GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Umsetzung von höheren Sicherheits-anforderungen mit Linux Bordmitteln

Unveränderliche Logdateien aller Server und privilegierte Admin Workstations

Sebastian Krey Freja Nordsiek Julian Kunkel

09.09.2025

Outline

1 GWDG and NHR

2 HPC Systems at GWDG

3 Privileged admin workstations

4 Central WORM logserver

5 Summary

About GWDG

- IT service center and data center operation for **University Göttingen** and **Max Planck Society** (MPG) since 1970
- Operating site of “North German Supercomputing Alliance” (**HLRN**) since 2018, since 2021 part of **NHR**
- AI Service Center **KISSKI** for critical infrastructure
- HPC operating site for the “German Aerospace Center” (**DLR**) since 2022

Network for National High-Performance Computing

NHR-NORD@GÖTTINGEN

UGOE + GWDG

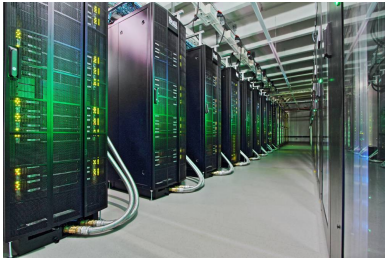

- Since 2021: Funding for national Tier-2 supercomputing (62.5M=C p.a.)
  - Nine centres
  - Annual funding 7,3M€ p.a.
- Usable for researchers at all German universities
  - Up to 1,200k CPU core/1500 GPU hours p.a. usable without application
  - Secure Workflow for processing sensitive data (medical, financial, etc.)
  - Larger projects require application
    → `https://docs.hpc.gwdg.de/start_here/nhr_application_process/index.html`

HPC systems at GWDG

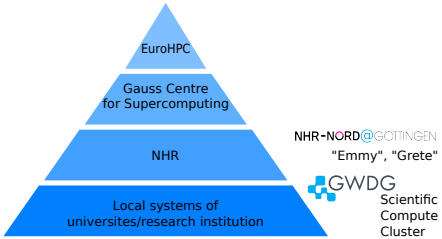

- Tier 2: **HLRN/NHR “Emmy”**
  Top500 #47 Nov. 2020, now #219
- Tier 2: **NHR/KISSKI “Grete”**
  Top500 #141 Nov. 2023, Green500 #22, now #226/#70
- Tier 2: **NHR/KISSKI “Grete Phase 3”**
  Top500 #274 Nov. 2024, Green500 #24, now #315/#35
- Tier 3: **Scientific Compute Cluster**
- **“CARO”** for **DLR**
  Top500 #135 Nov. 2021, now #335
- Several smaller systems for MPG and UGOE

Privileged admin workstations

What are privileged admin workstation

- Terminology from the Microsoft world
- Centrally managed devices for administration of critical infrastructure
- Maximize privilege separation (admin account of PAW has no admin privileges on the critical infrastructure and vice versa)
- Dedicated network segment (VPN profile)
- Restricted network access (no internet access)
- Restricted software availability (only tools necessary for the task)
- Logging of configuration changes and software installations

Why and for what use PAW?

- In MS world: All tasks requiring domain admin privileges (root on all nodes of the domain)
- Jumphost for external admin access
- Central logserver
- Key management systems
- Build infrastructure for OS images
- Network management

What is required and desired for setting up a PAW?

- Required
  - Automated deployment and configuration of PAW
  - Monitoring of configuration and software installations on the PAW
  - Central logging and analysis infrastructure for monitoring data
  - Network or VPN profiles for accessing only the critical infrastructure via PAW
  - Automated user setup
- Desired
  - No complicated device management infrastructure
  - Personalized work environment on PAW should be possible
  - Software installation from white listed pool

GWDG HPC solution

- Management:
  - Mainly based on Fedora and Kickstart
  - Installation ISO and repo config from trusted mirror
  - All configurations in internal Git repository
  - Git repository requires signed commits
  - List of approved keys for signature also in Git
- Configuration
  - UEFI Secure boot
  - Systemd-boot without kernel commandline editor
  - `root` account locked, LUKS, preconfigured `sudo`, `firewalld`, etc.
  - Removal of remote management tools (SSH server, cockpit, etc.)
  - `pam-u2f` for 2FA
  - User setup includes eduroam and eduVPN setup
  - Shell script for software installation kiosk
  - Shell script with systemd timer for monitoring

Central logging infrastructure

- Graylog based
- Daily Systemd timer executes monitoring script and creates JSON
- Logging of Laptop type, Serial number, BIOS information
- `rkhunter` results
- All software packages with version (package manager and flatpaks)
- All executeables with set SUID, SGID or capabilities

Central write once logserver

Why WORM logs?

Processing risk class D information (e.g. medical data) require auditable access logs. This means:

- Central log aggregation
- Verifiable integrity
- Redundant long term storage
- Prevent changes to log files
- No deletion from single person

How to solve the requirements?

Central: Redirect Journald to Rsyslog, send Rsyslog to central server, individual files for each server

Integrity: Daily logfiles, create checksums

Redundant storage: Offsite backups (incl. checksums) on tape

Deletion prevention : `root` is god $\rightarrow$ difficult (can even change SELinux settings)

File modificationdeletion prevention

Modificationdeletion prevention from normal users easy:

- Extended attributes `append-only` and `immutable`
- Creating new log files with apppend`-only` prevents deletion of older log entries
- Adding `immutable` attribute at the end of the day (before checksum creation) to prevent deletion
- Integrity check of scripts handling these operations with checksums

Superuser `root` still has permission to remove these extended attributes
$\rightarrow$ can modify or delete files and hide it by creating new checksums.

Further integrity enhancing measurements

Securing remote `root` logins:

- Changing the append-only and immutable attributes require `CAP_LINUX_IMMUTABLE`
- Removing this capability from sshd via (immutable) service override
  $\rightarrow$ `root` via SSH has lost the permission to change append-only and immutable attribute
- Remote administration via SSH still possible
- Prevent usage of IPMI remote and serial console, e.g. no network connection
- Prevent local IPMI usage/configuration change by disabling OS access to BMC (BIOS setting)

Further integrity enhancing measurements

Securing local accesss:

- Server location in accesss restricted area of data center
- Persons with access to data center area must not know password for root or sudo enabled user accounts
- Persons with local login permissions must not have permission to access the restricted data center area alone
- Additionally require 2FA for local login (e.g. FIDO Key via pam-u2f)

Summary

- Privileged admin workstations:
  - PAWs have sensible usage scenarios
  - No difficult device management infrastructure for setup necessary
  - Simple solutions based on known tools like Kickstart, Graylog and shell scripting allows easy review and good understanding of the concepts
  - Adaptable to different scenarios
- Central WORM logs:
  - Auditable logs good for access control as well as forensics
  - Most requirements easy to implement
  - Restricting delete permissions for `root` possible but not easy
  - Cost intensive special purpose hardware can be prevented