SIDROPS Agenda For IETF-110 (version 2)

Session: March 10th 2021, 12:00 - 14:00 (UTC)

1. Agenda bashing and Chair's slides - [5 minutes]

2. Ben Maddison - [10 minutes] RPKI MaxLength

draft-ietf-sidrops-rpkimaxlen

- Requests WGLC

- Comments/Questions

    - Alexander Asimov: Have different view on this draft. Why these changes: what kind of hijack/attack vector trying to address? or, "mistakes" ?
        - Ben: The specific attack vector is where you have a ROA issued, covering len up to eg /24 but in the ordinary course of ops all that is announced is some aggregate shorter than a 24. The attack vector we discussed is the trivial spoof AS Path, announce upstream to peers/transits and become bestpath for the prefix, because only longest-match in routing table. Longer prefix is auth for origination.
        - Asimov: so... speaking about intentional hijack not a mistake. Ben: yes, mostly. Making a mistake of this kind would be hard to do in most router implementations, but if made as a mistake would protect against as well. Asimov: helpful only in case of hijack. But, imagine two islands: one good, ISPs sign ROA, origin-validation happens, bad island, no ROA no verification, and hijackers there. In your scenario if such intentional hijack in bad island,

and there is somewhere inbetween, more specific, which will begin to “spread” from the island to other places, what you’re suggesting is the good island is protected against this activity, but all space between these islands have a high chance ot use the prefix. In response, the victim won’t be able to advertise its own more specific, because the ROA won’t be permitted. Partial adoption of ROV, these kind of “strengthened” ROA not using MaxLen will result in worse problems: space hijacked, and nothing in response at the moment.

- Ben if I understand correctly, by issuing only strict ROA, limit ability to respond to sub-prefix hijack, lose ability to propagate any de-agg except in your immediate neighbourhood. Asimov: if applying, egress router ROV, won’t even be able to pass from own ISP. can remove, but anyway Ben: probably a valid criticism. Unlikely scenario, because of where ROV is happening in the wild today. There aren’r readily distinguishable islands is/ isn’t happening. Increasingly the case large transits in topological centre are doing ROV. I think we are better off optimising for the case, the protection of existence of validating parties, rather than optimise for the case propagate defensive de-aggregation into non_ROV network. makes sense? Asimov: make security statement in doc, “you may have some problem because of …” so Operators aware, issuing this kind of ROA, they limit their defence of identity.
    - Ben: send text

- Rudiger Volk: (from older reading of the text) Do we tell, do we want to tell, that issuing any ROA for which we are not actually doing valid

announcement, or intend to do it in the short time, can be viewed as an “invitation” to an attacker who can take the AS mentioned in the ROA, to take over. The fact that the MaxLength is a way of doing the most dangerous type of attack vector, seems not really to be the most important thing.
        - Ben: Exactly what I was driving at on slide 5. Previous wording implied MaxLength in and of itself was the problem. we’ve clarified the text to say whenever you have a ROA which is not usually announced, longer than usual announce, the use of the maxlength attribute is a shortcut to create. Fundamental problem is covering a prefix not usually announced. MaxLength is a symptom
    - Sriram Kotikalapudi: even when a prefix is hijacked, not neccessarily maliciously, somewhere in the network away from you the hijack will be dropped, ROV being used, you are losing visibility when it is a normal prefix hijacked non maliciously, Want this kept in mind
3. Tim Bruijnzeels - [15 minutes] Prefer RRDP

draft-ietf-sidrops-prefer-rrdp

- Comments/Questions
    - Job Snijders: Not entirely sure what publication phase 2, end 2021 means to the community. RP s/w supports previous versions up to 1 year, looking at our logs, industry needs 18mo for 90% of RPs to upgrade. I think phase 2, whatever it means, should factor in a 1.5y lead time.
        - Tim: yes, appreciate. I think publication for phase 2 is requirement for current impl to support rrdp, doesn’t introduce rsync, so nothing should break at that point. even if you publish doc with normative MUST support rrdp, can still publish one that says “you really ought to use it (to RPs)” -still don’t say to RPS

"don't have to rsync any more".
        - Randy: Its a format change as well as a content change, future state has to cope in existing things
            - Tim: happy to take out implementation status report, phases beyond (really deprecate rsync) launch separate effort, focus on phase 1 & 2 Randy: want to go all the way, think Rsync deprecation is the goal, can stay in the single doc. lay out full plan, occasional reports on where we are, what has to be done etc. Tim: but the moment you update separate RFCs, is when you make a separate document. Currently we talk about a plan, but at some time you want to execute. At that time, you want a separate document to update whatever things we need to update (Randy: reasonable. eg in january implementation report or some -bis to an RFC. Want to separate the URI problem because thats messy)
- Rudiger Volk: Some "whataboutism" As we are writing down, a plan for successive protocols, would it make sense, help, but not only increase the complexity of the stuff, but target protocol specify preference, make the newer IP protocol "mandatory" v4/v6
    - Tim: make IPv6 mandatory for the repositories? I think thats a fine thing to do, but probably orthoganal to this
- Chris Morrow: put the doc split, implementation reports, please take to the ML, with the points. The implementation report thing is useful and the 'require running code before publish draft' thing would be helpful. (Randy: I did) Anything else? Tim:

nope. good. follow up next steps on list.

4. Job Snijders - [10 minutes] RPKI Signed Checklists draft-spaghetti-sidrops-rpki-rsc

- Comments/Questions
  - Tim Bruijnzeels: RSC in and by itself, has a place for the more simple use-case compared to RTA. RTA was not only about having multiple parties signing, the other thing it contained is the inclusion of the cryptographic material needed for validation inside the CMS which is not allowed inside the RPKI signed object RFC. There are at least two other use-cases in the world, could have a situation where multuple parties need to attest to shared resources, and, giving things to somebody else, without needing full RPKI validation in other regards. That being said, we need to think where we're heading with RTA specification. It may be we look at the RFC specification, wrap that, have multiple objects in one thing, in one go. if useful to have all the CA CRL shipped with it to make quick validation "right now" can look at in an enclosing structure. I can see the use-case for the simple case, think about RTA if we keep current spec of leverage this in some way.
    - Job: want to emphasise, RSC, RTA are not mutually exclusive. different semantics. multiple signers on a single SHA256 hash, the other is single attest to multiple SHA256 hashes so... fundamentally different, but from a "getting things done" perspective, industry has been waiting for RPKI community to delivery some technology which fits the workflows we both agree exist. Let the RSC effort proceed so simple case is covered, let RTA continue to be explored. I had trouble implementing RTA.
    - Job: Validation, intermediate objects present has not been finished. The APNIC demo of RTA with multiple signers under multiple trust

anchors, not confident its a robust strategy because it demands multiple instances of OpenSSL, not perfect robust fit for RTA but may be my abilities and understanding. Hope is, in RSC these complications don't exist and WG arrives at that point.

5. Job Snijders - [20 minutes] RPKI Validation Re-reconsidered

draft-spaghetti-sidrops-rpki-validation-update

- Comments/Questions
  - Chris: request codepoint early allocation, we can do that shortly.
    - Job: the IANA early codepoint allocation is for RSC not for this. This is a request for WG adoption.
      - Chris: mailed? Job: I will do so shortly.

6. Alexander Azimov - [10 minutes] ASPA drafts

draft-ietf-sidrops-aspa-profile

draft-ietf-sidrops-aspa-verification

- Comments/Questions
  - Randy Bush: reason you cannot find 8210bis, you're looking for a YMBK doc and its an IETF doc. its WGLC, waiting for you (Alex) to stop changing this draft.
    - Alex: thanks, noted

7. Sriram Kotikalapudi - [20 minutes] On the accuracy of algorithms for ASPA based route leak detection (joint work with Jakob Heitz)

- Comments/Questions
  - Alexander Asimov: (unintelligible)
    - Sriram: what I looked at, what we were trying to do, seems like the basic principles are similar/same, as described. Efficiency aside,

variants which can be implemented based on the outline

- Ben Maddison: To echo that, I think we have three algorithmic, in-code representations of the same set of ideas, more or less equivalent. The version in draft 07. There's a version I wrote based on the email Jacob sent to the ML a couple of weeks back, there's your version in slides. they all arive to same conclusions "under the hood" -two overriding considerations in std; readability and understandability, secondly, need much more clear line of reasoning, to go from what we understand as a route leak type trying to detect, what the algorithm looks like. setting out the logical steps for people not reading the iterations of the draft, to understand what runs on their routers.
  - Sriram: to add to what Ben said, in the draft, the g function defined, in slides, is quite helpful. just a matter in section 4 of the draft, instead of ASi ASg, invalid, give it a name like "g" P, NP, NA notation, if you use each hop check is valid/invalid/unknown you're saying the hop is invalid but the path is valid, that happens, confuses the reader, to make it dis-ambiguous, in addition to using a function like g, its a hop-check function, use P/NP notation dont mix it up with overall path validity when re-writing, making text more useful
    - Asimov: the function has already its name, its called ??? in separate section. Do you think the naming of these functions its not a problem

8. Christopher Morrow: running code, before we push drafts with implementation changes. Seems totally reasonable to me. At the very least, opportunity to see how changes actually work. Nobody disagrees with that in the threads. Chairs need to do work on a little Charter update, discuss with the IESG. Speak now if have comment

- Warren: Not sure it needs a charter update. can get charter update, sounds like a lot of “Faff” update the SIDROPS wiki, not going to progress docs until you have <x/> Happy to do the AD faffy stuff
- Rudiger: asking for running code, or running code and positive interop test?
- Keyur: having had an experience of IPR, 2 impl, strongly required deployed or no. can start with one. interop doc in some cases would be an added plus, otherwise start with one impl. don’t think we need interop in many cases.
    - Chris: don’t know particularly required but helpful to review/declare “hey, 3 of us … and nothing bad happened” or counter-case
- Job: to re-iterate, some of what was already said, think we should copy a lot of WG culture from IDR. depending on draft at hand, merely an implt report, normative term handling by s/w and potentially post, interop. not all IDRs require interop but the ones which do we should be testing. Proprietary can also do reports, interop. copy from how IDR handles it. would positively benefit this group
- Doug Montgomery: Interop in algo like ASPA, where systems interface, what we’re interested in is different impls behave consistently which is broader definition of interop
    - Chris: in my head, running code/interop are kind-of the same thing to me. convo earlier today, all kind of the same thing to me
    - Warren: more I think, more best to have “strong suggestion” for code requirements, not in charter can chat more
- Randy: its an ops group not a proto group: interoperability of *WHAT*
    - Job: for example, interop of RSC, RTA, or other newly defined objects in the RPKI. we can confirm if multiple impls can handle
        - Randy but they don’t belong here. This is SIDR-OPS not SIDR. if you want protocol

which needs interoperation it needs to be done somwhere else, or the group needs to move out of OPS and have a massive re-charter