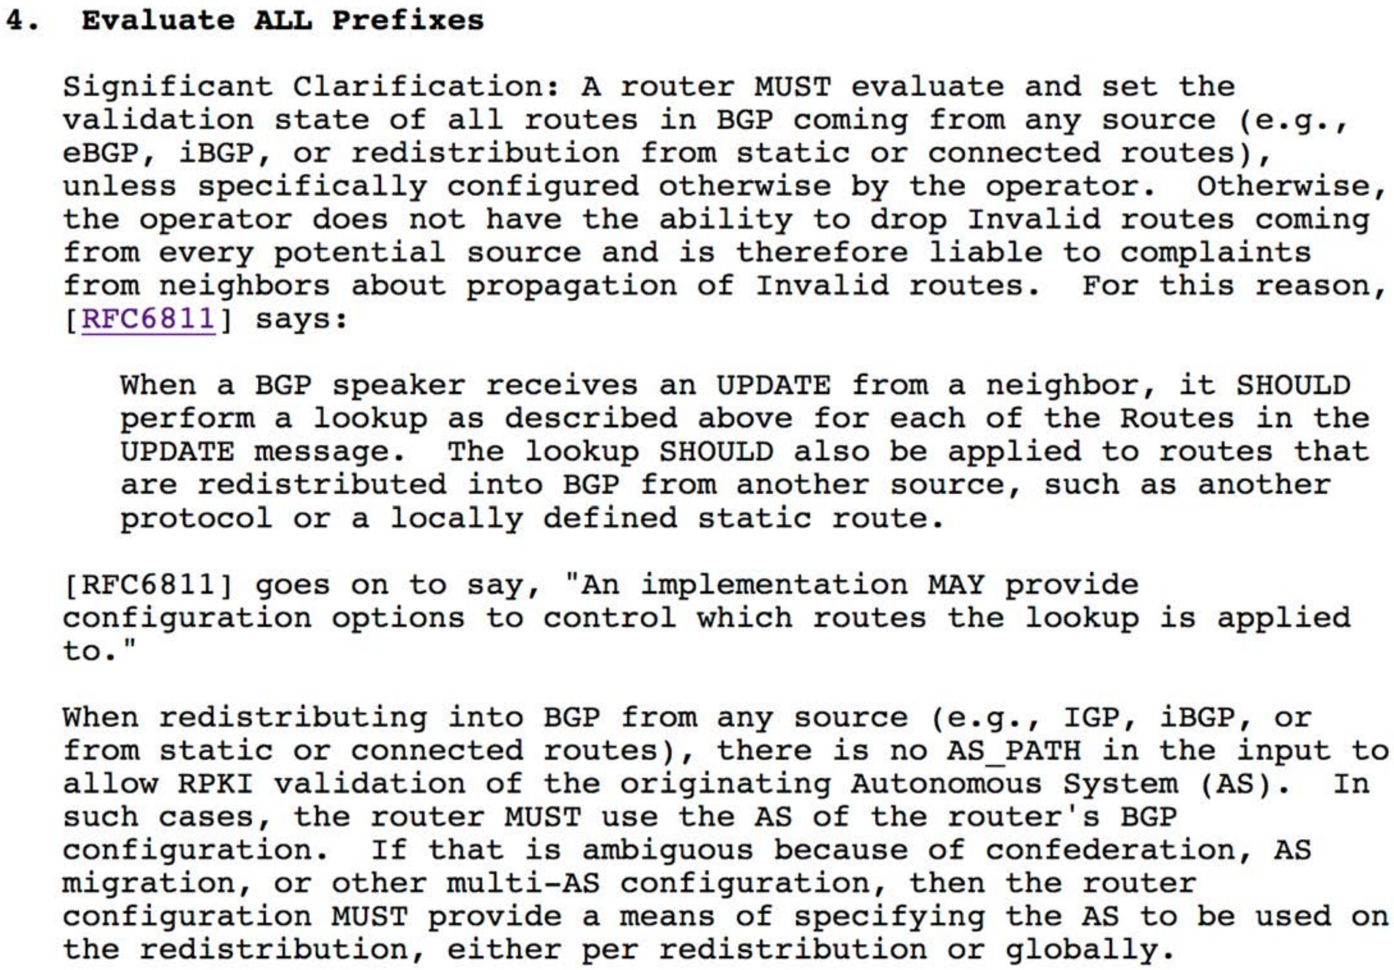

**4. Evaluate ALL Prefixes**

Significant Clarification: A router MUST evaluate and set the validation state of all routes in BGP coming from any source (e.g., eBGP, iBGP, or redistribution from static or connected routes), unless specifically configured otherwise by the operator. Otherwise, the operator does not have the ability to drop Invalid routes coming from every potential source and is therefore liable to complaints from neighbors about propagation of Invalid routes. For this reason, [RFC6811] says:

> When a BGP speaker receives an UPDATE from a neighbor, it SHOULD perform a lookup as described above for each of the Routes in the UPDATE message. The lookup SHOULD also be applied to routes that are redistributed into BGP from another source, such as another protocol or a locally defined static route.

[RFC6811] goes on to say, "An implementation MAY provide configuration options to control which routes the lookup is applied to."

When redistributing into BGP from any source (e.g., IGP, iBGP, or from static or connected routes), there is no AS_PATH in the input to allow RPKI validation of the originating Autonomous System (AS). In such cases, the router MUST use the AS of the router's BGP configuration. If that is ambiguous because of confederation, AS migration, or other multi-AS configuration, then the router configuration MUST provide a means of specifying the AS to be used on the redistribution, either per redistribution or globally.

# Things To Look Out For – IOS/IOS XE Doing Its Own Things

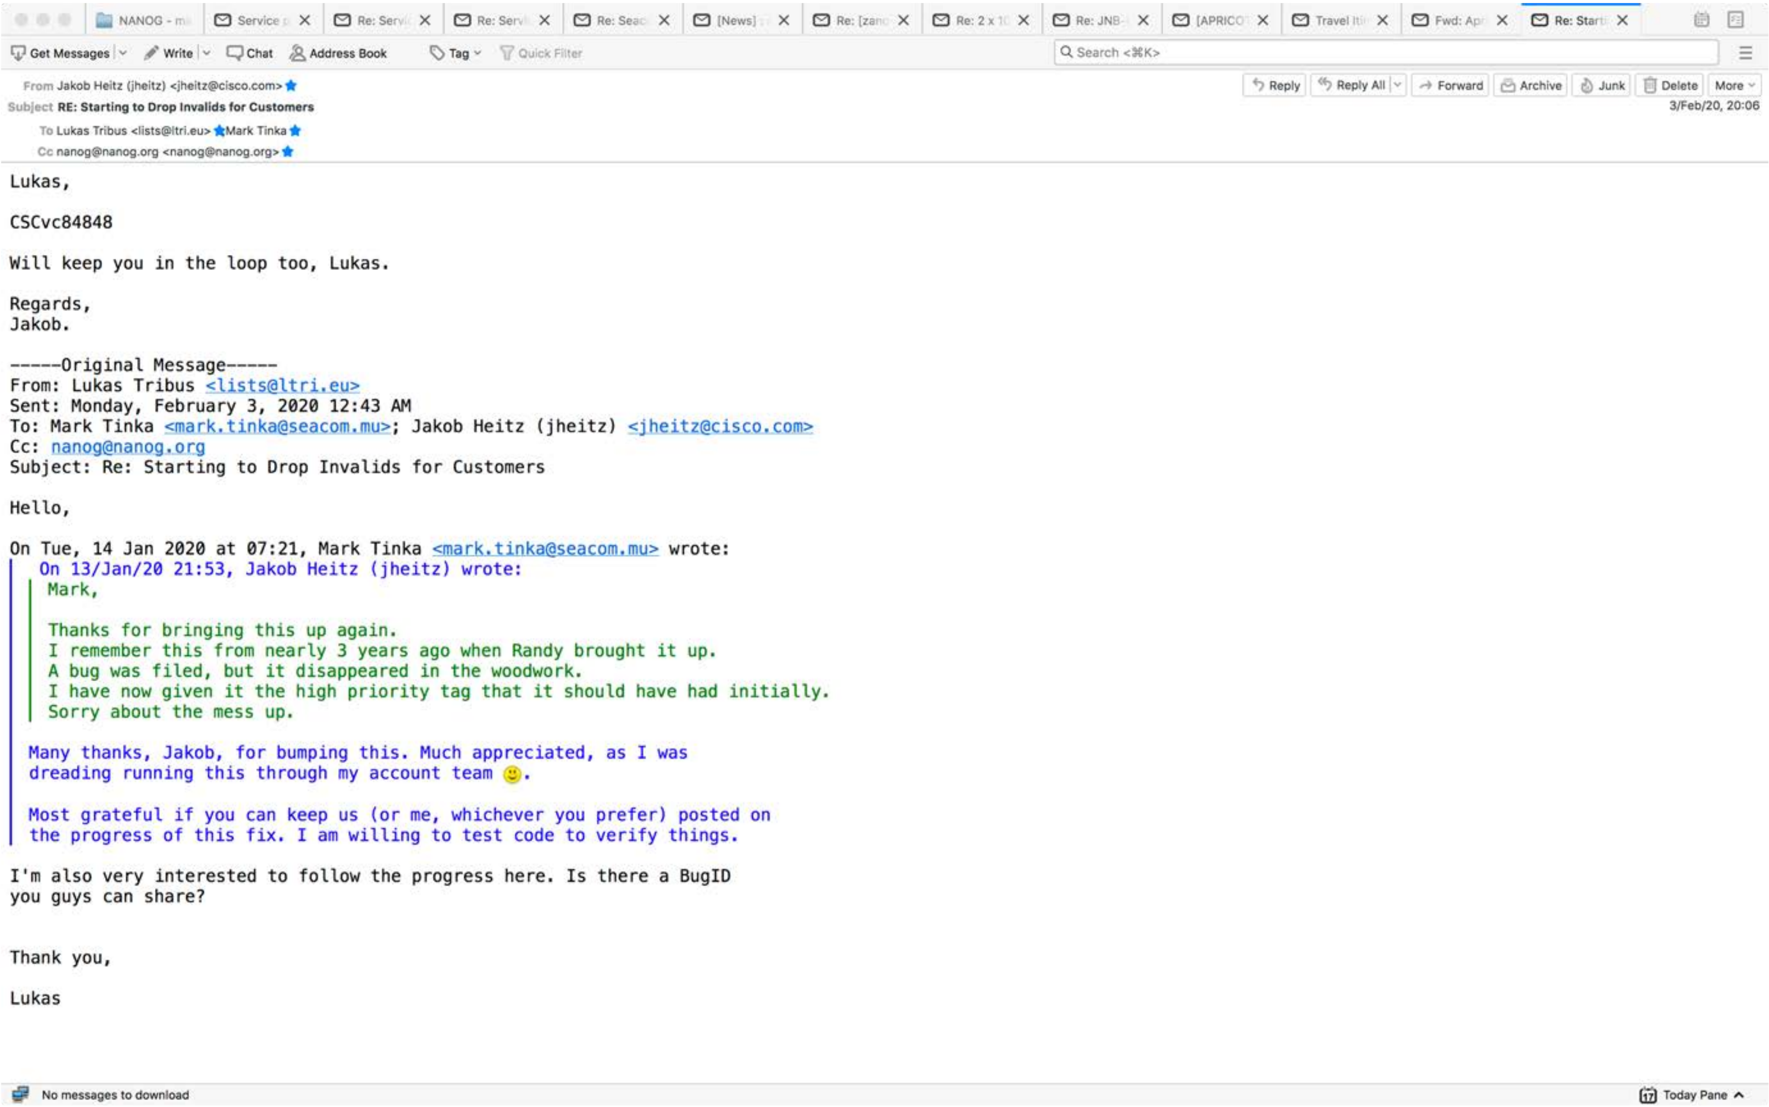

From Jakob Heitz (jheitz) <jheitz@cisco.com>
Subject **RE: Starting to Drop Invalids for Customers**
To Lukas Tribus <lists@ltri.eu> Mark Tinka
Cc nanog@nanog.org <nanog@nanog.org>

3/Feb/20, 20:06

Lukas,

CSCvc84848

Will keep you in the loop too, Lukas.

Regards,
Jakob.

-----Original Message-----
From: Lukas Tribus <lists@ltri.eu>
Sent: Monday, February 3, 2020 12:43 AM
To: Mark Tinka <mark.tinka@seacom.mu>; Jakob Heitz (jheitz) <jheitz@cisco.com>
Cc: nanog@nanog.org
Subject: Re: Starting to Drop Invalids for Customers

Hello,

On Tue, 14 Jan 2020 at 07:21, Mark Tinka <mark.tinka@seacom.mu> wrote:

> On 13/Jan/20 21:53, Jakob Heitz (jheitz) wrote:
>
>> Mark,
>>
>> Thanks for bringing this up again.
>> I remember this from nearly 3 years ago when Randy brought it up.
>> A bug was filed, but it disappeared in the woodwork.
>> I have now given it the high priority tag that it should have had initially.
>> Sorry about the mess up.
>
> Many thanks, Jakob, for bumping this. Much appreciated, as I was
> dreading running this through my account team 🙂.
>
> Most grateful if you can keep us (or me, whichever you prefer) posted on
> the progress of this fix. I am willing to test code to verify things.

I'm also very interested to follow the progress here. Is there a BugID
you guys can share?

Thank you,

Lukas

# Things To Look Out For – Junos Opaque Community Bug

- For use-cases where RPKI state is transmitted in BGP communities.
- Junos will not send the BGP community values correctly.
- The issue is fixed in the following releases:
    - 17.4R3.
    - 18.2R3.
    - 18.4R2.

# Things To Look Out For

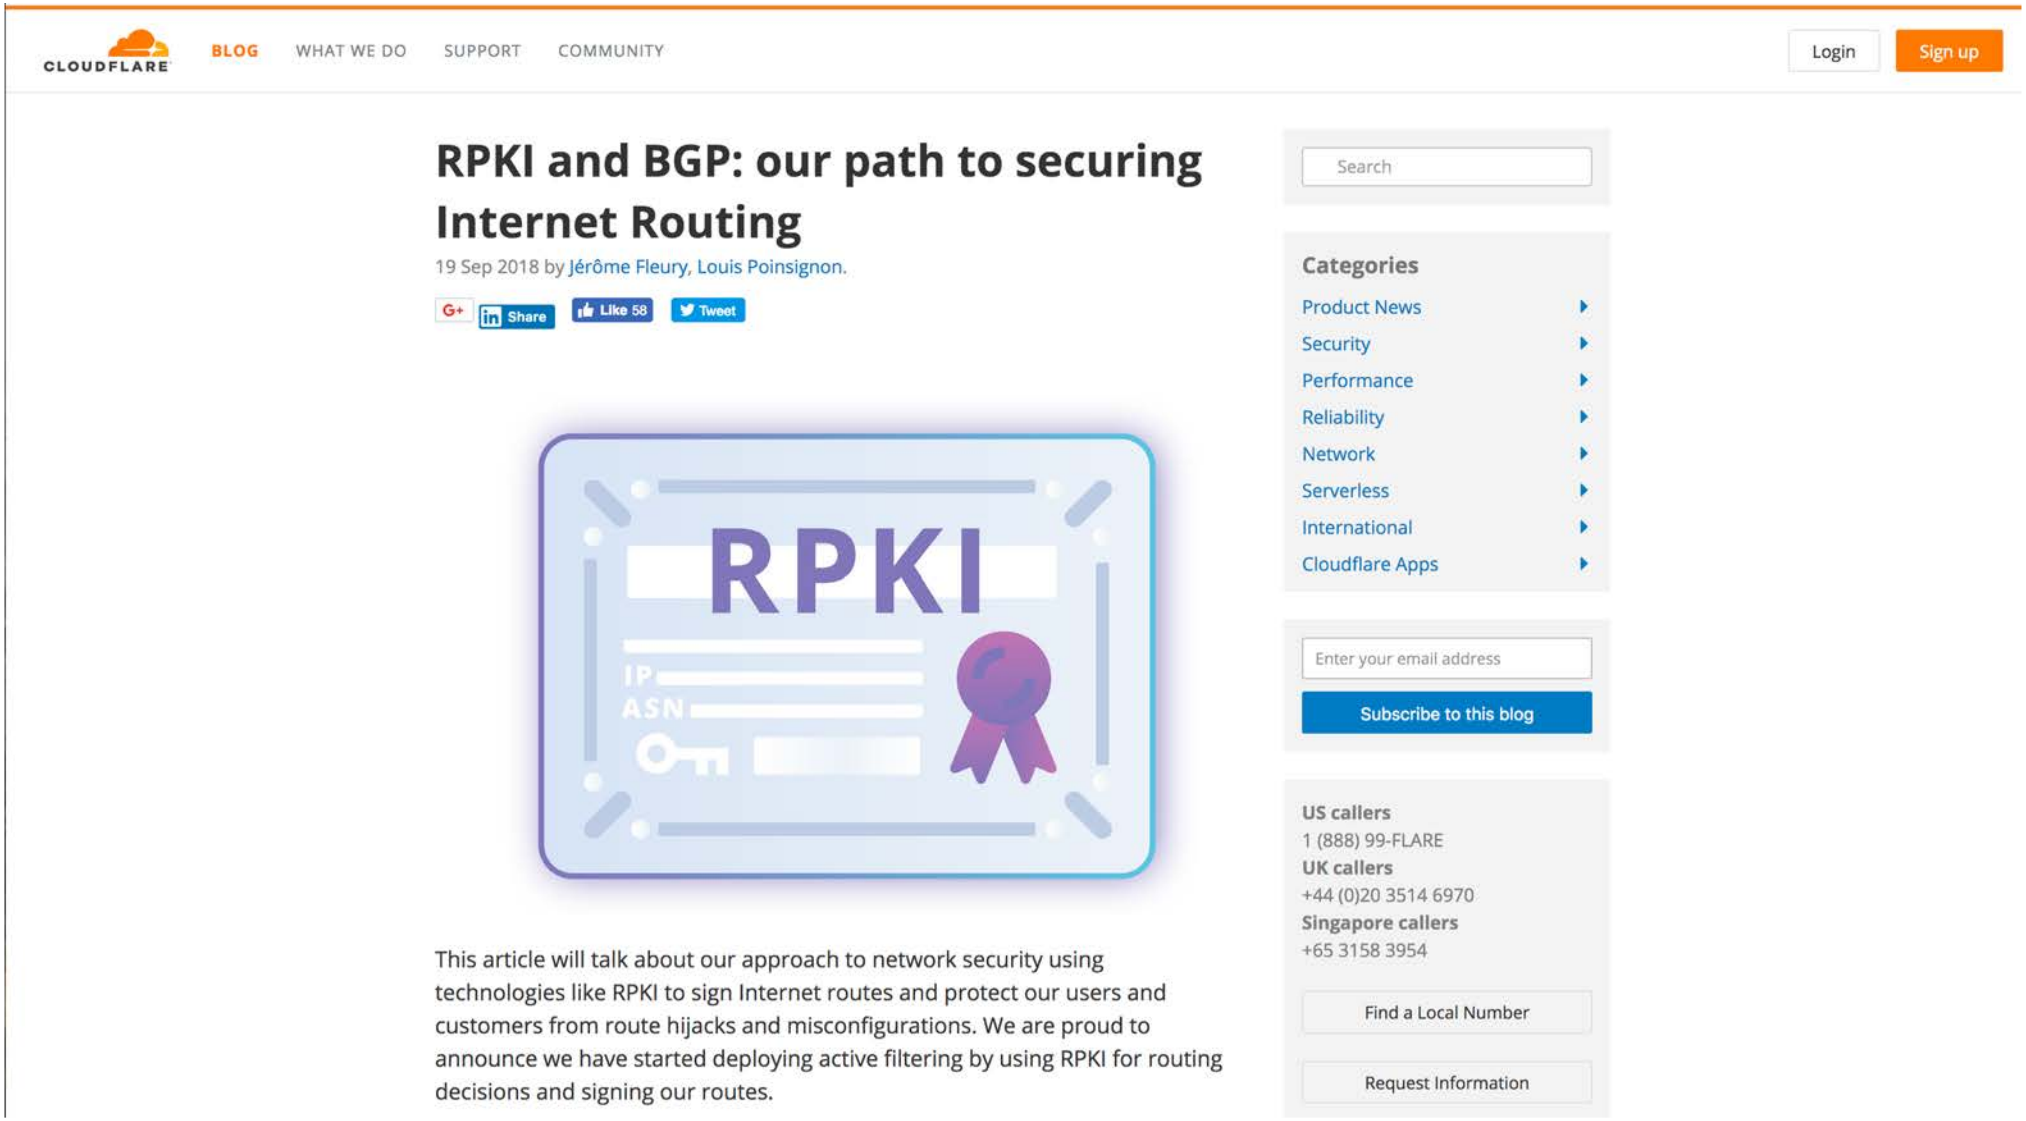

## RPKI and BGP: our path to securing Internet Routing

19 Sep 2018 by Jérôme Fleury, Louis Poinsignon.

This article will talk about our approach to network security using technologies like RPKI to sign Internet routes and protect our users and customers from route hijacks and misconfigurations. We are proud to announce we have started deploying active filtering by using RPKI for routing decisions and signing our routes.

**US callers**
1 (888) 99-FLARE
**UK callers**
+44 (0)20 3514 6970
**Singapore callers**
+65 3158 3954

# RPKI for Africa

LIQUID
TELECOM

SEACOM

workonline
communications

**Will Drop Invalids 1st April, 2019**

Thank You

Q&A

**mark.tinka@seacom.mu**