

The LSST Data Facility Service Monitoring Framework

J.D. Maloney | NCSA



Large Synoptic Survey Telescope



The Vision



Tools Chosen



Implementation of Tools




Looking Forward



The Vision

Increase Visibility:

- Have a view of all systems, services, and applications operated by and relied on by the LDF
- Be able to view historical trends for event correlation and capacity management
- Use available metrics data to project future needs
- Have rich alerts that gave more information
- In certain cases attempt to catch issues before they create trouble to mitigate a service interruption





Scalable Deployment:

- Rolling out new machines requires little to no additional effort
- Removing hardware and services is simple and quick
- Still leverage configuration management, but in a more simple fashion
- Greater portability between systems and environments
- All monitoring data aggregated in one place for ease of access and replication



Organization Integration:

- Multiple views of metrics geared toward different operational roles and use cases
- Federated access to different groups
- Integration with existing authentication/account management
- Allow viewers to build their own views of the data gathered



Tools Chosen

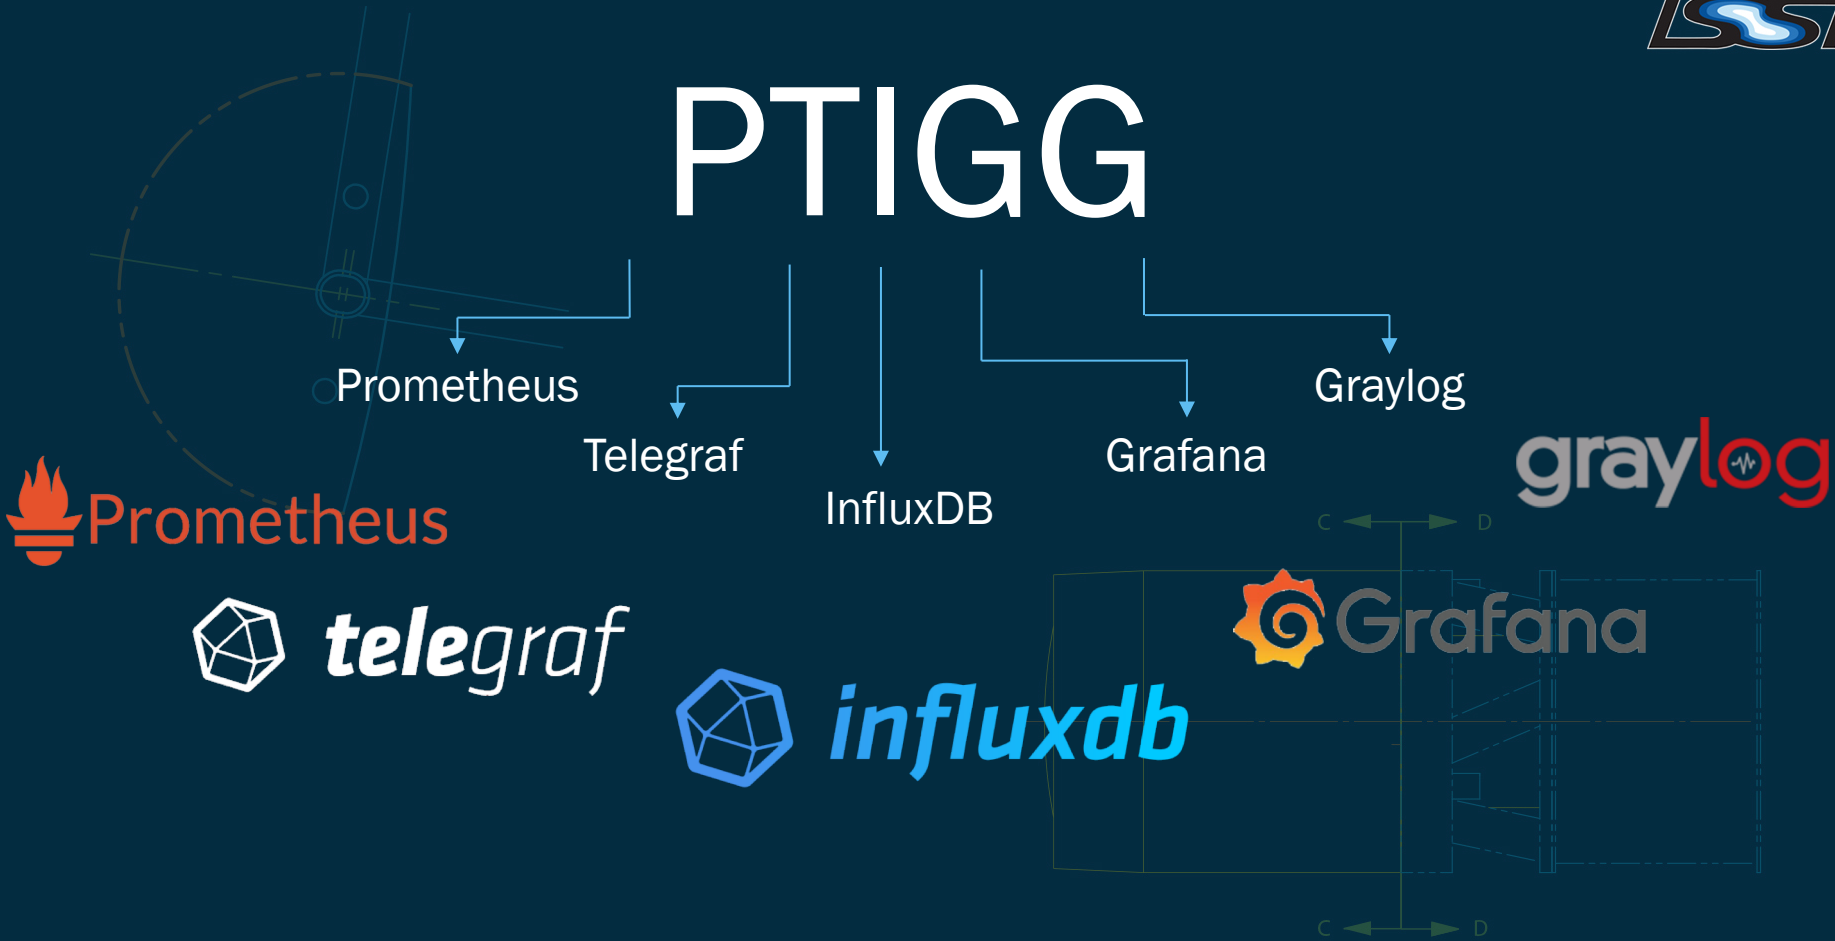


Broader Vision:

- This monitoring stack is a leveraging of common tools that are already in use across many of NCSA's other production systems and clusters
- NCSA's engineering sub teams (systems, storage, networking, and security) are all leveraging these tools in their environments
- Tools are all available free and open-source and have a strong community of users



PTIGG



Metric Collection:

- Telegraf gathers metrics on bare-metal systems and services
 - 1 RPM and 1 configuration file built by hiera
 - Dozens of plugins for various metrics and services
- Prometheus gathers metrics on containers and applications
 - Static binary that is easy to build and deploy across many environments
 - Very light weight



Metric Storage:

- Optimized for time-series data that streams in from thousands of sources
- Telegraf has built in compatibility
- Supports the Prometheus write API as of v1.4
 - Allows for long term storage of metrics on disk
- Supports down-sampling, replication, and continuous queries
- Clustering solution available for High-Availability
- Can easily scale to over 100,000 entries per second



Metric Visualization & Alerts:


- Powerful time-series visualization tools with plug-in support for further expansion
- Teams support to allow separate view points and dashboard folders for organization
 - LDAP support for team groups coming soon
- Alerting built in with many methods of broadcasting alerts
 - eg. Using Slack, email, PagerDuty
- Light-weight, super easy to deploy and maintain



Log Analysis & Alerting:

- Built on top of the popular and well established Elastic Search database
- Good out of the box configuration
- Built in support for authentication
- Built in alerting based on rules that can notify admins of issues
- Helps with event correlation and anticipation

graylog






Implementation of Tools

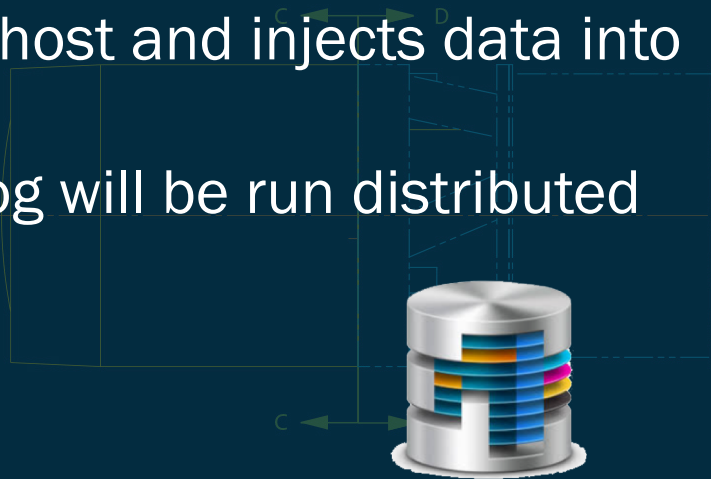
Collector Deployment:

- Telegraf is deployed via Puppet configuration management
 - Configuration file is built by hiera based on node role
 - Machines start sending metrics from installation onward
- Prometheus is deployed with containers to gather metrics
 - Pushes data into InfluxDB for longer term storage
- Developing guidelines and code examples for application-level integration with service monitoring



Metric Storage Deployment:

- InfluxDB deployed via puppet on bare-metal host
 - Databases live on local ZFS file system
- Metrics are backed up nightly to parallel file system
- Database will be replicated to Base Site for geo-redundancy and availability
- Prometheus server runs on same host and injects data into InfluxDB
- Elastic Search database for Graylog will be run distributed across multiple VMs



Metric Visualization/Alert Deployment:

- Multiple Grafana instances are run in VMs
 - All access same backend databases
 - Allow for development of dashboards outside of production environment
 - Can test future Grafana versions and features
 - Can scale resources based on demand
- Graylog will also run in the VM environment for many of the same reasons as above
- These VMs are also under puppet control for easy provisioning



Looking Forward:



- This monitoring framework is up and running on current LDF services and systems
- As services and systems are brought online they will be incorporated
 - Additional dashboards and views will be added
 - Further expansions of developer facing alerts
 - Data gathered will be used to project future need to guide acquisitions



Available Today

Check out dashboards that are available today!

<https://monitor-ncsa.lsst.org>



Large Synoptic Survey Telescope



Questions?



Large Synoptic Survey Telescope

