

# GitLab: update on status and plans

ASDF 4-Nov-2021

Alex Lossent

# Meet the GitLab team

Alexandre Lossent as service manager. Core team:

— **Emmanouil Fokas**


— **Ismael Posada Trobo**

— **Carina Antunes**

---



Manos is handling the day to day infrastructure operations and the setup of the new infrastructure.



Ismael leads the deployment of the new infrastructure.



Carina is the newest addition to the team and investigates alternatives to Premium features.

# De facto platform for development at CERN

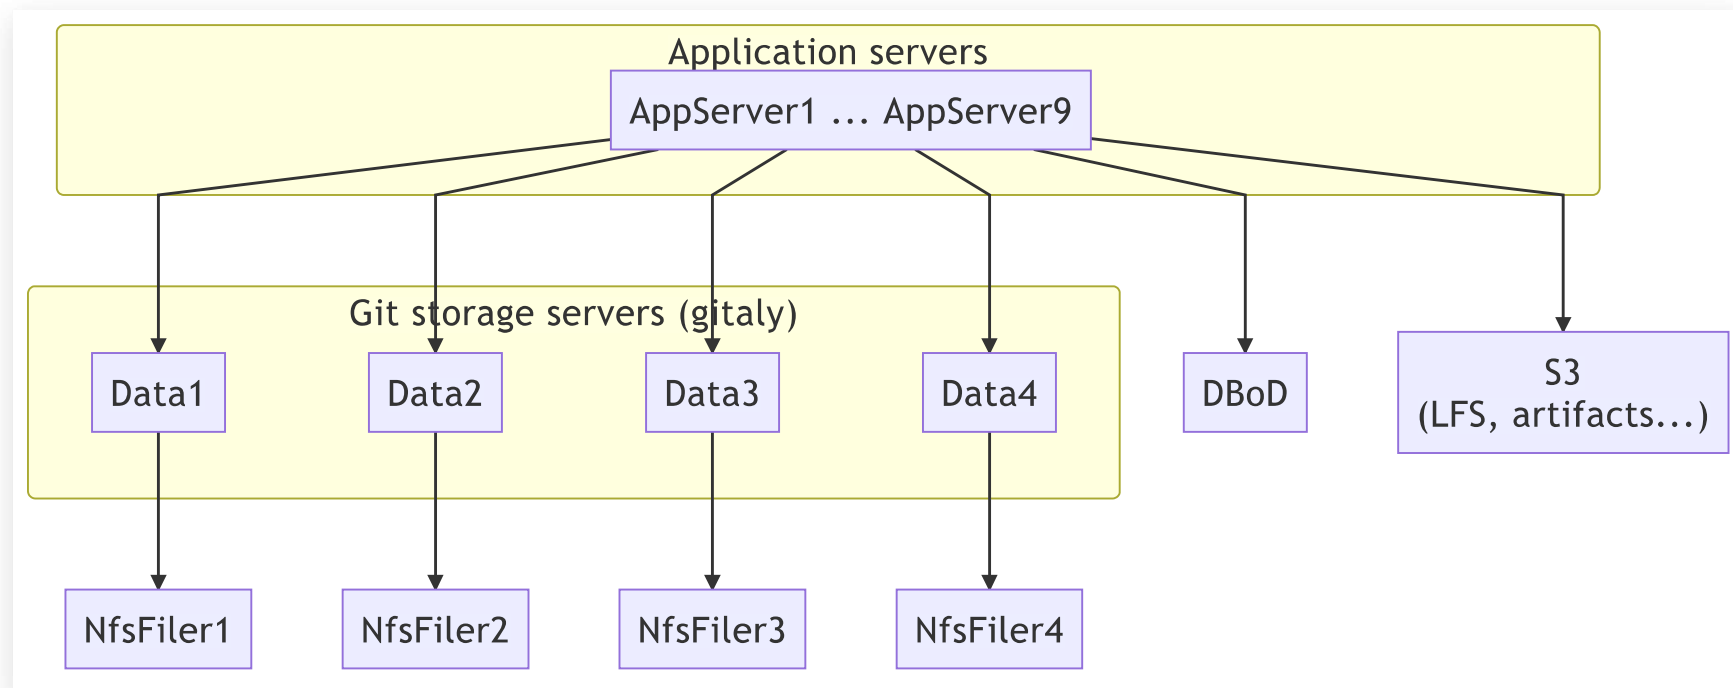
- Attract and consolidate development workflows
  - 103k projects (git repos 7.5TB, other: 100+TB)
  - 16k users (+ 8k archived CERN accounts)
  - ~450 req/s
- Enable widespread tests and automation
  - 16k CI jobs/day

# Recent highlights

- Package repositories:
  - Conan, Maven, NPM, NuGet, PyPI, RubyGems
- Unicorn to Puma: 150->100ms avg response
- Prototype of new shared CI runners by IT/CM
  - GPU, TPU, ARM

# Deployment

- Home-made containerized deployment in OKD3



# What's next

Q1 2022

- New GitLab deployment
- Switch to new SSO

2022

- Prepare for switching to Free GitLab

# New deployment

- Why?
  - Openshift 3 EOL
  - NFS filers deprecated in favor of CephFS
  - *standardize*: upstream Helm chart on Magnum Kubernetes

# New deployment

- What to expect? (as GitLab user)
  - Probably a couple hours downtime to switch
  - Otherwise, aiming at *identical functionality/feature set*
  - Possible initial downsides, should improve later on:
    - <1min downtime for GitLab updates
    - Downtime for Kubernetes cluster updates
    - IPv6: unsure yet
- Benefit: facilitate long-term maintenance



# GitLab Licenses

- Historically, CERN subscribed to GitLab Starter license
  - LDAP groups (e-groups)
  - Support
- Forced upgrade to Premium in 2020
  - Starter discontinued
  - New license deal running until March 2023
  - List price is ~50x higher

# Preparing for GitLab Free

- Why?
  - Unbearable risk of exploding license costs
- What to expect? (as GitLab user)
  - No loss of information
  - No change to core workflows: git repos, code review and CI
  - Premium features cannot be used
- Benefit: ensure the perennity of development workflows built around GitLab

# What exactly would stop working

- <https://about.gitlab.com/pricing/self-managed/feature-comparison/>
- Very widely used:
  - LDAP sync to use e-groups for permissions
  - Kerberos for passwordless git operations
  - Git repository mirroring

# Use e-groups for permissions

- Must-have
- Early design of custom sync solution - promising
- Opportunity: transition from LDAP/e-groups to Grappa
- Estimated effort: 2-3 person-months

# Passwordless git operations

- Kerberos seems less relevant today
  - At odds with 2FA, federated accounts
- Authentication tokens are now common practice
  - Replacement for Kerberos
- Opportunity:
  - Users: align with common practice (GitHub...)
  - Infra: greatly simplify deployment
- Users: one-time git configuration change
- Infra: communication

# Git repo mirroring

- Very popular, e.g. GitHub projects
- Seems to justify a dedicated service
- Opportunity: mirror whole GitHub orgs
  - Archive and improve visibility/discoverability of CERN code in GitHub
- Estimated effort: ~3 person-months

# Other Premium features

- Most Premium features are not used (never enabled)
- A lot have workarounds/alternatives
  - e.g. Push Rules => “code quality” CI jobs
- The number of projects in need of significant workflow changes seems limited
  - e.g. Merge Trains, Run fork pipelines in parent project...
- A few person-months to document alternatives, help users transition workflows

# Service delivery: opportunities

- Some user communities have expressed interest in purchasing Premium licenses for themselves
- New deployment enables segmentation of the service
  - E.g. provide a dedicated second instance
  - Moving projects may have a nontrivial cost for users (update all references in tool chain...)



# Support from GitLab

- Not a concern nowadays
  - Product maturity: no need to contact GitLab support for the last couple years
- Powerless at pushing forward feature requests important to CERN user community

# Timeline

- Q4 2021: raise awareness of the plan (ITUM, coordination meetings...)
- Q1 2022: new deployment, new SSO
- H1 2022: prototype LDAP/Krb/Mirror replacements
- 2022: align with department strategy/decision process
- End 2022 (probably?): license negotiation, last call for decision to downgrade, progressively shut down Premium features
- March 2023: end of current license deal

# Conclusion

- Technically, we have a plan to mitigate the risk and ensure the perennity of workflows built around GitLab
- Obtain sign-off by the user community

