Large-Scale Measurement of Broadband Performance (LMAP)
MINUTES

Meeting: IETF88, 2 Nov 2013 - 8 Nov 2013, Vancouver BC

Location: Hyatt Regency Vancouver, Georgia A, Wednesday 6 Nov 2013, 13:00-15:00 (Afternoon Session 1)

Chairs: Dan Romascanu <dromasca@avaya.com>

Jason Weil <jason.weil@twcable.com>

Minutes: Steve Miller <sdmiller@verisign.com>
version: 0.1

==========
AGENDA:

1. Note Well, Note Takers, Jabber Scribes, Agenda Bashing - Chairs (5 min)

2. WG Status - Chairs (5 min)

3. Liaisons Status - Chairs and Benoit (5 min)

4. LMAP Use Cases - Marc (20 min)

5. LMAP Framework - Phil (20 min)

6. IPPM work on Metrics Registry - Bill (20 min)

7. LMAP Information Model - Juergen (20 min)

8. HTTP-based Protocol Proposal - Marcelo (5 min)

9. Open Mic and Next Steps (remaining time)

Reading List

https://datatracker.ietf.org/doc/draft-ietf-lmap-use-cases/
https://datatracker.ietf.org/doc/draft-ietf-lmap-framework/
https://datatracker.ietf.org/doc/draft-claise-ippm-perf-metric-registry/
https://datatracker.ietf.org/doc/draft-mornulo-ippm-

registry/
https://datatracker.ietf.org/doc/draft-mornulo-ippm-registry-columns/
https://datatracker.ietf.org/doc/draft-burbridge-lmap-information-model/
https://datatracker.ietf.org/doc/draft-bagnulo-lmap-http/

==========
MEETING REPORT

1. Note Well, Note Takers, Jabber Scribes, Agenda Bashing - Chairs

Steve Miller and Sam Aldrin volunteered as minutes scribes.

Bill Cerveny volunteered as jabber scribe.

- Dan: This is not privacy BOF :D. If you are here for BOF, it is happening in other room.

2. WG Status - Chairs

Almost on track with respect to milestones.

Two I-Ds published: LMAP Use Cases and LMAP Framework

Performance metric registry work happening in IPPM, as that seems more appropriate.

Next dates are ambitious and we need to stay focused in order to hit those dates.

Information Model I-D is next in the charter and we have a draft; will discuss in this meeting.

3. Liaisons Status - Chairs and Benoit Claise

There is interest outside the IETF in this work, as other standards organizations are targeting similar work; we may be useful to them and/or vice versa.

One formal communication from IEEE: 802.16 - Performance Management Task Force; no real ongoing engagement.

Chairs and Benoit are in contact with other organizations, who are asking for info on how to provide inputs and the like (IETF process, how to

access documents, how to access email list, the fact that we do most of our work via email, how we prefer to receive information and contributions).

Benoit said that ARCEP (a French regulator) was in contact.

Benoit also said that BEREC (an organization similar to the FCC for EU telecommunication policy) was in touch -- not a liaison but reaching out to people with knowledge.

If we don't take input into account, our RFCs may not do what they want.

We got feedback from the FCC initially, which was good.

We've also received feedback from EU regulators.

- Goal is to have right people discussing in the room
  - To continue the discussion about regulators from Berlin mtg, as regulators are going to use it, need to consider them and their input
  - If you know any regulators, please contact chairs or AD, to ensure that the docs generated here will not go unused.

Dan: we need to continue to exchange information. Please feel free to talk about LMAP in our companies, academic events, etc. If you need help with presentations, contact Dan and he'll try to help.

When we're getting near Last Call, we should also circulate documents and try to elicit input from these other groups at that point, that would be a good time.

We did receive multiple formal liaisons from (Broadband Forum), to which we have responded.

Dave Sinicrope (BBF Liaison Manager) from Ericcson: Liaison response was reviewed by the BBF (WT304 is their equivalent to LMAP) fairly quickly. He emphasized to them that if they want to use out work, they should get involved here, and that if they want our feedback, they should liaise early and often. They meet in Rome in two weeks; it's likely afterwards that we'll see docs and that they'll ask us for feedback on

those. Dave will be at that meeting but not in those discussions specifically.

Jason Livingood asked about WebEx access. Dan and Jason responded: it's one-way only, so it would only help for document sharing, please use the audio feed plus Jabber.

4. LMAP Use Cases - Marc Linsner presenting

- After Berlin WG, we resolved and merged to one draft.
- This contain two use cases

Consolidated three use-case drafts; now have ISP and regulatory use cases.

End-user use case moved to appendix.

ISP use case hasn't changed substantially since IETF 87.

- ISP use case is to finding, isolate and fixing the problems in the network
- Second one is regulator use case.
- Regardless of technical details/differences regulators need to use it

Plan to turn doc around by end of November.

There have been comments on the list about regulator, other use cases; will review and integrate into next draft as appropriate.

- Add clarification, what does QOE mean
- Frode submitted extensive text for regulator use case
- Understanding is, this use case is from EU regulators.

Dan : Is it diff use case or diff aspects of regulator use case, which we have it now

Marc : it is not much of a difference, in my opinion.

Outstanding issues: We don't really have a requirements section. Do we need one, or is the constraint discussion in this doc adequate?

Phil Eardley: on requirements, having a formal set of requirements doesn't work all that well in the IETF, as it's not really top-down unlike other standards

bodies. Would prefer where we have discussion and reach agreement on what we're trying to do, identify constraints and assumptions, and at a high level identify the protocol model. That seems more fruitful than a more formal requirements list.

Benoit: Basically agreed with Phil. Also: security/privacy needs more meat.

Marcelo Bagnulo: We still need to do more with v4 vs v6, multi-interface, and security/privacy. Concerned about having wars over formal requirements and getting bogged down there. Feels the group is largely converging and that we don't need a formal list of requirements.

Al Morton: multipath TCP group had numbered requirements; he observed that there were only a few of them. Earlier LMAP draft had nearly 100 requirements, which was far too many; maybe it'd be useful to have a generally-stated, small number of requirements.

Phil: consensus is no formal list of requirements.
Dan: concurs.

Jason: make sure that the framework matches the use-case doc.

Dan (as contributor): Need to start being careful. With the focus on privacy in IETF 88, we may be in a situation where we'll be under scrutiny for such aspects. Since talking about monitoring in a provider environment, with end users, the issue of privacy can't be avoided. Discussing this is fine; we need to see where to put it. Could write use case for privacy; put in security-considerations section; put in privacy-considerations section.

Marc: good feedback, thanks.

Outstanding issues: (contd)
-End user use case – to have a formal section in FW draft or present state is good enough

Dan : With limitation within charter, end user can start subset of tests and see subset of results, not sure if it is separate use case. This is

Mike : It appears as a way out but

Marc : There is a goal to have small number of use cases to get work done. That is what drove this

Mike : we might be

Steve : What is missing is, multi people telling what to do. i.e. number of MA's involved etc. How this is different to what is in the play?

Benoit: don't design with the end-user case in mind, e.g., do my MAs discover each other and synchronize somehow, etc.? As an end user, might install the software to be a MA but not looking at it so much as a bigger requestor of tests.

Marcelo: One guy with MA wanting to do stuff. We are doing a control protocol and a report protocol but those aren't needed for the end-user case, the end user just tells the MA to go do something and emit the results. End user just needs to be able to activate tests, seems like an implementation detail. From LMAP perspective, seems like it's covered, nothing's really different.

Dan : Just another i/f in the hands of End user

Nalini Elkins: the end-user experience is very important and needs to be somewhere. We have been working on performance metrics built into the protocol that don't require an agent. A perspective she wanted to talk about to see if it's of interest to the group.

Marcelo: "no agent" means "no active agent but passive agent".

Dan: may be passive-agent stuff already with metrics built in (e.g., RTP). An issue but somewhat orthogonal.

Alan Clark: If LMAP only intends for residential BB, then EU need not be considered. But if it is over all BB, then EU should be considered. You can't just do with simple protocol. Who is allowed to see what is very important.

-if business services, end users are highly likely to want to do tests. Service providers and enterprises want to cooperate; with end-user/multi-party testing, we need to worry about who gets to see

which results.

Al Morton: User requested testing is in scope as part of SP use case.

Dan: yes, but as part of ISP use case.

Marcelo: important but not in scope.

Dan: tends to agree, scope is large-scale access. Some businesses may benefit from this kind of measurement but want they require may not be large-scale and may not be in scope of working group.

Klaus (Nieminen?), FICORA (a Finnish regulator): not sure what service provider means here.
Dan: SP is domain manager not just ISP.

Klaus: ISP may be offering the service.
Marc or Phil: no one will go to jail for using this stuff for other purposes.

Marc or Phil: cellular/mobile devices are in scope, so are there other things we need to put in here to cover that?
Dan: send mail to list
Andrea (possibly Soppera): Mobile devices are capped on usage, and we need to think about the impact of testing on those caps.
Marc or Phil: FCC is launching mobile testing currently. Test suite is 10Kbytes, low volume, ad-hoc.

5. LMAP Framework - Phil

Slide 2:

Milestone to submit this in December: needs reviews! He's had three people send comments but needs more. Will try to react to those comments and rev the doc and send for last call to get more comments.

- Merged the three individual drafts into one WG doc
- Some aspects were redundant. Many were unchanged.
- Introduced high level protocol model. Without any specifics.

Slide 3 (Framework):

Per slide

Slide 4 (Bootstrapping):

Not defining protocol but defining process, may be access-style/device/technology-specific.

Need to get MA to register and to get its ID.

Slide 5 (Bootstrapping):

Per slide

Slide 6 (Control Protocol):

Controller to tell MA how and when to run what tests and how to report.

- Control protocol is to instruct how to start/stop etc, to perform measurement task.

Slide 7 (Control Protocol):

No signaling about progress of measurement task: no feedback, it just does what it's told.

Dan: controller is free to run the send clock.

Slide 8 (Control Protocol):

MA does what it's told, no negotiation.

Can be told to report to more than one collector.

Dan: when talking about reporting to more than one controller, do we report the same test to all, or some subset to one and some to another? If the latter, control protocol needs to allow telling who gets the results for each test.

Phil : reports could be sent for multiple controllers

(name missing): thinks the lack of a way for the MA to say no is a problem.

Marcelo: intent is that the capability request is done first, so the controller will know what each MA will do and thus won't ask a MA to do a thing it can't do.

Slide 9 (Control Protocol):

Controller can ask the MA to stop tests (e.g., if there's a network issue, maybe ask the agents to stop in order to not make it worse).

Details are open, how complicated is that? Will discuss on list.

Alan Clark: can an agent run more than one test at

a time? Some attempt to clarify that.

-Phil - No.

Marcelo — That is not what he is asking.

Andrea- You could issue latency and performance test at the same time.

Alan: dispersion test might require exclusive access, while other tests can play well with others. If agent is out of service or moves to a different ISP or whatever, how to deal with telling it to stop doing something forever? There are a bunch of exception conditions to consider.

Marcelo: some stuff Alan mentioned is covered in the rest of the slides or in the draft. You can suppress tests and you can get failure reports back.

Marc: you're not thinking there is a session open between MA and controller all the time, right? Suppression would only occur once the MA contacts the controller again. Not clear, details haven't been worked out yet.

(name missing): What about auto-suppression? Some things that need to get box into auto-suppression state.

Marcelo: had some discussions about suppression. When a MA gets a new schedule, it should toss its old schedule of tests; the new schedule needs to contain anything from the old that still needed to keep going. So to suppress permanently, just don't include something in the newly-downloaded schedule. Temporary-suppress is a specific message ("shut up!").

Alan Clark: one more exception -- what if the MA is power-cycled? Does it start where it left off? Does it try to catch up? Or does it not do anything except for getting the schedule again from the controller?

Marcelo: If it reboots, it does nothing until it talks to the controller, there's no history.

Alan: We need to have a backoff algorithm in the MA reboot sequence so we don't kill the controller when

100K devices reboot at once.

Marcelo: that's why we put some randomness in the protocol.

Slide 10 (Control Protocol):

MA might need to tell controller it failed: might not be able to do what it's told because the request was incomplete or because it is capable of doing what it's told but there's an unexpected resource issue ("task OK but I can't execute it now").

Slide 11 (Report Protocol):

Sends the measurement that it did as part of the report so that the controller has the full context.

Slide 12 (Report Protocol):

How to report if there's potential impact from cross traffic?

Slide 13 (Privacy):

Largely text from Al Morton

Needs review

Move to separate draft?

- A major new section in the new version is about privacy.

- Tries to follow the RFC6973.

Nalini: question on reporting. Pervasive paranoia: is there incentive for the MA to cheat or report certain results (e..g, to make SLA look artificially good or artificially bad)?

Matt Mathis: is there bias here, such as the results not looking bad because the MA couldn't submit the bad results because it couldn't reach its controller then? -I am worried about the loss of samples of data upon failure.

Marcelo: asked for clarification

Phil: we did talk about cheating and regulators would be concerned. The FCC view is that this is solved at the policy level (e.g., ISPs must sign a will-not-cheat contract).

Matt Mathis: are cases where it's easy to skew results by accident that can't be fixed by policy (see his previous comment).

Dan: why move privacy to separate draft?

Marcelo: because half these docs have big

sections on privacy, maybe it makes sense to move it and just say it once.

Dan: reflection of the times! (-: His instinct: separate section but not separate doc.

Benoit: you got 15 pages of contribution, keep it!

Alan Clark: can fingerprint users by knowing where cellphones go (FIXME I didn't get the gist of the question)

Dan: certainly an important issue

Dan: idea of providing aggregates in privacy-protecting way?

Marcelo: MAs belong to group ID, no separate MA ID, so you get location and group ID but that's it.

Marc — They keep IP address.

Al Morton — Which one?

Steve: If you have a device persistently crappy, with group-ID, there is no way to figure out which MA is causing it.

Dan — We are going deeper into the specifics. Still open to comments about how useful group-ID is.

Dan: group ID probably useful but still open to comments.

Slide 14 (Privacy TOC):

per slide

Slide 15:

Discussion between Dan and Phil, wasn't able to note it.

Slide 16:

Dan: Can you commit to the timeline for the I-D?

Phil: will try to but will be away for a whie.

Dan: would like to be capable of collecting input and communicating to external forums that we're in last call in month of December.

Marcelo: said we can take care of all the reviews, the workload will depend on what's in those.

6. IPPM work on Metrics Registry - Bill

Bill Ceverny -- co-chair of IPPM (with Brian Trammel, who was in a different WG meeting during the LMAP meeting)

Slight delay in swapping machines, because Bill had done the right thing and submitted PDFs rather than

PowerPoint. (-:

Slide 2:

IPPM got two proposals in Berlin: one for active, more IPPM-like measurements, for which there is more prior art, and one for passive, which depends on flow keys, less "standard".

Slide 3:

Convened a design team, which talked/emailed and then met on Monday of IETF 88.

Slide 4:

Highlighted that a design that is implementable would result in multiple implementors getting similar results, and one that is deployable is one that could be configured and rolled out by someone who's not a data scientist. Otherwise per slides.

Slide 5:

Matt: are these IANA registries?

Bill: yes.

Slide 6:

Dan (as contributor): Clarification. XRBLOCK has its own registry, why do we maybe need another?

Al: Slightly different set of metrics, passive to some degree but embedded in endpoints, not same as usual passive metering; endpoint-enabled knowledge of what flows between endpoints so may have more data that isn't available in the more usual, purely-passive context.

Marcelo: would like to have registry of all metrics.

Sam Aldrin: registry is within IPPM working group. Just focused on IP metrics, or including things like L2 or application metrics?

Marcelo: not just focused on IP metrics.

Slide 7:

Per slide

Slide 8:

will also present at IPPM tomorrow but otherwise per slide

Benoit: IPPM chairs asked to come up with list of open issues for discussion, are some in there that need more discussion.

Dan: do you have committed authors for all

docs? (answer: yes) Would it be good to share list of open issues here but disuss on IPPM or LMAP lists, or both?

Matt: doc should say in the doc title where it should be discussed, so if filename says IPPM, do in IPPM.

Benoit : LMAP should discuss if this fulfills the needs of LMAP or not?

Dan: WRT call for adoption: can't do that til it meets needs of LMAP.

7. LMAP Information Model - Juergen

- Goal is to instruct MA to measure and results are posted
- IM is split into multiple sections rather than to deal as a whole
- Instructions contains tasks, reference to metrics, options, cycle-id etc.
- Instruction for task start, stop etc.
- Timing – periodic, immediate etc instructions
- Reporting – What is the data that is flowing from MA to the collector.

Slides 2-5: per slide

Slide 6:

pseudocode easier to deal with than the chart he'd used in a previous draft, which is why he did it.

Tasks:

names, registry, options, cycle ID (tag to associate measurements across activities)

Report:

set of channels

Schedules:

names, tasknames to correlate back to tasks, list of report channels, timing object to say when to do this

Suppression:

start/end date, name of tasks to be suppressed

Current doc has just one suppression thing, he decided we should change that while putting the slide together, and will revise doc to reflect the change.

Slide 7:

Channel has name, collector/controller URL, maybe credentials, and channel-specific timing info

Slide 8:

Timing has name and is periodic, calendar, immediate, or one-off

Slide 9:

Calendar time close to what cron does and an active/inactive pair of times

Randomness: type, upper and lower bounds, spreads.

Slide 10 (Reporting):

When

MA ID (optional) and group ID (also optional) -- need at least one

report context: info on context in which measurement was carried out (not yet fleshed out)

Set of tasks:

task config (not yet fleshed out)

report headers and a list of (time, cross-traffic, values) tuples

Kevin: how does randomness play with resolution of types?

Juergen: you have a time when you should do something and you use the randomizing to advance or retard that. Not fully worked out.

Marcelo: avoiding synchronization effects, so randomization is likely to be a smallish range of seconds.

Some discussion of stochastic intervals (specify range, then "flip a coin" during the range periodically to see whether or not we do the test).

Marcelo: one test might be (say) a Poisson sample; Marcelo's take on randomness just moves that around to avoid synchronization.

Slide 11 (Config/Pre-Config, Logging):

Report-ID is a boolean (the "use MA ID versus use group ID" selector)

May have separate channels for logging vs reporting

Alan Clark: assume pushed or pulled?

Juergen: Likely pushed.

Alan: may regret requiring it to be pushed.

Marcelo: why are you asking this?

Alan: Engineering implications of doing one

versus the other, doesn't want to lock us in.

Marcelo: for information model, though, not for protocol how will that affect things?

Alan: having trouble separating protocol from information model.

Juergen: only assumption is that the MA initiates the connection.

MA needs to have some info preconfigured so it can make initial contact; might be updated on first fetch.

Al Morton: MA-preconfig might be the place to put in the randomness to avoid the bootstrapping "100K things all check in at once" problem.

Juergen: yup thanks.

Logging: datetime + event. Event not fully fleshed out, not clear how fleshed out it needs to be.

Slide 12 (Status info):

per slide

Slide 13:

Status info about interfaces

routing and DNS server are optional.

Marcelo: idea is that this would support request made previously.

Dan: What was it?

Marcelo: If controller has multiple interfaces, we find out so we can instruct it which to use.

Dan: where are we?

Juergen: is this notation useful? Seems so to him but wants feedback.

Phil: much easier to understand, thanks. Also: is this time to make it a working group doc?

Dan: also asking that.

Juergen: is that an admin question?

Dan: yeah, but it has to be ready.

Juergen: more writing needed but it's getting there.

Hum: is this ready to become a WG document? Multiple hums yes, only one said no, so consensus is yes, will ask on the mailing list in a few days.

Marcelo: is this isn't ready, why, what needs to happen?

Benoit (the one who hummed no): use case doc isn't ready so how can this reflect that? How does multi-interface stuff plug into the IPFIX/YANG model?

Dan (as contributor): From org/language point of view, seems good, and have given time for other submissions to come in (and they either didn't or did but weren't as good). Trying to gain some time by not serializing the two documents, and it can adjust to match the use case as it evolves.

Benoit: as soon as one is accepted, there is a lack of energy, so he wants to keep the energy there for these docs.

Marcelo: thinks the energy won't drop off til the document is done.

Benoit: yes, but everyone would say that!

8. HTTP-based Protocol Proposal - Marcelo

Dan: this is work in the second stage of the WG charter. Encouraging people to start down this path but these are individual contributions or solutions at this time.

Marcelo: this focuses on the new stuff but that's integrated into the previous (IETF 87) presentation.

Slide 2:

Since last meeting: added logging and status.

Slides 3-6:

Per slide.

Slide 7:

Per slide.

Some discussion (Steve Miller plus Marcelo and Juergen) about whether or not this is REST or just RESTful; Miller to provide better citation or example as to just what he's thinking in this regard.

Slides 8-12:

Per slides.

Steve Miller: How would we pass back aggregate results in this model, such as a histogram of all latencies versus a single result)?

Marcelo: That's an IPPM issue.

Nalini Elkins: having issues with no-cross-traffic

and duration stuff. How do you know no cross-traffic? What if traffic starts at some point during your test?

Al: will specify in protocol but implementation details may have to be taken up in next phase. There is some language in the framework; if that's not good enough they can write more words.

Nalini: if the MA is sitting at an end device, are you plugged into the stack so you know nothing is going on?

(FIXME no name) yes: if at demarc, it can see cross traffic, though it might not be at the demarc). If cross-traffic happens during test, test might be invalid and have to be tossed, which is OK.

Juergen: A lot will be defined as part of the test in IPPM. Will have to work out how to report when cross-traffic starts (don't report, report with flag, report with info on amount of cross-traffic?)

Marcelo: thinks we have a good definition of no cross traffic somewhere.

Juergen: potentially delay if detect cross traffic at start.

9. Open Mic and Next Steps

Dan: that completes the agenda, there's two minutes available for discussion and comments.

Al: would it help if we have a viewgraph to lock down things already agreed upon?

Steve Miller: yes!

Dan: that would be a good thing. Didn't do that this time but did it last time, might have been good to continue that. Last time, did a breakfast meeting for newcomers, which gave them time to ask questions and to come up to speed; there were few newcomers last time (which is why he didn't do that this time).

Dan: Done. Blue sheets!