

dCache - data access optimization in a hybrid cloud

Wednesday, 3 April 2019 16:30 (30 minutes)

For over a decade commercial providers offer their computer resources as public clouds. With a couple of clicks one can build and run a full data center without any local infrastructure. While this is a very attractive approach, high costs and/or legal aspects force sites to run their own infrastructure, often as a private cloud, though. However, such local resources may be insufficient if user demand can't be foreseen in advance.

One option to cope with such peak loads is to build a hybrid cloud by combining the local instance with external resources or public clouds. However, jobs running on additional CPU resources may provide low efficiency when remote data access is required. Moreover, additional network traffic may produce higher network utilization, thus higher overall costs or even network bandwidth starvation.

For years, dCache.ORG has provided robust software, called dCache, that is used at more than 80 Universities and research institutes around the world, allowing these sites to provide reliable storage services for the WLCG experiments and many other scientific communities. The flexible architecture of dCache allows running it in a wide variety of configurations and platforms. In this presentation we will show how to deploy a distributed dCache instance, that dynamically adds caching nodes to provide local data access and an optimized resource usage in a hybrid cloud.

Primary authors: Mr ROSSY, Albert (FNAL); Mr LITVINTSEV, Dmitry (FNAL); Mr STAREK, Jürgen (DESY); Dr SAHAKYAN, Marina (DESY); Dr ADEYEMI, Olufemi (DESY); Dr FUHRMANN, Patrick (DESY/dCache.org); Dr MILLAR, Paul (DESY); Ms YASAR, Sibel (desy); Mr MKRTCHYAN, Tigran (DESY)

Presenter: Mr SCHUH, Michael (DESY)

Session Classification: Data Management & Big Data

Track Classification: Data Management & Big Data