



Berner Fachhochschule
Haute école spécialisée bernoise
Bern University of Applied Sciences

Geo Data Storage

ITS Meeting

Thursday, 0th June 2015 16:00-17:45

BFH – Conference Room C.O.R.A. (Dammweg 3, 1. 1st Floor, Bern)

Daniel Baumann <daniel.baumann@bfh.ch>

IT System Engineer, Infrastructure Team

Overview

- ▶ Storage System
- ▶ Geo Data as a Service
- ▶ Future Plans?

Storage System

Storage (1)

Legacy Storage

- ▶ single storage systems
- ▶ replication to other single storage systems

Disadvantages

- ▶ **scale-up** possible only to a certain point
- ▶ is not flexible, often mixes up all the layers (blackbox instead of KISS)
- ▶ hardware and software vendor lock-in

Storage (2)

Distributed Storage

- ▶ many individual systems forming a cluster
- ▶ self-healing and self-managing

Advantages

- ▶ **scale-out** to infinity and beyond :)
- ▶ whitebox systems
- ▶ can be geographically distributed

Storage (3)

Software Defined Storage

- ▶ storage independent of hardware
- ▶ provisioning and management of storage through software (APIs)

Ceph (<https://ceph.com>)

- ▶ default scale-out filesystem on Linux
- ▶ basis for OpenStack Cloud
- ▶ broad Cloud and HPC user community
- ▶ scale-out up to 30PB without problems, afterwards tuning needed

Ceph OSD Nodes

Parts

- ▶ 1x Supermicro 2U Server
- ▶ 4x 10Gbit Fiber
- ▶ 1x Intel 2.4GHz Octo-Core CPU
- ▶ 1x Kingston 64GB RAM
- ▶ 2x Intel 200GB SSDs (Operating System)
- ▶ 3x Seagate 4TB HDDs (Data)
- ▶ 3x Western Digital 4 TB HDDs (Data)

Capacity

- ▶ initial 24TB brutto, 12TB netto per OSD (Ceph replica=2)
- ▶ initial 6 OSDs with total 144TB brutto, 72TB netto



Ceph MON Nodes

Parts

- ▶ 1x Supermicro 1U Server
- ▶ 2x 10Gbit Fiber
- ▶ 2x Intel 2.4GHz Octo-Core CPUs
- ▶ 1x Kingston 64GB RAM
- ▶ 2x Intel 100GB SSDs (Operating System)

Capacity

- ▶ Ceph MON/MDS/RWG for at least 3PB of OSDs



Timeline Ceph Storage Cluster

August 2015

- ▶ Hardware Setup

September 2015

- ▶ Software Setup
- ▶ initial data import via NFS and SMB Fileserver on existing VMware ESX Cluster

October/November 2015

- ▶ OpenStack Setup

December 2015

- ▶ Replacing NFS and SMB Fileservers on VMware ESX with Fileservers on new OpenStack Cluster

Geo Data as a Service

Service

Reason

- ▶ Geo Data are large files
- ▶ storage is not expensive but not free either

Service

Reason

- ▶ Geo Data are large files
- ▶ storage is not expensive but not free either

Objectives

- ▶ ensure timely data updates and imports
- ▶ ensure access to everyone
- ▶ ensure scientific usage and application
- ▶ ensure vendor neutral access to information (ArcGis, QGis, etc.)
- ▶ ensure scalability and „cheap“ extension (8'500 CHF/24 TB netto for 5 years)

Service

Reason

- ▶ Geo Data are large files
- ▶ storage is not expensive but not free either

Objectives

- ▶ ensure timely data updates and imports
- ▶ ensure access to everyone
- ▶ ensure scientific usage and application
- ▶ ensure vendor neutral access to information (ArcGis, QGis, etc.)
- ▶ ensure scalability and „cheap“ extension (8'500 CHF/24 TB netto for 5 years)

Preconditions

- ▶ centrally managed, large storage system
- ▶ procedure to organize Geo data on the filesystem

Procedure

Initial Import

- ▶ Geo Data share(s) are read-only for users on purpose
- ▶ ITS takes care about the import from TI and HAFL sources
- ▶ ITS imports „normalized“ data only

Procedure

Initial Import

- ▶ Geo Data share(s) are read-only for users on purpose
- ▶ ITS takes care about the import from TI and HAFL sources
- ▶ ITS imports „normalized“ data only

Regular Imports and Updates

- ▶ via OTRS Ticket
- ▶ ITS imports the „normalized“ data

Procedure

Initial Import

- ▶ Geo Data share(s) are read-only for users on purpose
- ▶ ITS takes care about the import from TI and HAFL sources
- ▶ ITS imports „normalized“ data only

Regular Imports and Updates

- ▶ via OTRS Ticket
- ▶ ITS imports the „normalized“ data

Data Schema

- ▶ Proposal from 2015-05-25 for data schema on the filesystem (handout)

Discussion?

- ▶ How much new data in the future?
- ▶ Update intervals of existing data sets?

Future Plans?

Future

We...

- ▶ increase performance in Spring 2016
- ▶ double the netto capacity from 72TB to 144TB in Spring 2016

Future

We...

- ▶ increase performance in Spring 2016
- ▶ double the netto capacity from 72TB to 144TB in Spring 2016

You...

- ▶ Plans for HAFL?
- ▶ Plans for TI?
- ▶ Plans for Fachstelle GIS?

Future

We...

- ▶ increase performance in Spring 2016
- ▶ double the netto capacity from 72TB to 144TB in Spring 2016

You...

- ▶ Plans for HAFL?
- ▶ Plans for TI?
- ▶ Plans for Fachstelle GIS?

Further Steps

- ▶ Setup Storage System
- ▶ Reviewing and normalizing data for import together (~end of September)

Thank You for Your Attention.

♥ Source Code is freely available

```
git clone git://git.bfh.ch/git/staff/bad9/other/talks.git
```