Data storage at the RIPE NCC

Robert Kisteleki

RIPE NCC R&D

RIPE NCC

Data collection exercises

We run multiple measurement systems:

- Test Traffic Measurements (TTM)
  - To be decommissioned soon
- DNSMON
  - DNS root and TLD monitoring
  - “Powered by TTM”, will be “powered by Atlas”
- Routing Information System (RIS)
  - BGP information from ~12 collectors, ~700 peers
- RIPE Atlas
  - Distributed measurements from tiny devices (and more)

In RIPE Atlas

In RIPE Atlas:

- 2500+ probes active as of now
- Supplying ~60M data points a day
- We expect to double-triple that this year:
  - DNSMON -> Atlas migration
  - Atlas Anchors as targets
- User Defined Measurements available since 2012-03
- The probes use ~1% their capacity
- Atlas Anchors are coming

In RIPE Atlas

Daily fixed

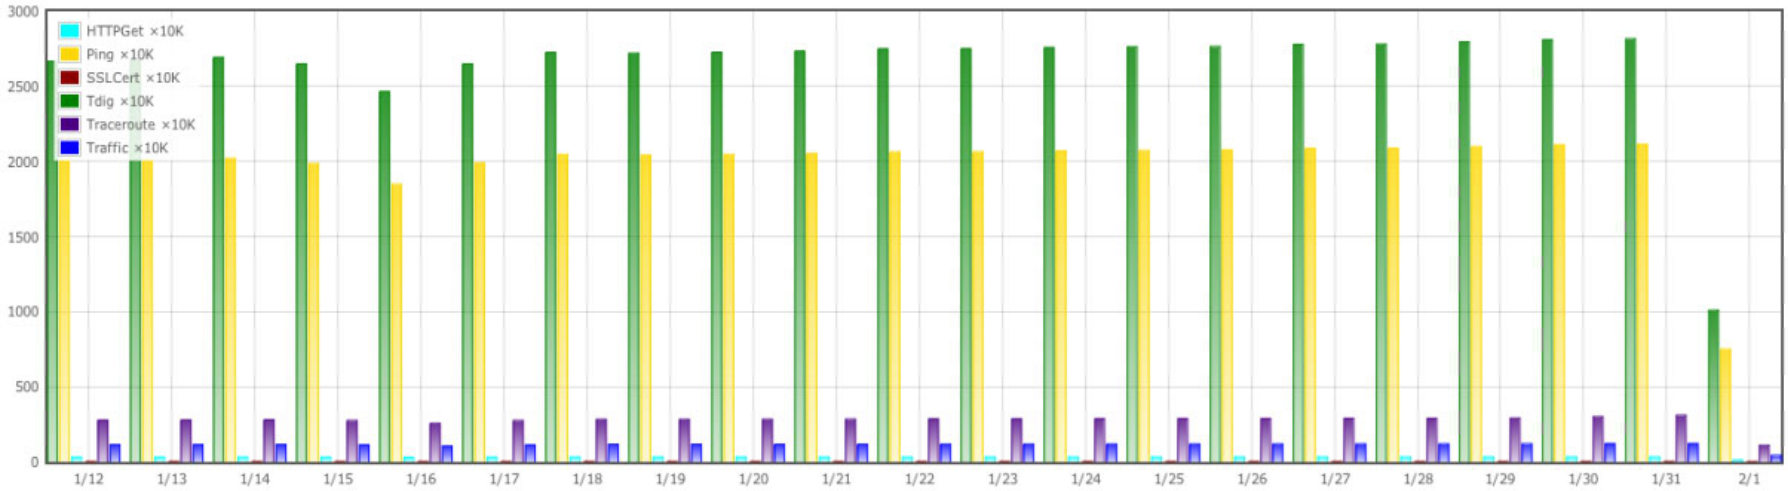

Daily UDM

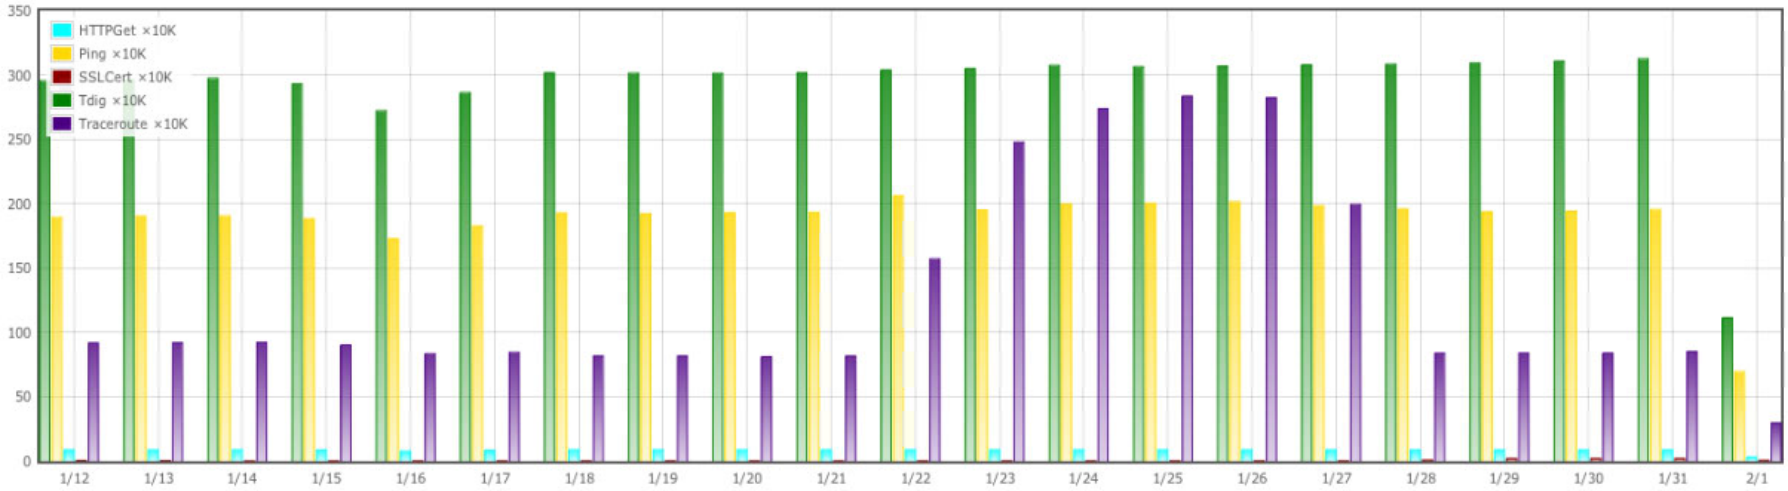

In RIPE Atlas

The difficulty is to store/retrieve this data.

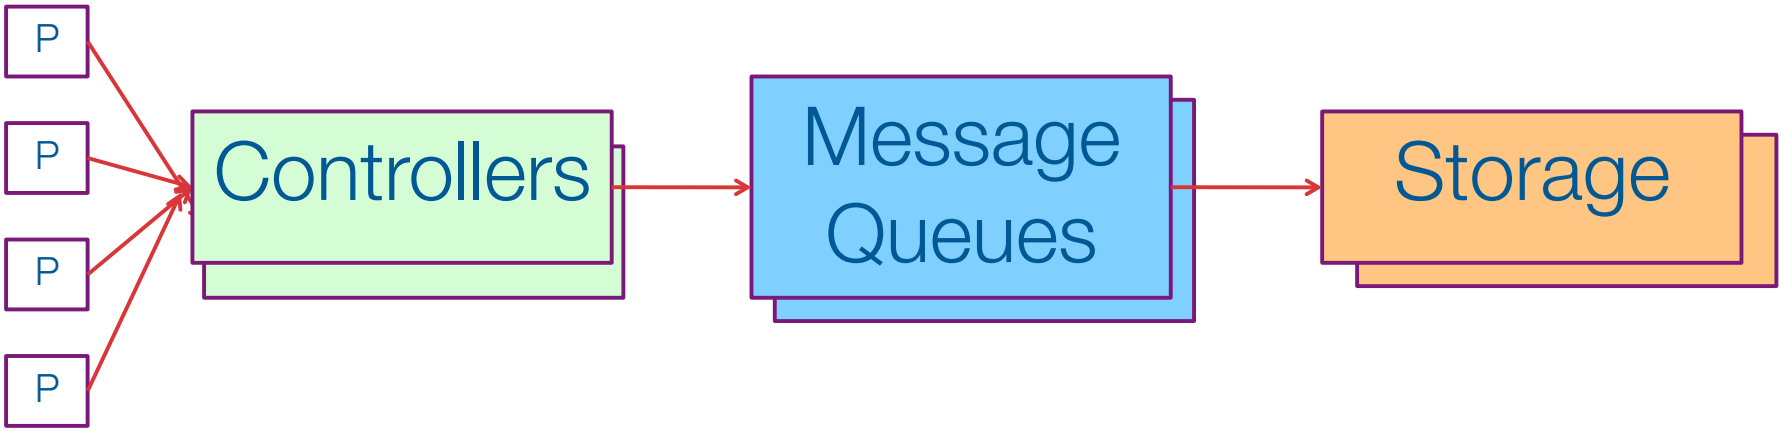

In RIPE Atlas

- Probes supply JSON data
  - We needed a {key->value} format
  - Not so compact but compresses well
  - For the lookups you need indexing anyway, so parsing performance is not an issue
  - JSON has very good tool support

Components we use

On the storage side:

- A bunch of regular machines
- Hadoop/HDFS as infrastructure
- HBase for storage
- RabbitMQ + Flume for transferring/inserting data
- Thrift for retrieval
- Map/Reduce jobs and Hive for number crunching

Components we use

All of these have their own pros/cons

- There's a steep learning curve
- You need (some of) these for big data
  - Unless you're Google

Most of them are bleeding edge

- Memory leaks (-> crashes) and "random events" do happen
- Once you tame them, they work well

In RIPE Atlas

Internally we serve:

- “data downloads”
  - Full result data for a specific time period
  - You get results in full detail
  - Slow for large result sets
- “latest X” results per probe, measurement
  - What’s the latest result for a measurement?
  - Can specify certain fields
  - (or the latest X results, cached)
- Coming: multi-resolution aggregates
  - To facilitate visualising long term trends

In RIPE Atlas

Interacting with the system

- We're introducing various APIs:
  - Searching in existing measurements ✔
  - Looking up meta info ✔
  - Downloading data ✔
  - Searching for vantage points ✔
  - Specifying / modifying / stopping measurements (coming)
- We'd love to open all data to the public, that needs more work
  - Good news: most if it is already public

Bottom line

Some takeaway messages:

- On this scale you need a solution that scales automatically
- Off the shelf components exist, but you do need to tailor them to your needs
  - That can be tricky