

TSD

The Norwegian node runs the TSD 2.0 system, developed at the University of Oslo (UiO), for their secure data processing. Users connect via Remote Desktop to virtual machines (VMs, Windows or Linux) with substantial computing power. If even more power is needed, jobs can be run on the TSD high-performance computing (HPC) resource. A storage area accessible from the VM is designated per project. Storage for TSD consists of a part of UiO's Astrastore, which is logically separated from the rest of Astrastore, delivered partly via a dedicated fibre channel (for VM storage and database), partly via NFS (the main bulk).

The system is firewalled and all incoming traffic goes through two (for failover) gateway servers, which route traffic per project to the respective VM. Authentication is 2-factor, users have to use a Google Authenticator App on their smartphone, or a Yubikey physical device to obtain the one-time code.

Projects are on separate VLANs, and no data can be transferred between projects, even if they have the same owner. Data can be transferred into a project by any project member, but only a project leader can decide to export data from a project. Import/export is done via a third gateway machine termed filesluice. Users can via 2-factor ssh login to the filesluice see their export and import folders of the project. Import/exporting is done via scp, and the filesluice NFS-mounts import/export folders from an inside VM which itself NFS mounts them from block storage. The inside sluice VM runs an automatic job that moves files from import folders to projects and from project export folders to the sluice export folder.

Heavy lifting is done on the HPC resource, which operates without a login node. There is a frontend VM that runs a DHCP server and cluster manager. The HPC runs a SLURM queuing system and since there is no node login, jobs can be submitted into the SLURM queue.

The entire TSD is provisioned by the inhouse designed Cerebrum system, running inside the firewall, which controls users, projects, storage, and VMs.

Mosler

The Swedish contribution to the Tryggve collaboration is the Mosler system, developed by BILS. It is modeled after the Norwegian TSD system, and has undergone some adaptations to suit local needs. Like TSD, it represents a PaaS delivering virtual machines to the user. Currently the basic allocation consists of a service node and an eight-core compute node. The storage and compute architecture are based at the Swedish HPC centre UPPMAX.

Like on TSD, Mosler project environments are isolated from each other and have no Internet access. Unlike TSD, Mosler uses sftp instead of scp to load data in to project storage via a filesuice. The sftp method is considered safer and provides a chrooting server for the sftp-server so one does not need to build a chroot environment for the scp process.

Mosler uses Openstack for provisioning the system, and user accounts are imported from UPPMAX. Access to the system is provided via a 2-factor authentication based on LinOTP, which enables access through the Thinlinc remote desktop server system to the users.

ePouta

ePouta is Finland's part of Tryggve, and it is designed to deliver virtual cloud capacity to customer organisations (universities, research institutes, companies). As it only delivers virtual computing nodes and storage, ePouta is seen as an infrastructure-as-a-service (IaaS).

To connect to ePouta research institutes will use an optical private network (OPN) that each respective institute must set up towards ePouta. After this, the institutes will be responsible for running their virtual cloud environments. Like in TSD and Mosler, there will not be communication between these environments and also not between an environment and the Internet.

Since the institutes themselves are responsible for the environments, user access, management and data import/export is not governed by the ePouta system.

The provisioning system used by ePouta is Openstack, which opens for possible resource sharing with the Mosler system.