

# PromCon 2017 Dev Summit

Saturday August 19 2017 9am to 5pm

Google 6th Floor  
33 Erika-Mann-Straße  
München, Germany

## Attendees

- Richard Hartmann - Organizer, Lead
- Murali Suriar <[murali@suriar.com](mailto:murali@suriar.com)>
- Deepak Kakadia <[dkakadia@google.com](mailto:dkakadia@google.com)>
- Goutham Veeramachaneni <[cs14btech11014@iith.ac.in](mailto:cs14btech11014@iith.ac.in)>
- Ben Kochie <[superq@gmail.com](mailto:superq@gmail.com)>
- Brian Brazil <[brian.brazil@robustperception.io](mailto:brian.brazil@robustperception.io)>
- Conor Broderick@robustperception.io
- Aaron Kirkbride <[aaron@weave.works](mailto:aaron@weave.works)>
- Julius Volz <[julius.volz@gmail.com](mailto:julius.volz@gmail.com)>
- Frederic Branczyk <[frederic.branczyk@coreos.com](mailto:frederic.branczyk@coreos.com)>
- Martín Ferrari <[tincho@tincho.org](mailto:tincho@tincho.org)>
- David Kaltschmidt <[david.kaltschmidt@gmail.com](mailto:david.kaltschmidt@gmail.com)>
- Nathan Haugo, [nathan@influxdb.com](mailto:nathan@influxdb.com)
- Paul Dix <[paul@influxdb.com](mailto:paul@influxdb.com)>
- Matt Bostock <[matt@mattbostock.com](mailto:matt@mattbostock.com)>
- Nicolai von Neudeck <[nicolai@vonneudeck.com](mailto:nicolai@vonneudeck.com)>
- Kevin Lyda <[kevin@phrve.com](mailto:kevin@phrve.com)>

## Discussion Topics

- Alertmanager event sink? –RichiH
  - Could link to other org-specific tools
  - Already has group id to discern events
  - Label store should suffice for this
  - Assuming 10k events per year over 10 years
    - 100k label sets are near trivial
    - [various] Discussion of Google's outalator
    - Does this fit in alertmanager, or should it be an external integration?
    - Iso have a log of all notifications and silences so that people can
- SLOG / Annotations (Ben)
- Alertmanager search through Prometheus reverse index? – RichiH / Frederick
- Design for the versioned docs
- Adhoc scraping via the Prometheus UI. – Goutham

- Label normalization - [https://github.com/prometheus/snmp\\_exporter/issues/180](https://github.com/prometheus/snmp_exporter/issues/180) (Ben/Richard)
- Alertmanager exposing its metrics
- Default retention time for 2.0
  - 1 year or 10 years?
- Further SD discussions/planning? - Tincho
- Venue next year
  - Amount of people?
  - People?
  - Etc
- Remote storage stuff, read/write - Brian
- Event Correlation - Deepak
- OpenMetrics - next steps? - Paul
- PromQL as a standard (e.g. people may use it with TSDB library)? What should its limits be? - Matt
- 5 minute staleness vs generic data engine -- richih
- Service discovery as a generic standard? -- richih
- Onboarding process -- goutham
- Relying on Hashing & Timing of scrapes -- RichiH
  - SpaceNet rely on this
- Local time as standard process metric -- RichiH

## Meeting Minutes

- EVENT SINK
  - RichiH Want to persist alerts
  - Bbrazil Like an outalator
  - [frederick] This would be a useful example integration with the alertmanager webhooks.
    - [bbrazil] This should be in pagerduty.
    - [julius] u --
    - [bbrazil] We should a
    - [RichiH] Gossip for sync?
      - Jrv: not all info
      - Frederick: there's some info
      - Bbrazil: stability risk
  - [bjk] Manual timeline/eventlog thing. Taggable company event database that happened in the system. INFO level alerts. Few hundred events per day. Every manual or automated change got documented
    - Eg. for annotations in Grafana

- Deploy
    - Alerts
    - Who is on call
  - Frederick we have a retry mechanism, could use that
  - RESULT: We want this, it should not lose events; TBD
- Alertmanager search with inverted index
  - Frederick search happens in UI
    - Not very great
    - Re-use inverted index for searching, filtering, etc; in server
    - No persistency
  - [jrv] should perform well
  - [Frederick] would need to keep memory usage in mind
  - [bbrazil] people must not depend on am having history for reporting
  - [RichiH] event sink will take care of that
  - [matt bostock] what's the use case for this?
  - [frederick] we need to do this eventually anyway
  - [matt] should we test load at Cloudflare
    - [frederick] yes
  - [frederick] no load testing yet
    - [bbrazil] propose 1 million active alerts
      - More, we need sharding at some point
    - [frederick] plan to submit a talk for kubecon, so i need to do this
- Design for the versioned docs
  - [goutham] everything is in one line
    - Nice to have is sequential
    - But makes it hard to make version
    - Some docs need to combine different versions of software
  - [bbrazil] current layout is assuming that's not an issue, need to redo
  - [paul] we break things up by products
    - Makes it harder for new people
      - Need landing area
    - We have versions specifically indexed
    - When people search, they frequently land on old versions
  - [RichiH] brackets of compatible versions, links within always link to latest
    - Also carry /latest/
    - [murali] What if features change / flags are different, etc?
      - [RichiH] system-inherent, but at least no one statically links to old stuff
    - [Nicolai] we need something to help build a mental model
  - [David] Organize docs around use cases, not components, but keep reference
  - [all] split up guides and reference; mixing them will end up annoying everyone
  - [Paul] Guides should always be for latest

- [Richih] what if they need to understand older versions? Don't throw away knowledge
  - [David] how about tagging "minimum version" for functions, like in python docs etc
  - [goutham] we need something actionable really quickly
  - [bbrazil] do what was already suggested on issue suggested
  - [goutham]
  - [bbrazil] mark functions etc as "this has been added in version x.y"
- Where do ordering functions etc belong
  - PromQL
  - API
  - UI
  - [Tincho] I'd argue to keep useful functionality and not assume that everyone is using grafana.
  - [bbrazil] should we add functions to promql?
  - [matt] can we have promql as a library?
  -
- PromQL as a standard?
  - [bbrazil] what about feature requests which we don't need in prometheus
    - PromQL tied to Prometheus
  - [RichiH] No, groundwater, personal weight, etc are more than 5 minute delta
  - [matt] added functions like covariance and others, this could extend the language
  - [paul] abstracts languages away from execution engines, individual langs can get features independently
  - [ben] promql abstraction layer could help when integrating with other systems like innodb
  - [Tincho] to what extent?
  - [bbrazil] maybe some functions are too heavy, need guidelines on what we allow to become supported functions
  - [matt] discussions on where to draw the line continues. Holt\_winters features prominently (for marketing)
  - [paul] take hints from graphite, was optimised to easily add functions, led to plenty.
  - [bbrazil] more constraints by what results can be represented by vector/matrix, jupyter being mentioned, and R was powerful, but can get out of hand if things get too niche. Users need to understand statistics ultimately
  - [all] general debate about handing users the gun of statistics in the first place
  - [bbrazil] inject functions like go templating can easily work
  - [bbrazil] a promql spec should define the general way functions/operators/aggregators work for things like label matching - but not the exact operators/functions
  - [paul] have our own promql parser, could benefit from promql lib, but be happy to adapt parser to spec, but ultimately need to keep in mind we support other query langs too

- [bbrazil] we could switch PromQL to nanoseconds internally and then lose precision for Prometheus; small performance hit; no objection
- [matt] should start with separate lib?
- [julius] exists already but as part of a repo
- [bbrazil] should stay close to prometheus, as close as possible, for control
- [bbrazil] paul, how do dataframes align across multiple query langs??
- RESULT:
  - Paul will evaluate if they prefer to re-implement or collaborate
    - Possibly switch to nanoseconds internally
  - Keep it in Prometheus repo
  - allow users to extend functions
  - We can always split it to another repo later if we ever need to
- Service Discovery as a lib as well?
  - [all] random bits about SD
  - [frederick] how it's very targeted how targets are consumed
  - [bbrazil] should we have relabeling as a lib?
    - [frederick] i hacked that once
  - [richi] a library is thinking to short; we should have this as a standard to enable SD across multiple systems
    - [tincho] Just use file format
  - [bbrazil] if we break it out, it will gain the ability to pull via http; ok with that trade-off
  - [kevin] so many systems out there, would another standard really help?
  - [bbrazil,ben] standards make adoptions easier (they don't have to understand it)
  - [goutham, Tincho] what's next for SD, for people to maintain it, it could turn into something like exporters, people could write their own. [File sd](#) as the only mainline service discovery.
  - [bbrazil] we should clean our house first, need unit tests, clean up the bugs
  - [all] SD is headline feature of prometheus, should not be jeopardized
  - [some] SD as a sidecar?
  - [bbrazil] limit on SDs prometheus can accept? Limited by maintainers
  - [ben] how about the plugin model?
    - [bbrazil] with go? can of worms
    - [frederick] with explosion of containers and plugins every binary will be different
    - [bbrazil,matt] reading file from disk seems to be the clean demarcation line
  - Testing
    - [frederick] we had things broken on two point releases
    - [matt] AWS etc force people to test against their systems
    - [goutham] that means someone adding marathon needs to give us marathon servers
    - [matt] do you want to rely on external services for your CI?
    - [frederick] CNCF has a CI group, Max will bring that up in the next call
  - RESULT:

- Write an adapter
  - Show interdependencies, interfaces
- Need to clean up before we invite contributors
  - Testing
  - Fix known bugs

COME/GO Matt left

INFO Lunch 12:30 - 13:24

- Adhoc scraping via the Prometheus UI. – Goutham
  - [goutham] for testing, it would be nice to define a temp target via the UI
  - [frederick / bbrazil] security relevant
    - DOS vector
  - [RichiH] make it a config / command line option
    - [bbrazil] overengineering, debug should be always on if it exists
  - [David] use case?
    - [goutham] debugging, see what targets return
  - [many] could do file\_sd and cURL new targets
  - [bbrazil] Better error messages are a good idea generally.
  - [some] who needs this?
    - [bbrazil] people who don't know how their network works
    - [ben] we want to get away from people needing ssh [tincho] wat!?
  - RESULT:
    - Do the math, then simply cache the last scrape result or 1kB of it
    - Also improve error messages & logging
- /targets endpoint/page
  - RESULT
    - Searchable, groupable, foldable, filterable, option to show only errors “no up”
    - Job & target specific page which also shows last scrape output
- Label normalization - [https://github.com/prometheus/snmp\\_exporter/issues/180](https://github.com/prometheus/snmp_exporter/issues/180)
  - [ben] pattern that labels are minimized, in the mysql example: we dont need to the normalized (?) one because we're not a mysql DB
  - [bbrazil] snmp example: timeseries can break when switch ops changes the interface id on the switch itself, so `without()` wont work, you get discontinuity on the data
  - [richi] do we want to store the normalized or the denormalized one?
  - [bbrazil] revives instance vs host discussion (motivation: share rules between people should be supported)
    - [richi] i eliminate the port, so instance is the host
    - [bbrazil] only works if jobs are different?, can make it all work with relabeling
  - [bbrazil] promql needs to support the most general case, period
    - [ben] snmp example (ifIndex) joins and withouts make use of promql too difficult, massively impractical “if i have to do this on every query”
      - [bbrazil] maybe we adapt promql to make it easier?

- [ben] passionate appeal: we dont want to block our power users, but we currently are, I wanna use the generator
        - [julius] other people write their own snmp exporters because of this
        - [richi] I need both `if` labels in there, so I need to hand-write my configs
      - [ben] we should recommend normalizing labels, but should allow denormalized ones
    - [julius] reigning in the debate: make it optional in the generator
    - RESULT:
      - Denormalizing is OK in some cases
      - Multiple label lookups will be added to the generator
      - Add docs section on when/how to denormalize data
- Alertmanager exposing its metrics
  - RESULT: ???
- Default retention time for 2.0
  - [RichiH] this was a joke, but why not increase?
  - Marketing material
  - Should we get user feedback?
  - [ben] what do other systems have?
    - [richi] some forever, cacti 1 year
  - [murali] what's the failure mode? If the server goes down, it's down
  - [nikolai] whatever we do, dont choose 1year sharp, because some events happen exactly yearly
    - [some] what does honeycomb use? we just choose not that!
    - [paul] we have infinite, like a DB
  - [ben] how about Infinite, plus a promtool for cleanup
    - [kevin] the faster we write the tool, the safer it is for the user, general LOLs
  - [bbrazil] don't pick infinity, but just below
  - RESULT:
    - Retention is +Inf time & bytes
    - Promtool will get ability to truncate
    - Flag for retention max in both time & bytes
    - Expose metrics for retention time & bytes, plus current retention time & bytes
- Further SD discussions/planning? - Tincho
  - [goutham] how to onboard new devs like Tincho? How to get new devs?
  - [paul] they just appear for Telegraph at least, we wrote a guide on how to write a Telegraf plugin, that helped a lot
    - [frederick] goutham, you did this, any feedback?
    - [goutham] wrote a doc how to review
    - [bbrazil] see readme in SD repo

- [julius] need bandwidth to review existing PRs, because e.g., we don't know how zookeeper works, or different cloud providers, may be certain devs can own those SD parts?
  - [paul] how do you maintain quality?
    - [bbrazil,julius] we don't, exporter maintainers are in charge, some have integration tests
- [some] so far fabian, grobie and bbrazil have been doing PR reviews
  - [ben] we need to knock on corporate doors (coreos, gitlab, etc) to allocate a FTE, new project with build-kite to help out with integration tests
  - [paul] if you ask corps you need to limit the scope and tasks to not be bottomless, and then this person has to have the skills
  - [goutham,paul,julius] we need a harness to make writing tests easier
  - [julius] need to make SD devving sexy again, give out praise and conference talks
- [murali] stepping back, could we benefit from a generic test harness?
  - [jrv] yes, but there's different scope as we need Azure, AWS, etc accounts
- RESULT:
  - [murali] look up what OpenStack is doing
  - [frederick] CI WG can help with that
- Venue next year
  - [richi] current format should be maintained, CNCF operated is more expensive
    - General agreement
  - Amount of people?
    - [richi] we did 80 last year, 220 this year; could have done 300-350
      - How many next year? 500 could be doable?
      - [murali] do we want to?
    - [bbrazil] inflection points: <300 it stays a community conference. Should stay on order 300-400 max
    - [frederick] single track is great
      - General agreement
      - [bbrazil] maybe consider more tracks in the future
  - People?
    - [richi] whoever becomes overlord gets to choose the city, should be a transport hub
    - [paul] how about NYC
      - [richi] cant do USA, had issues with visas
    - [bbrazil,richi] if i do it -> dublin, if i do it -> munich
  - [richi] how big are google conference rooms in dublin?
    - [murali] looking at the foundry: 360 people
  - [bbrazil] are we targeting ourselves? Had lots of use cases this year.
    - [richi] how about panels?
      - [julius] gophercon had community days, kinda informal panels, people could listen in, with a hackathon in the background



- [ben] could help with SD recruiting
      - [bbrazil] keep ticket price at around 100.
    - RESULT:
      - Munich or dublin seems to be finalists
      - 200-300 ish is the mark
      - Consider having use-case and design day, panel discussion for dev summit, but still some private, relaxed thinking time
- OpenMetrics - next steps? - Paul
  - [murali] group has become too big, things take forever so there is no immediate value, there is no running code
  - [richi] i'll weed out some issues after the summer holidays, and reach out to contributors (google/stackdriver, openTSDB, influx, weaveworks, etc), then hammer out design (in protobuf)
  - [bbrazil] could use first mover advantage, hammer something out and implement something
  - [paul] then we write some code, and we have 2 ref implementations, job done!
  - RESULT:
    - RichiH to start closing tickets towards the end of summer holidays
      - Rest will be done with rough consensus
      - Get buy-in from Prometheus, Influx, Tom Wilkie, Solarwinds, Stackdriver, Grafana, etc
- 5 minute staleness vs generic data engine -- richih
  - [bbrazil Prometheus implementation detail
  - [RichiH] Yes, libTSDB & PromQL are all that's needed
    - [goutham] yes
  - RESULT:
    - We want this, but not urgent
- Onboarding process -- goutham
  - [goutham] huge descriptions helped me understand issues, but hard to get
  - [goutham] github labels might help, e.g. lowhanging-fruit/help-wanted, not sure what the requirements are to apply the label
    - [frederick] there is a benefit in adding big descriptions on those, then the committer is likely to come back for more
  - [RichiH] cant be too aggressive about onboarding; onboard too many and you burn out if people don't stay; been there
  - [Tincho] prom IRC channel is too busy
    - [julius] yes, shall create new one for devs only (keep old users in original channel)
    - [bbrazil] some issues with long messages (Matrix long messages are annoying)
  - RESULT:
    - We tweet easy issues at people
    - Try to be extra nice to the ones GitHub marks as first-time users

- Created new channel “#prometheus-dev” on IRC/matrix

COME/GO: Frederick, Paul and Nathan left

INFO: Nicolai was caught napping

- Relying on Hashing & Timing of scrapes -- RichiH
  - SpaceNet & Cloudflare rely on this for snmp\_exporter because of slow hardware
    - [RichiH] Changing this now would take Clouflare offline  
The underlying reason is HA: two proms hit the target at the same time, current fix is tweaking the exporter
  - [bbrazil] you shouldn't be caching, this is a capacity problem, also people shouldn't rely on this implementation detail, but if you really need it, caching might help, but caching messes up timestamps
  - [bbrazil] are there devices that cannot deal with interleaving walks?
    - [RichiH] I tested, none of them did this well
  - [murali] we have to be careful not to build a request death machine
  - [RichiH] perhaps it must live in the exporter (not prometheus)
    - [murali] makes sense because this is where the pain is
    - [RichiH] caching messes up my rate spikes
  - RESULT:
    - Hashing will change, but we will talk about mitigations well before
      - If the cache exposes timestamp, we need to take new staleness into account
        - OpenMetrics support for “this has gone away”?
    - Document that scrape phase is unspecified
      - Also add this to changelog when it changes
- System time as standard process metric -- RichiH
  - [RichiH] I need this to detect clock skew
  - [bbrazil] you can get this in v2 with node\_exporters clock node\_time - timestamp(node\_time)
    - [RichiH] but i need this on machines where node exporter is not installed
  - [murali] you've got a use case but I have not heard a principled argument on why we should include time, and not any other node level metric in the client libs
  - [bjk] Even Google didn't need that
    - [RichiH] hmm, cant argue against that, even though I don't buy the “none of the big corps have that problem” argument
  - RESULT:
    - Clear resistance from most sides, no action
- RichiH's tactical dashboard
  - Shows off Sorted table by urgency, alerts are associated to tickets via ticketing system
    - Silenced alerts show up in gray for 12h after a silence, then disappear, reappear 6h in gray before the silence ends
    - Hyperlinks Alerts to runbooks

- Hyperlinks assigned users to the tickets
    - Can filter by teams or show global state
      - Shows critical events globally, ignoring filters
  - Also shows planned & emergency maintenance notices for context
  - Also shows new, untaken, stale, high prio tickets
    - With hyperlinks
  - Keeps SREs sane and shows what to do `_now_`
  - If anyone wants it, ask RichiH
  - Frederick might be interested in implementing something similar upstream
- Remote storage-y stuff, read/write - Brian
  - [bbrazil] we have influxDB interested, and others. Things are being worked on, generally going well
  - [goutham] feedback from paul and tom is they prefer blocks being shipped
  - [bbrazil] current implementation to write is simple but limiting, if we only send blocks we can do proper snapshotting too
  - [RichiH] looks like we end up with 3 interfaces: remote-write-batch (blocks), remote-write (realtime), remote-read
  - [RichiH] long-term storage might be libTSDb with PromQL
    - Can be used for data analysis (R/Ooohrrr/Romeo)
    - Can be long-time storage
    - Has eventual consistency
    - Does not need to have the 5 minute staleness
  - [bbrazil] implications to HA story: ideally prom runs in k8s with an external storage, if prom restarts it can use remote-read to get the data from external storage
  - [RichiH] what about compression/downsampling?
    - [bbrazil] this is a hard problem, prefer if the storage engines deal with it. Problem is maintaining correctness
  - [bbrazil] What do tell to people who ask for long-term storage?
    - [julius] I just tell people about the ones I know
    - [bbrazil] there is a PR about integrations, the storage solutions should be listed there
  - [bbrazil] ultimately prometheus is ephemeral. Getting data back in after a restart could become essential
    - [kevin] my customers think of TSDBs as forever databases, but the current prometheus can't deliver on that
  - RESULT:
    - Just discussion
- Event Correlation - Deepak
  - Direction for Prometheus for Data Science:
    - Functions in Prometheus
    - External Functions Data out, process, Data In
  - Recommended Interface to pull data from prometheus to External Analytics Tools( Jupyter, Colab ) then return result

- Synchronous
- Asynchronous - Correlation ( requires ingestion from external data sources)
- Prometheus -> Pub/Sub -> OpenTSDB & BigQuery
- FCC requirement of 3 years retention
  - [bbrazil] as soon as laws are involved, Prometheus is out for storage
  - [Ben/Brian] Use another database
- How to point to events?
  - [bjk] Link from Alerts to ELK etc with the correct timestamp
- RESULT:
- Simple multitenancy with forced labels - Tincho
  - [Tincho] how can this be done?
    - [ben] improbably is doing this, and Gitlab has an issue out for this
    - [julius] need to ensure the label is there
    - [ben] how about treating it at a k8s level: one prom per customer, and then it becomes a routing issue, and federation can help with cross-instance metrics
  - [bbrazil] not sure if label\_replace() and vector() are supported
    - [julius] probably safe because no data is retrieved by them
- Current UI gripes - David
  -