Webex recording:
https://ietf.webex.com/ietf/lsr.php?AT=pb&SP=EC&rID=6714127&rKey=0ad5dd297cf4b13e

Jabber Log: http://www.ietf.org/jabber/logs/sidr/2012-04-30.html

Meeting delayed due to technical difficulties – getting remote participation tools working together (teleconference call, webex session, webex integrated conference call, conference phone, feedback, etc.)

Chris Morrow and Sandy Murphy serving as chairs

Intro slides (note well, housekeeping tasks, agenda) viewed

(notes slide 2)
Basic Incremental Deployment Strategy

- Origin Validation
  - Deploy code
  - Start doing some metrics
  - Start turning it on
- Bgpsec
  - Deploy code
  - Start doing some metrics
  - Start checking – logging
  - Don't need to worry about local originated route

(notes slide 3)
Operational Issues in Incremental Deployment

- Bgpsec incremental deployment
  - Randy: if I am on this router, I want origin validation checked now on this router – don't want to find out from neighbors that things are bad
  - Chris says want to check signatures to make sure the origin is validated and authentic
  - Chris:If you don't have RRs, you have full mesh, simplest case
  - Randy: if a RR and its clients have no externals, then those RRs don't need to speak bgpsec
  - John: to repeat: in general RRs have to be bgpsec aware because bgpsec attrbs are non-transitive. But only those RRs in a control path betw two ebgp borders need to have bgpsec enabled.
  - Chris: in the case of bgpsec capability negotiated – don't send attrbs. So if RR is not bgpsec aware will not receive attrbs and they don't get to other side of AS. So RRs need to be upgraded when they are in the control path between ebgp routers.

  - Randy: there are routers in your topology that will not need to be bgpsec aware for your AS to be bgpsec capable. Some routers will be bgpsec capable but may not validate and may sign anyway.
  - For locally originated routes – you only apply bgpsec attrbs (signatures) when you are sending the update over an ebgp session

(notes slide 4)
[Sidebar: Issues in Protocol Deployment]

- Stack of to be discussed
  - Confeds
  - Add-path
  - AS aliases (replace-as in juniper speak)
  - Forward policy (eg route servers) –brief discussion pointed to need for firm definition (is route server example only case? Etc.)
  - Brian's AH-HA to be explained later

(notes slide 5)
Operational Issues in Incremental Deployment (continued)

- In bgpsec we only do signatures and validations of signatures at the borders
  - What about RRs doing best path selection
  - A:Edge router passes on results of its signature validation in some AS specific policy knob (community, localpref, etc)
  - Q: is there any requirement that RR consider signatures in best path selection
  - A: (Ruediger) RR should do its best path selection based on this local policy knob
  - A: (Jason) the local policy knob is not bgpsec attribute so any router can make a decision based on that. Operational quirk – have to set up RR to make decisions on locally originated routes as if they were signed.
  - A: Randy: routes is marked somehow – and the mark is available to your policy.
  - A: John: you should mark the route as if it had been signature checked and marked (you trust yourself to have locally originated)
  - Q: Sandy: you have to be careful to make sure that there's no strange way to make the AS_PATH null so an external path doesn't look local (as an attack)
  - A: Brian: if you redistribute from bgp to igp and back to bgp, origin validation should catch that.

(notes slide 6)

- Another inflection point:
  - Going from not doing bgpsec to doing bgpsec
  - Could change preference of best path
    - Could make big changes in traffic movement in your AS, traffic swings
  - Randy: at layer 9, we are going to a more secure network, that's a feature!
  - Chris: hopefully you do some modeling first

        - In case of 701, as-path wins, not peer
    - Brian: much bigger than that. When change of policy of two bgpsec capable peers, to change to prefer signed routes, you may have difficulties making that an incremental change
    - Randy: you are changing your policy, you know there will be consequences, these are consequences you want to have happen.
    - Wes: need tools Randy: have tools now to model policy change
    - Brian: but there may not be a way to tell who goes first – chicken and egg
    - Chris: internally, you may choose to prefer signed, but process will be:
        - Deploy code
        - First mark routes with potential bgpsec result
        - Look and see what is right and what is wrong
        - Still believe this will be multi-year process (5? 7?)
        - Not a flag day.
    - Wes: in documentation of simple case or somewhere
        - Update code
        - Watch what goes by and see what changes – dump somewhere else
        - Right policy based on that, before you pull trigger on policy change
        - Don't drop invalid initially, just downpref it.
        - Maybe this goes in a secops drop

(notes slide 7)

- Randy: ops do these things every day, we have tools (not great), we do deployments, upgrade links, etc.
- Brian: impact is minimized by making policy change late
- Chris: every network will make a different choice. And a deployment time frame of years. But not a knife edge deployment – takes a couple years of
- Chris/John: do we need an RFC, or a presentation, John: but confeds may need protocol change, Chris/John: some common deployment+network models
- Warren: isn't this just another knob in policy? But then there are presentations about how you tune meds, so maybe.
- Chris: don't think this is RFC, maybe vendor docs.
- Brian: think RFC is right documentation model just simple case with RRs. (Haven't proved it is scalable if you can't show you've considered that.) Brian volunteers to write up simple case with RRs.
- Ruediger: remember 4byte AS attribute – there was no deployment draft. Chris: to Ruediger – are you saying that went well and is proof we don't need a draft, or that is a warning that we DO need one.
- Randy: can't write docs for all the possible things it would be good to know. Where do you draw a line. Need a decision of what MUST go in. Heather: docs can't hurt – what is the down side?

(notes slide 8)

After Break – Next Topics?: Forward Policy Signing

- John: votes to discuss things that change protocol spec. Eg confeds

- Brian begins: need some way for signer sending update to express what it was authorizing the recipient to do. eg.: authorizing route server to remove itself from as-path
  - If the recipient might do something other than just add its as to the as-path, there should be somewhere that this authorization should be expressed.
- John: interesting, but old bgpsec goal was to protect what is currently in bgp without changing the spec'd model. You are talking about changing the bgp spec'd behavior. Might be a non-trivial change.
- Rob: you might want to think about whether this needs to be in the bgpsec protocol – maybe a new rpki object?
- Warren: you have one suggestion – but there are lots of others. Carrying policy around
- Randy: not carrying policy – carrying a policy of what YOU can do – a very deep hole
- John: is there good semantics in what we have now? Is there room for later added semantics? Those should be the decision points
- Brian: but bgpsec is authorizing the path validation that comes to you
- Randy: but we forward on destination, we accept as-path and propagate
- Brian: but we are validating the path, right?
- Randy/John: yes, path we got, you are talking about path forward.

(notes slide 9)

Route Server

- Y does not know that RS is really a route server
- X and Y should authorize and check the role of the route server
- Randy: Y can validate that someone doing pcount=0 is a route server to it
- How does Z know that Y was correct in what it did

  X -----> RS ----> Y ---> Z
- Randy: X saying that RS is an authorized route server is forward policy restriction
- And is a big hole
- Randy: we are trying to protect the protocol from being violated, not the business relationship
- John: kinda analgous to route leaks – you want Z to second guess if what Y did was in accordance with what X wanted it to do. Given that we have the option for Y to check the use of pcount and right now Z must trust that. Suppose you stuck something in RPKI – who would be making the authorization
- Brian: all members of the route server would authorize route server
- Jason: ebgp multi-hop and gre tunnels should be considered as well.

(notes slide 10)

Discussion of Confeds

- Is there some way to indicate that the as-path is only for in the boundary of a confed?
- Is there a way to indicate semantics not already covered in the bgpsec attrb?
- Issue with confeds:
  - In current real world – you have some sub-as's, adding themselves to the path as as-confed sequences, get removed at boundaries.

- But no way to add these as-confed type things in the as-path. Needed for loop detection.
- Need some flag in sig attrb? Mark each – throw away all of the confed-marked sig attrb.
- Warren: don't you know all the as's used inside the boundary? A: you don't know the enumeration
- Randy: spec says it can be arbitrary topologies, but deployed topologies are more simple – a core AS with some stub ASs attached to core.
- Q:Sandy – external links come in where?
- Randy: the simple case – only possible loops in this case would be solved by adding normal sig attrbs and then dropping at the boundaries.
- Rob: think John says: mark confeds inside confed so there is an easy way to strip the confed sigs. Randy says: inside confed, add sigs as usual and strip at border.
- John: current bgp implementation strips all confed sequences and does not consider the as number
- John: this is part of discussion of can we get rid of AS-PATH entirely
- Randy: external AS neighbor of stub AS in the confed topolgy knows only the core AS number. Externally only core AS number is seen.
- Brian: we could do this using pcount=0, using some local AS

<<< picture at board was large circle labeled "coreAS" with smaller circles adjoining at the boundary being the stub ASs>>>

(notes slide 11)

- John: by defn: you don't use core AS as the public AS. Randy: oh, no, not true. John: ooo, did not know that could happen. Wes: that's how we do it also.
- Chris: external to ebgp router in subAS, cbgp from subAS to core, External forward signs to ebgp customer facign router, cbgps forward signs to core, cbgp by core to subAS, then subAS ebgp router wants to strip the first subAS number and does not know that subAS number. (Hard to do this in text).
- John: One way to tell what to strip: start at origin. When you find a forward sign to your public AS, strip everything past that.
- (External neighbor thinks it is connecting to the public AS. Because some people use the public AS as the core AS, can't use the core AS as the marker of the confed topology border.)
- One thing that this hack does break: right now in the protocol, my AS can appear in the protocol more than once – the allow loops (allow ownAS in (in cisco), loops (in juniper))
- Heather: lies are bad, so dropping lies are perfectly fine.
- Warren: what about path poisoning? A: path poisoning will not work in bgpsec. End story.
- (Work at board emphasizing that customer external AS configs itself as peering with the public AS. Wes wants to confirm with his network.)

(notes slide 12)

More Agenda Bashing

- Some early departures want to influence the agenda ordering.
  - Confeds (cont'd)
  - Add-path
  - AS aliasing (replace-AS, local-AS, etc.)
  - Revisit the incremental intra-AS deployment – anything more there?
  - AS-PATH – in the protocol
- Brian took five min to discuss a out-of-the box way to do different crypto (maybe shared secret based) – on the list. (See Brian's AH-HA in notes slide 3)

(notes slide 13)

Discussion of Confeds (contd)

- Confeds solution:
  - Use pcount=0 everywhere and strip all pcount=0 that have your AS in them.
    - Problem – you may be usign a sub-as to peer with a RS
- Have an attribute that has an attrib that carried the confed nature of paths. John proposes that that attrb is the as-path, but that attrb is not currently in the bgpsec protocol
  - Sriram: the as's are in the bgpsec attrb A: but we are speaking about the real bgp as-path attrb
  - Warren: was not thinking of using as-path attrb, a brand new attrb.
- John: Augment current bgpsec sig attrb with a flag field that says "confed".
- Rob: likes having a separate attrb, so people who want this pay the pain, not the whole world. John: but bgp attrb space is limited, would rather not use another bit for this. Small point, but.
- Sirram: what distinguishes confeds in current bgp? John: type code. We don't have that in bgpsec protocol. Maybe we can add that. (randy: why have extra complexity – people will think up ways to use it – can't predict future)
- John's summation
  - 1 rewind (start from origin to first signing mentioning public as – relies on external session going to public AS)
  - 2 Flag in sig attr
  - 3 Child-of-AS-path (some other attribute that carries "confed" characteristic
- Warren: can we just use a marker in the sig attrb?
  - Trust model: ski length 0 , sig length 0, --- marker – don't need to sign over it.
  - Rob: repurpose whole bgpsec sig attrb to carry this John: that's what we bit into when we got rid of as-path
- John adds
  - 4 explicit "entered confed" marker
- Me: is this a problem with someone entering the marker who should not do so?
  - A: entry point into confed knows that it is in a confed and knows it is the entry point
  - And can be required to check incoming data for entered confed marker – and throw error
- John: likes marker, only thing better about flag is it fits existing structure

- Warren: doesn't like flag because it gives 7 extra bits Matt: additional bits would allow some coloring
- Wes: additional bits in attrb – does that get signed? A: yes.
- Warren: that means all routers have to sign on entry and exit from confeds – incremental deployment issue
- Randy: learning anything more? John: probably not. Go off, think about it, discuss on email?

(notes slide 14)

Add-Paths

- Issue: only sign your best path
  - But that breaks add-path
  - Sriram had text and will send it to the sidr list.

(notes slide15)

AS aliasing Replace AS Local AS

- Wes: channeling Shane – what about AS merging?
- And path poisoning
- John: in my implementation
  - There's something that lets me be any one of a set of ASs when I talk to an external peer
  - Strip private ASs – ok if they are all on the right side of the path
    - Warren: Don't you have to rewind until you find the first that is not private and not in your member list
    - Would that not work for confeds?
  - Replace arbitrary AS in the middle of the path – that's never going to work, give it up

(notes slide 16)

Replace as – AS migration

- AS 1 is connected to 2 and 3. but 3 thinks 1 is 5

  2 ----- 1 ----- 3
  - One way is use pcount=0.
  - This exposes the fact that 1 exists to 3
  - Randy: what does 3 see from 1? Jason: It might see 5 1 2 or 5 2, depends on the config. From 3 to 2, 2 sees 3 1 or 3 5 1.
  - Going to have to be knob.
- Documented? (keeps coming up, so should be documented) (Matt: doesn't think it belongs in core spec, Geoff agrees on jabber, Randy agrees to go into bgpsec ops)
- Remove private hack: what about forward signing problem if external peer signs to private
  - Wes: can you transit a private AS?
  - Randy: keys used to sign for router with private AS decend from a trust anchor that anyone wanting to validate the signatures must have that same private trust anchor (operationally quite messy)
  - Geoff on jabber: that would mean that if a private AS the rest of the Internet must have that private trust anchor
  - Randy: can't transit a private AS in bgpsec

    - Wes: but what about external person peering with the private AS?
- John: but we have 4 byte AS – why do we need private ASs
- Ruediger: private AS as stub is OK – remove private AS and originate at upstream where private AS is stripped.
- Wes: is it ok to make it impossible to transit a private AS – because you run into all these problems.
- Brian: use public IP address for the router? (but how to establish relationship between router IP addr and AS number)

(notes slide 17)

There Might Be 50 Ways to Leave Your Lover but 150 Ways to Mess Yourself Up in BGP

- Ruediger: don't add complexity for those 5000 AS numbers. Do not propagate private ASs into the public internet.
- Ruediger: some customer is sending 200 /32s with private ASs.
- Geoff: I see about 10 of them.
- Rob: but these are supposed to be used in trust boundaries and stripped outside – transiting is different and hard to define trust model.
- Wes: if you have a reason you need to transit the AS that is using a private AS, use a public AS (maybe one that is shared for this purpose).
- Chris: need to address this if there is a use case where it woulod be very painful to use any other mechanism.
- Brian: people are using BGP but their engineering was done by consultants who have been gone – and want to add bgpsec.
- So until we find the painful can't-do-this-any-other-way case, we will move on.
- Jason: when people are using private AS, do they sign? Always? Never? Sometimes?
- Brad NTT: some customer who share a public AS number. [RFC2270 sort of ASN] – (701 uses one of these things as a shared cutomer AS number) – 701 could sign the prefix as [RFC2270 ASN] and 701. The [RFC2270 ASN] could be stripped (Wes: for Sprint this does get stripped.) If the IP addr belongs to the customer, the [RFC2270 ASN] is not stripped. Allows one customer to sign an announcement originating another [RFC2270 AS] customer's space (because each customer would allow [RFC 2270 AS] to originate their IP prefix)
- Discussion of DNS analogy that I missed about primaries and secondaries and muddying the waters.
- When there's a routing issue because a /24 is flapping, you want it to be obvious your customer's fault not yours and that is why you do not strip their use of RFC2270 AS
- Brian: if the [RFC2270 AS] is doing the signing on a router that is not under the customer control
- Can't tell the AS just go get your own public AS, because they are not multi-homed and do not qualify under RIR policy.
    - Just fix the RIR policy

(notes slide 18)

Path Poisoning

- Just can not work

- Deal!

(notes slide 19)

Intra-as deployment

- Are we done?
- Chris: we have a path forward
  - brian volunteered to write a simple RR case
  - And then things will go forward
- Randy: do we understand the technology and are comfortable with it.
  - Simple example seems to work.
    - Brian but what about case where AS wants to sign the origin
      Randy: when injection occurs from another protocol into my bgp, there is no AS-PATH. Inside my AS, I know it is mine, so I know what the AS number is, so I can add that first signature.

(notes slide 20)

Discussion of Signing over IBGP Sessions

- John: what is use case for signing over ibgp session? Don't trust router?
- Brian: you are leasing routers, leasing bandwidth, people can put themselves in MITM position. Some people may want to use bgpsec on ibgp sessions rather than transport on their sessions.
- Warren: how many people are using transport security on their ibgp sessions? A: most of the big guys will do that!!!
- Rob: hard to say there may be a need for this someday so lets add this complexity now.
- Geoff: hard to address unless we understand threat model.
- Brian: was raised on list as comment on threat model document and "shouted down by chairs".
- Ruediger: do you want the thousands of ibgp routes I carry internally in my network to have this protection?
- Ruediger: Brian, what you are interested in is protecting the other attrbs in the update, not the empty as-path.
- John: since I don't understand the use case: (a) transport security will not solve problem? or (b) I don't want to use transport security?
- Brian: I don't want to use transport security. Don't want to have to deploy transport - two security solutions at the same time.
- Sandy: but transport security protects against other vulnerabilities that bgpsec of ibgp updates would leave unprotected.
- Brian: understand that, presenting this as a incremental way to get bgpsec deployed.
- Rob: thinks security ADs would object to a solution that would leave such an obvious vulnerability uncovered.
- Ruediger: does your proposal apply only to originated routes? Are other attrbs that are not used externally (if you don't your own network) also to be signed. Interested when mitm can inject routes. Think if you can't trust your own internal network, someone can inject something bad

(notes slide 21)

Prepending Issues

- <anonymous> reports from message from Jeff Haas. In bgp today, upstreams can add prepends of my AS number (with or without my permission) and those prepends stay when upstream propagtes the update.
  - John: some implementations do loop detection on ibgp as well as ebgp, which means prepended looks like a loop.
- <anonymous> prepend a private AS the needed number of times and translate to my public AS number
- (Brian reports a solution to his own concern about untrustworthy internal networks: use private AS internally to sign with, then strip on exit)

(notes slide 22)

Protecting Other Attributes

- Brian: Ruediger brought up a point – is there any way to protect any attribute than AS_PATH
  - Randy: unsure of trust model of other attrb
  - Randy: two possibilities presently
    - Business relationship (route leaks)
    - Signature expiration time
- Geoff on jabber: Brian needs to go back to SteveKent's analysis of what needs to be signed.

(notes slide 23)

Freshness

- SteveBellovin: still concerned that we don't have a way in a short time frame to roll a router signing key.
  - Randy: concern is security compromise? Yes
  - Beaconing does not seem to work
- One suggestion (Steve Kent's ??) flood a "refresh your rpki" in bgp. Randy: inband? The one that wants to do this is compromised router SteveB/John: or the AS/NOC
  - Beaconing has problem that there is incentive to beacon too often, this does not
- A: does this accomplish this B: does it work
- Rob: how do you authenticate this? A: what you want to do is rate limit this. Q: how to sign A: AS number-ish
- <Geoff Huston> a broadcast message to everyone else that says "generate refresh queries to the repositories" is a great DOS attack
- Warren: someone owns my router and issues enough of these to make rate limiter throw alarm and router's key gets rolled. (if noc gets compromised, game over) Wes: or employee leaves. There are methods to manage this.
- Brian: if you could flood info and identify alternate paths, could signal with route leak material to ensure people who need the new info get it. Chris: not talking about BGP data, talking about RPKI data like CRLs. Brian: using BGP to signal

(notes slide 24)

Freshness (Cont'd)

- Back to freshness and stolen keys…
- Chris: problem is that we have all data in same place and so same fetch process applies (poll from repository)

  - Can we have a separate communication mechanism for some objects?
  - Rob: there is a protocol for retrieving CRLs in real time, but we are trying to avoid it!
  - Chris: Danny asks why twitter can communicate in near real time
- Brian: if 13 routers are compromised… warren: if someone owns 13 routers they could just turn them off.
- steveB has left but he said he would be happier if there was a way to get the key change distributed faster. Rob: me too, I just don't know how to do it cheaply.
- Chris: a way for repository to say "ok, ok, ok, new info"
- Rob: poll or flood, take your pick
- Chris: just trying to say that there are different classes of data and some have different freshness concerns. Rob: but just talking about different protocols. Would like to talk to this with Tim and SteveB in the room. Chris: besides different protocols can we concern ourselves with different requirements for different objects? Need DATA. Tim might be interested in some long term measurements. Rob: Randy says we don't have a long term baseline for how this would work in the future – we don't know what the expected behavior would be. Should be.
- Reports that rsync has some features that don't work well in validators. Rob: like maybe an atomic publication of a bunch of data.
- Chris: freshness requirement – some time to live – cert expiration time gets reached.
- Rob: DNSSEC has signature lifetimes AND ttls. Should not be combined. Signatures are end-end, ttls are cache behavior. Different animals.

(notes slide 25)

Repository Discussion

- Chris: if I am gathering from the world but you won't give me this one object so I keep using the object I have but then the sig expires and OOPS. A: yep.
- Rob: for any publication mechanism, you have to worry about how long it takes for data to get through the system.
- SamWeiler: any real need to have a push of data? (Actually a notify – go get new info). IETF has an old protocol 2244 called ACAP and ask for notifies for anything that would meet that search criteria. <<something about it does not scale>> Brian: but you can't block a notify Rob: oh, yes you can. SamWeiler: acap has authentication per user.

(notes slide 26)

Repository Discussion (Cont'd)

- Chris: rqmts about repository – how many objects, how many retrievals, how often, etc.
  - Rob: George Michaelson in RIR group looking at this long ago – 10K-100K of objects in the rpki. So nowadays, maybe 100K-1M. Originally retrievals once per day, now I poll once per hour, could poll several times a day (4-6hrs) but 1/day is a bit slow.
  - If you are doing a make before break, everyone needs to get the make before you do the break, so any slower than 1/day meant you deserved to lose.

    - ROAs are most frequent data change. Manifests change when any data changes. Crls change ..<missed>
- Chris: we have 1M objects now, in future bounded by (multiple of the) size of routing system. Currently 405K+8K routes (v4+v6). Geoff reports his extreme outside view is upper bound 25% growth per year.
- Churn is low, number of objects is large.
- Geoff: v6 AS advertises 2 routes, v4 AS advertises 4 routes.
- Warren: deaggreages my routes is an advantage, deaggregates my ROAs is not. Geoff: but you want to have every alternative covered. Would pre-provision all 5 neighbors even though currently announcing to just 2.

(note slide 27)

Discussion of Repository Size

- Chris: are there other things that I might want to certify that I might want to use the rpki system for?
    - Geoff: may want own system for each
    - Rob: there are router keys, crls, manifests, etc.
    - Chris: size of routing system times multiple of 2? For right now. That seems to give lots of room and a safe upper bound.
    - Geoff: 600K in five years Chris: so you will need 1.1-1.2K objects in five years. If the change rate is 20% then you need to extranct 250K objects each interval.
    - Geoff: you are needing a new object, not a return to previously announced state.
    - Geoff: 10-12 new ASs a day, 30 new routes.
    - Rob: use new keys each time, so a return to previous state still looks the same.
    - Chris: need to understand the size and the speed of change. Need to collect the data or extract from routing history
    - Geoff: troll through RV and RIS and presume it is all certified.
    - Chris: that is what Tim and (RIPE guy from Canada) are interested in this (Canada guy = andre took) and are trying to prepare for ietf vancouver.
    - Rob: if you are running this with rsync that is binary only and uses inodes, you might want to be looking at database.
    - Chris: rsync works (Rob: has a lot of the nice properties) but having more protocols is OK. (people troubled by having a large number). Need 99% uptime. Rob: no SLA for DNS servers. Brian: there is for root servers. Chris: for RIRs, this might be like root servers, but for me, it is more like my local dns server

(note slide 28)

Repository Discussion (Cont'd)

- Chris: stuff you retrieve needs to be high rate, stuff you serve may not need to be updated at lower rate. Rob: probably more necessary to be available for people to retrieve than for me to be able to update quickly
- Wes: this may change if we are dropping invalids. I recall calls to noc that says “you must update your dns server immediately”. I don’t think that it is valid to assume those calls won’t come in. Rob: distinction that your customers will want

a immediate reaction and the chance that the call will come in when you have the system down for maintenance.

- Brian: you want your servers available to your neighbors than to the people 12 zones away. Rob/Chris: don't buy that – global need for this data. Everyone wants a copy of the worldwide system.
- Warren: how much data is 1M objects? Rob: 1K per object, max. Warren: round trip to himalyas is 10 sec. Rob: worried about pounding primary servers into the ground.
- Sriram: if we are doing BrianW/Roque's proposal from last IETF for replay protection, you could increase churn significantly.
- Brian: you could do it that frequently, but we weren't proposing that you do it that quickly. If you are worried about some particular event , you could increase rate.
- RussH: during algorithm rollover, you need another factor of 2 in that – multiplier is 2 at the peak and builds up and falls off from there.
- Chris: also think we need to capture the entire size of the repository
- Rob: CRLS are potentially unbounded because people might have too long expiration times, which means nothing ever expires from CRL. (story of very very very large CRLS because of long expiration times)