



perfSONAR 5.0

Preparing for the perfSONAR 5.0 Release

Lætitia A Delvaux, Szymon Trocha

Poznań Supercomputing and Networking Center (PL)

WP6 T3

GÉANT Infoshare, VC

31 March 2023

Public (PU)

perfSONAR

perf5.0SONAR

Preparing for the perfSONAR 5.0 Release

perfSONAR is developed by a partnership of



ESnet



perfSONAR 5.0 will be released

April 17th



What we'll cover today

- What's in the release?
- Why upgrade to perfSONAR 5.0?
- What will happen to my old data?
- What other information do I need to know?
- What can I do between now and April 17th?
- What is coming after 5.0?




The path to perfSONAR 5.0

- perfSONAR 4.0 was 6 years ago...exactly (April 17, 2017)
- perfSONAR 5.0 Beta first released on May 31, 2022
- The largest change is the replacement of the default archive where measurements are stored

The Old: 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
 - Many old libraries that need large rewrites to update
- Community successes from which we can learn
 - WLCG, NetSage, Stardust, others



The New: OpenSearch

- Derivative of Elasticsearch with friendlier licensing
- Default perfSONAR bundles will rely on OpenSearch
- Will maintain compatibility with Elasticsearch for those with existing installations

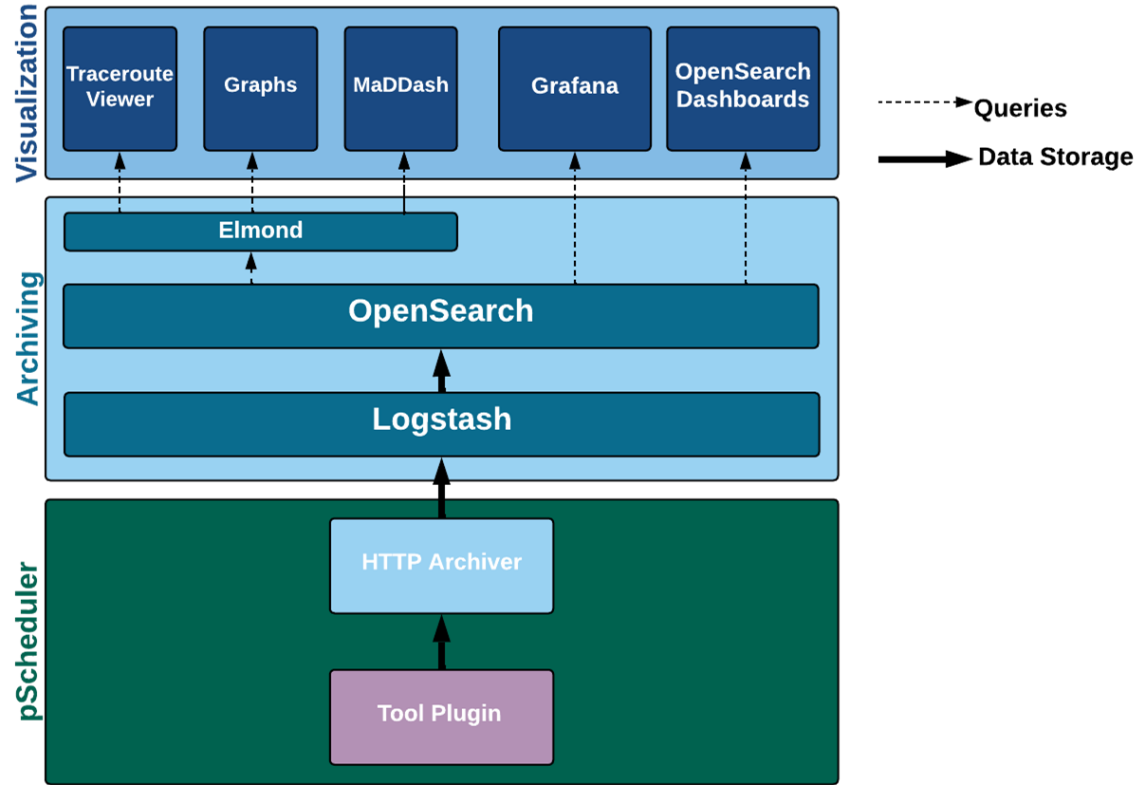


The New: How it fits together

Elmond: Backward compatibility interface for queries

OpenSearch: Stores the data

Logstash: Adds additional metadata and formats for storage



What are the benefits of upgrading?

- OpenSearch enables some new capabilities
 - You can now integrate perfSONAR data in interface like Grafana directly. Try our MaDDash-like plugin <https://grafana.com/grafana/plugins/esnet-matrix-panel/>
 - New API options including a SQL interface for getting at data
 - Foundation for many new features in future releases like better viz, alarming and more
- Ubuntu 20 support!
 - More on other OSes later
- Convenience features
 - configure-archives option now enabled by default in pSConfig
 - Host info JSON service now available on testpoint - if you are in WLCG dashboards this make testpoint a viable option
- Toolkit UI Improvements
 - pScheduler Plot command integrated in UI
 - Minor adjustments such as a logout button and more



What are the benefits of upgrading (continued)?

- New pScheduler Plug-ins
 - **Tests:**
 - MTU (mtu)
 - pScheduler Response Time (psresponse)
 - WiFi BSSID Scan (wifibssid)
 - 802.1x Authentication (dot1x)
 - DHCP Response Time (dhcp)
 - **Tools:**
 - halfping - tests latency by halving the return value of the round-trip time as determined by ping.
 - **Archives:**
 - postgresql - Inserts results into a column in a table
 - tcp - Open a connection to a TCP port, send result, disconnect
 - udp - Sends results in datagrams (those that will fit in them)



Example: Grafana

TODO: Delete this text

See <https://34.28.56.245/>

- This is a Grafana instance pointing directly at OpenSearch on a 5.0 toolkit.
- Highly customizable using standard Grafana plugins
- Show:
 - Data Source: <https://34.28.56.245/datasources/edit/gKB9-sBVk>
 - Overview: <https://34.28.56.245/d/tqaA0QfVvk/perfsonar-overview?orgId=1>
 - Bring up edit panel on MaDDash-like panel, show how you can customize things like stats used, threshold, box size, etc.
 - Click on one of MaDDash-like box and show some graphs
- Be clear this is not something installed by pS packages, but that anyone can do and cookbook will be provided. 5.1 and future versions will look to further streamline these types of activities and look forward to what people come up with on their own.



What will happen to my old data?

- The answer depends on which perfSONAR bundle you are running
- **Testpoint:** Good news! Testpoints don't store data so these hosts are unaffected.
- **Toolkit: Your host will be auto-updated and data will NOT be migrated.** It will remain on disk and you will have the option to move to a dedicated esmond instance on another host.
- **Central Management/Stand-alone Esmond:** Your host will not be automatically updated to OpenSearch. Clients can continue writing and you can control when you migrate.



Why aren't you migrating my Toolkit data?

- Esmond since its inception has aged-out data after 6 months, so we are talking about a relatively small set of data
- Running Esmond and OpenSearch in parallel on the same host is more memory and process intensive than the typical Toolkit can handle...in addition to having to support a theoretical migration process
- We will provide instructions for moving Esmond data to a new host and bringing up a standalone Esmond if you really need it.
- **If you are not comfortable with the above, you have the option to disable auto-updates**



What about my central Esmond instance?

- These packages will NOT be automatically upgraded (i.e. your install will be untouched)
- We realize the typical central Esmond instance has more clients to serve and may require greater control of the migration process.
- While you can't migrate data, you can run Opensearch and esmond in parallel on different hosts for a period of time. Simply include both archives in pSConfig templates.
- **When you are ready to move, update your pSConfig definition to only point at the HTTP archiver writing to Logstash/Opensearch**



pScheduler Limits Changes

- **This is an important one that may be easy to miss!**
- Support for “test” type limit will be dropped.
- If you are using the default Toolkit limits file, then there is no action required
- If you have made any changes to your limits file, then you will need to manually update them.
- **pScheduler will stop running on update if your limits are not valid!**
- Run the helper script here to check your limits:

<https://github.com/perfsonar/pscheduler-limit-checker-for-5.0>



ISO Image No Longer Supported

- We will no longer support the 5.0 ISO image
- Users can install the base operating system of their choice and use yum or apt to install perfSONAR
- The ISO image requires special care and feeding for installs to work in its boot environment
- The perfSONAR packages have evolved to the point that installation is relatively straight-forward without the ISO



Packaging Changes

- **perfsnar-centralmanagement retired:** This was the bundle that installed esmond and maddash on same host
- **New perfsnar-archive package:** This installs and configures OpenSearch and Logstash for you. If you want a replacement for centralmanagement then install *perfsnar-archive* and *maddash*
- **New optional perfsnar-dashboards package:** Installs OpenSearch Dashboards (equivalent of Kibana) for browsing data.



What can I do between now and April 17th?

- **Check pScheduler limits file** - See <https://github.com/perfsonar/pscheduler-limit-checker-for-5.0>
- **Review the data strategy proposed** - If you think you'll want a standalone esmond instance, you can start on that now. While not recommended, if you feel the need then you can also disable auto-updates
- **Install the beta** - If you want some more experience with 5.0, try our beta https://docs.perfsonar.net/release_candidates/5.0.0/install_rcs.html
- **Remember we're here to help** - The perfSONAR development team understands this is a big change and every scenario may be slightly different. Let us know if you have questions and we'll do our best to answer.
- **Check out the new Lookup Service UI** - <https://stats.perfsonar.net>



What's after 5.0?


- Additional OSES by Early Summer:
 - Alma 8, Alma 9 (works with other flavors of EL as well like Rocky)
 - Ubuntu 22
 - Debian 11
- Beyond Summer:
 - Grafana visualization and MaDDash Replacement (In Progress)
 - See <https://grafana.com/grafana/plugins/esnet-matrix-panel/>
 - Threaded iperf3 (In Progress)
 - Threaded powstream (In Progress)
 - pSConfig in python (In Progress)
 - Revamped PWA (Early Design Phase)



Summary

- April 17th is the 5.0 Release.
- Toolkit data will NOT be migrated to Opensearch
- If you run a central Esmond instance, you can start thinking about how you want to upgrade but don't have to have it figured out by April 17th
- Check your pscheduler limits
- Additional OS support will be available not long after April 17th
- We are here to help you





Questions and Answers

Question and answer icon by iconsphere from The Noun Project





Thank You

www.geant.org



The scientific work is published for the realization of the international project co-financed by Polish Ministry of Science and Higher Education from financial resources of the programme entitled "PMW"