Karlsruher Institut für Technologie

DARWIN

A Lightweight Analysis and Grid Facility for the DARWIN Experiment

CHEP 2024 – 19.10.2024 - 25.10.2024

Robin Hofsaess, Sebastian Brommer, Florian von Cube, Manuel Giffels, Markus Klute, Benedikt Maier | 21. October 2024

DARWIN: Dark Matter Wimp Search with Liquid Xenon

The Experiment

- Future astroparticle observatory in planning and R&D phase
- Evolving from the XENON experiment
- Growing international collaboration with currently around 200 people; now part of XLZD

Measuring Principle

- Multi-ton target TPC with liquid xenon for the direct detection of particle dark matter over the full WIMP-parameter space
- Measures WIMP-induced nuclear recoil spectra
- Also suited for other rare event searches, like neutrinoless double-beta decays, solar axions, or galactic supernovae

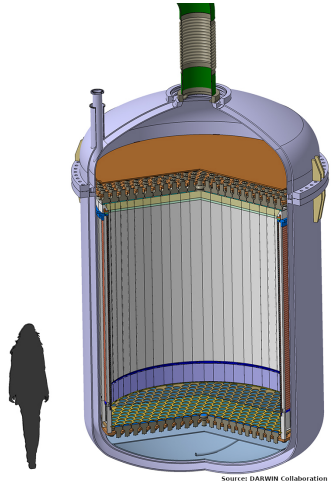

Computing in a New(-ish) Collaboration

Situation

- Typically **no central IT services** available yet
- **No dedicated** infrastructure or IT personnel
- Often **no joint computing platform**, but self-made solutions for different sub-groups

Requirements

- **Main computing needs** in R&D phase: simulations and framework development
- Accessible for **all members**
- **Lightweight** setup with simple deployment
- Good **scalability** according to the computing need

Computing in a New(-ish) Collaboration

Situation

- Typically **no central IT services** available yet
- **No dedicated** infrastructure or IT personnel
- Often **no joint computing platform**, but self-made solutions for different sub-groups

Requirements

- **Main computing needs** in R&D phase: simulations and framework development
- Accessible for **all members**
- **Lightweight** setup with simple deployment
- Good **scalability** according to the computing need

Our Solution: A Lightweight Analysis and Grid Facility

A **unified**, **future-proof** concept for **users** and **central production** needs of an experiment

Designing an Analysis and Grid Facility for DARWIN

Our Philosophy

- Has to be **easy to deploy** and **easy to use**
- Easily **extendable** for growing collaboration needs
- Rely on **existing and established** tools and experience gained from LHC Computing
- Use **modern**, **open source**, industry **standard** technologies

Designing an Analysis and Grid Facility for DARWIN

Our Philosophy

- Has to be **easy to deploy** and **easy to use**
- Easily **extendable** for growing collaboration needs
- Rely on **existing and established** tools and experience gained from LHC Computing
- Use **modern**, **open source**, industry **standard** technologies

Requirements from the User Point-of-View

To provide a common facility to support the full set of users (analysts and production), a flexible approach is required:

Requirements

- One common entrypoint, Single-Sign On
- Interactive development and analysis (e.g., as JupyterHub)
- Classic **SSH login** and **batch system** for users and central production
- Common **storage** entry point

Central Token-Based User Management for Accessing

Karlsruher Institut für Technologie

Interactive (web-based) and user access via CLI (SSH) is all handled with tokens!

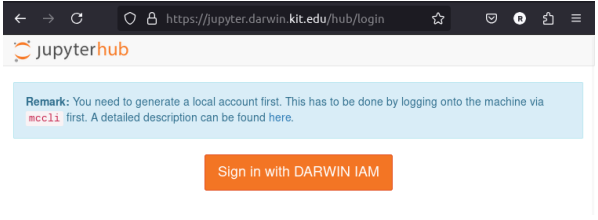

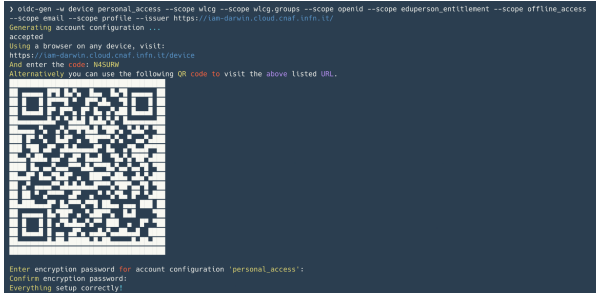

Token Authentication

- Detailed permissions handled via group memberships and protected scopes
- Approval of new users and user management done in IAM instance by a **collaboration manager**

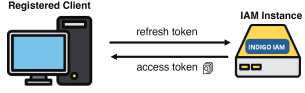

SSH with Tokens

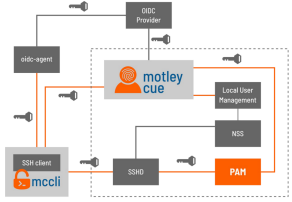

Client: mccli* and oidc-agent*

- **oidc-agent** allows users to obtain access tokens on CLI by registering the client as a device with an IAM instance
- **mccli** as a wrapper around the regular SSH client

Server: motley cue*

- Validate an access token
- Automatically create new local users
- Map the token to a **local unix user** and handle PAM authentication
- **LDAP** instance as backend for local user login

**Tools developed as part of the European Open Science Cloud*

Batch System and Opportunistic Resources

For larger scale **analysis** and **production**, resources are provided by GridKa via an **HTCondor batch system**

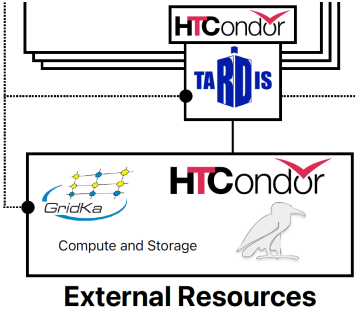

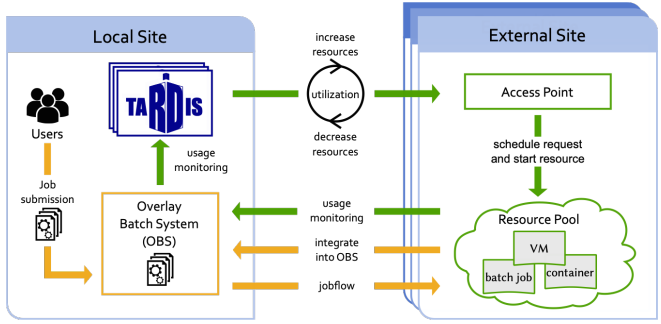

On-Demand Dynamic Allocation of Additional Resources

- Realized with our meta-scheduler COBalD/TARDIS
- Used in production at the German WLCG Tier-1 since **several years**

Scalability by Dynamic Integration of Distributed Resources

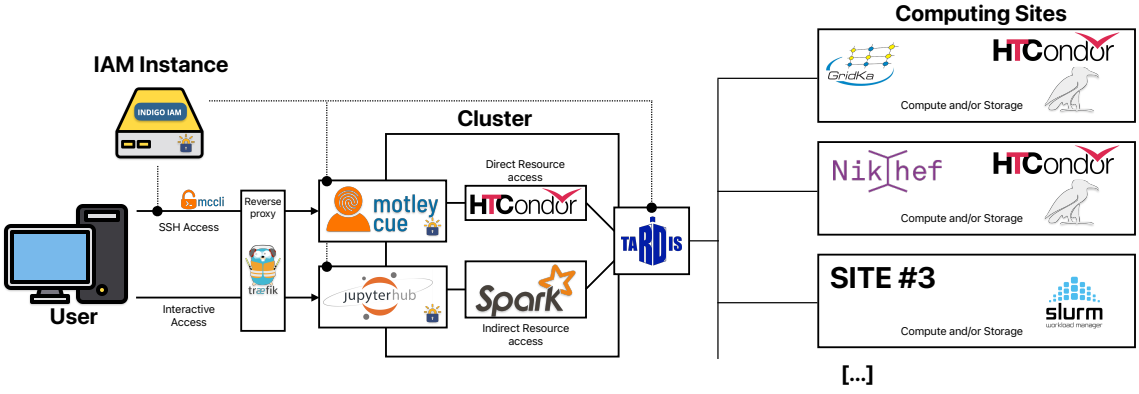

Further compute resources can easily be integrated via COBaID/TARDIS and are accessible via the OBS!

Management and Monitoring

Management

- The setup is managed by an additional **service node**
- Minimal effort setup with the provided **ansible** roles
- Small maintenance effort thanks to **containerization**

Monitoring

- A full stack with standard tools is used for monitoring
- Host and container monitoring
- User monitoring, e.g. local disk usage
- Self-made site tests

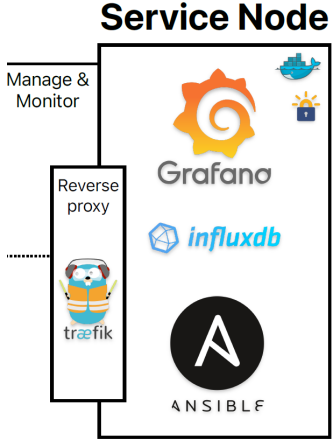

Prototype for an Analysis and Grid Facility for DARWIN

Full Prototype Setup:

- One **cluster node** as entry point (RHEL9, 96-Core EPYC)
- **Service node** for management and deployment
- IAM Instance for DARWIN provided by CNAF
- External resources can be dynamically integrated with COBalD/TARDIS
- Storage provided by GridKa

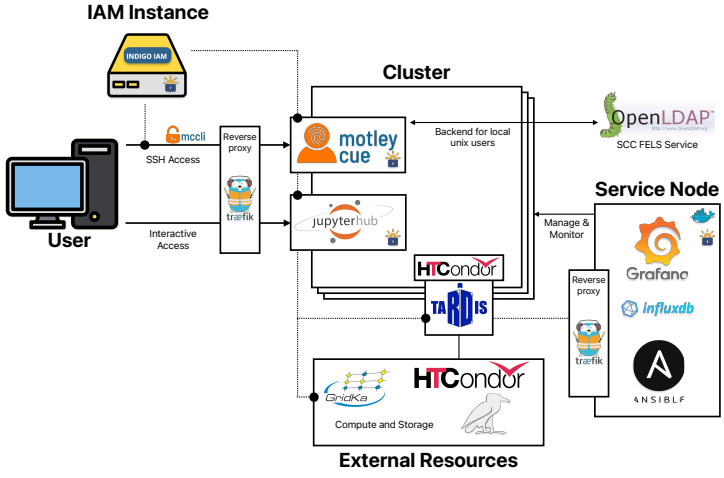

Conclusion and Outlook

- We provide a comprehensive, lightweight analysis and grid facility concept that can be easily adapted
- The concept covers all typical use-cases for a young collaboration

Future-Proof

- Fully **token** based setup with SSO
- High **scalability** for growing resource demands
- Leveraging modern **open source** technologies with broad community support

Management and Monitoring

- Easy, automatic deployment with **ansible** and containerization
- Comprehensive **monitoring** for a secure operation
- Administration effort is kept minimal

Outlook

- We are able to integrate further resources, e.g. from NIKHEF, when required
- The full ansible setup repository will be released **soon**

Prototype

- A running prototype instance for the DARWIN collaboration is **available**
- First users are testing the facility
- **We are ready for more!** ;-)

BACKUP

Storage

Local Storage:

- For software, code, etc
- Currently, 20TB NVMe on the machine itself
- CVMFS avaialble for access to software stacks and analysis containers

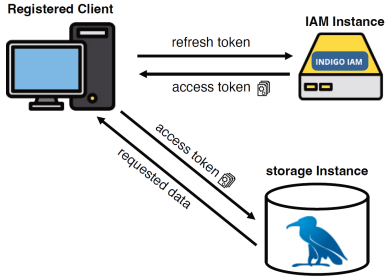

Remote Storage:

- **dCache** instance at GridKa for users and central production
- Data transfers via common grid tools, e.g. XRootD
- Storage access fully handled via tokens

Remote Storage, Jobs, and Tokens

- **Problem:** Access tokens are short-lived, refresh tokens are not meant to leave a registered device
- **Solution:** mytoken integration in HTCondor
- Automatic renewal via HTCondor mechanisms, mytoken does not leave the submit Node