



Virtualization of Grid Services



A. Gellrich for the Grid Team at DESY*, Germany

DESY

DESY, a member of the German Helmholtz Association (HGF), is one of the world-wide leading centers for research with particle accelerators and synchrotron light. DESY is a WLCG Tier-2 center for LHC experiments ATLAS and CMS and participates in the EU-projects EGI, the successor of EGEE, in the federation NGI-DE. DESY was founding partner of the German Grid initiative D-GRID and played a leading role in the HEP community project (HEPCG) and in the integration project (DGI-2).

Grid at DESY

The Grid site DESY-HH, which is the home to 10 VOs and supports in total 20 VOs, incl. ATLAS and CMS. All VOs are using one common Grid infrastructure. In addition to the 4784 job slots (2GB mem/slot, 15GB scratch/slot) with a total of 38kHS06 and the 3 dCache-SEs with a total of 4PB of disk space, all Grid services which make up a complete Grid infrastructure are provided, incl. multiple instances of BDII, LB, LFC, PX, SCAS, VOMS, and WMS. Most of the Grid services run in virtual machines.

Concept


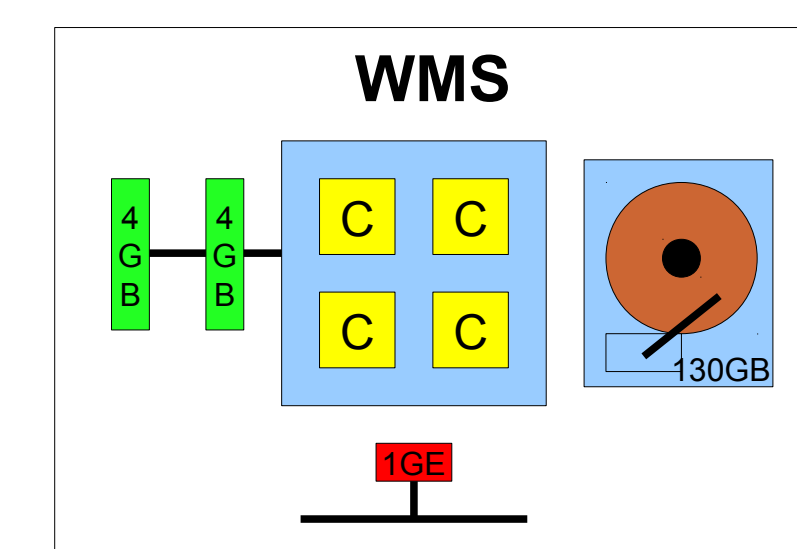
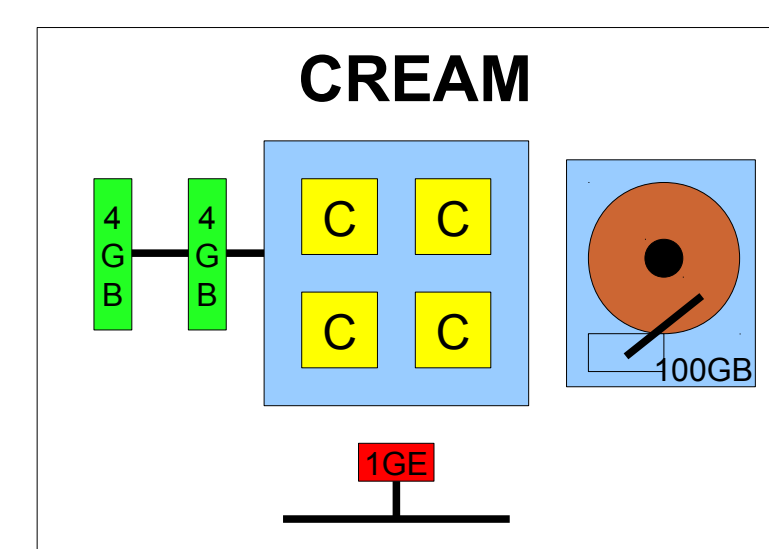
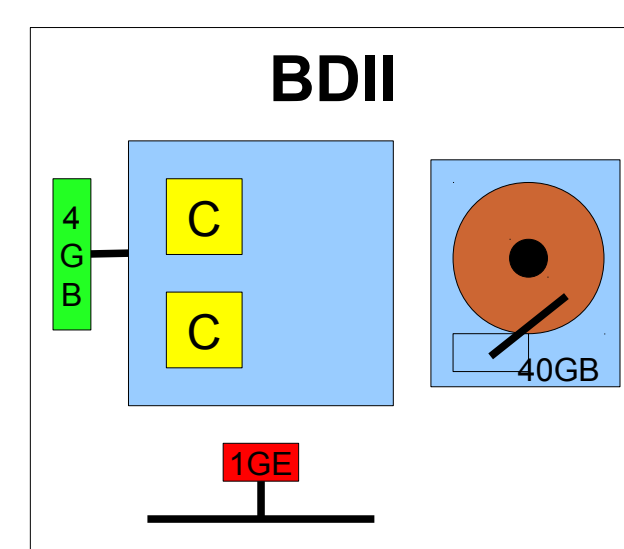
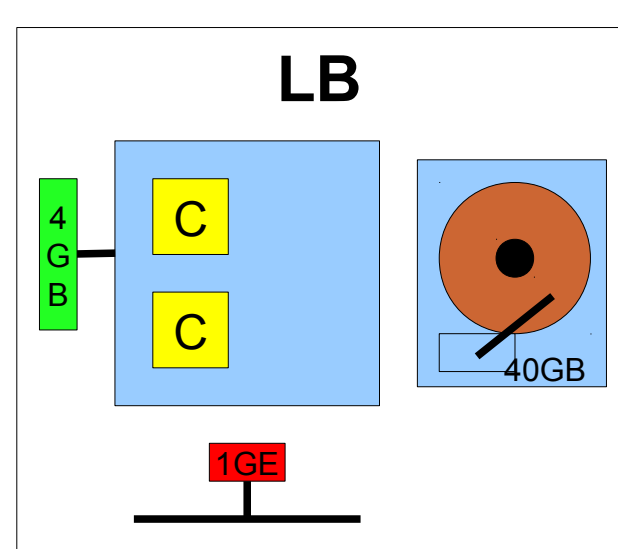
The DESY-HH Grid site supports 20 global VOs. 10 of the VOs are hosted at DESY-HH and all necessary Grid services are provided. In order to guarantee stable and performant operations of the Grid services, not only appropriate hardware which meets the resource requirements of the services, is needed, but also scenarios to migrate servers from out-dated to new hardware must be foreseen. Virtualization technologies, allowing for automatic or semi-automatic migration of services between hardware set-ups, were studied and are used in production at DESY-HH.

Technologies

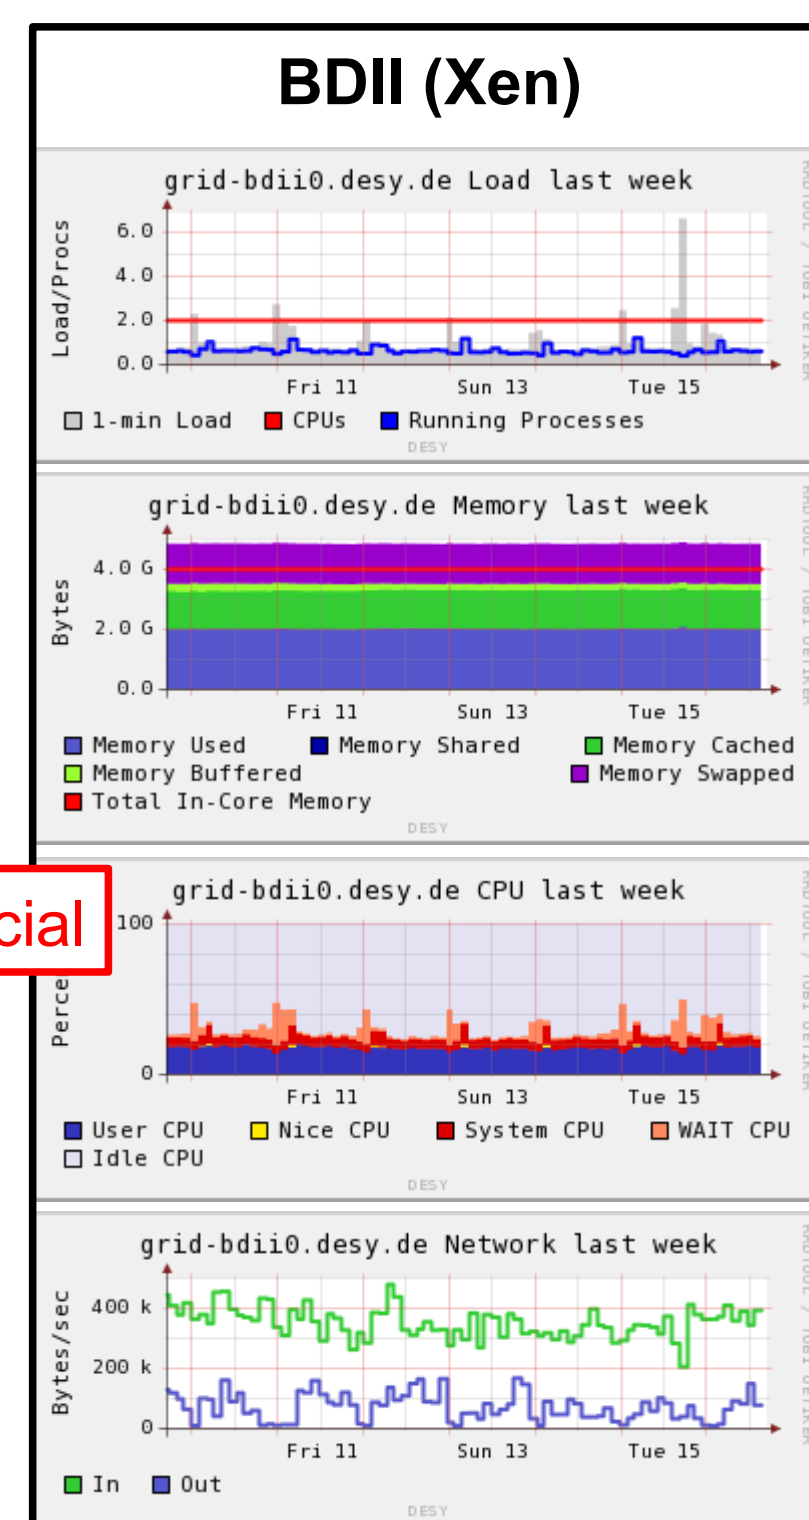
At the Grid site DESY-HH two virtualization systems are in use for Grid services; cloud technologies are being investigated to host worker nodes (WN):

XenCenter: shared NetApp-based storage, incl. snap-shoting and online migration; Windows-based
KVM: uses virt-manager; host-based; no online action; needs manual copying of images

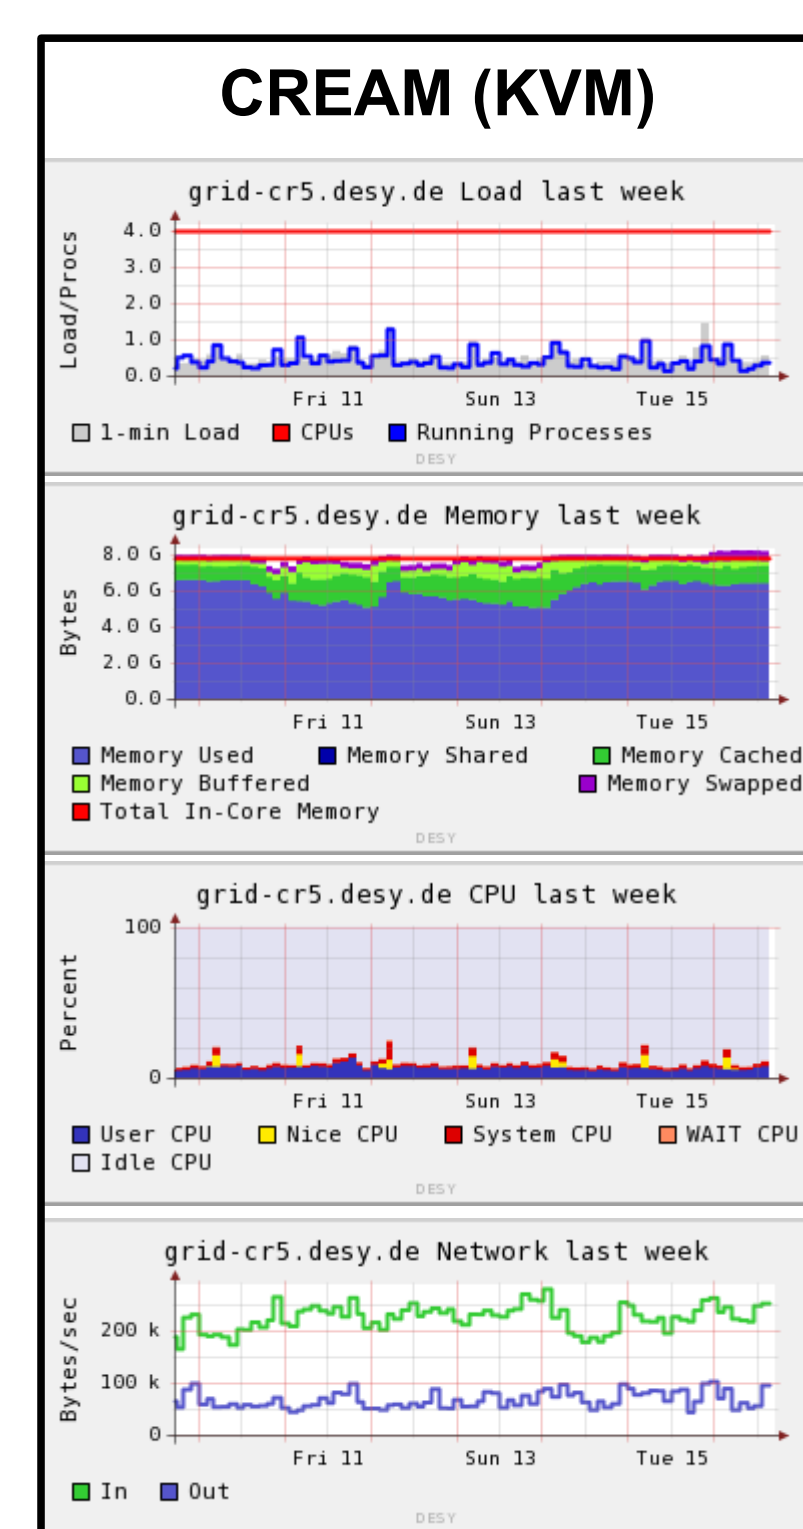


Light-weight Grid services (APEL, BDII, LB, SCAS, VOMS) are running in XenCenter. Heavy-weight Grid services (CREAM, WMS) remain on their original hardware as Virtual Machines (VM), using KVM.



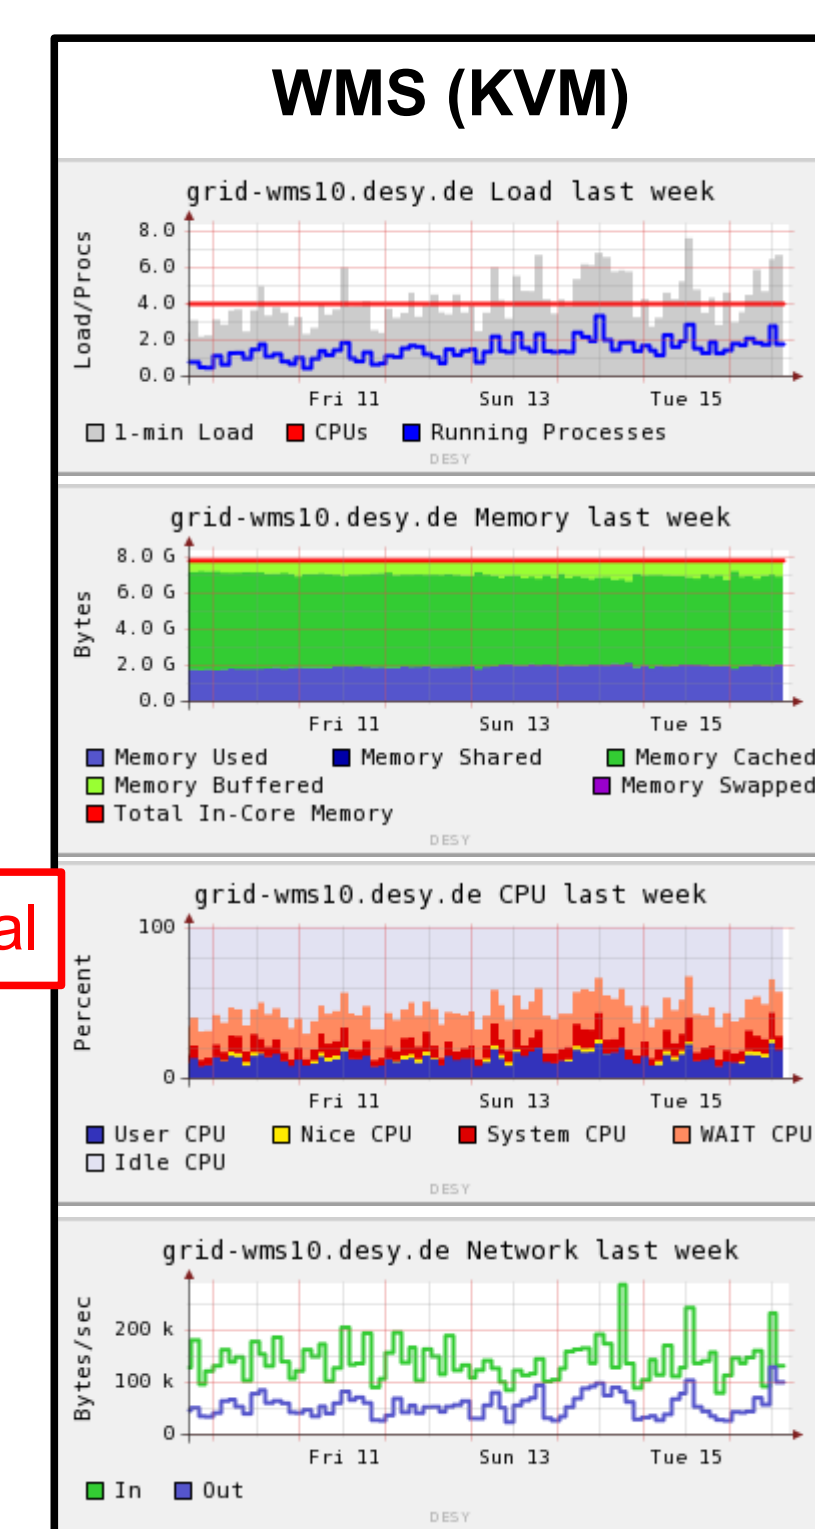
crucial



crucial



crucial



Experiences

Long term operational experiences showed that the server machines, which host Grid services, are stable and reliable. The utilization of mirroring RAID-system for system and data disks has prevented DESY-HH from service interruptions and data losses for many years. Therefore online migration capabilities incl. snap-shots are not the main design criteria for virtualizing Grid servers. KVM offers a typical HEP-like Unix/Linux environment with scripting capabilities. XenCenter is Windows-based and has a more sophisticated GUI. For heavy-weight performance-critical services such as WMS and CREAM, in XenCenter the underlying storage system in its current set-up is a bottleneck. Already light-weight Grid services show big cpu_wait_io values.

Outlook

The most crucial aspect is the performance of heavy-weight services such as WMS and CREAM in the storage access. From the operational point of view a solution, which is offered as a central service by the computer center, incl. maintenance, is favored and most promising. XenCenter-based solutions would need as of now major improvements in the storage access. KVM-based solutions are typical for HEP-Unix/Linux environments but currently lack convenient migration tools. Similar tools are needed for XenCenter. With respect to hardware, the usage of SSDs could increase the I/O performance and decrease cpu_wait_io. Tests are on-going. Also for Grid services a cloud-like infrastructure could be envisaged.