



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



New use cases and usage models for Cray DataWarp

Maxime Martinasso, CSCS

Cray DataWarp

Idea

Bring SSDs closer to the compute nodes

DataWarp modes

- cache between application and parallel FS
checkpoint/restart, periodic output, libraries
- scratch or additional /tmp
- shared storage: multiple jobs
- swap: over commit memory

DataWarp issues

No client-side cache, low performance

Data Science use cases

Data exploration

- manipulate and visualize a large set of unstructured data
- interactivity
- stream data (in/out)
- multiple users/systems
- Spark, Jupyter, streaming, many other tools
- HDFS, Alluxio, HBase, ...

Criterion

- performance: interactivity, small delay response
- functionality: deploy own file system or database
- connectivity: stream, data updates from outside
- connectivity: web enabled access to storage

Deep Learning use cases

Training Neural Network

- provide fast access to common training data sets (ImageNet, YouTube-8M, ...)
- data is persistent and read-only
- data sets are large but often made of small files (images)
- multiple users/systems
- CNTK, TensorFlow, Caffe, MxNet, ...

Criterion

- performance: optimize for read bandwidth and meta-data accesses (not common HPC FS requirements)
- persistence: data sets remain available for months/years

Complex scientific workflows

Large projects such as the Human Brain Project

- users become virtual organizations (VO)
- data is generated by different sources
- manipulate large structured data, periodically select+save
- pre-analysis on subsets of data applied to the entire data set
- simulation is driven by data analysis
- data storage integrated in the workflow tool (web portal)

Criterion

- functionality: Posix limits the usage of VO, needs special FS
- connectivity: data are shared among VOs
- performance: data should remain close to compute
- performance: workflow restart time should be minimum (save)

Requirements (1)

performance

- data remains close to compute
- high read/write bandwidth and meta data accesses
- interactivity, small delay response

functionality

- deploy own file system or database

Requirements (2)

connectivity

- web enabled data access
- data updates from outside (stream)

persistence

- data sets remain available for large period of time
- permanent point of access for sharing data among “users”

Vision on Datawarp

DataWarp as a Data enabler for HPC

- Deploy your own data manager (fs, database, . . .)
- Connect remotely and share the data access
- Manipulate data interactively and get fast response

