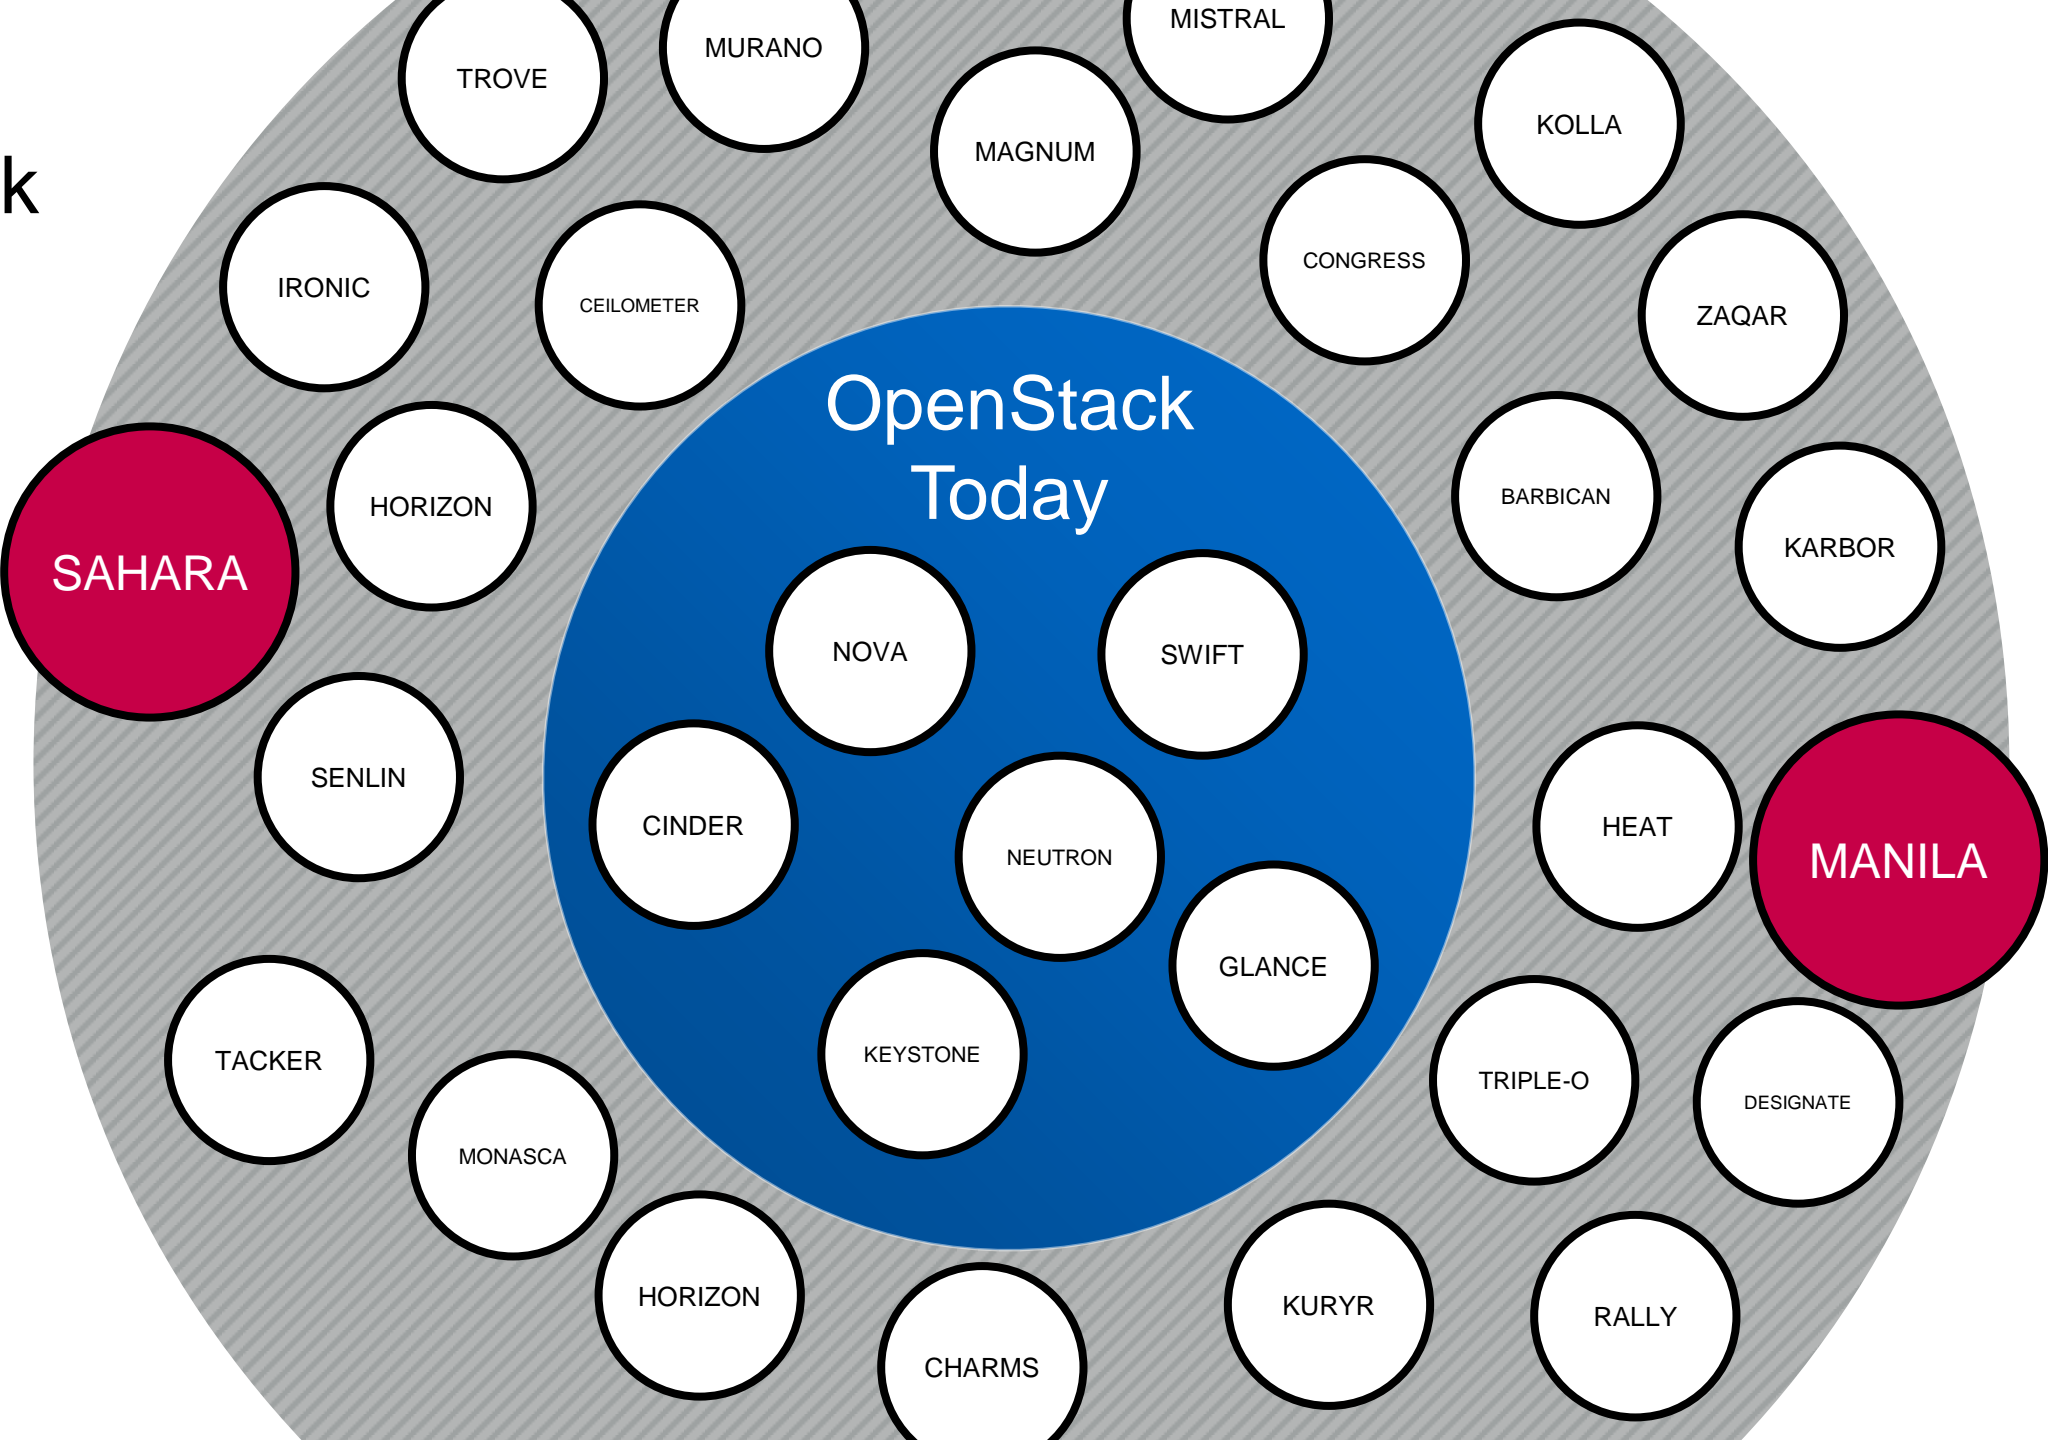


Using OpenStack Sahara & Manila to run Analytics against NFS Datastores

January 2017

Presenter: Chris Madden
NetApp

OpenStack
Keeps Growing!



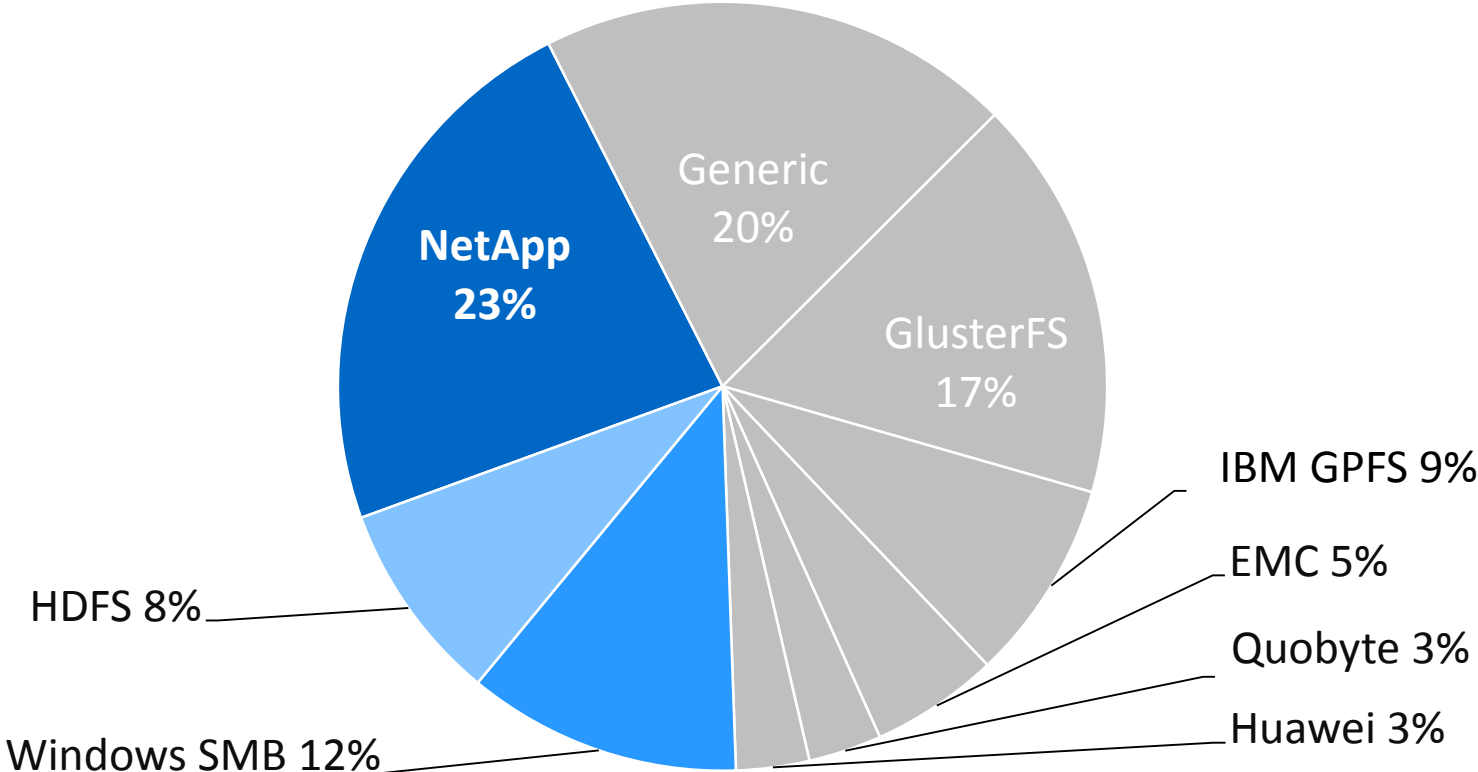
Lots of data in shared file systems



storage sold
for Shared File Systems
(IDC)

OpenStack Foundation User Survey Results

Shared File System Drivers in Use (April 2016)

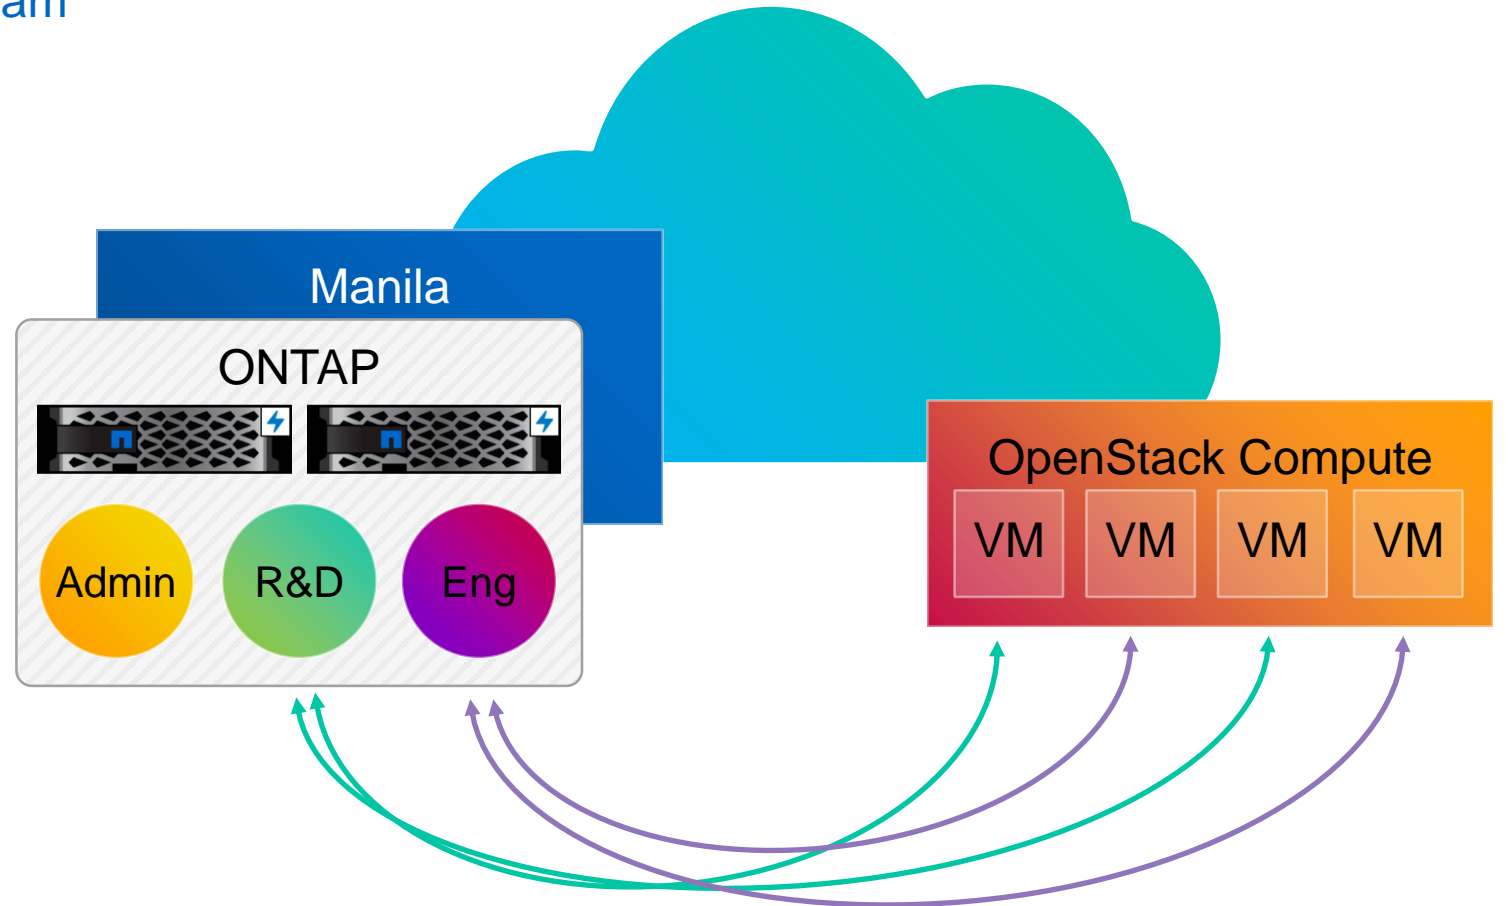


- NetApp is the founder of, leading contributor to, and most widely deployed option for Manila

OpenStack Manila

The OpenStack Shared File Service Program

- Manila delivers file systems as a service:
 - NFS, CIFS, HDFS, and more
- Production-ready in Liberty
- Supports a wide range of network topologies



Hadoop Clusters are not easy!

Same for Spark, Storm, etc...

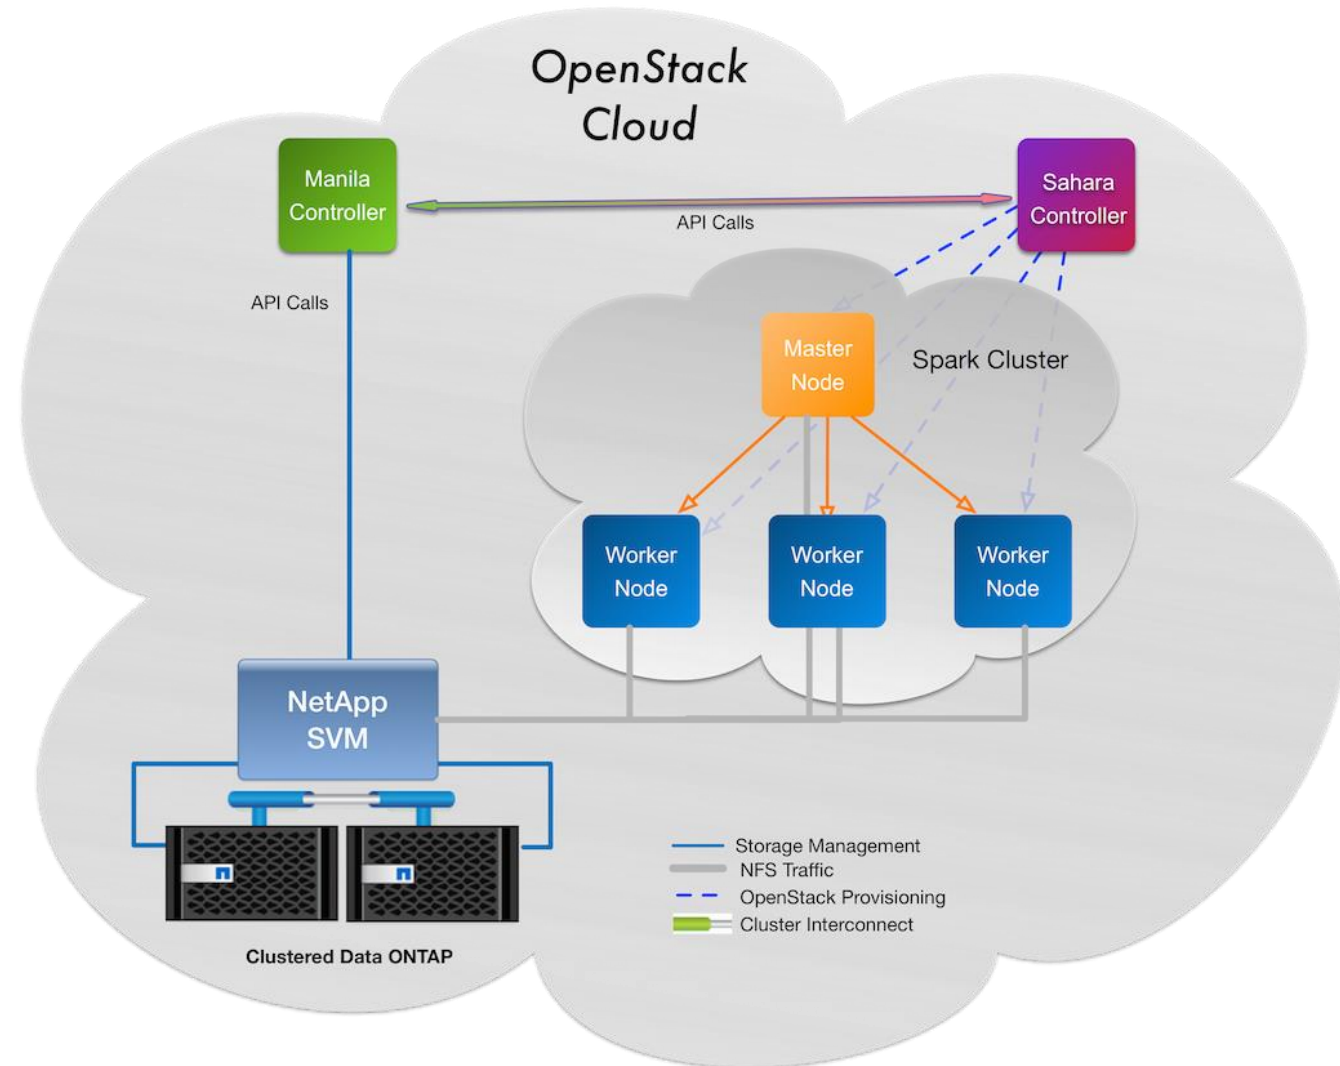
- Difficult to configure & maintain
- Adaption to changing demands
- Hardware headaches
 - Commodity hardware is cheap but requires more work
 - Reliable hardware is expensive
- Hadoop dev is very difficult without a real cluster



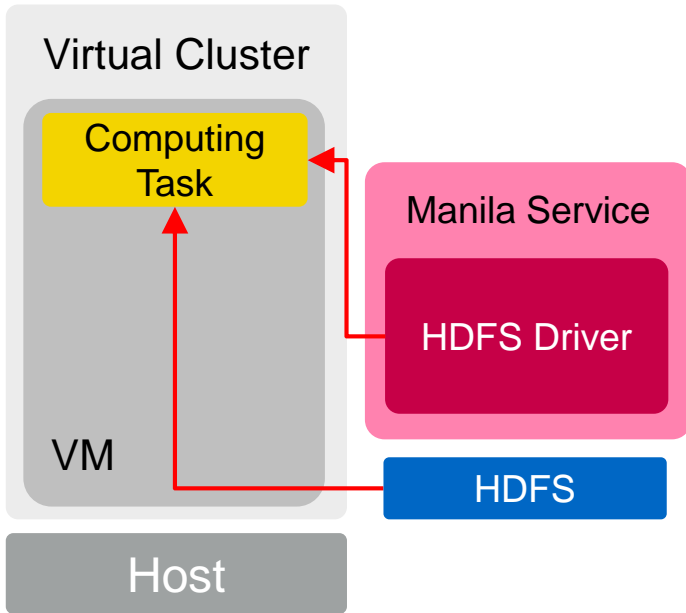
OpenStack Sahara

The OpenStack Data Processing Program

- Put Hadoop and alike in a OpenStack cloud!
- Then have easy-to-use, standardized interfaces:
 - To create clusters (reliably and repeatedly)
 - To scale clusters
 - To run data processing jobs
 - On any popular data processing framework
 - With sensible defaults that just work
 - And sophisticated configuration management for expert users
- That's OpenStack Sahara!

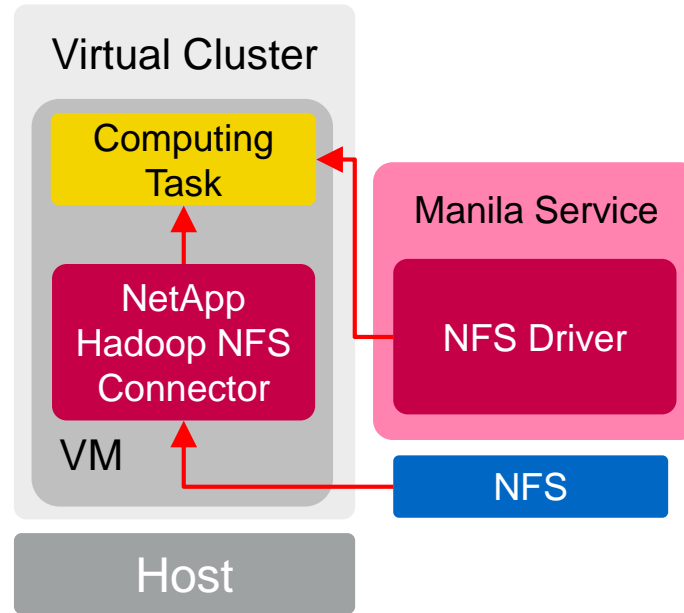


Technical Overview



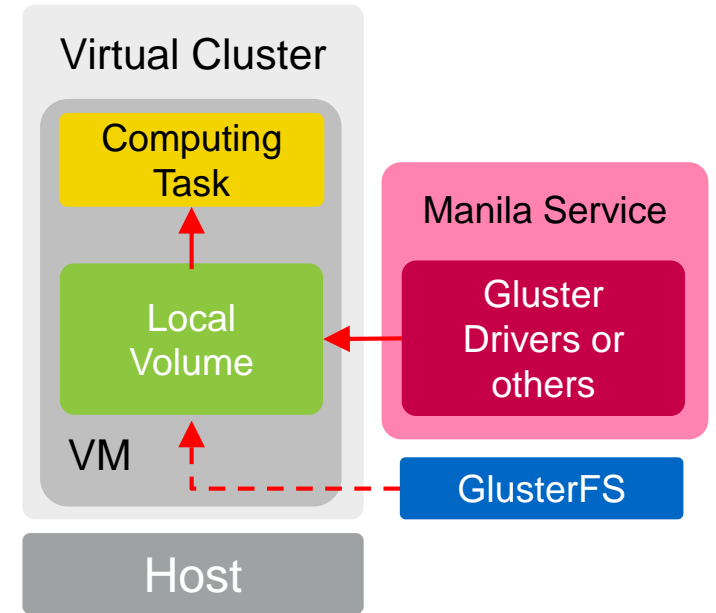
External HDFS via Manila

- Sahara can support external HDFS by using the HDFS driver in Manila



NFS directly via Manila

- NetApp Hadoop NFS Connector brings the NFS capability into Hadoop



Local Storage in Manila

- Use local storage in Hadoop to remote mount any type of storage in Manila

Summary

- Manila adds Fileshares-aaS to OpenStack
- Sahara adds Analytics-aaS to OpenStack
- NetApp NFS Connector for HDFS allows access of NFS data through HDFS protocol
- All together:
Easy analytics against existing datasets in file shares
- Resources
 - https://www.youtube.com/watch?v=Tf_8v0OvR80
 - <https://github.com/NetApp/NetApp-Hadoop-NFS-Connector>
 - <https://www.netapp.com/us/media/tr-4464.pdf>

