Uncovering Lustre Performance Issues in Operational Weather Forecasting at DMI with View for ClusterStor

The Danish Meteorological Institute

Uncovering Lustre Performance Issues in Operational Weather Forecasting at DMI with View for ClusterStor

- Keywords
  - Reliable and timely production
  - System wide I/O accounting for identifying hot spots
  - Merging PBSPro and ALPS events with lustre jobstats
  - Evaluating View for ClusterStor
- Presenters
  - Thomas Lorenzen, <tl@dmi.dk>
    Systems Analyst, Danish Meteorological Institute
  - Torben Kling Petersen, <tpetersen@cray.com>
    Principal Engineer, Cray

Agenda and intro

- The task of a meteorological service
- Supercomputer and I/O overview
- History of I/O monitoring and accounting
- View experiences : Installation and operation
- View experiences : Usability and cases
- Evaluation and future outlook

The task of a meteorological service

- The Danish Meteorological Institute obligates to forecast evolution of weather and ocean for the Kingdom of Denmark
- The Kingdom of Denmark is not just Denmark but includes Greenland as well
- Timely forecasts require supercomputing power
- Resiliency mandate redundancy, either ”2*N” or ”N+1”
- Two XC50 systems, one for production and one for research and development, with shared Sonexion and Netapp
- Hosted at the Icelandic Meteorological Office in a joint partnership

The task of a meteorological service
Deterministic models

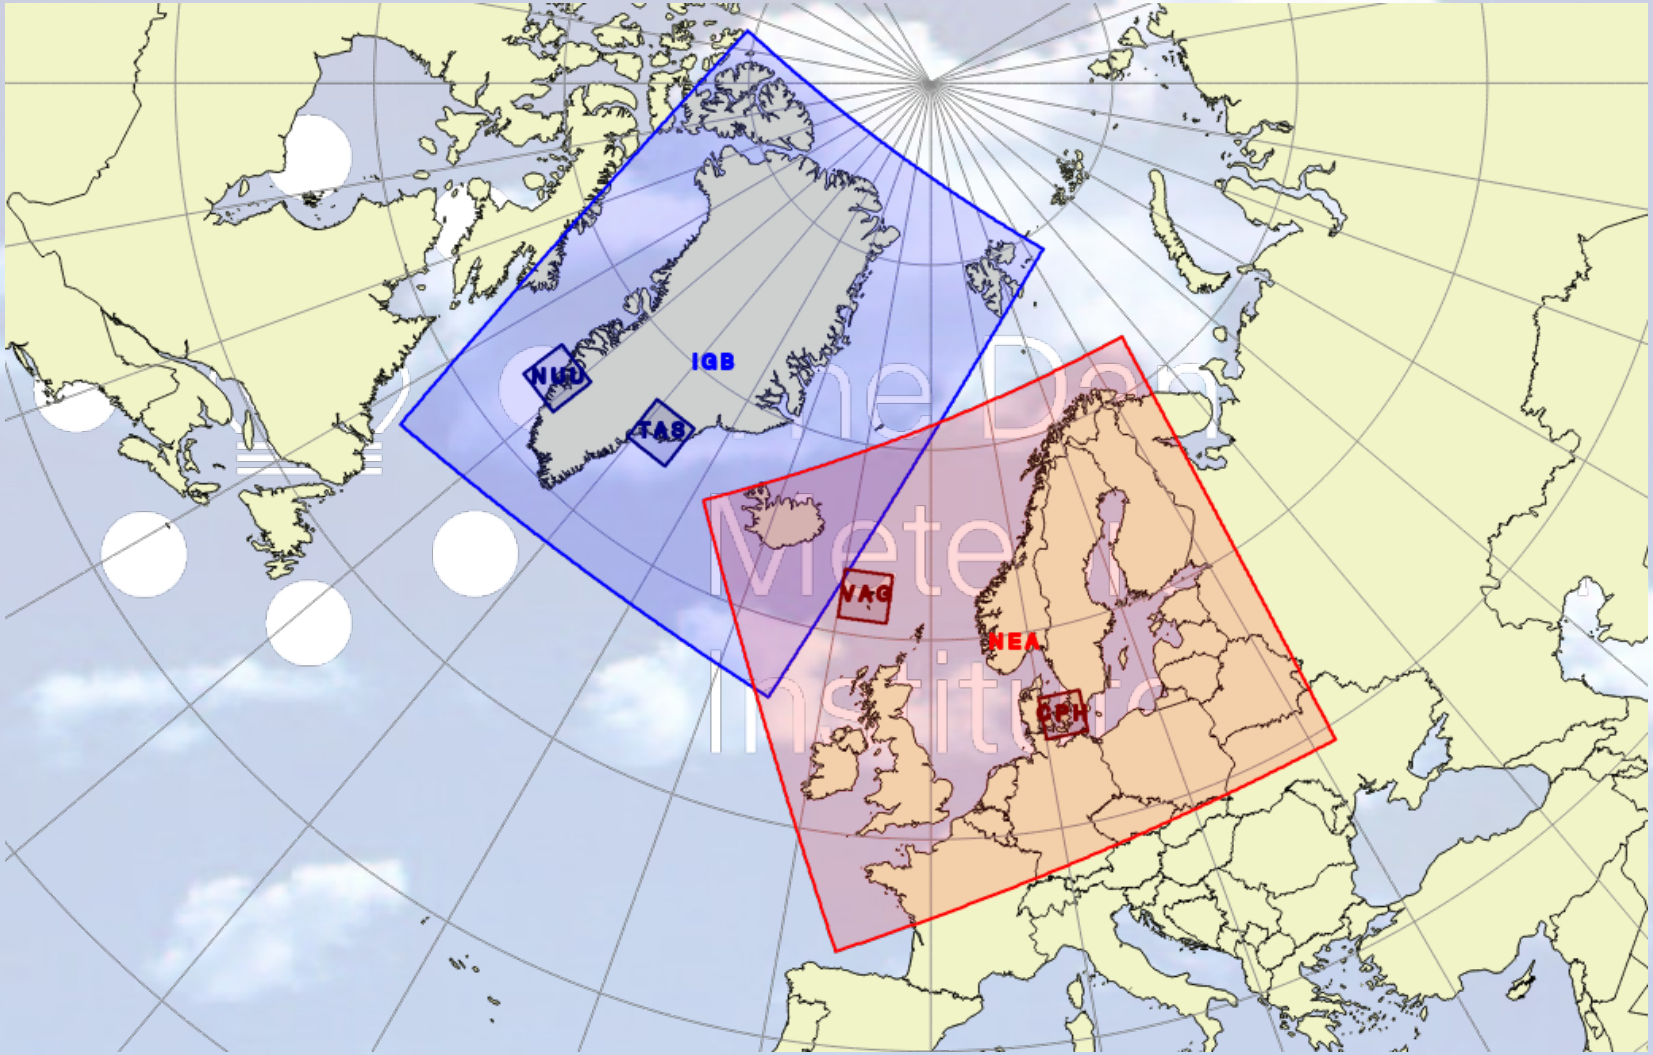

The task of a meteorological service
Ensemble models

The task of a meteorological service
The site

The task of a meteorological service
System utilization

Supercomputer and I/O overview

- Two XC50 systems and their shared storage
  - hpcopr : System for timecritical operational production
  - hpcdev : System for research and development
  - sonexion2000 : A single shared lustre file system
- Each XC50 system has three different user faced node types
  - Traditional compute nodes : Scheduled via ALPS via PBSPro
  - Repurposed compute nodes : Scheduled via PBSPro
  - Login nodes : Interactive and scheduled via PBSPro
- Sonexion mds shared between opr and dev
- Sonexion osses separated between opr and dev via lustre pools

Supercomputer and I/O overview

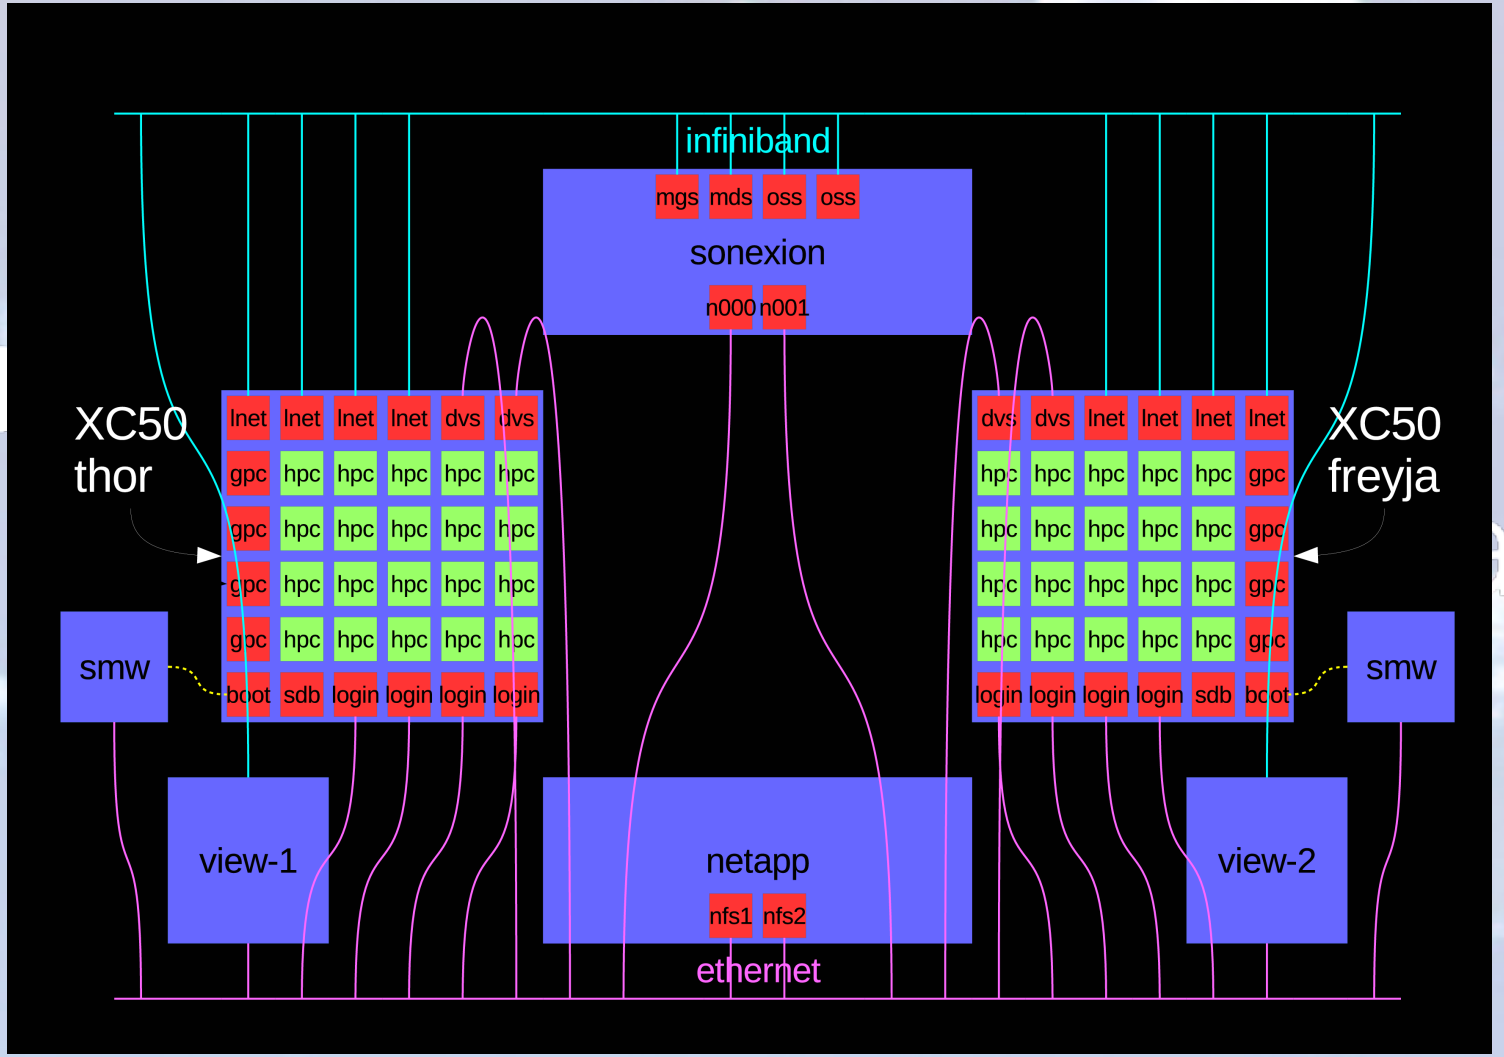

Supercomputer and I/O overview

- Potential pitfalls of the lustre setup
  - Shared resources between opr and dev
  - Insufficient resources
- Monthly raid consistency checks incur a bit of production delay
  - Implies that resource usage is maxed out already
  - Either acquire more or assess on the applications

Supercomputer and I/O overview

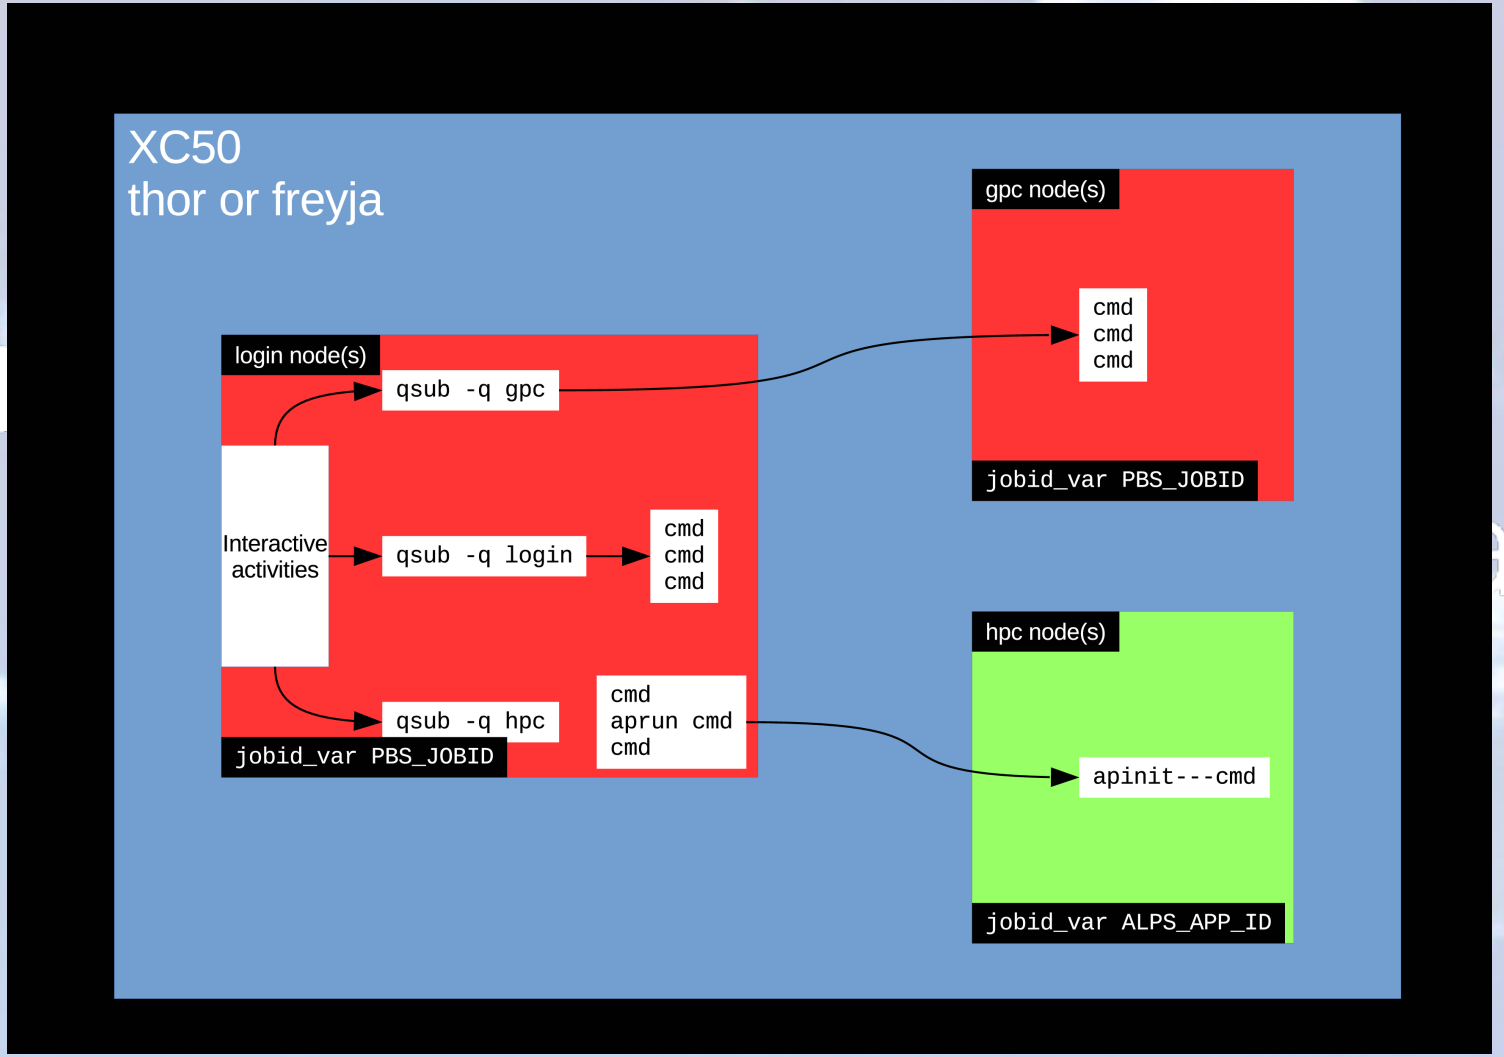

Supercomputer and I/O overview

- Desire to map out the I/O usage profile of all user activities
- Two purposes served
  - User approach : Asses the I/O characteristics of a particular job
  - Sysadm approach : Spot jobs with bad I/O characteristics potentially impacting other jobs
- The latter approach is the focus of this presentation
- The I/O usage profiles may be explicitly or implicitly useful
  - Some outliers may be spotted directly from the sampled metrics
  - Other outliers surface when coupled with knowledge and expectation of job characteristics

History of I/O monitoring and accounting

- System installed end 2015 with SU10 and lustre 2.5.x
- Setup supported LMT and web interface for real time monitor
- However, SU10 not immediately supporting jobstats despite jobstats being available in lustre 2.5.x
- Also, LMT carried a limited number of lustre usage counters with desire for more
- The Cstream aka SeaStream application programming interface announced to be avaiable in SU23
- Planning challenges due to only a single sonexion being available precluded upgrade until SU31

History of I/O monitoring and accounting

- Fall 2016 and spring 2017 featured a joint effort with DMI and Cray on an alternative approach
- The Open XDMoD was crafted to integrate to XC systems
- By means of rur plugins PBSPro account logs were amended with lustre stats for ALPS jobs
- Plain PBSPro jobs on repurposed compute nodes more difficult, since such nodes are used by several jobs
- Also, going the way of Open XDMoD required a significant amount of work for local adaption
- Better have a purpose built tool, if only such existed

History of I/O monitoring and accounting

- At CUG-2017, when DMI and Cray presented their Open XDMoD experiences, Caribou was also presented
- Many similarities were seen, but whereas Open XDMoD needed tweaking, Caribou was born tightly integrated to XC
- At CUG-2018 the renamed View for ClusterStor product was presented
- In autumn 2018 DMI and Cray made an agreement with DMI evaluating View for ClusterStor and providing feedback to Cray

View experiences : Installation and operation

- Fetched View tar ball via crayport.cray.com and follow the installation instruction as found on pubs.cray.com
- Minimal centos sufficed, since tar ball contains all dependencies for docker containers, which View makes heavy use of
- Installation instructions are comprehensive and easy to follow
- No unpleasant surprises
- For ALPS (jobeventd on smw), things worked mostly out of the box
- For PBSPro (jobeventd on sdb), some additional tweakings were needed

View experiences : Installation and operation

- By default there is no external network on sdb
- Proof of concept established via ssh port forwarding via login
- Version 13.0.408 of PBSPro still uses ancient python 2.5.1
- The jobevent prologue uses the runjob hook
- The hooks framework has the runjob hook as an obvious prologue hook on the pbs server. Unfortunately no obvious epilogue hook exists, which executes of the pbs server
- Hence jobevent epilogue is essentially a polling mechanism looking for finished jobs, which details are then fetched via the tracejob tool. This carries overhead and is not scalable

View experiences : Usability and cases

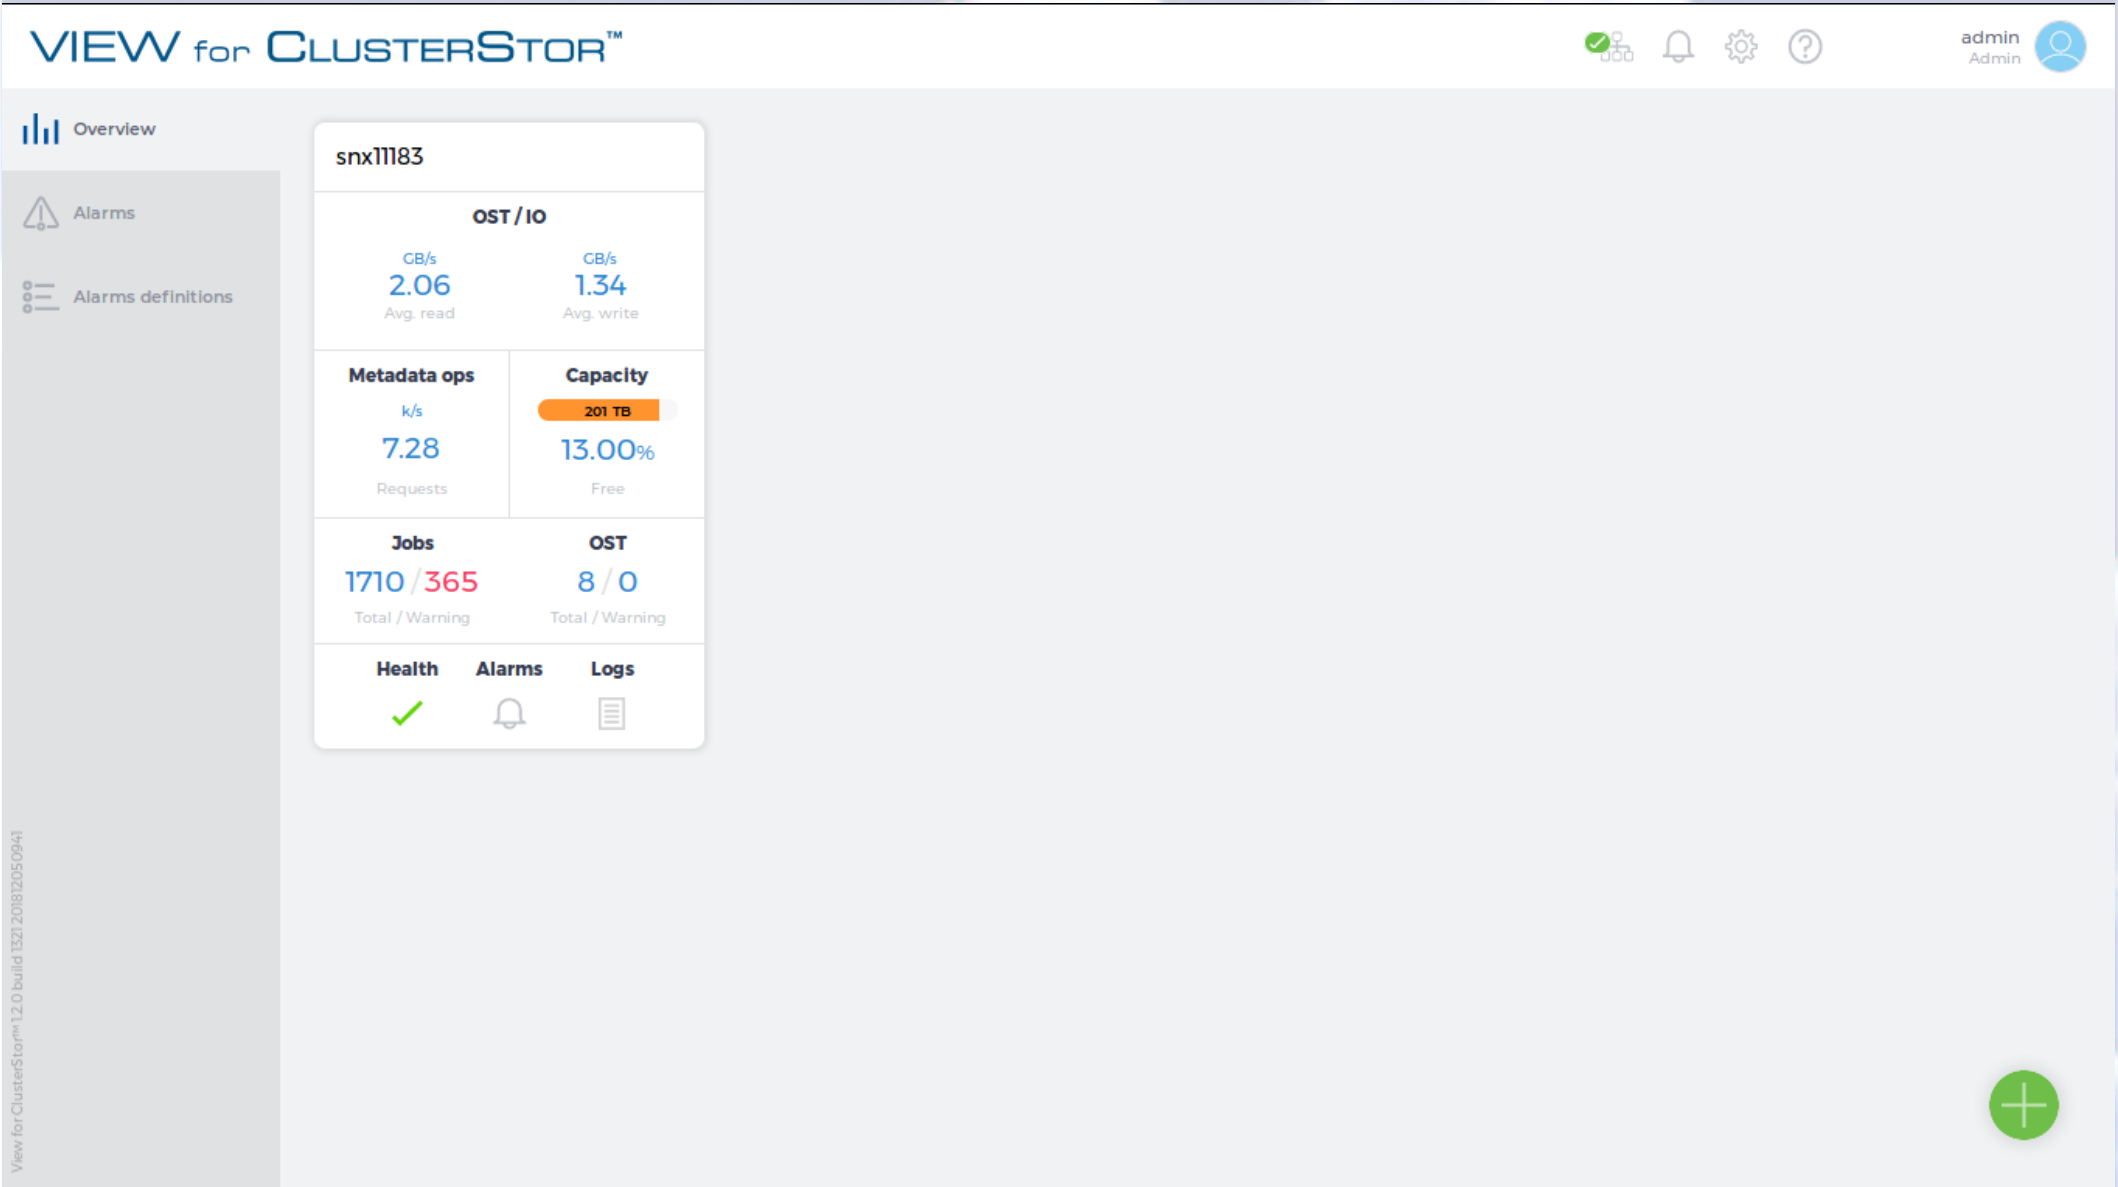

View experiences : Usability and cases

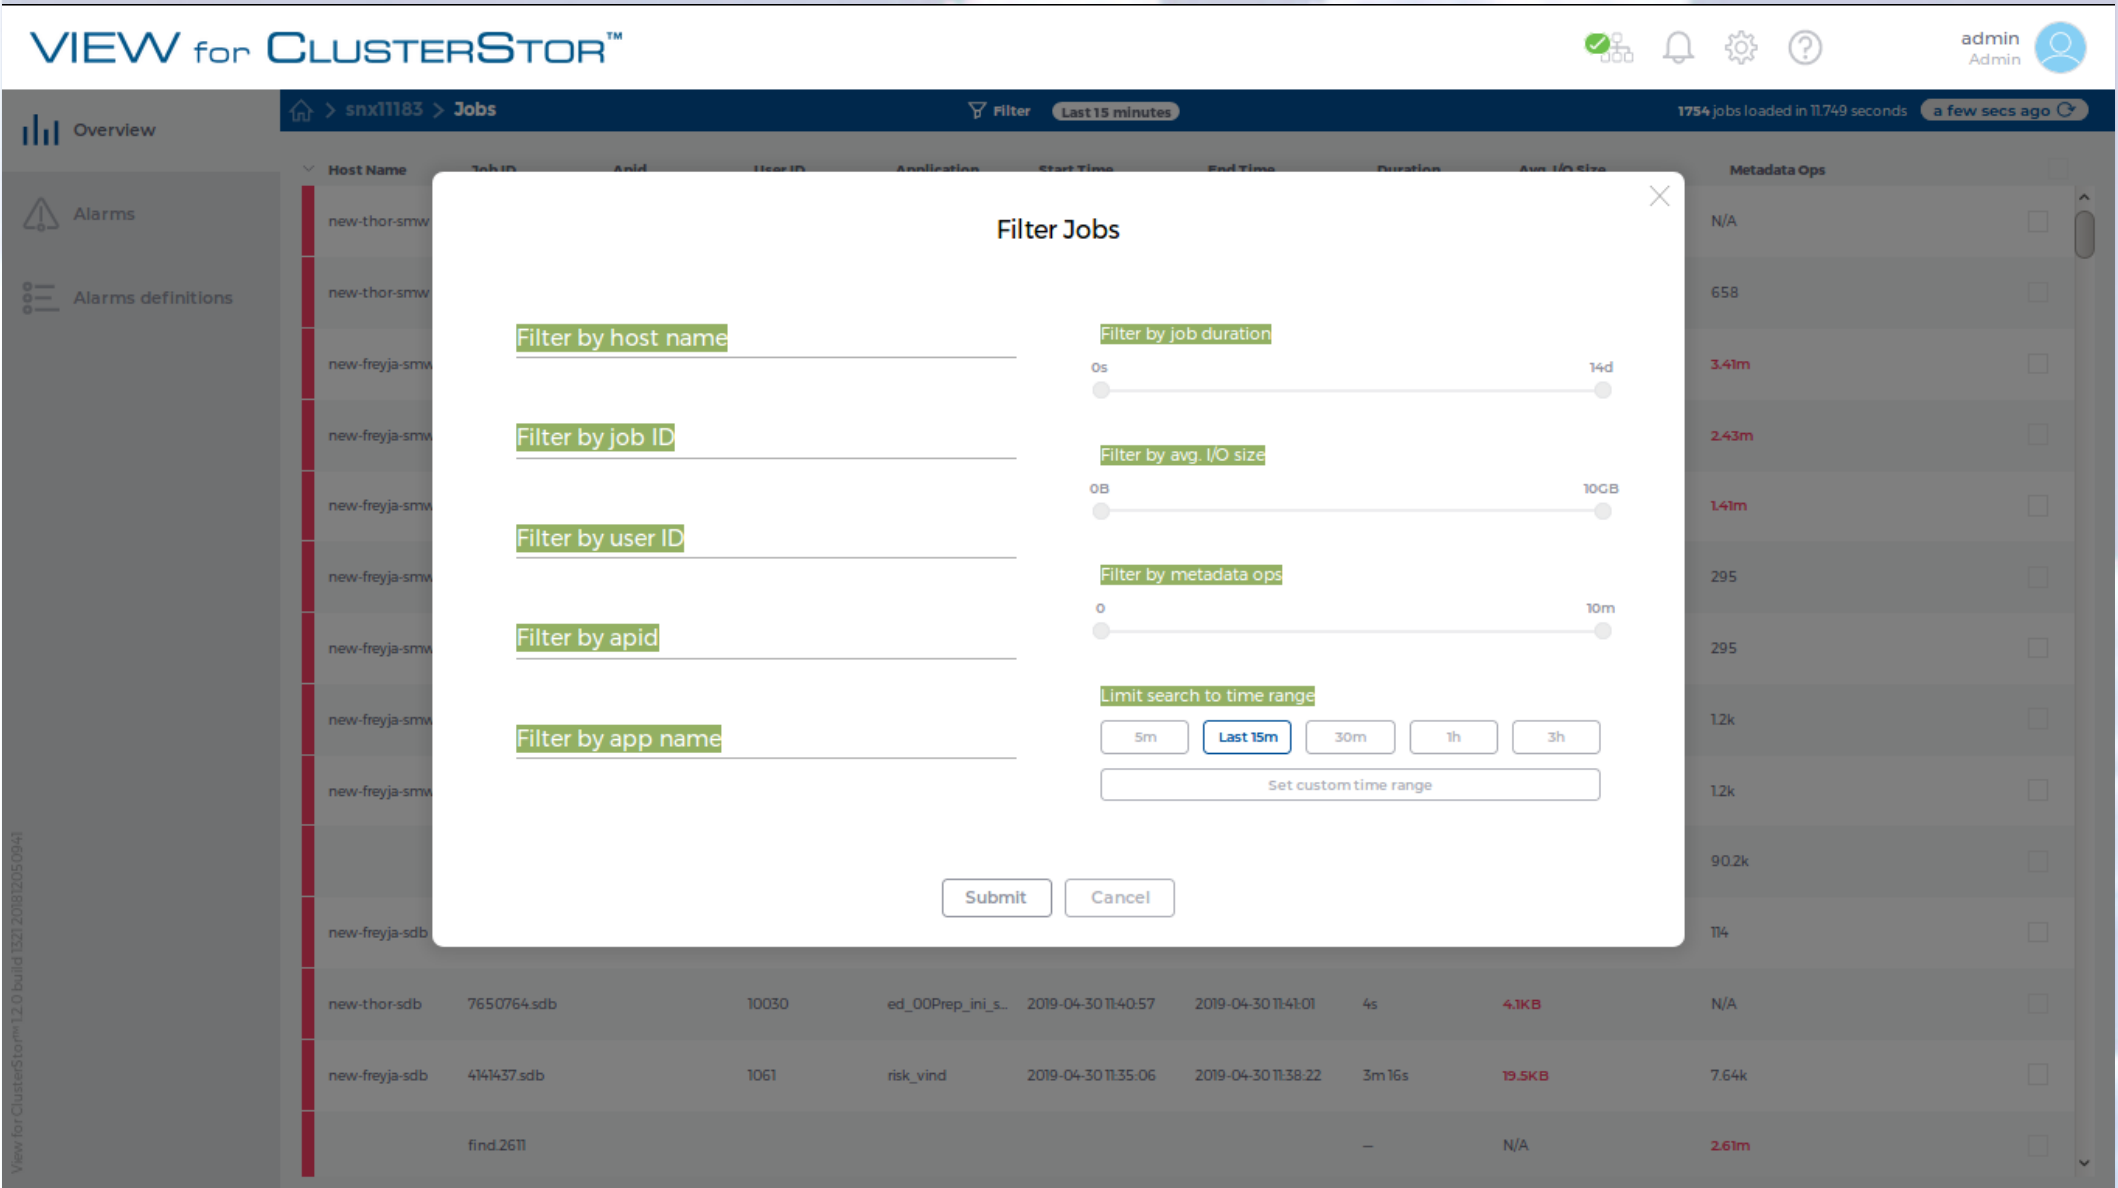

View experiences : Usability and cases

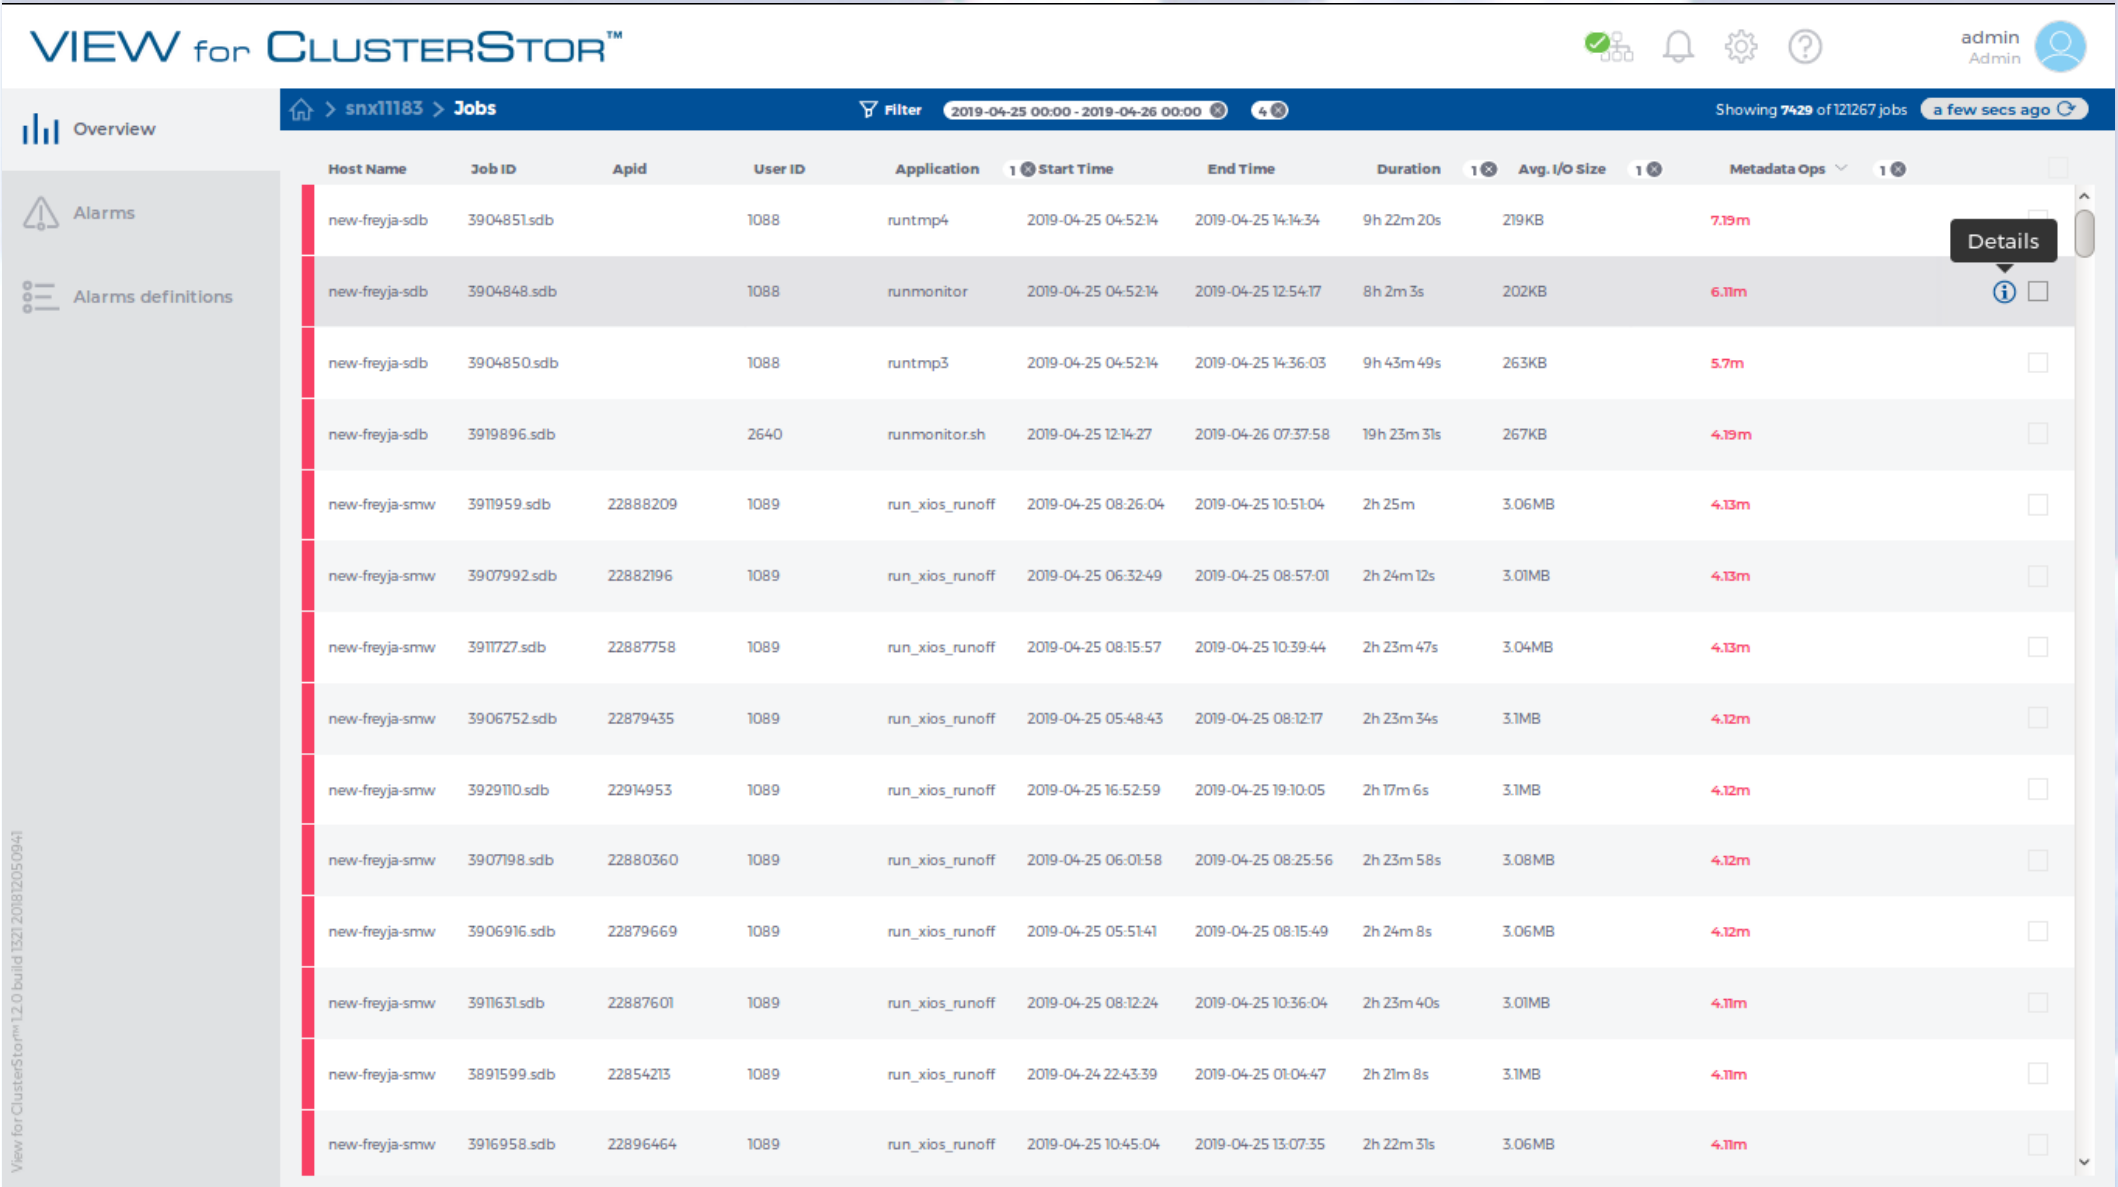

View experiences : Usability and cases

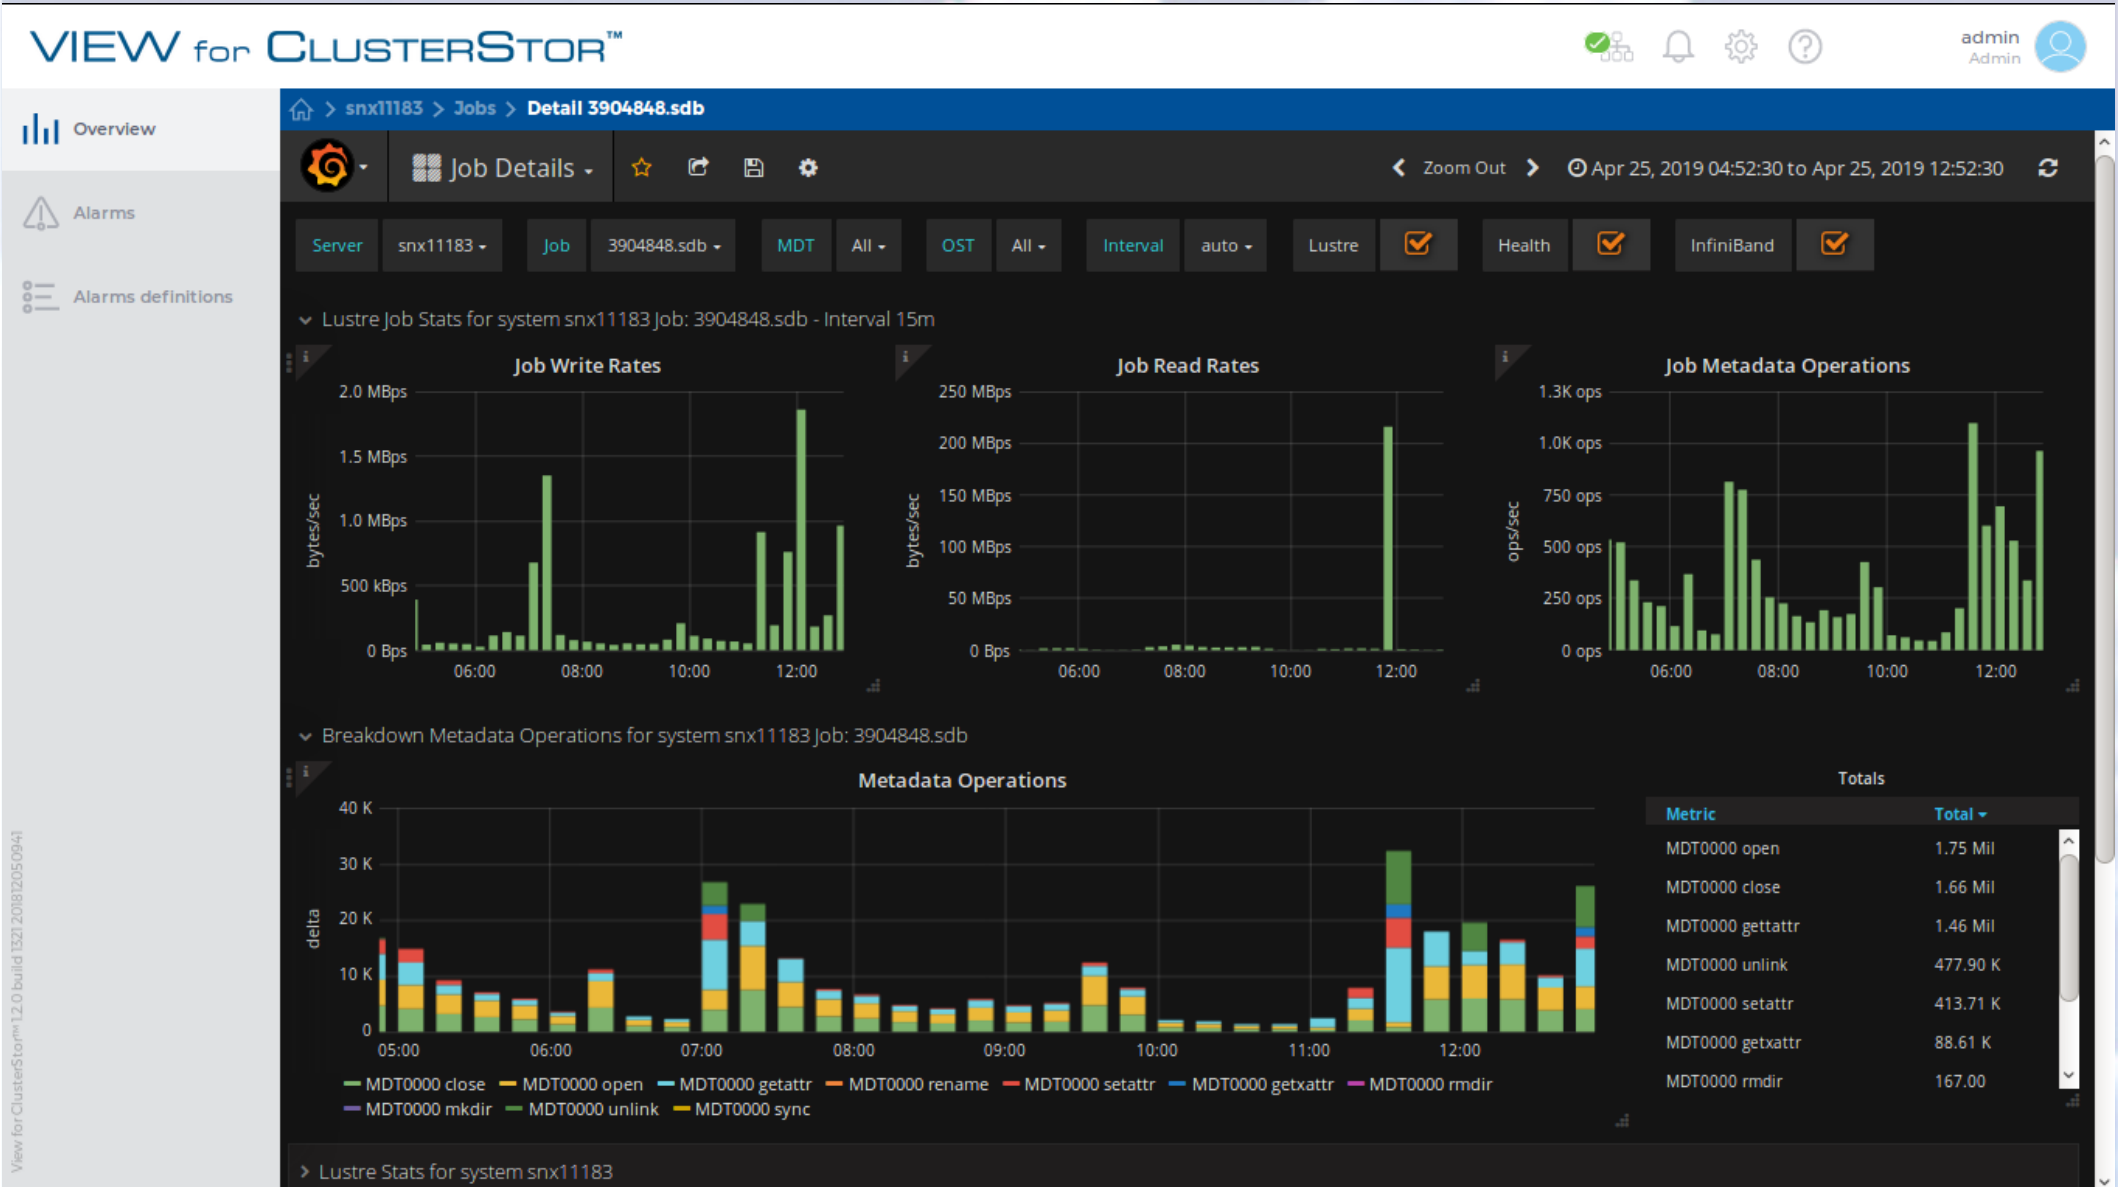

View experiences : Usability and cases

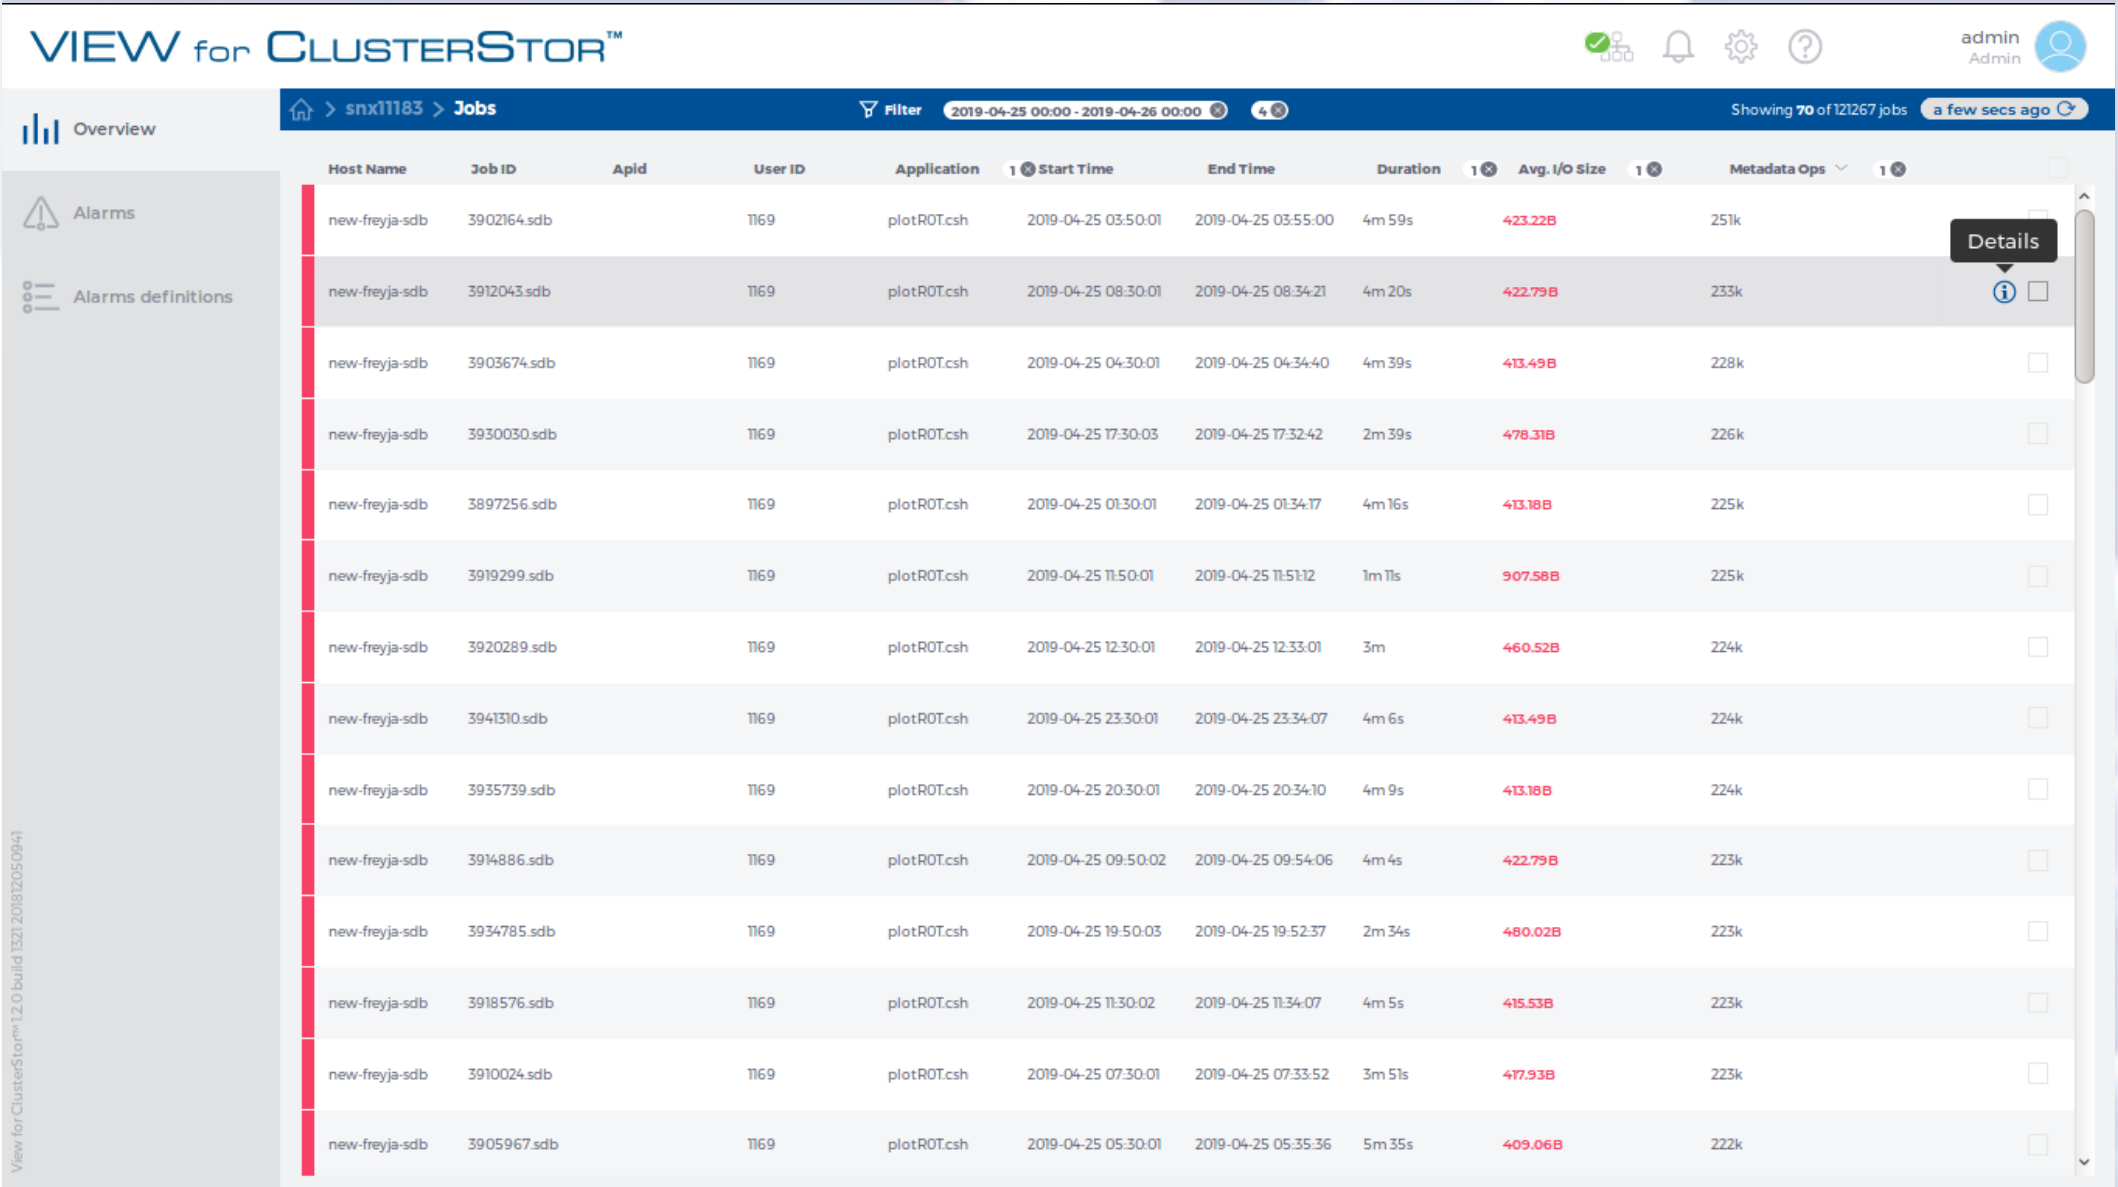

View experiences : Usability and cases

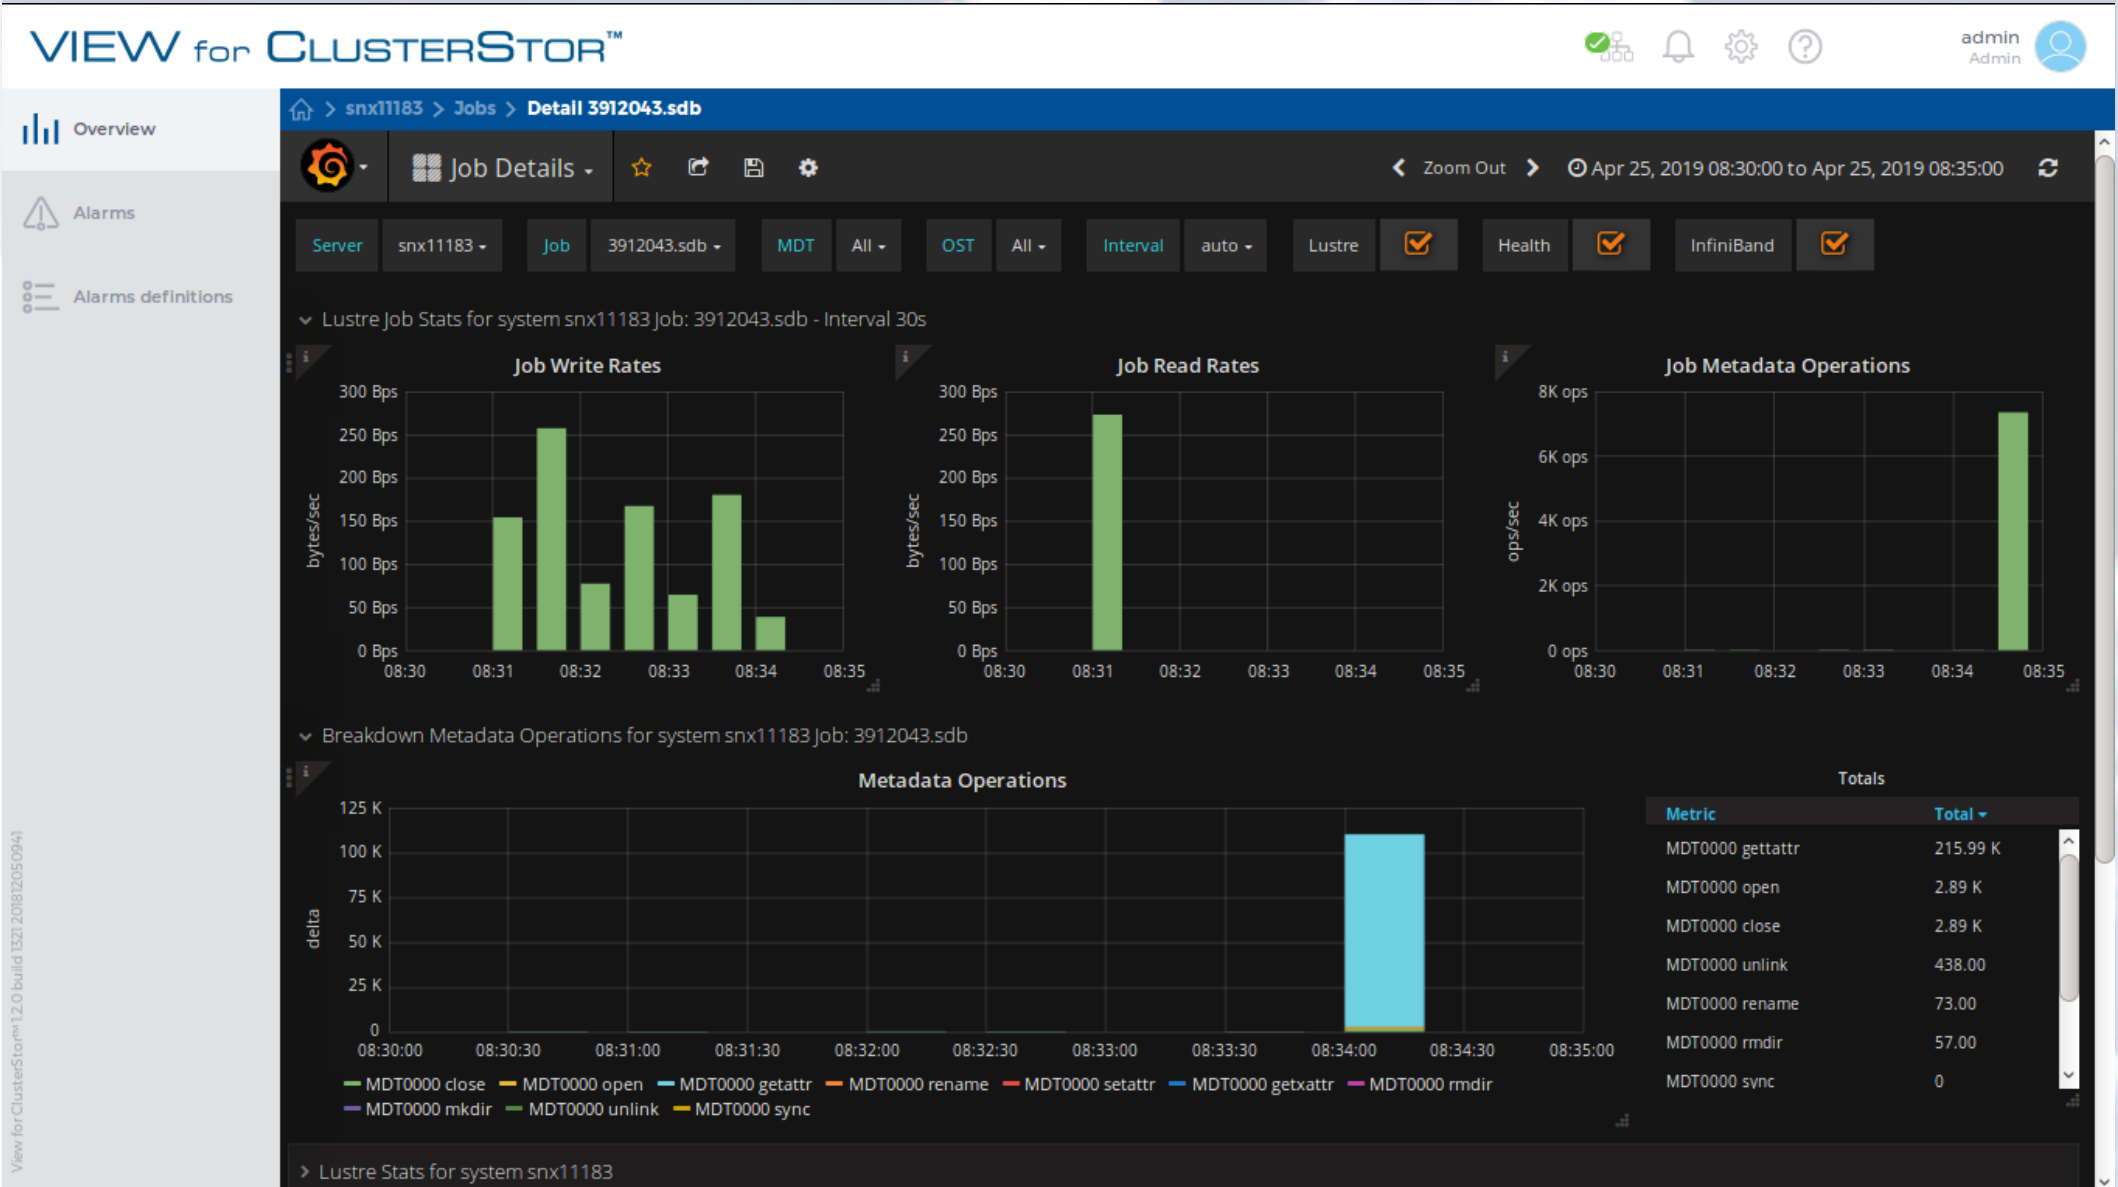

View experiences : Usability and cases

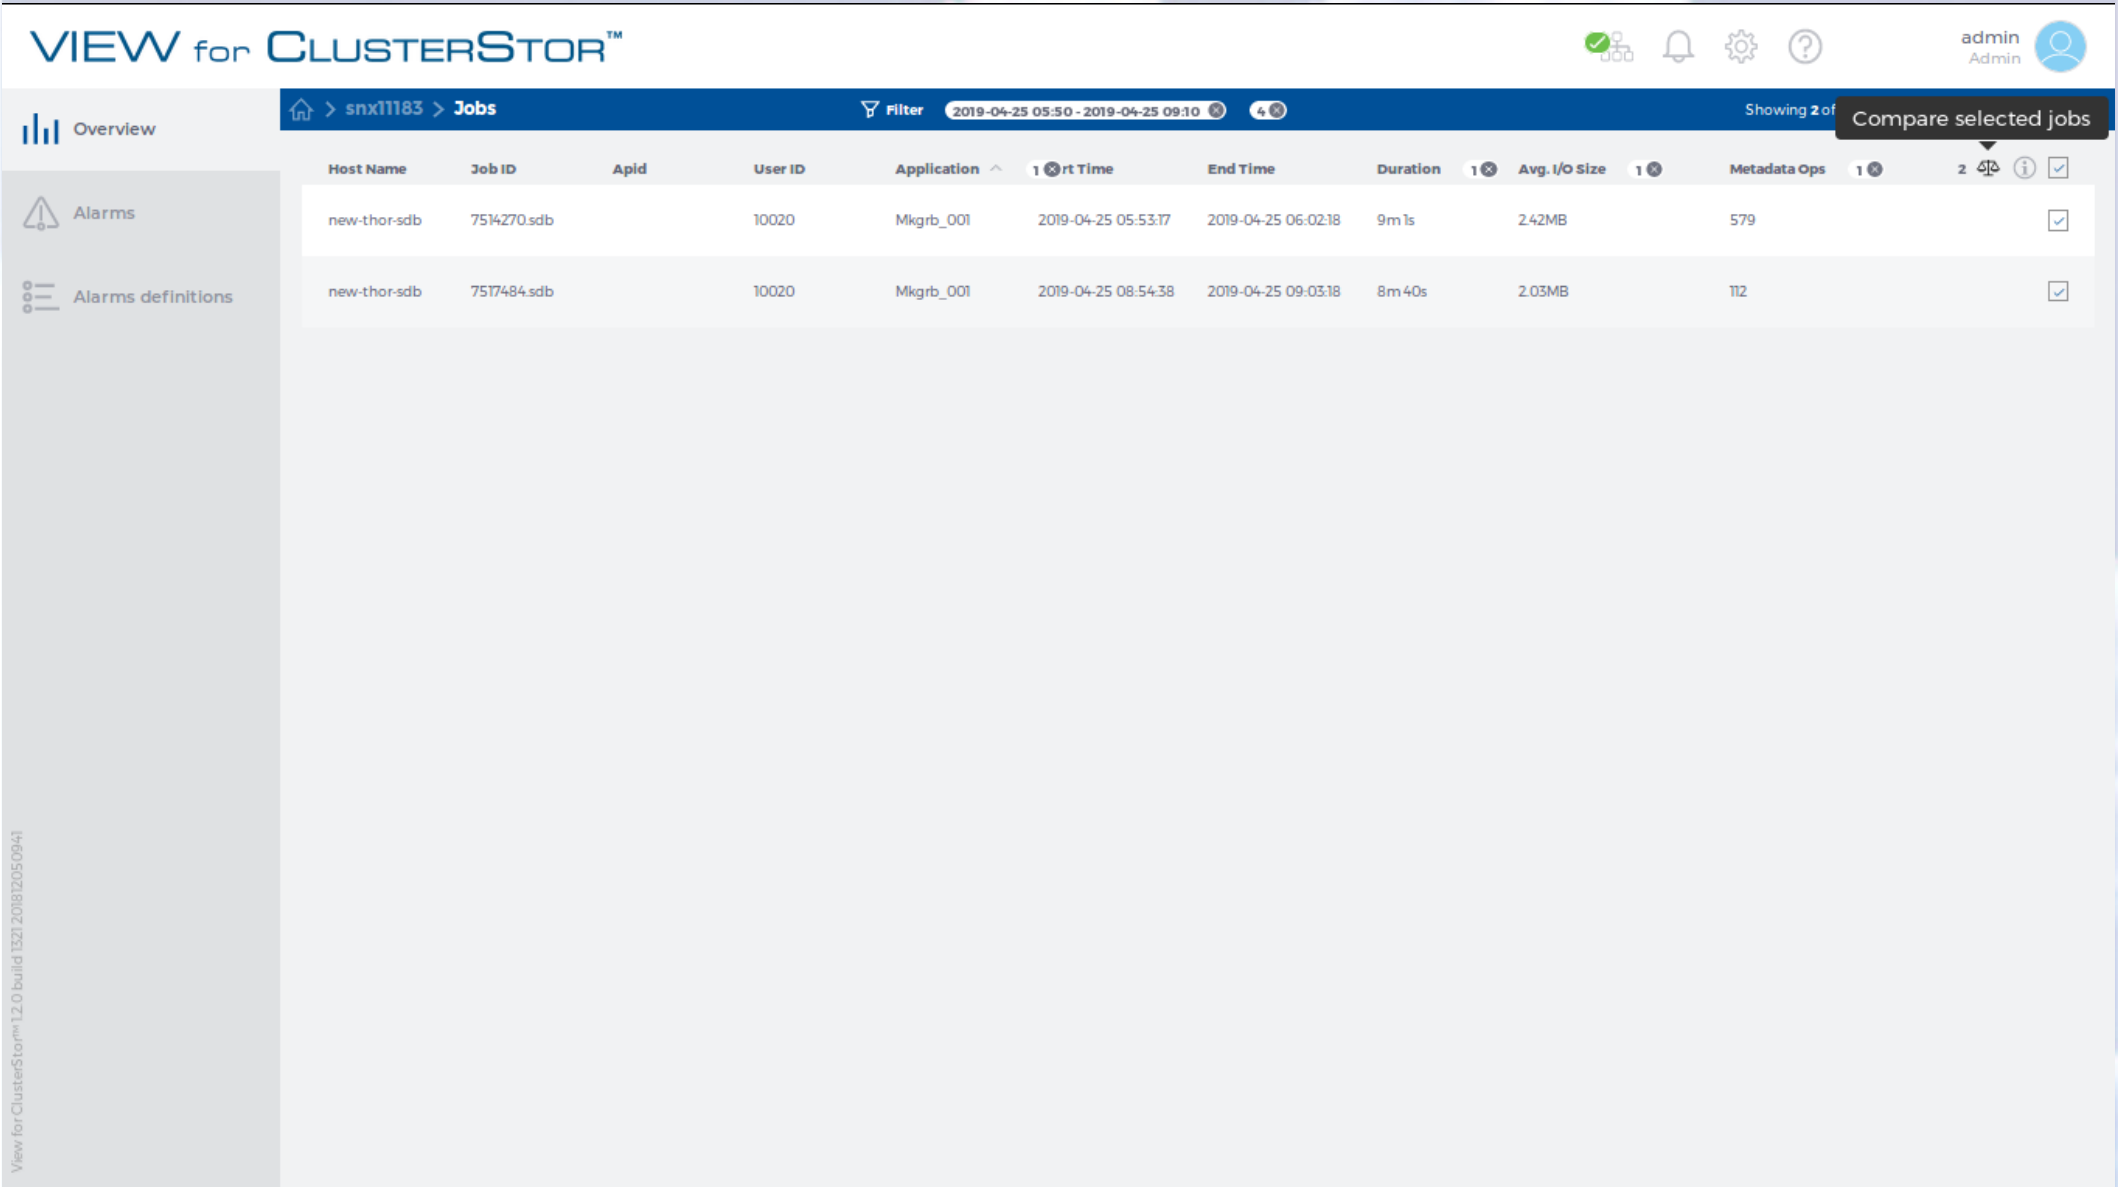

View experiences : Usability and cases

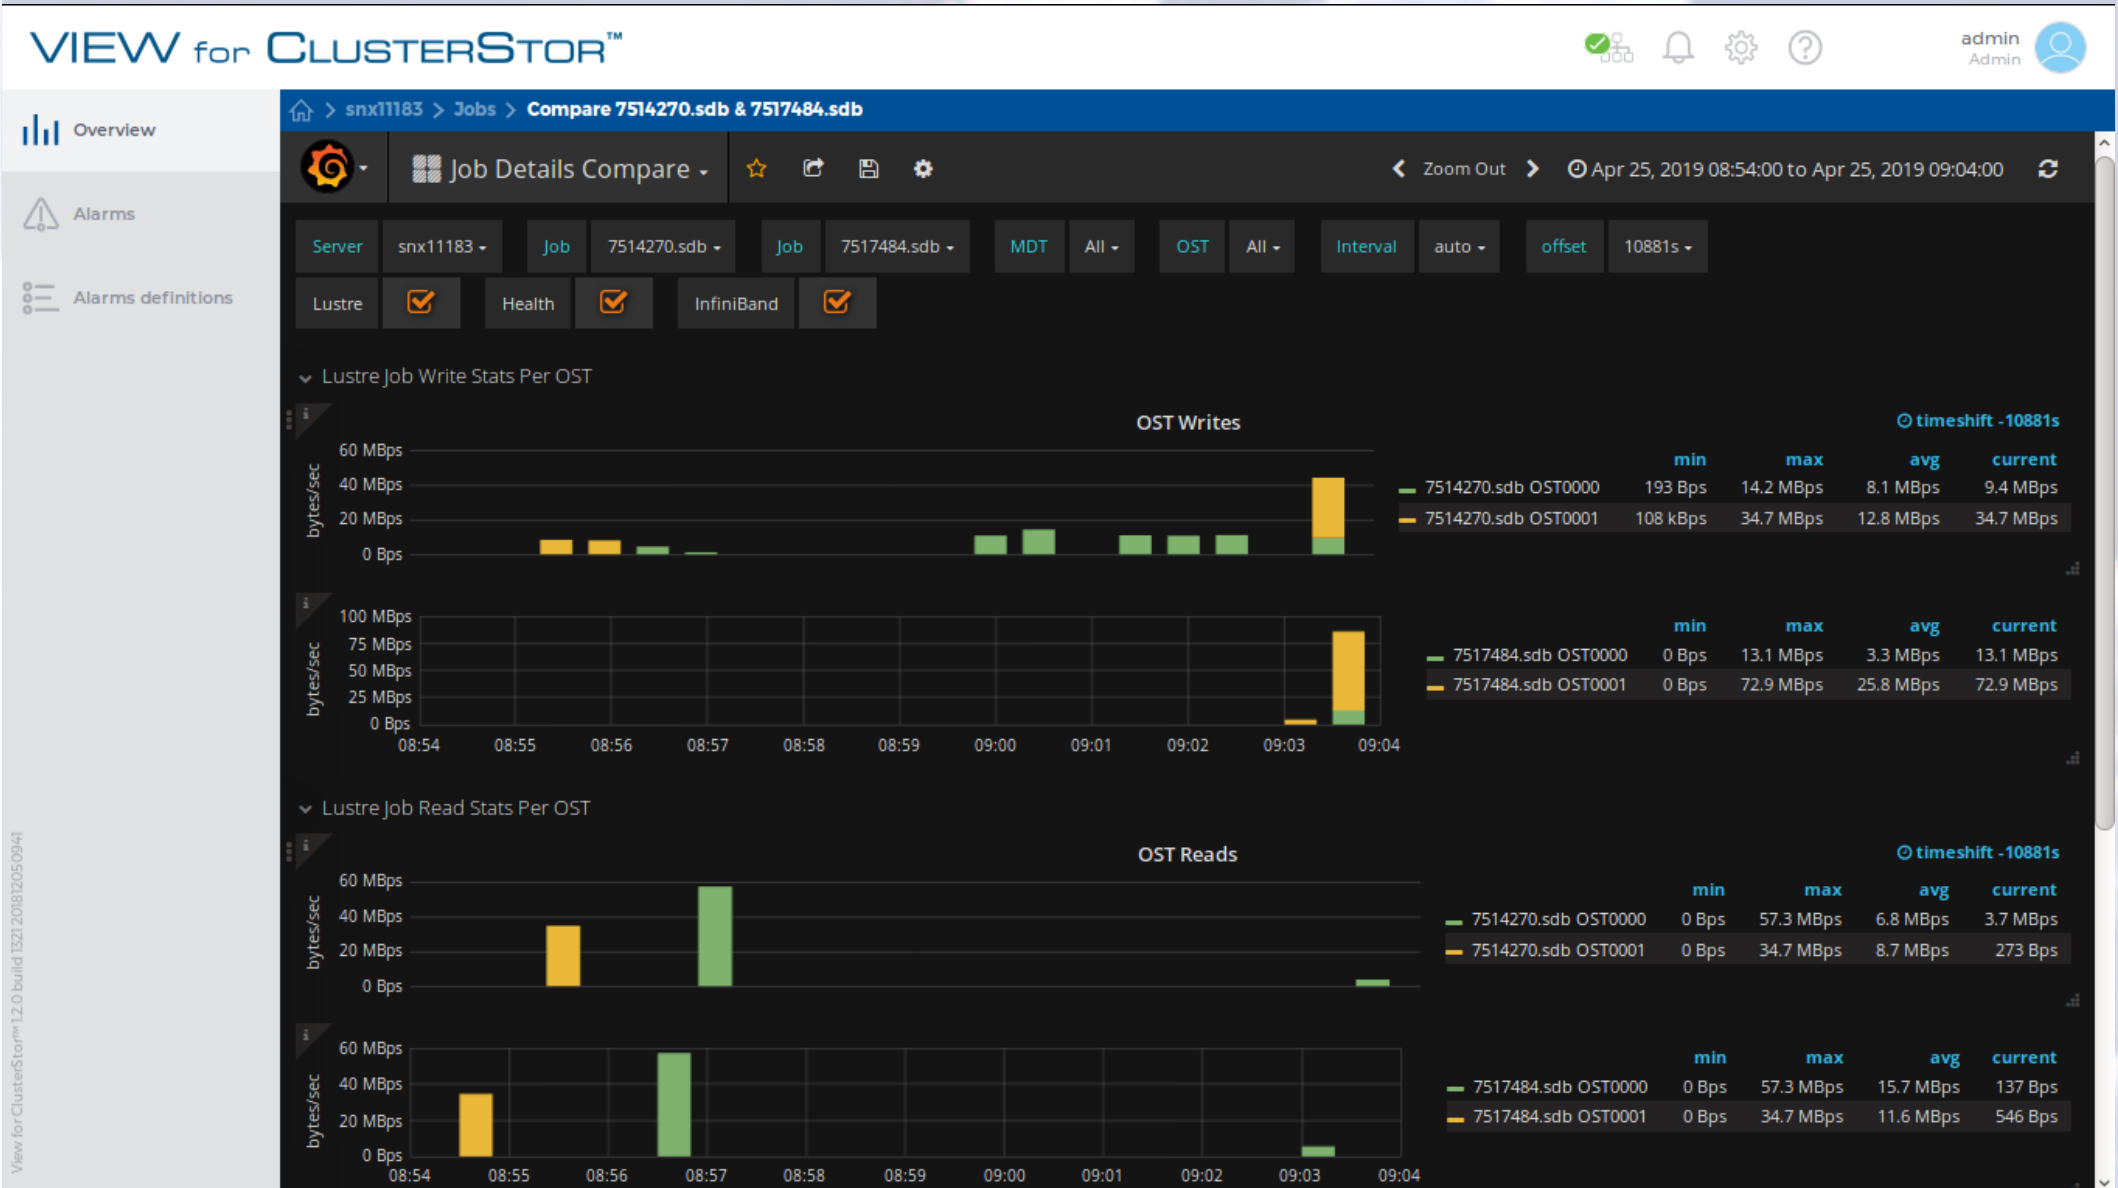

Evaluation and future outlook

- View glues lustre stats and jobstats with PBSPro and ALPS job events to give a comprehensive system wide I/O account
- View is an out of the box experience for ALPS, but for PBSPro workarounds have been needed, and the setup is less scalable
- The dashboard of the detailed job view is customizable, but the overall job view table lacks this ability
- Apparent misses of jobs or account for I/O activities in the web interface have been spotted and reported
- However, the View tool seems good and intuitive in getting overview of I/O activities, both system wide and per job
- Hence View can make a positive difference for sites in need of doing detailed I/O analysis at the level of individual jobs

Summary and outro

- The task of a meteorological service
- Supercomputer and I/O overview
- History of I/O monitoring and accounting
- View experiences : Installation and operation
- View experiences : Usability and cases
- Evaluation and future outlook