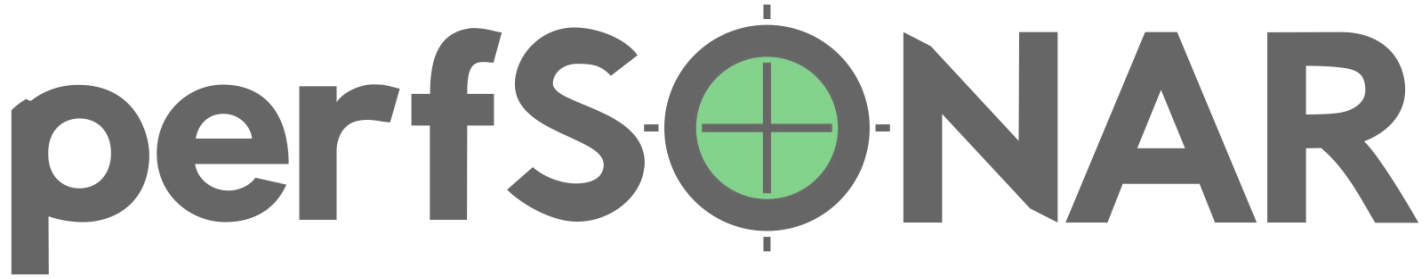

perfSONAR new Opensearch backend

Luan Rios (RNP)

3rd European perfSONAR User Workshop organized by GÉANT Project

perfSONAR is developed by a partnership of

INDIANA UNIVERSITY

perfSONAR 5.0 Context

- Bump from 3.X to 4.X was almost 5 years ago
- Enough of a change we thought was right time to go from 4.X to 5.X
- Multiple things going into perfSONAR 5.0, but this talk is mostly going to be focused on the archiving of measurements since that's the biggest change

Motivation for backend refactoring

- At some point problems get complicated enough you want to be able to easily integrate with other data such as:
  - Interfaces stats (e.g. SNMP)
  - Flow
  - Optical
  - More…
- These are hard problems, but we think we can better position perfSONAR for this type of integration
- It starts with the metadata and data that we *archive*

What do we mean by an archive?

- **Archive** is where we store measurements long-term

- Archive is where visualizations get their data

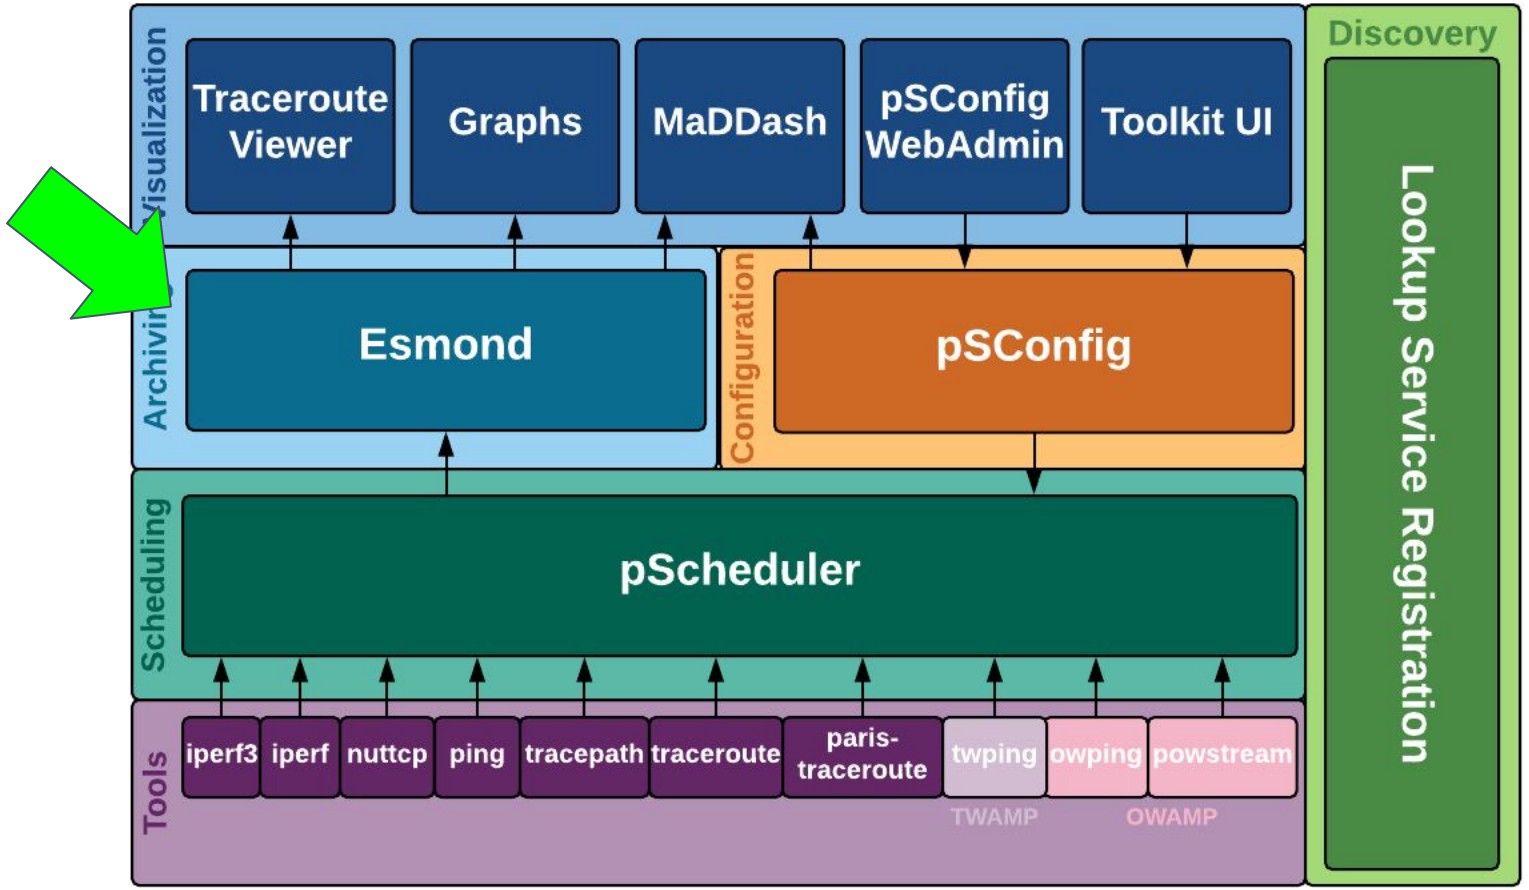

What is Esmond?

- Default archive that most users run
- Django app with custom REST API
- Use two backend databases
  - PostgreSQL
  - Cassandra

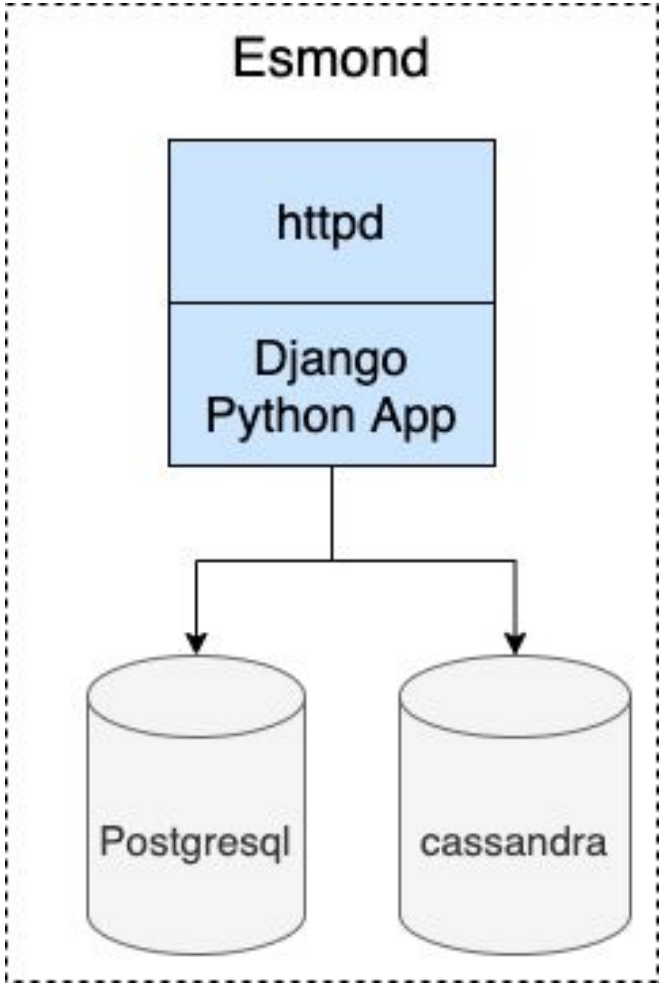

Why replace esmond?

- Many good open source options for time-series storage
  - Rich query languages
  - Integration with off-the-shelf visualization platforms
  - Better support for backups, scaling, etc
  - Cloud vs On-Prem Deployments
- Stability
  - Cassandra one of the main source of issues on user list
- Maintainability
  - Less custom code
- Community successes from which we can learn
  - WLCG
  - NetSage

Elasticsearch, Logstash, Kibana (ELK) and Grafana

- **Elasticsearch** - Stores and indexes documents and lets you do searches
- **Logstash -** Accepts input from lots of different sources, enriches with location data and more, can output it to different places (like Elasticsearch)
- **Kibana -** Visualizes data in elasticsearch
- **Grafana* -** Visualization platform for ElasticSearch and more

Elastic Licensing Changes

- Elastic has always had a tiered licensing model:
  - The core of Elasticsearch and Kibana were Apache 2.0 - just like perfSONAR
  - Some of the features fell under their their “Basic” license which was free but they didn’t necessarily share the source. This includes many of their auth features, rollups, etc.
  - Advanced features such as more alerting plugins, field level access control, built-in anomaly detection using ML fel under their commercial license.
- Enter Amazon’s OpenDistro (now named OpenSearch)
  - In 2019, Amazon decided to build a service out of Elasticsearch using the bits of code under Apache 2.0 and re-implement the feature themselves in the other tiers. They called this implementation OpenDistro.
  - Since 2019, whenever Elastic did a release, OpenDistro would merge in the changes. They’ve maintainers API compatibility and have tried to make things compatible.
- Earlier this year Elastic decided they didn’t like Amazon making money off their product without them getting a cut, so switched Elasticsearch and Kibana to Server Side Public License (SSPL)
  - Further added license checks to Logstash to make things more challenging

Moving to OpenSearch

- Default perfSONAR bundles will rely on OpenSearch

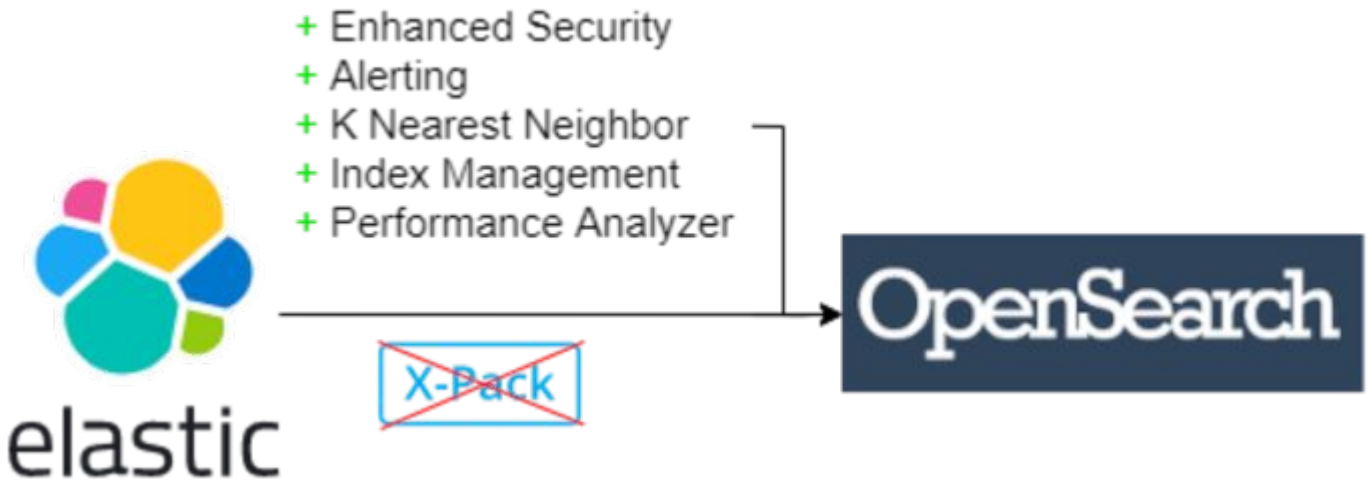

The Software Pieces

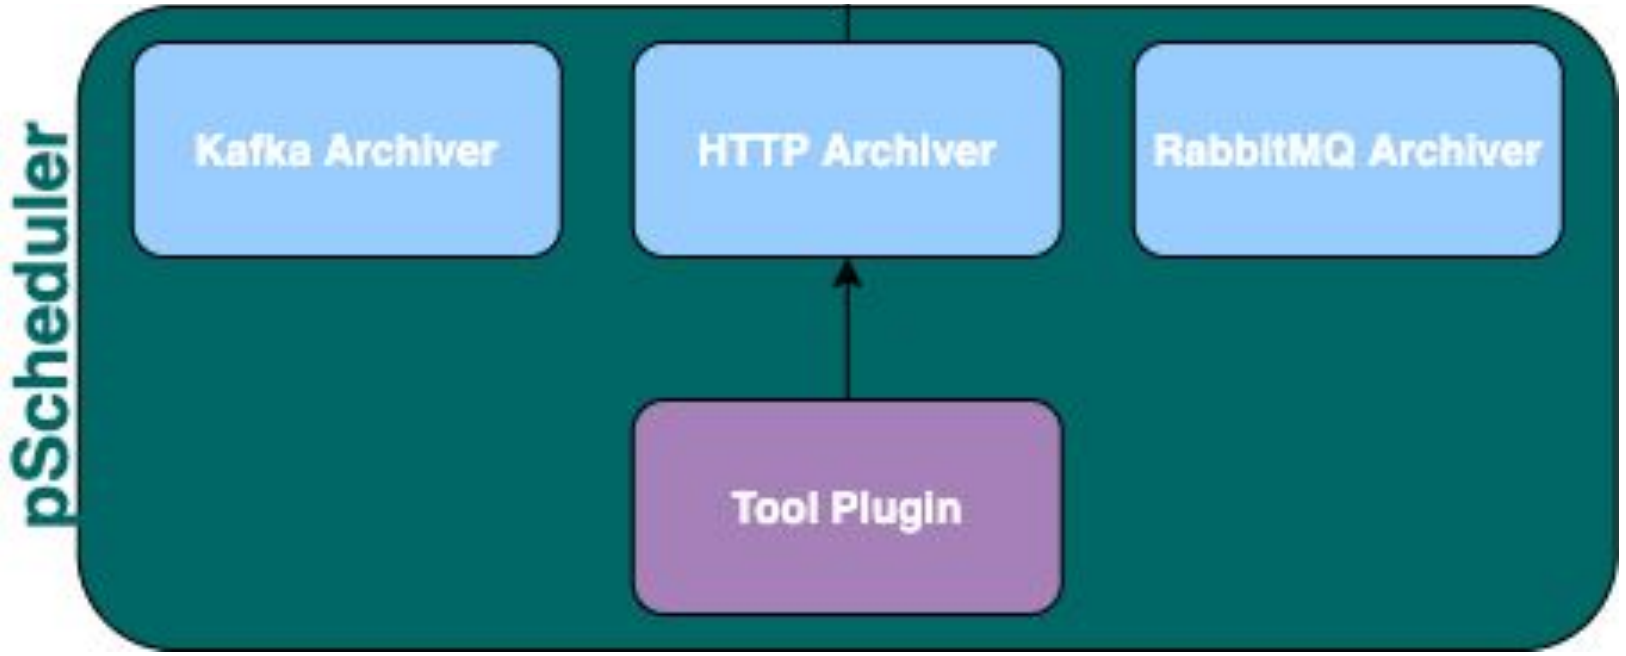

The Software Pieces

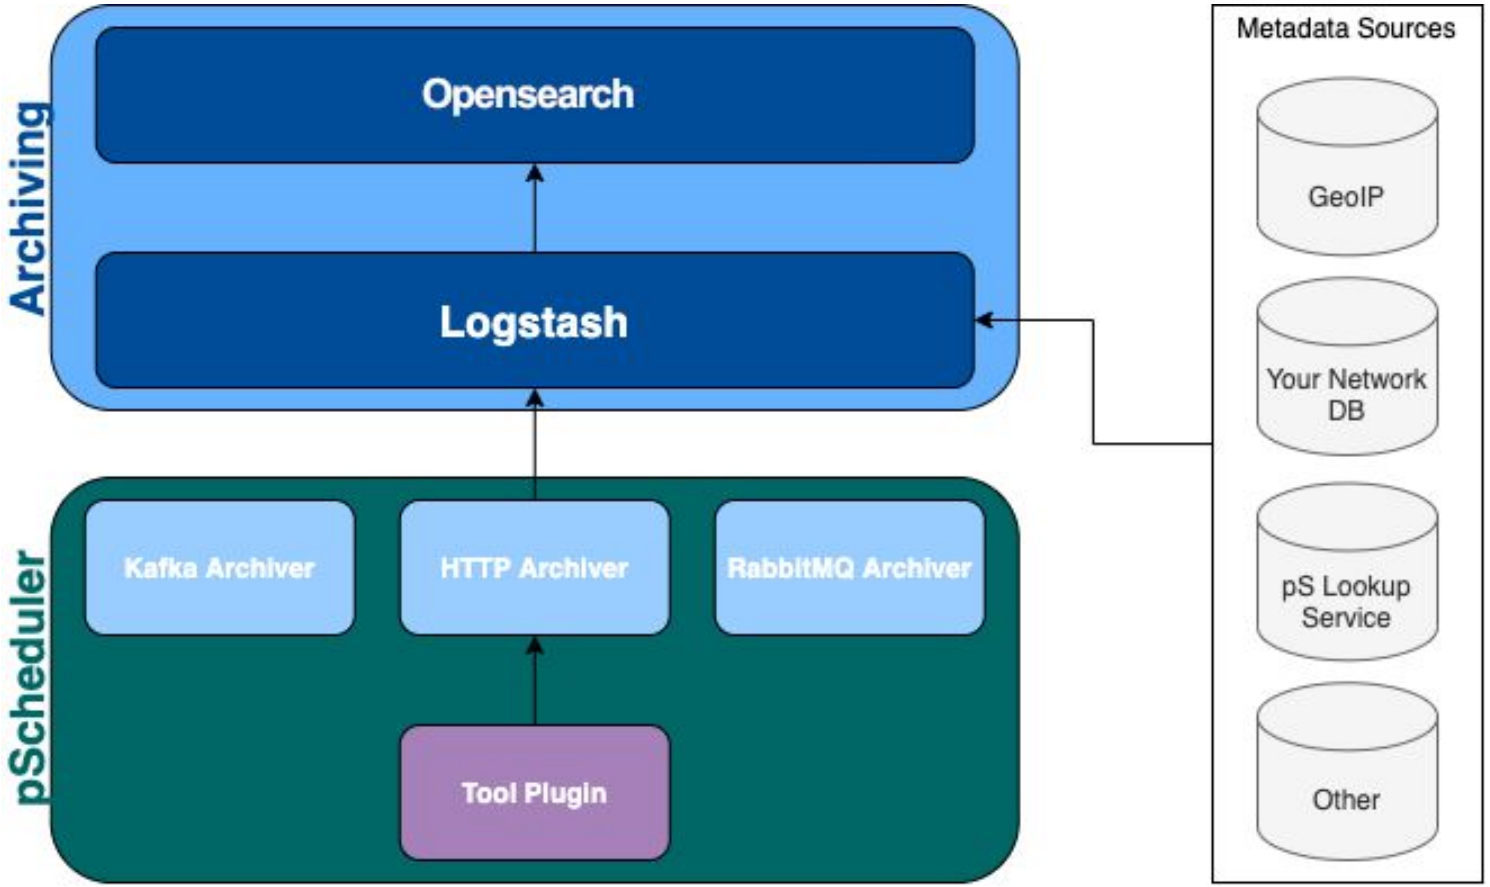

The Software Pieces

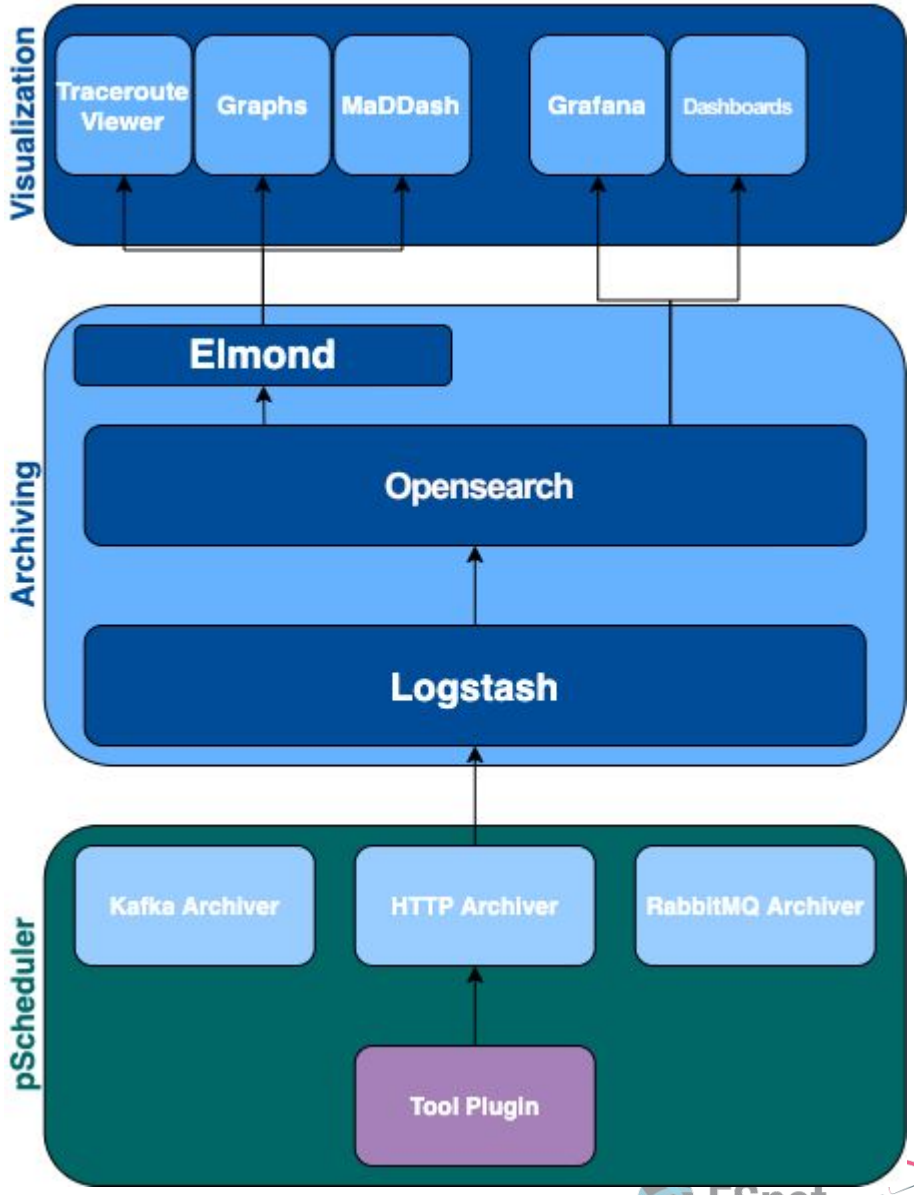

How does this fit together?

- Multiple new parts, but not a radical shift in architecture
- New components map to existing bundles

Packages relationship 4.x

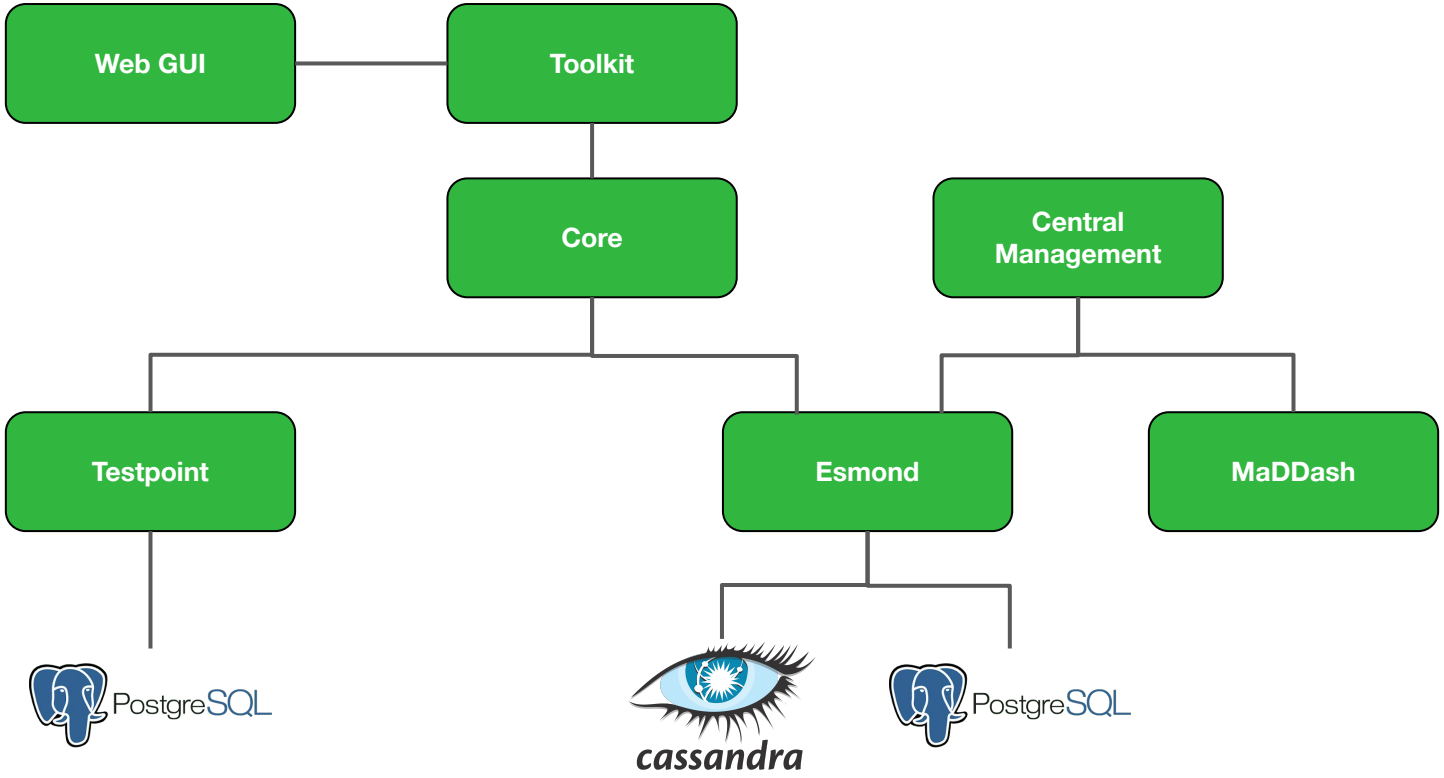

Packages relationship 5.x

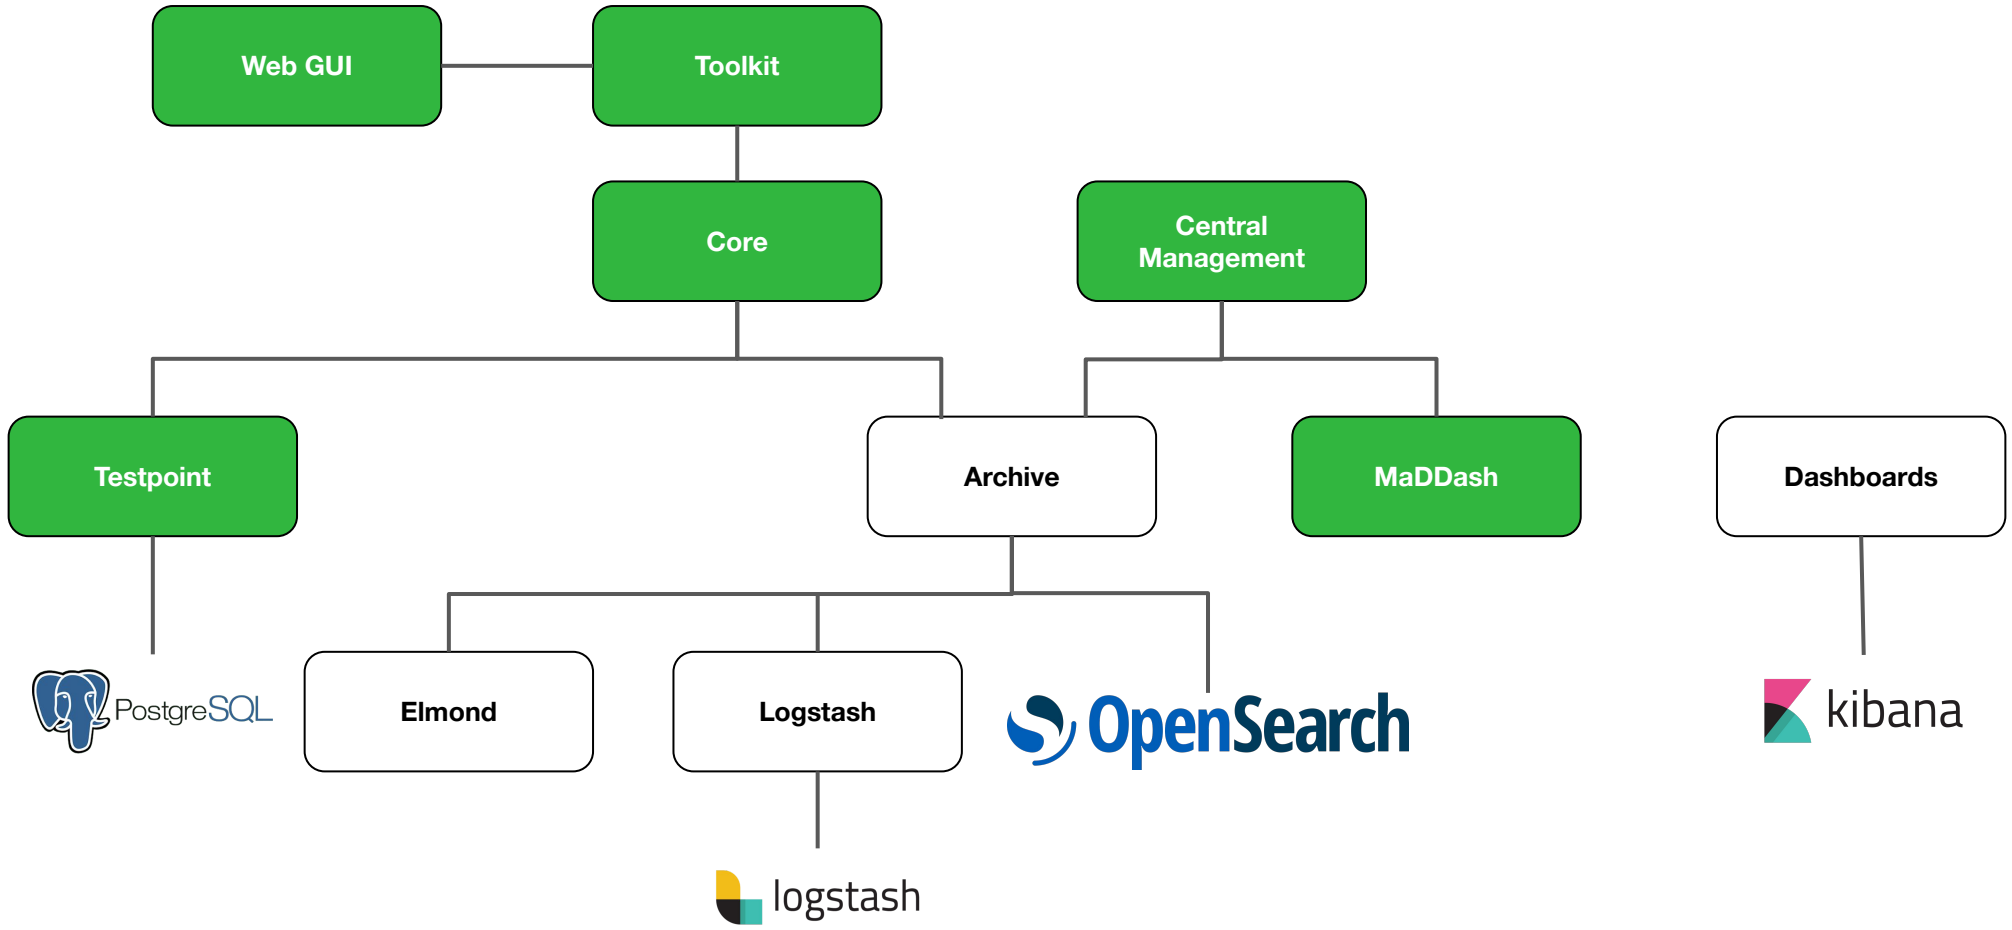

Elmond

- Converts Esmond queries to queries understood by Elastic
- Python 3 Flask application

```
{
  "ELASTIC_HOSTS": [
    "https://admin:5y3...MsI@localhost:9200"
  ],
  "ELASTIC_PARAMS": {
    "use_ssl": true,
    "ca_certs": "/etc/elasticsearch/root-ca.pem",
    "client_cert": "/etc/elasticsearch/admin.pem",
    "client_key": "/etc/elasticsearch/admin-key.pem"
  },
  "PROXY_PATH": "/esmond/perfsonar/archive",
  "FORCE_HTTPS_URLS": true,
  [...]
}
```

`/etc/perfsonar/elmond/elmond.conf`

Logstash

- The perfSONAR Logstash pipeline used to enrich data before archiving

```
## Logstash environment variables.
log_level=info
opensearch_output_host=https://localhost:9200
opensearch_output_user=pscheduler_logstash
opensearch_output_password=pscheduler_logstash
```

`/etc/perfsonar/logstash/logstash_sysconfig`

1. Input (IP/port)
2. Build pscheduler object
3. Normalize IP addresses
4. Convert ISO8601 durations to seconds
5. Lookup GeoIP information
6. Process each type of task
7. Output (index template)

```
input {
  http {
    host => "localhost"
    port => "11283" # ACII 112=p, 83=S
  }
}
```

`/usr/lib/perfsonar/logstash/pipeline/*`

Archive

- The perfSONAR Measurement Archive based on OpenSearch
- Configures Elmond to connect to OpenSearch
- Configures OpenSearch security
  - certificates and user permissions
- Configures ISM policy for pscheduler index

Question and answer icon by iconosphere from The Noun Project

Questions and Answers

perfSONAR

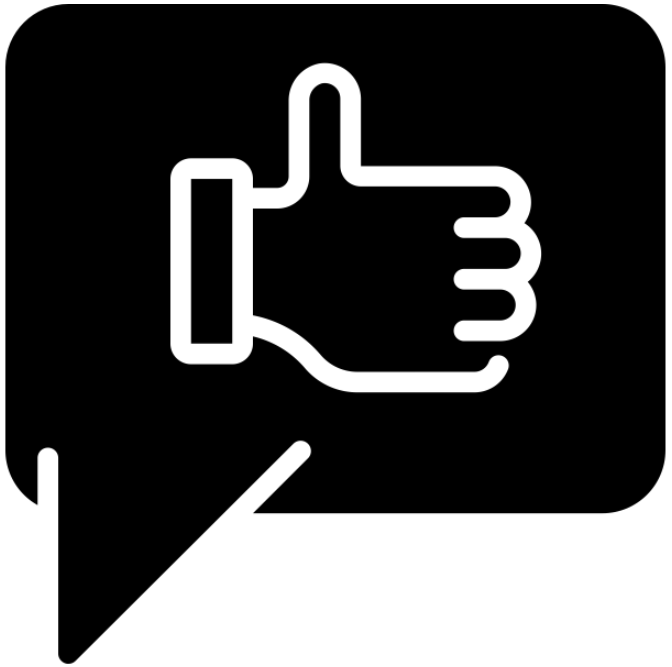

Thanks!

For more information,
please visit our web site:
https://www.perfsonar.net

Thanks icon by priyanka from The Noun Project

perfSONAR is developed by a partnership of ESnet GÉANT INDIANA UNIVERSITY INTERNET2 RNP ORGANIZAÇÃO SOCIAL DO MCTI UNIVERSITY OF MICHIGAN