



Dynamic Computing Resource Extension using COBalD/TARDIS

R. Florian von Cube, René Caspart, Max Fischer, Manuel Giffels, Eileen Kühn,

Günter Quast, Matthias J. Schnepf

Data Life Cycle in Physics - June 8, 2020

INSTITUTE OF EXPERIMENTAL PARTICLE PHYSICS (ETP), KIT



www.kit.edu

Future Computing in Physics



- Data volume will increase with future experiments
- Modern analysis techniques will require dedicated hardware
- Cross-community resources are favored by funding agencies

Results in a need for more computing power, while having access to a more heterogeneous infrastructure

Computing Infrastructure at ETP



ETP: Institute of Experimental Particle Physics

- Local resources
 - Desktop cluster, worker nodes
 - Adapted to our needs, but small
 - Require maintenance

ForHLRII

- HPC cluster at SCC, KIT
- Shared with other communities
- Low priority user to fill up unused resources

- BWForCluster "NEMO"
 - HPC Cluster at Uni Freiburg
 - O(1000) cores
 - Shared with other communities
- TOpAS
 - Throughtput optimized
 - Joined operation by ETP and GridKa
 - GPU nodes

All made available through single point of entry

Dynamic and transparent integration

- Dynamically integrate multiple, heterogeneous resources into overlay batch system (OBS)
 - Transparent to the user
 - Single point of entry
 - Allows for hassle-free maintenance
- Provision resources using container and virtualization technologies
 - Provide specific software stacks
 - Maintain well-defined environment





Opportunistic Resources



Opportunistic resource in the HEP context

Resource, not specifically designed, but temporarily made available for HEP workflows.

This may include

- HPC centers
- Resources at cloud providers



Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers





Resources are dynamically allocated on available cloud providers



Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





Dynamic, Opportunistic Resource Scheduling

TARDIS and COBALD

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





TARDIS and COBALD Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD

Assesses the suitability of resources to the current job mix with metrics *allocation* and utilization.

Utilization Allocation not used





Dynamic, Opportunistic Resource Scheduling

TARDIS

Dynamically provisions and integrates resources into overlay batch system.

COBALD





Computing Infrastructure at ETP



Local resources

- Desktop cluster, worker nodes
- Adapted to our needs, but small
- Require maintenance

ForHLRII

- HPC cluster at SCC, KIT
- Shared with other communities
- Low priority user to fill up unused resources

- BWForCluster "NEMO"
 - HPC Cluster at Uni Freiburg
 - O(1000) cores
 - Shared with other communities
- TOpAS
 - Throughtput optimized
 - Joined operation by ETP and GridKa
 - GPU nodes

All made available through single point of entry

8 Data Life Cycle in Physics – June 8, 2020



Computing Infrastructure at ETP



Data Life Cycle in Physics – June 8, 2020

Dynamic Computing Resource Extension using COBalD/TARDIS R. Florian von Cube (ETP, KIT)

Computing Infrastructure at ETP

Local resources

Desk

8000 Adap ForHI R2 igh Throughput WNs (static) 7000 Requ nmunities Desktops (static) Dedicated WNs (static 6000 ForHLR 5000¹ 40 AD HPC Share FTP and 3000 Low Dynamically integrated 2000 used using COBALD/TARDIS 1000 All made av 19 0 15 Aug 2019

Easily scales to larger and more heterogeneous infrastructures!

Used CPU cores per site



reiburg

BWForCluster "NEMO"

Opportunistic resources at GridKa



- Provides resources for multiple VOs through single CE
- Dynamically integrates resources for the VOs at their respective various computing centers using COBALD/TARDIS
- Transparent usage for the VOs



Allows for dynamic integration of dedicated hardware, e.g. GPUs

COBALD/TARDIS: Adapts to your needs



COBALD/TARDIS can be adapted to fit your setup.

Supported Overlay Batch Systems

- HTCondor
- SLURM (Currently in development by Uni Freiburg)

Supported (Opportunistic) Sites

- HTCONDOR OPENSTACK SLURM
- CLOUDSTACK MOAB

Join us on GitHub and Gitter!

Usage of TARDIS Outside Physics



The Tier 1 grid center GridKa and ETP use COBALD/TARDIS to support the fight against COVID-19:

- Resources are dynamically scheduled to run "work units" for Folding@home and Rosetta@home
- Statistics of provided resources
- KIT press release

Summary



- COBALD/TARDIS allow for dynamic integration of opportunistic resources and dedicated hardware, e.g. GPUs
- Usage of the opportunistic resource transparent for the user
- Application by no means limited to HEP:

Join us on GitHub and Gitter!