IETF 111 SIDROPS 28 July 2021 19:00 - 21:00 UTC

Chat log: https://jabber.ietf.org/jabber/logs/sidrops/2021-07-28.html

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

* Status Reviewed of WG drafts.
    * Job Snijders: Signed-TAL update is still being discussed, Can we hold for implementation reports?
        * Chris: Yes, will check mail (Keyur thought it was with IESG)
    * Warren Kumari: Hadn't seen it in IESG, thinks it may still be in WGLC, Doc hasn't been discussed
        * Chris: Can resolve this after Implementation report

2) Tom Harrison - [10 minutes] Signed TAL [draft-ietf-sidrops-signed-tal](https://datatracker.ietf.org/meeting/111/materials/slides-111-sidrops-signed-tal-00)
    * Maarten Aertsen: What can an RP do with a successor key?
        * Tom: It can carry out two validation runs and compare results. Build trust, by showing the user what happens IF we switch to the sucessor key. But it's important for validation it continues with the original key. Can't just "take the second key" and use it for the results it produces for offline/RTR use
    * Job Snijders: Architecture of some RPs (such as openbsd rpki-client) does not permit the parsing process to write/delete the system's TAL files: updating TALs is a privileged operation. Job can probably get revokation going, but building an arc to the next (successor key) probably demands a privileged, out of band access. Some RPs have a restricted access pattern which does not permit directly updating TAL.
    * Job Snijders: Is it clear which CA signs the TAK? can it be any subordinate CA, or is it explicit in the draft only signed by the TA?
        * Tom: an EE certificate produced by the TA, yes. Job: Clarify please
    * Randy: Phases depend on software upgrade in the field. Do you have measurements of upgrade rates and percentages we can base this on?
        * Tom: no, we don't have, but we have data from things like RRDP rates we can look at
    * Ben Maddison: If the revokation state is only held in TAK1, then the PK material in Key 1 needs to be privileged, it can "unrevoke" itself after a key roll
        * Tom: The idea is when an RP gets to phase 4 it deletes. Ben: Presupposes a "bad actor" doesn't hold the old key in their pocket to use to revoke later Tom: will note
        * Chris: can this go to the mailing list to be kept track of please?
    * Geoff: If a key revokes itself, there is no further work needed: they can walk away and the attack is done. There is an operational reason to e.g. regularly roll. The process seems awfully vulnerable to points during the roll, the roll itself becomes the

problem. There's no way of phasing in, it's a complete A/B switch but you allow a key to kill itself. What happens if it revokes without a successor? Bootstrap the new TAL because continuity lost. Replace one vuln, key compromise with second vuln: Fragility during roll. Security Review needs more detailed discussion to risks of self-revokation, fragility. A lot left to guess. Don't think this is ready yet to describe a robust rolling process.

    * Tom: To the current key revoked without successor: If a client gets to the point of no successor, the process aborts and tries again later. It's up to the TA to fix things
    * Geoff: But, to be in this situation is either a stuffup (security) or a stuffup (operation mistake) -the external basis of trust isn't there. It would not be rational to proceed.
* Maarten Aertsen: Ben stated the point I was trying to say earlier, succinctly. Good to either include in security considerations not neccessary to signal trust in successor key (mechanism kept simple)-explicit its not part of the protocol, or good to have some way to signal trust in the successor, help mitigate the scenario compromises a TA and there is no way to signal successor published after that time is controlled by the TA operator.
* Chris: comments should go to the list, so they can be worked on.

3) Job Snijders - [10 minutes] RPKI Signed Checklist [draft-ietf-sidrops-rpki-rsc](https://datatracker.ietf.org/meeting/111/materials/slides-111-sidrops-rsc-draft-ietf-sidrops-rpki-rsc-july-2021-update-01)

* George Michaelson: I like the normative requirement implementation table on the IETF side. Good use of tech.
* Randy Bush: If RSC are detached from RPKI datam then where do I look for revocation?
* Chris: in what CRL is it published?
    * Job: the RSC file contains an EE cert, the CRLdp in the EE cert shows where the specific CRL for the RSC is located. There must be a CRL: if it can't be found the signed object is meaningless.
* Tom Harrison: Server side implementation is planned for Q1/2 2022
* Ben Maddison: Procedural question: in early allocation there was an OID for content-type but not the module. doesn't impact on-the-wire or interop but does make code unmergeable. "TBD" has to be stripped out in code compiling. Is there early allocation for module subtree?
    * Russ Housely: Yes, can be given. Most implementors don't really care, I know your tooling does. A request for that arc will be approved if asked.
    * Job Snijders: 4 things we want from IANA: content-type, update to filename extensions, sub-registry and media type. content-type is the key one to keep moving forward, the other 3 can wait until RFC publication. IETF doesn't like recycling security sensitive codepoints, filenames .. future application may want the one we propose.
* Ben Maddison: Signer implementations: AfriNIC is keen to code. No timelines. So hopefully there will be two implementations in 2022

    * Warren Kumari: This is worth presenting in SEC area related WG so people aren't surprised when this shows up (present before WGLC)
    * Ties de Kock: Any key sign an RSC or a key usage flag to show its only being used below the arc.
        * Job Snijders: not sure key usage can be used in this context. e.g. Go lib code makes key usage flagging inoperable. Not sure how the ecosystem works, can verify. I know some of the RPs do verify key usage. Not sure how we would restrict this without making RSC difficult to use. I think the resource-subordination can naturally produce objects like RSC but if you have a specific suggestion we can look at it
        * Ben Maddison: X509 profile for Resource Certs is pretty specific what is allowed. The only divergence is the missing SIA. The content of the blobs you put a sig over, and in RSC are "not in scope" -even concievable a TA at the root might want a signature over something e.g. mailing verifiable copies of the CPS. Absent a specific attack vector don't see the utility in restricting it.
        * Ties de Kock: Take to list. Job: good to talk, now (in draft timescale?) eg reject TA (selfsign) -but would be good to understand what we want to prevent under certain circumstances, exclude either by normative terms or X509 extensions/settings

(Job: This is the applicable section on Key Usage https://datatracker.ietf.org/doc/html/rfc6487#section-4.8.4)

4) Haibo Wang - [10 minutes] Region Verification [draft-shen-sidrops-region-verification](https://datatracker.ietf.org/meeting/111/materials/slides-111-sidrops-verification-of-routes-using-region-authorization-00)
    * Randy Bush: This model assumes in region "no routers are compromised" - if this is true with neighbors then no body on the internet is compromised.

    * Ben Maddison: There are a lot of problems with this approach. The two most problemmatic are: The first one is the one Randy just called out, no operator in the DFZ is trusted enough not to be capable of making a configuration error, this is unworkable since few operators will establish a regional confederation and work on the assumption of trust inside the region. The second problem is that it makes the internet topology as a whole extremely fragile: the rules as enumerated would cause valid routes, transiting to be rejected as well. If peering broke, the "outside" routes would be rejected. ASPA objects in system to confirm the AS is a legit. transit provider, then there seems to be a very corner-case attack vector. I don't think this approach is worth pursuing.
        * Haibo: we don't show the solution here. Traditional route policy. Some providers have deployed local policy like this. Our solution is to make this easier. re-use policy, must set one-by-one, specify how to filter the routes. I don't think its 'not useful' -because some people have deployed like this. Can we take this to the list? I did not catch everything.
    * Rudiger Volk: problem scenario was "route propagation between various ASN not secure" then create fairly complex scheme of data,

checks to secure things. not questioning, what checks are done and what checks would be reasonable to do, at the various places, is missing. I have trouble to believe you really assume within one ISP 2 dozen AS are fully interconnected without any checks and thats reasonable policy. Randy was succinct, but please explain exactly what checks and guarantees you are assuming or dropping within your starting scenario

- * Haibo: did not get all of this, please take it to the list.

* Job Snijders: it appears to me your users are not like those in Europe, you are needing to form "aggregates" of some kind. I think you might need to go IDR to discuss. Encourage you to specify in ASN.1 how you envision the objects, to facilitate the discussion.

- * Haibo: yes, we're not like Europe (our ISPS)

4) Chris Morrow - [15 minutes] [Running code requirement](https://datatracker.ietf.org/meeting/111/materials/slides-111-sidrops-running-code-sidrops-00)

* Randy Bush: There was interoperable running code, for about half the things you mention, but not for newer inventions. This is not a routing WG. Not saying there doesn't need to be one, This is supposedly an OPS working group and shouldn't be developing protocol by the rules of the game

- * Chris: agree we did have some implementation interop/testing. As a group we fell down on that, But I think the fact we were able to build something a while back because of the implementation testing made a lot of that "go" -without that we're realising things which aren't working as optimally as we expect. Interop testing would be useful. I agree we're not a routing or protocol dev group. (Randy: we have an OPS and a Routing AD) lets hear them.

* Rudiger Volk: Hate to see emphasis on ensuring interop in the IETF protocol development process being lost, so I applaud trying to raise it where applicable. For the incidents that seemed to have kicked up the discussion here, I would also raise the question: Were some of those incidents observations which were unlikely to show up in an interop because they depended to some degree on having large-scale operation, including the fact that race conditions are really something hard to "force" finding the problems which may result in race coditions are also very hard. I agree with Randy, we did have small-scale interop testing in the past. We have actually done one RFC/standards track that quite certainly has no way of being interop tested: Just demands something on a single system (RFC8893)

* Warren Kumari: Just because a WG is an ops WG, it doesnt mean there is never a need for code, or implementation. As an example: protocol-type development in DNSOPS, not massive sets of changes, but part of ops/maintenance involves code & protocol updates. If SIDROPS wants a path of "implementation before publication" it should be carefully worded, including escape clauses so that for example BCP where there is no implementation, so the "must be done by 3" thing isn't trap

* John Scudder: I'm not upset about where work gets done, if it's work which needs to be done. I do agree with Rudiger: don't over-state the benefit. Its not a substitute for a comprehensive

test suite. In IDR the main experience and benefit is in knowing you've produced a spec somebody else can read/implement "at all" -can point to multiple spec where its resulted in updates post WGLC. implementors say "don't understand what to do here" and authors say "oh my.. we need to change" In IDR we don't have "rules" for implementation requirements. Its a tradition rather than a written rule. This has good and bad points, but they give the chairs lattitude.

     * Joel Jaggli: If we're making changes because we think things dont work well or could work better thats close to our charter. Formal or informal, it doesn't seem to preclude this kind of activity.

     * Chris: "This shit's broke y'all... let's fix it.." IDR has a process that seems to work. Job shows some methods in his RSC presentation which I think look useful. There is work for chairs: A process needs to be gone through. Warren's point hard/soft process requirements, It's good to have an exception process. A process in datatracker would be great too.

     * Warren Kumari: If you do have something like "waiting for implementations" tag on drafts, it requires followup work to check if the correct tag is set, and removed when multiple implementations exist. so, it involves work.

  * Job Snijders: lets take ASPA as an example that could be put into this process (continues with slides)

    * Rudiger Volk: Regarding the principles we were discussing before: One simple suggestion for preparing operational testbeds and so on, we should ask the RIRs, one or more, to simply support straight delegation of just AS resources of a member to a member CA and allow de-coupling the member CAs for IP addresses and ASNs, so address space operation is stable. The AS part is not yet populated and will be the cruical playing field for ASPA.

        * Job: Good news: You can use RSC with AS only, we have structure which can work here to test.

    * Rudiger Volk: Appreciate the work figuring out a testbed for ASPA, but to the general question, we have been discussing before interoperation testing but what we see here is far beyond that: "testbedding" -one should clearly understand the differences, mark if working on one or the other. They may overlap but they are very different things. The question for the general interoperability requirements: I wonder if we should actually try get consensus about what the primary goal of the exercise is. In times past, "official" IETF statement was 'interop is for quality' -making sure a few different people interpret the text in a conforming manner. We don't want to let that go, but back in the old times there were other interpretations that were not really the same: "This is a way to ensure free reference implementations available", which was not true. Yes, that happened, but was not the primary goal.

        * Job: running code is not open-source code. The licence has nothing to do with this. Interop is not about open source.

    * Rudiger Volk: When setting up a formal requirement, I think it makes sense if the primary goal of the exercise is spelt out.

        * Job: helps uncover if something is implementable at all.

* Keyur Patel: Chris and I have discussed this, the example here, ASPA is pretty narrowed down, but one difference between IDR and SIDROPS is the maturity we accept docs. In IDR the draft goes through considerable amount of discussion where it matures enough to allow implementation. In SIDROPS we accept first, then progress to last call. This has direct implications on the cost to design software. Churning becomes a problem, you will see the issues. Think about if this general approach is accepted as implementation needed, view requirements of WG adoption moving forward, have more discussions a priori before issuing a WG call.

* Ben Maddison: The application of requirement to implement document to document will vary so much, chairs should make a call on a doc-by-doc basis. Require source code to change implementations, should result in work which is done, tested in a lab or field so we're not in an endless cycle of revising published RFCs. To Job's point: There are two categories of things to worry about from interop perspective: "does my signer create a DER stream your RP can read" is one category. The ASPA testbed is in a different category, has some of that. More importantly: "does ASPA do what we want to the DFZ" and it goes beyond the testbed/interop outcome. I Agree with Job mostly. do ourselves a disservice if we try and do this on completely sandboxed separate interconnection over tunnels. I think we want the complexity of the real DFZ to find the problems. Run separate CAs and RPs but validate real routes.
    * Job: I agree would want to mimic the shape of the DFZ and we can take real data, or alike.

* Geoff Huston: This is nothing to do with ASPA in the first instance. its an IETF wide problem. the IETF talking about running code & rough consensus was the viability of the spec, implementable as they stood. The references to OSI producing non-interoperable implementations was directly the topic. "the spec is good enough for implementors to do it and work" but this is different: works at scale, works at all.. thats a mis-characterisation of what a standards do. Its a mischaracterisation of what a standards body should do. The test about viability at scale and market acceptance has gone beyond garages, its about enormous cost at scale. is this the IETF's business? We are not here to tell people how to spend their money: thats the job of the market not the IETFs jobs. We won't solve it in SIDROPS, we have AD here and this is for an IETF plenary. What does rough consensus and running code mean for the whole IETF? Please can we punt this back to the IESG, IETF-wide discussion might be better/focussed what is achievable and what we mean.
    * Job: I did not mean market acceptance. I meant we as WG can see a demonstration that it works at all. Before we publish a doc saying "possible for implementation from spec independently to work" there has to be some software which can work, we need proof-of-concepts. I'm not asking for commercial quality.
* Alex Asimov: Thanks for slides. Gaining operational experience is good. Will try to put some implementations status in open-source so more ISPs can validate ASPA behaviour before deploying, seeking vendor support.

* Job: to conclude: Dear chairs, we would like some codepoints for ASPA.
    * Chris: Please take to list and we can action.

Notes from the chat, requested to be included in the minutes
    * Tim Bruynzeels: We have not determined timeline yet, but want to implement ASPA support at some point in Krill. When we do we can provide a testbed - I would rather not commit to 24/7 support for it but for testing it should be fine.
    * Sriram Kotikalapudi: Just FYI. We (NIST) have an implementation of ASPA verification in Quagga as part of our BGP-SRx. The implementation includes a correction to the downstream detection algorithm that was presented/discussed at SIDROPS IETF 110. @Nathalie You may kindly include this.
    * Geoff Huston: its not the IETF's task to evaluate viability or not - thats more or less up to the market to determine - the role of the standards body is more prosaic than that - its put publish specs that are complete and useful to guide implementaitons

5) Chris Morrow - [10 minutes] 6486-bis [draft-ietf-sidrops-6486bis](https://datatracker.ietf.org/meeting/111/materials/slides-111-sidrops-6486-bis-infoz-00)
    * Nothing substantive noted here, we were out of time. Chris: I will ask on the list about 6486bis.