



DPM: Future Proof Storage



Alejandro Alvarez, Alexandre Beche, Fabrizio Furano, Martin Hellmich, Oliver Keeble, Ricardo Rocha

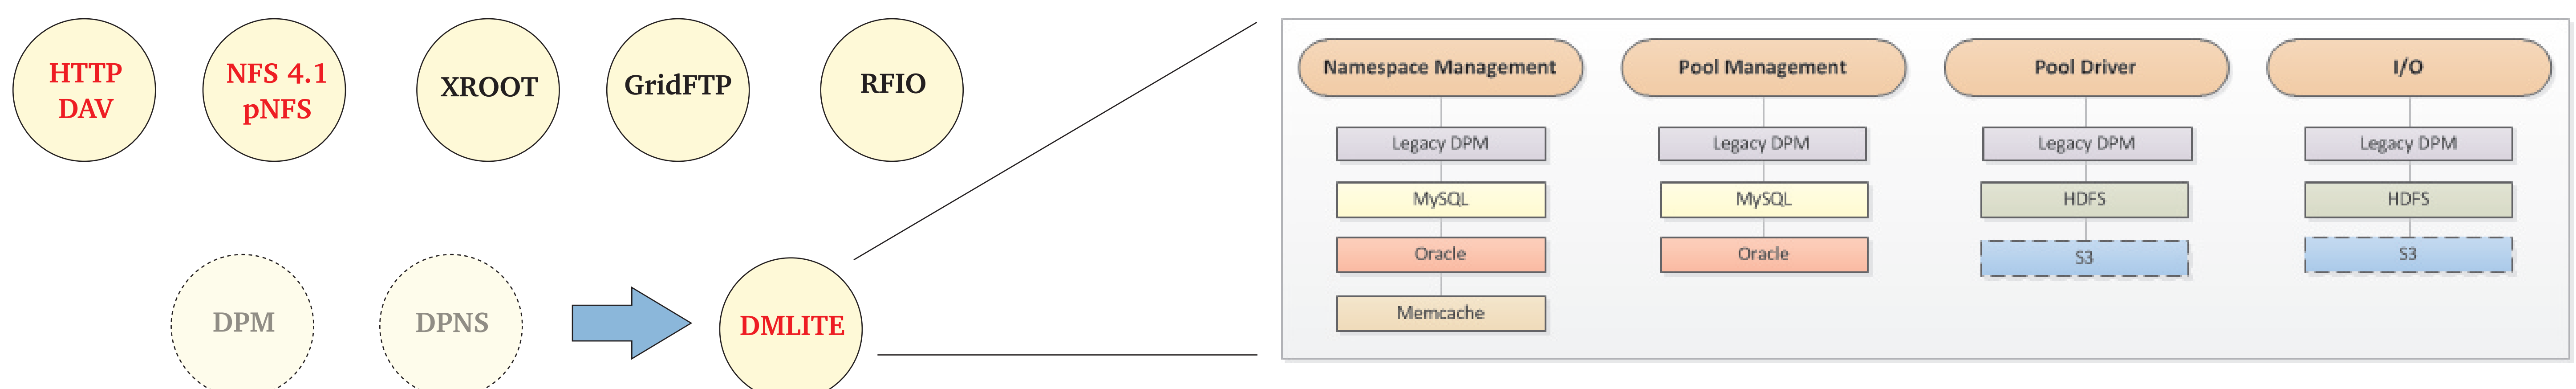
Lightweight Scalable Standards Based Extensible

Summary

The main objective of the Disk Pool Manager (DPM) is to provide a lightweight, standards based and grid-aware data storage solution. It simplifies the life of users and administrators, while offering good performance and allowing easy integration with new tools and systems.

Preparing the future: standards and dmlite

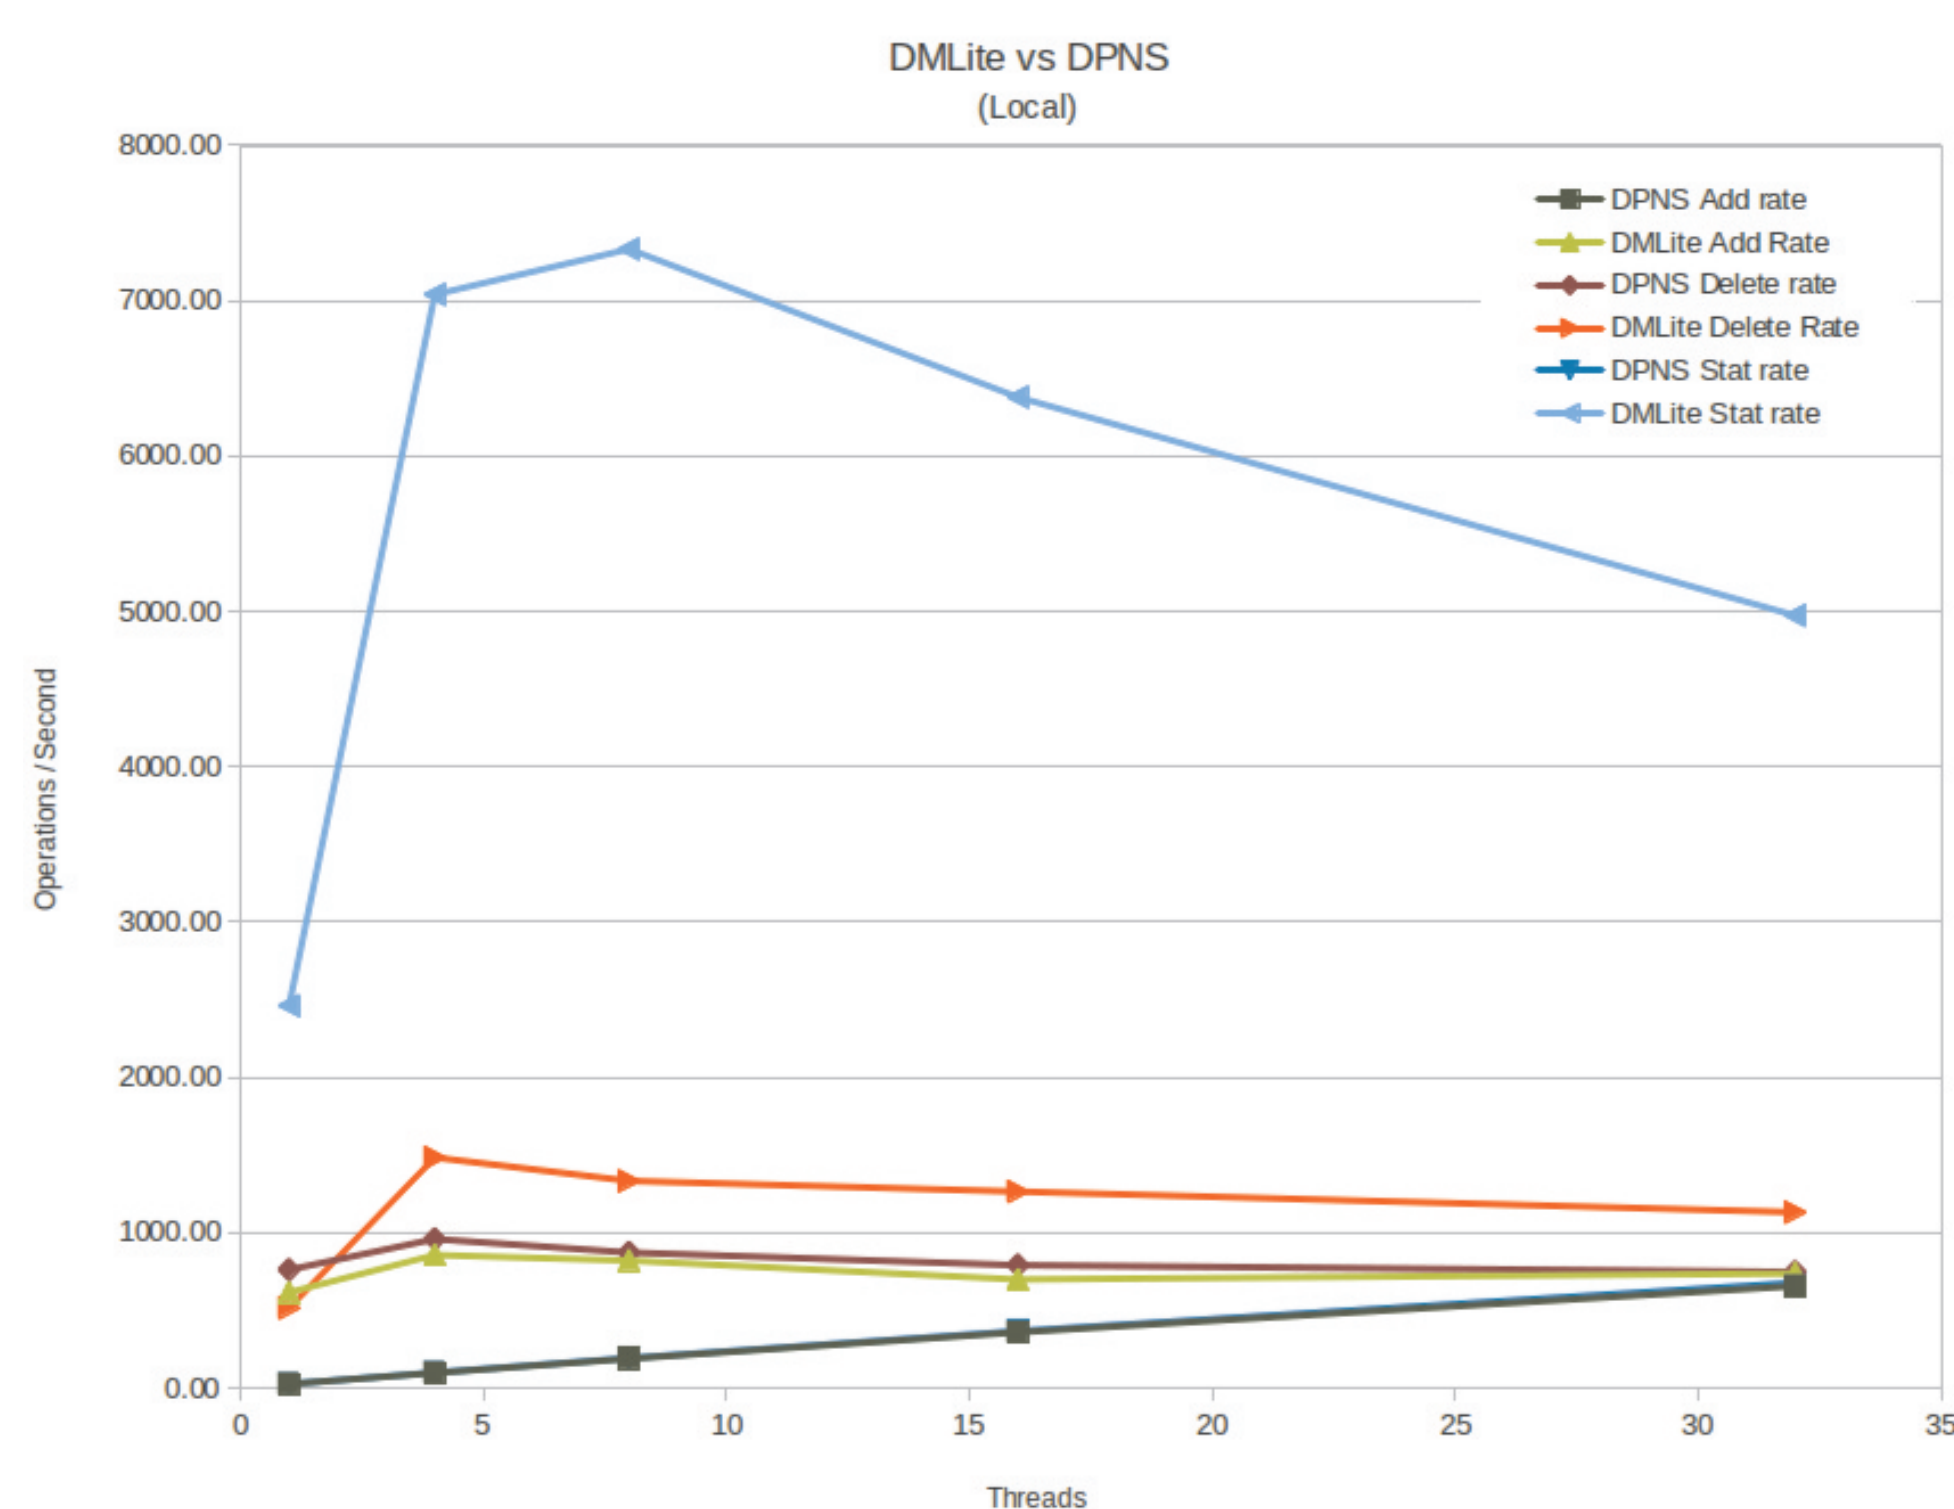
We have added standard protocol compliant frontends to the previous entry points of our system, basing them on standard building blocks as much as possible. For adequate performance, we have refactored the native daemons. The result is a plugin-based library called dmlite, which also opens our system to other existing tools and storage backends.



Plugins can be stacked at will. Existing implementations include Native DPM, MySQL, Oracle, **Memcache** and **Hadoop/HDFS**. **S3/cloud** coming very soon.


Performance Comparison

Offering the frontends direct database access (no daemon contact) has given us an improvement of up to 7 times in stat operations - the most frequent.



Metadata Performance with memcache

Plugin implementation for Namespace Management, offering even greater metadata access performance.



Best single node performance with a local memcached instance
Higher performance possible with multiple instances (horizontal scalability)

Monitoring and Configuration

DPM offers an extensive set of **Nagios** probes to track service status, metadata queries, system usage, and much more. We also provide **Puppet** manifests for easier deployment and service configuration.



Testing and Documentation

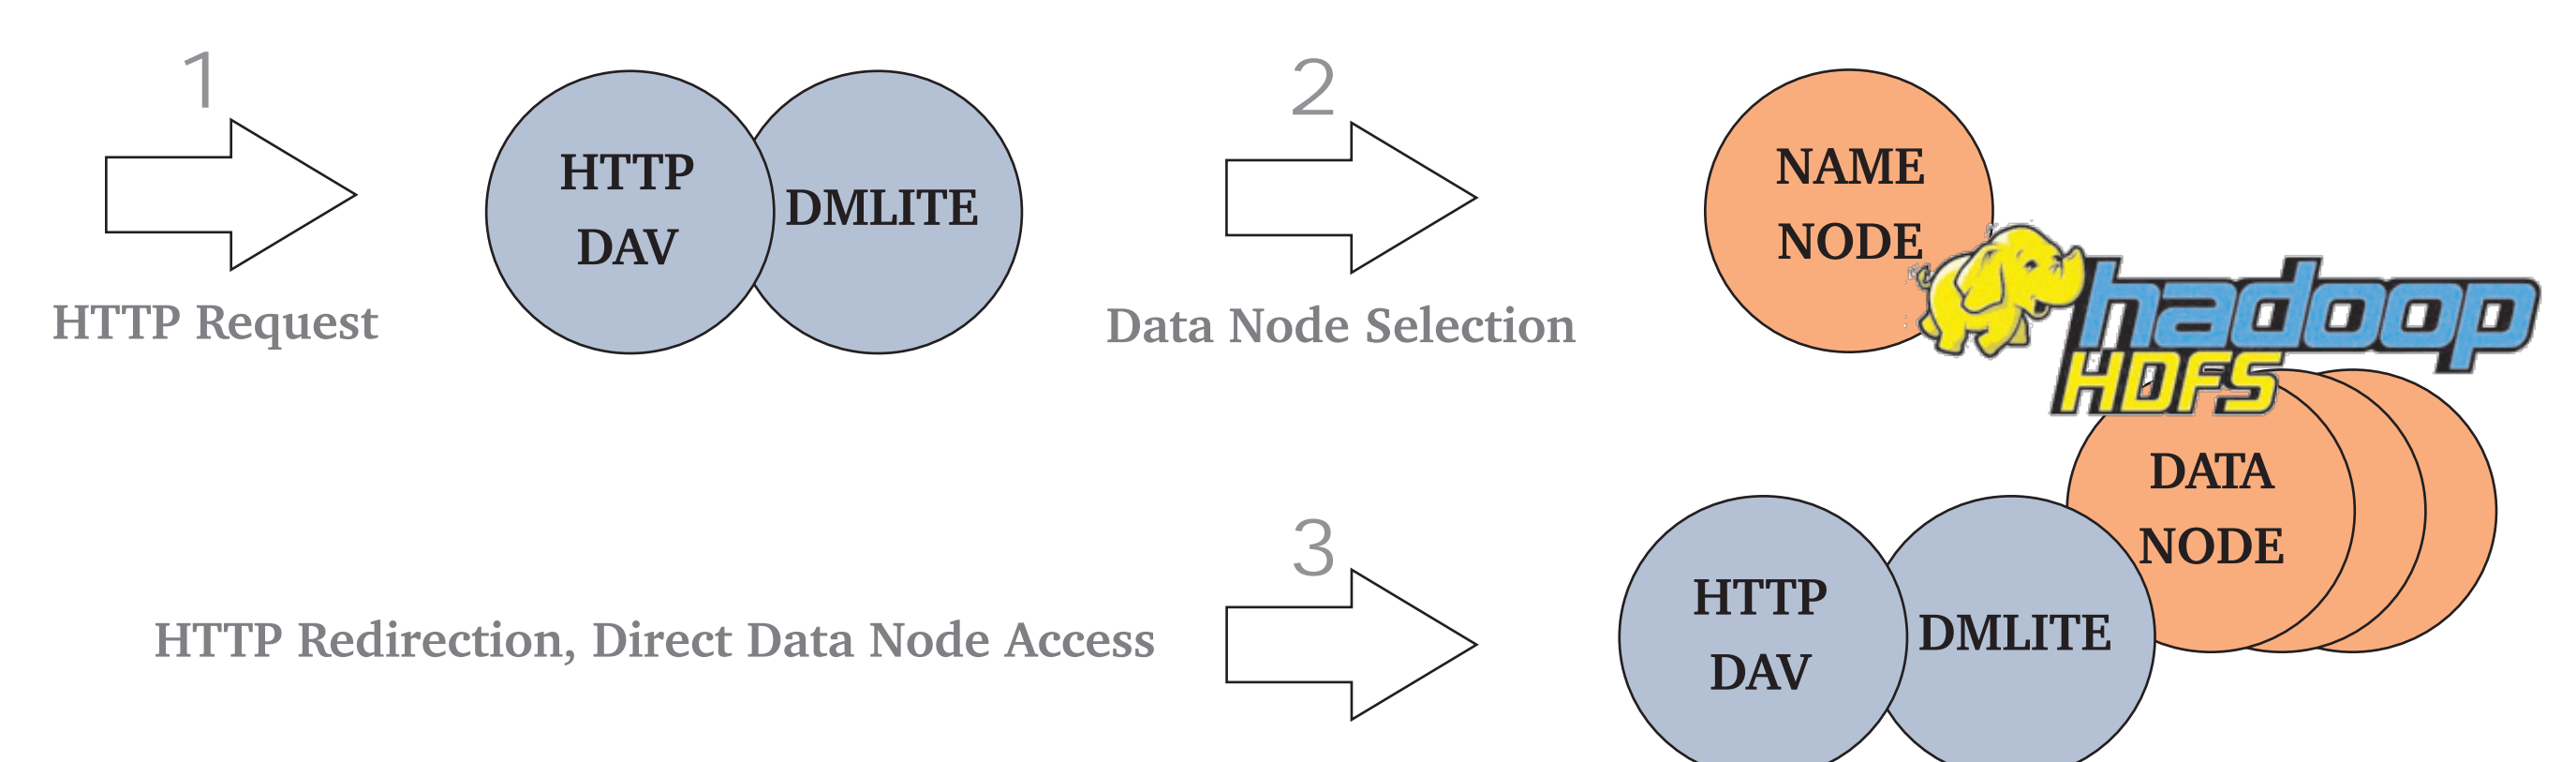
All new components follow a "release early, release often" principle - help us testing by using our beta repository. Check our website for further documentation, webinars and tutorials.

Acknowledgements

We would like to thank GridPP and ASGC for their extensive validation and testing efforts, and for their software contributions.

Extensibility, integration - HDFS, S3, ...

With dmlite sites can deploy new pools with distinct storage backends, under a single namespace. Currently HDFS and an early S3/cloud prototype.



Thanks to dmlite and the nice HDFS APIs, the implementation took 2 weeks
We envision integration with other popular storage systems

Conclusions and Future Work

DPM has come a long way, and it stays as the most widely used grid storage system. With the most recent developments we have greatly improved performance and reduced the load of packaging, deploying and configuring a complete storage system. Future work will continue focusing on improved performance and the introduction of new features by reusing standard and well supported building blocks.

