

Using OpenStack Sahara and Manila to run Analytics against NFS datastores

Tuesday 31 January 2017 10:10 (15 minutes)

The OpenStack Shared File Systems project (Manila) provides basic provisioning and management of file shares to users and services in an OpenStack cloud. The OpenStack Data Processing project (Sahara) provides a framework for exposing big data services, such as Spark and Hadoop, within an OpenStack cloud. Natural synergy and popular demand led the two project teams to develop a joint solution that exposes Manila file shares within the Sahara construct to solve real-world big data challenges. This contribution examines common workflows for how a Sahara user can access big data that resides in Hadoop, Swift, and Manila NFS shares. It also demonstrates how you can replicate Manila fileshares on-prem (in your private OpenStack cloud) to public clouds (such as AWS, or Azure),

Author: STUIJT, Tom (NetApp)

Co-author: Mr SIEBLER, Clemens (NetApp)

Presenter: STUIJT, Tom (NetApp)

Session Classification: Storage technologies