

# perfSONAR

# perf5.0NAR

## What to Expect from the Next Major Evolution of perfSONAR

Antoine Delvaux (GÉANT / PSNC)  
(thanks to Andy Lake (ESnet), Daniel Neto (RNP), Luan Rios (RNP))

3<sup>rd</sup> European perfSONAR User Workshop ▪ 24-25 May 2022

*perfSONAR is developed by a partnership of*



ESnet



INDIANA UNIVERSITY

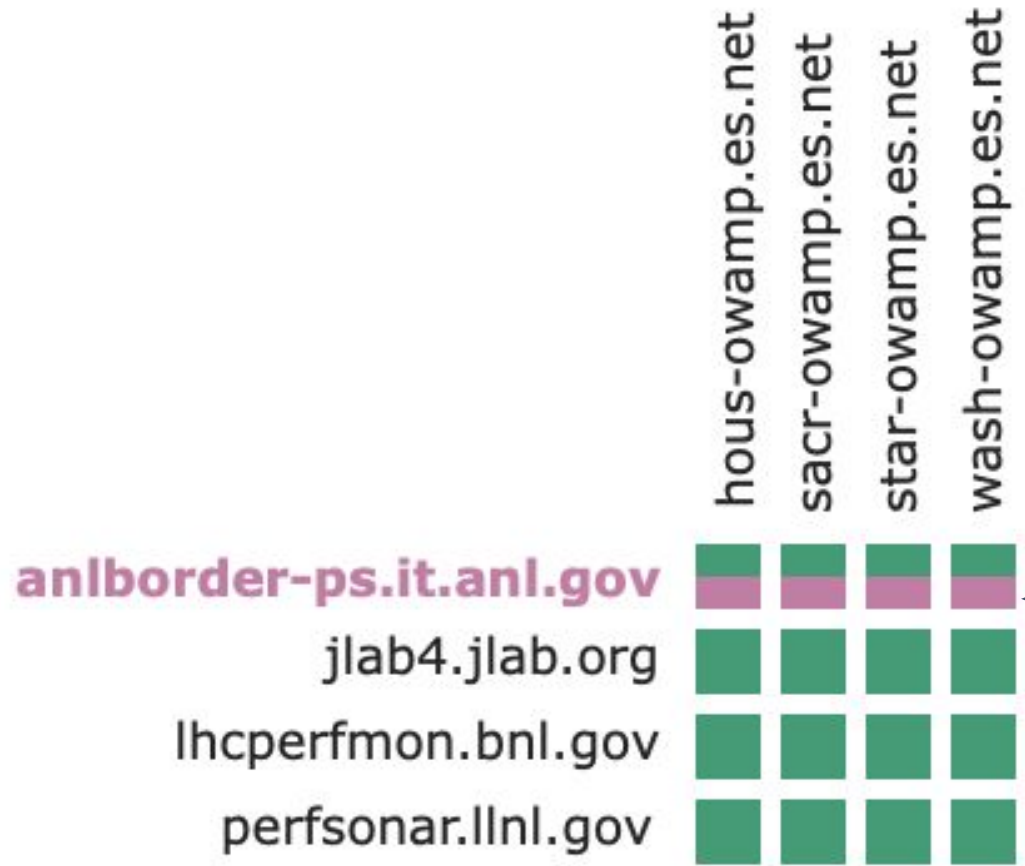


# What's in perfSONAR 5.0

- 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

# perfSONAR Today: MaDDash

- Provide a dashboard of “grids” called **MaDDash**
- Highlight problems when metrics like throughput or packet loss fall below a certain threshold
- ESnet’s lives at <http://ps-dashboard.es.net>



ANL seeing packet loss for traffic entering ANL from ESnet



# perfSONAR Today: Graphs

perfSONAR test results - [documentation](#)

[Share/open in new window](#)

<b>Source</b> hous-owamp.es.net 198.129.254.66 <a href="#">Host info</a>	<b>Destination</b> anlborder-ps.it.anl.gov 130.202.222.58 <a href="#">Host info</a>	<b>Report range</b> ← Choose → From To <b>Submit</b> Mon, 15 Jun 2020 15:16:45 GMT to Wed, 15 Jul 2020 15:16:45 GMT
---	--	--

Show/hide chart rows  Throughput  Packet Loss  Latency



← One direction has much higher throughput than the other

← Low throughput direction has quite a bit of loss

← “Bad direction” has choppy latency



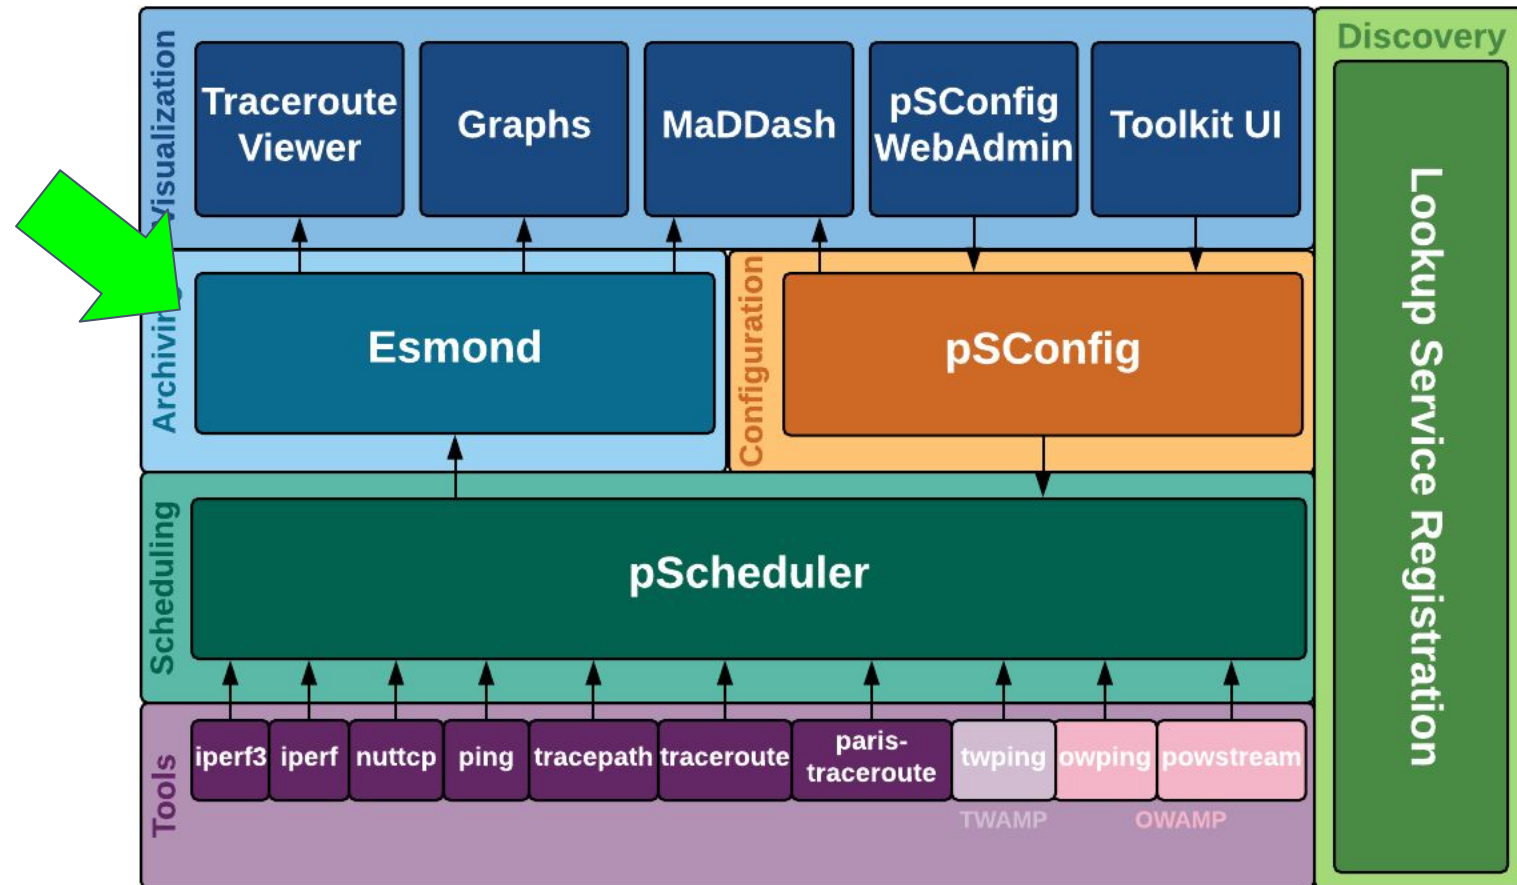
# Looking closer at network problems

- 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

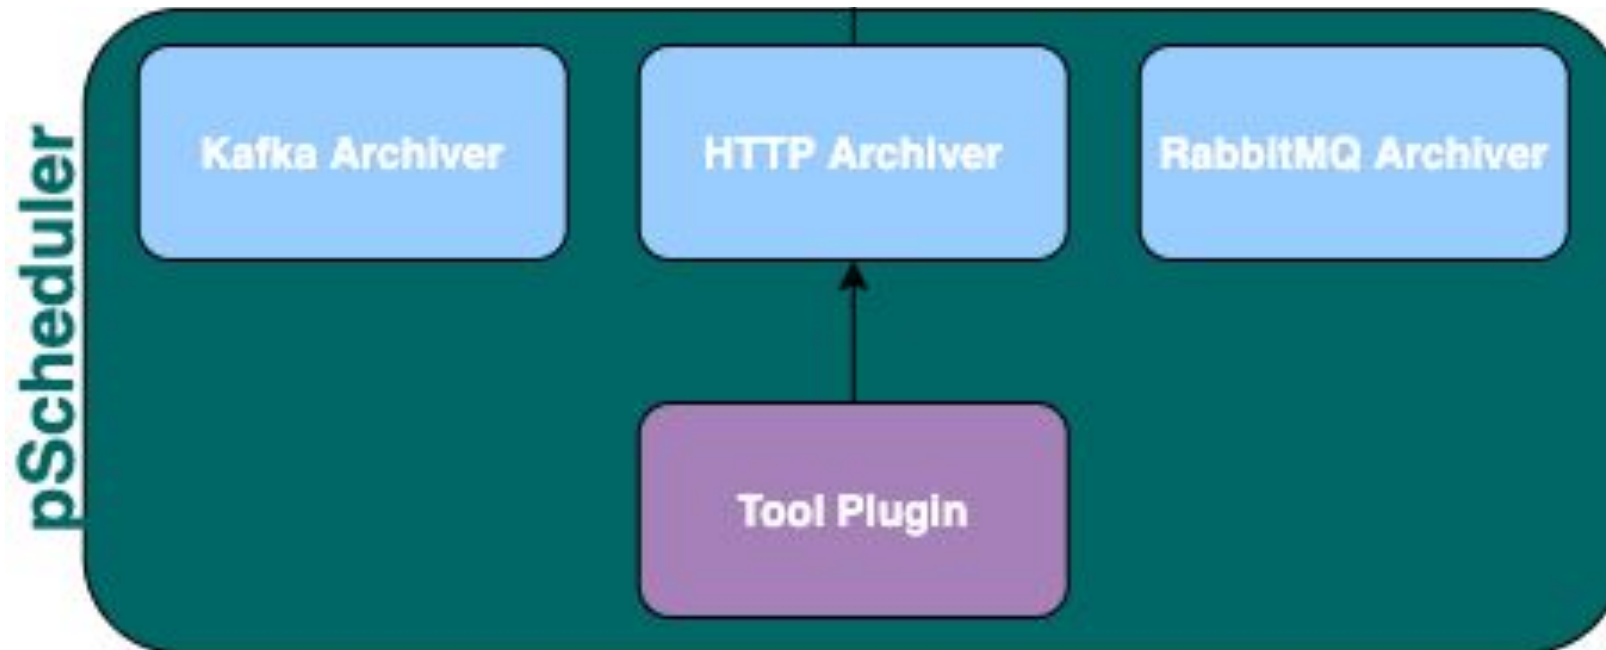


# Moving to OpenSearch

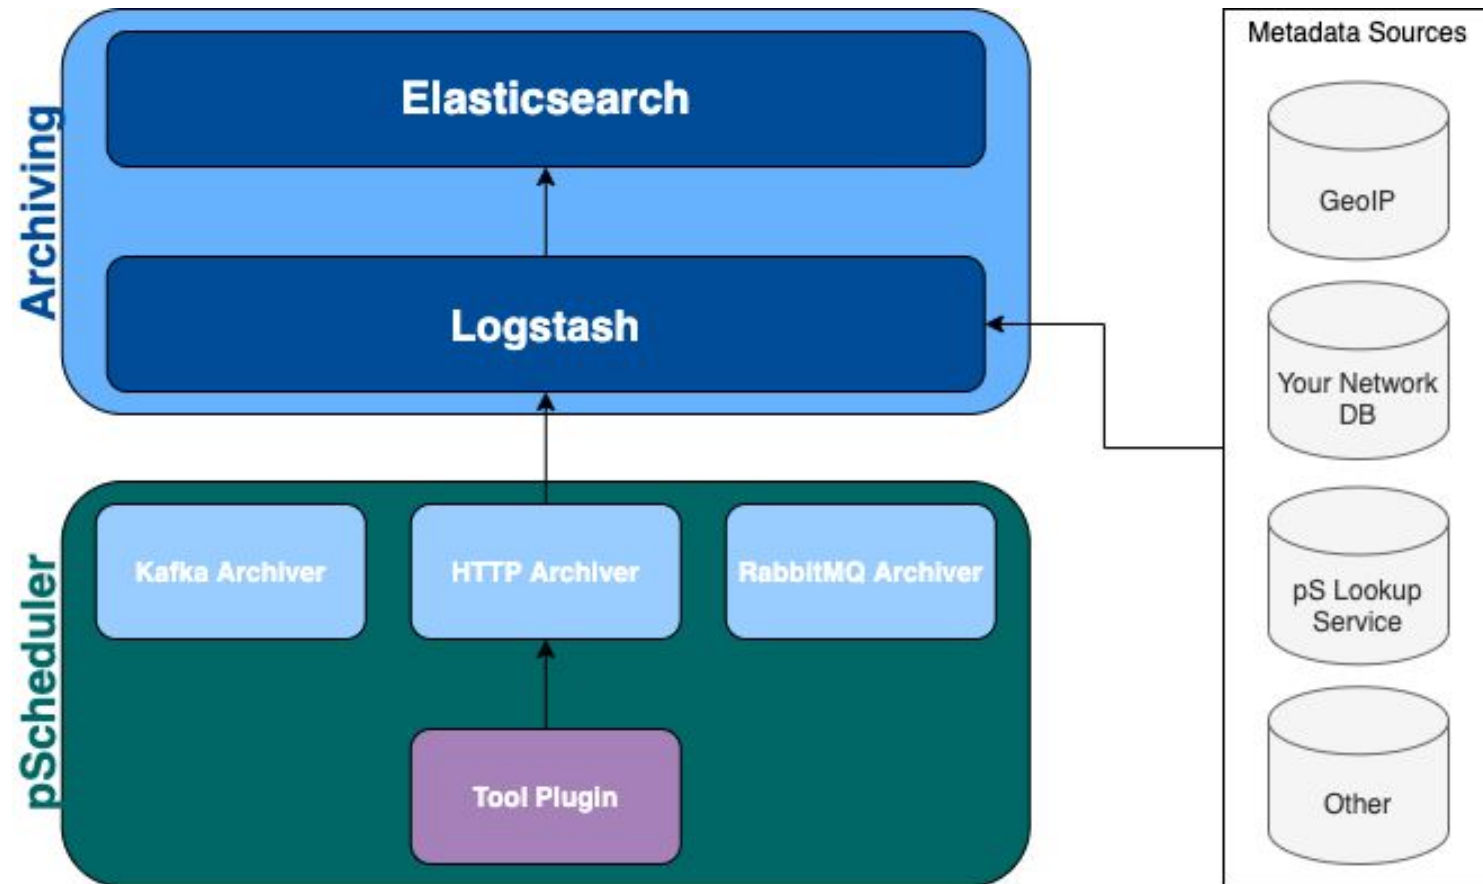
- Because of licensing changes
- Default perfSONAR bundles will rely on OpenSearch
  - Might be OpenDistro on Debian
- Will maintain compatibility with Elasticsearch for those with existing installations



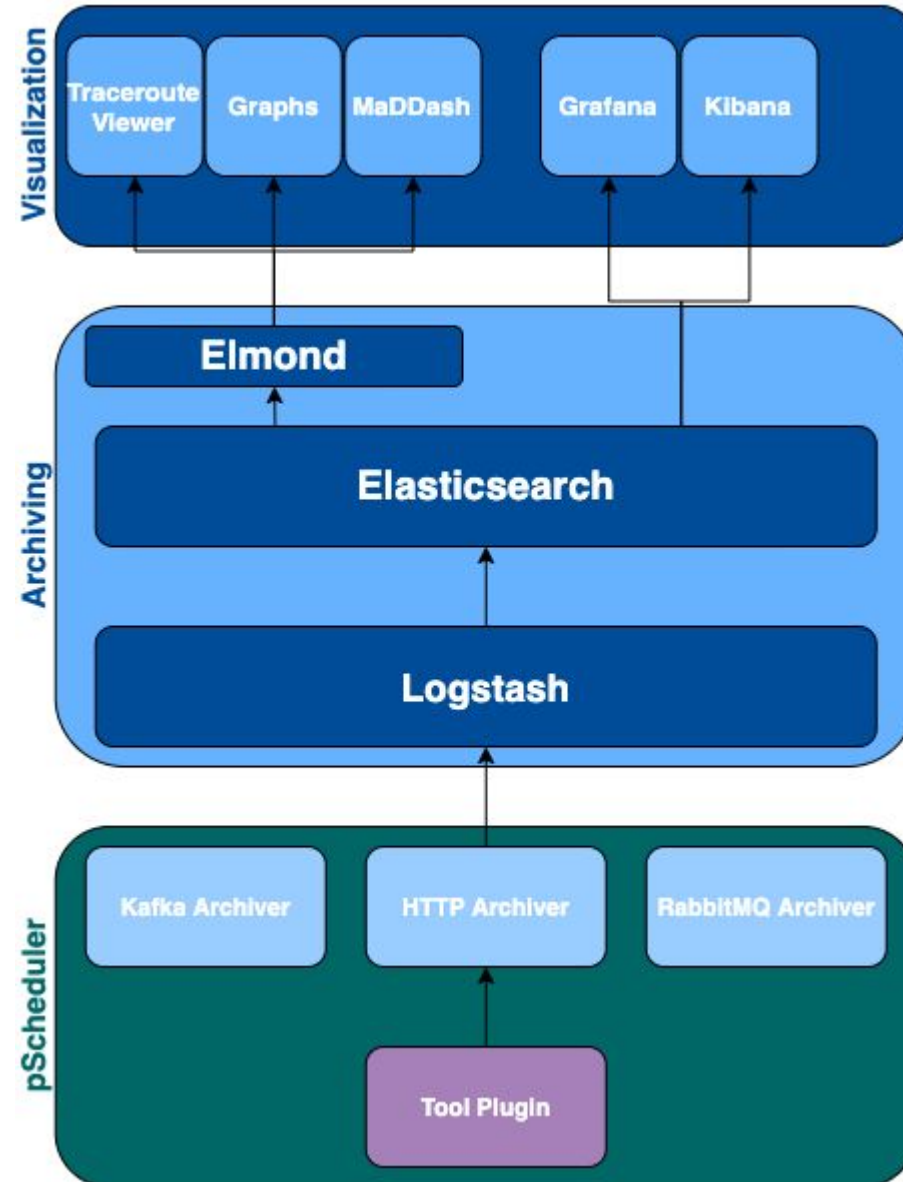
# The Software Pieces



# The Software Pieces



# The Software Pieces




# How does this fit together?

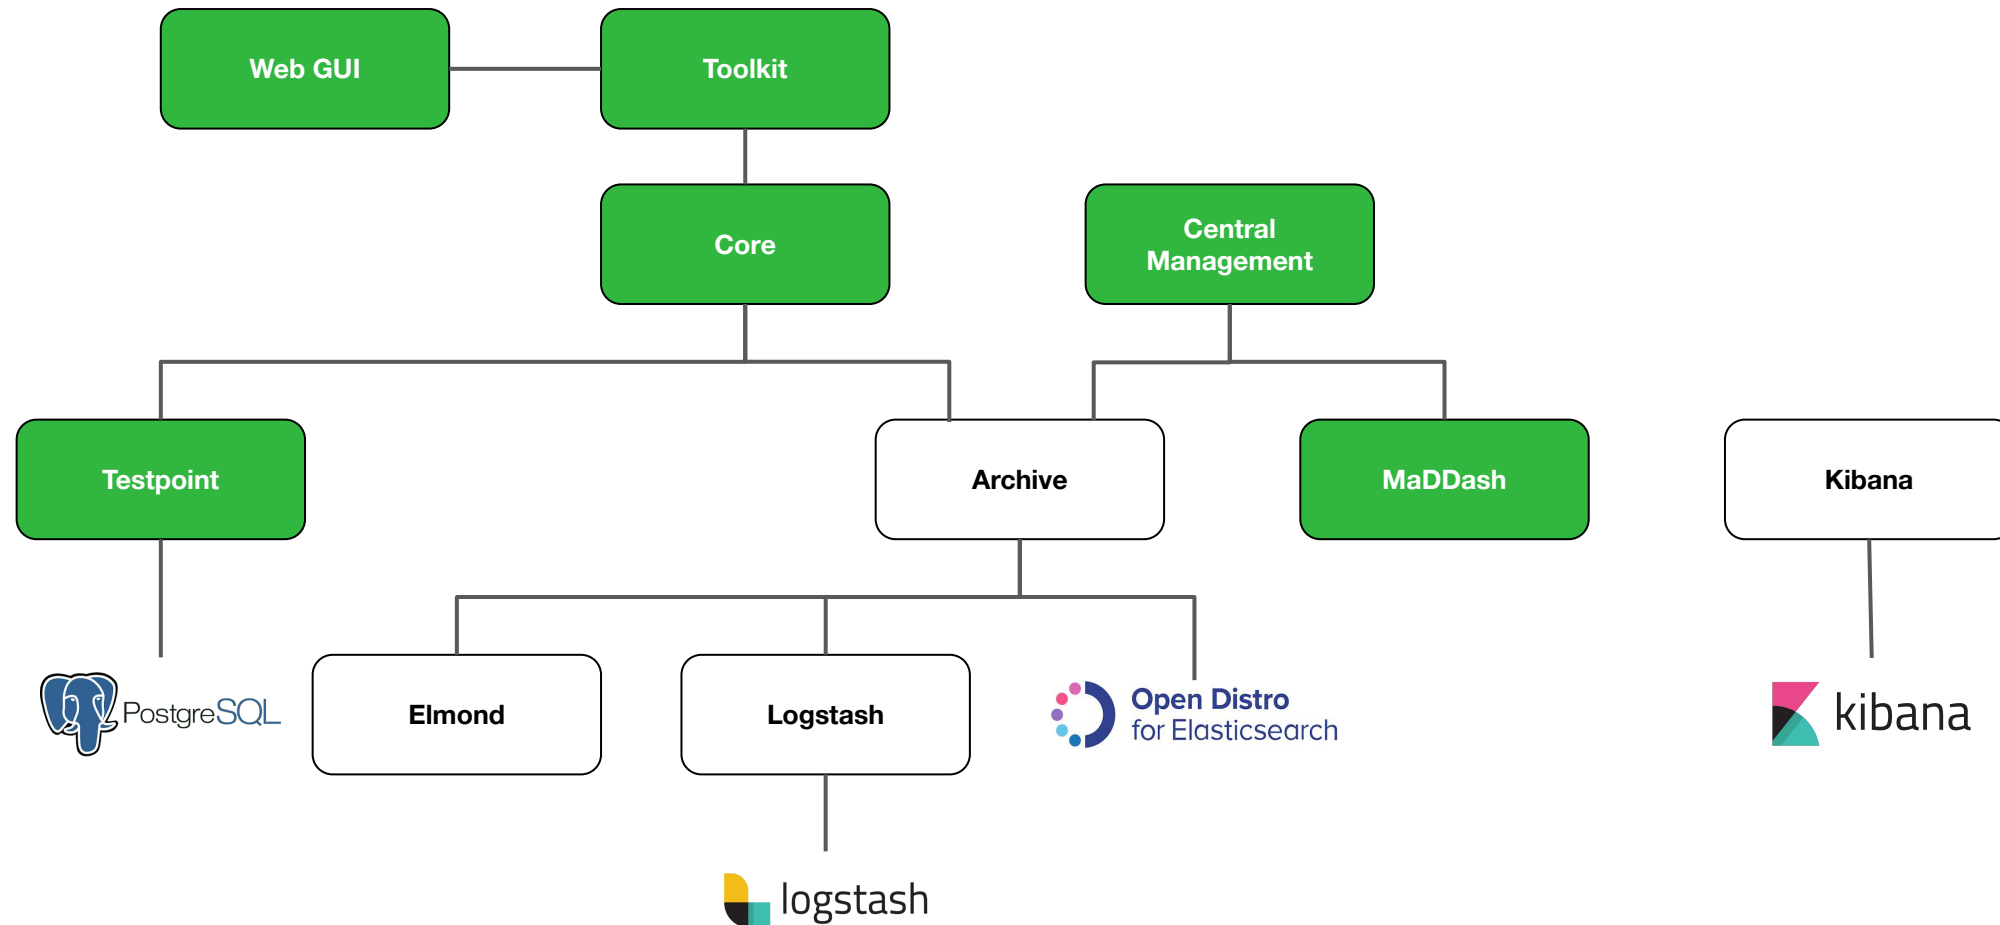
- Multiple new parts, but not a radical shift in architecture
- New components map to existing bundles



# 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
  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)

```
## Logstash environment variables.  
log_level=info  
elastic_output_host=https://localhost:9200/  
elastic_output_user=pscheduler_logstash  
elastic_output_password=pscheduler_logstash
```

`/etc/perfsonar/logstash/logstash_sysconfig`

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

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


## 5.0 Drawbacks

- Unfortunately, there are some
- Data history challenge
  - Very difficult to move data from Esmond to OpenSearch
  - Possibility to keep both backends up, during a transition period

# Other features of 5.0

- Various pScheduler improvements and some new plugins
- Some Toolkit UI improvements
- pSConfig Web Admin (PWA) changes
- Optional packages if you want
  - to keep Esmond
  - to use Kibana

# perfSONAR



Thanks icon by priyanka from The Noun Project

## Thanks!

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

*perfSONAR is developed by a partnership of*



ESnet



INDIANA UNIVERSITY

