

RADU POPESCU

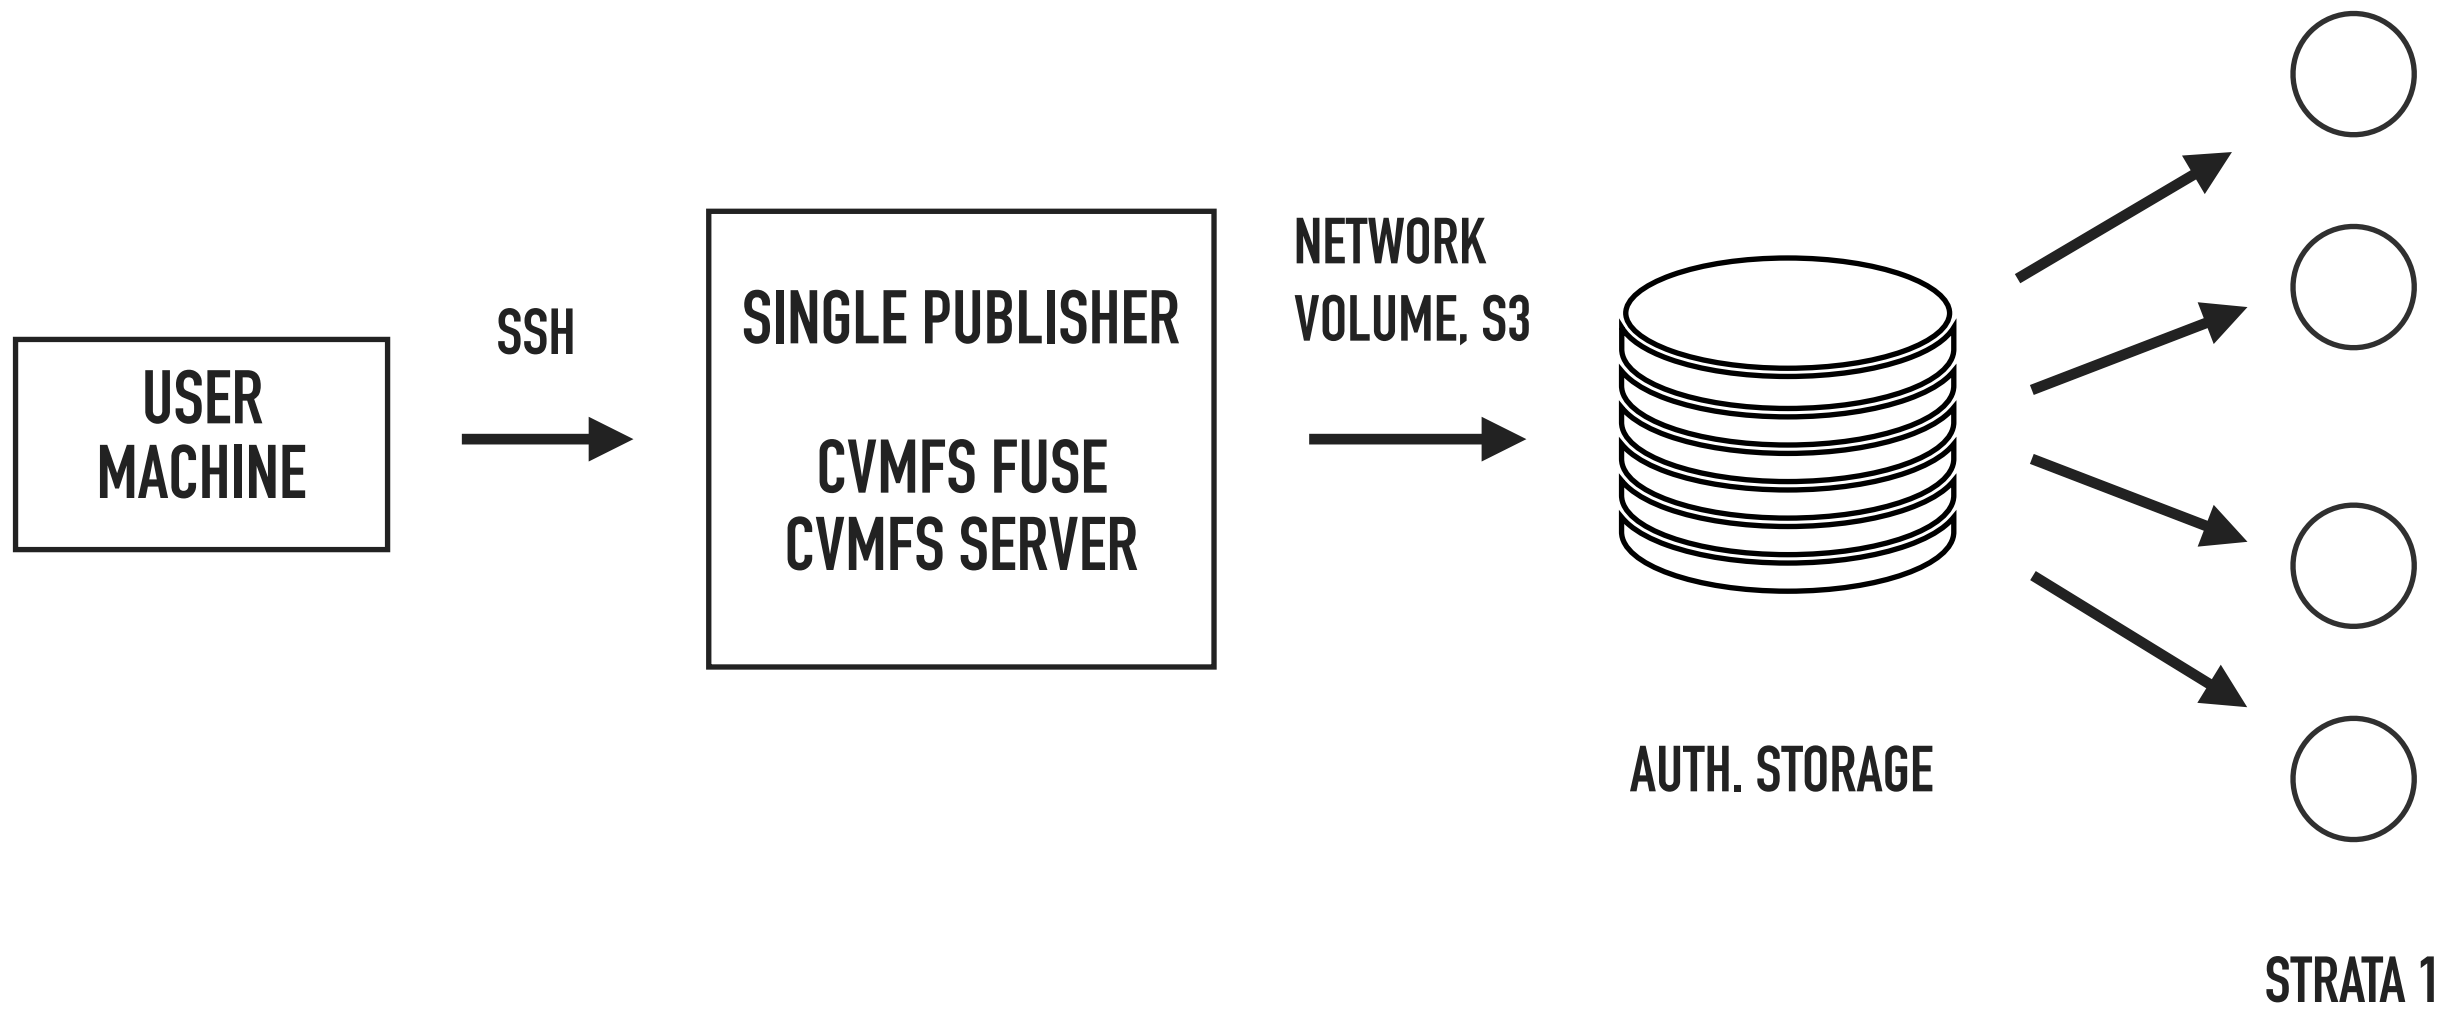
EVOLUTION OF THE GATEWAY SERVICES

CERNVM WORKSHOP, JUNE 2019

DESIGN OF THE CERNVM FILE SYSTEM

- CernVM-FS is not a general purpose file system
- Optimised for scalable reading of content
- Publish/subscribe model with a single source of new content per repository and many concurrent clients
- Command line publisher tools

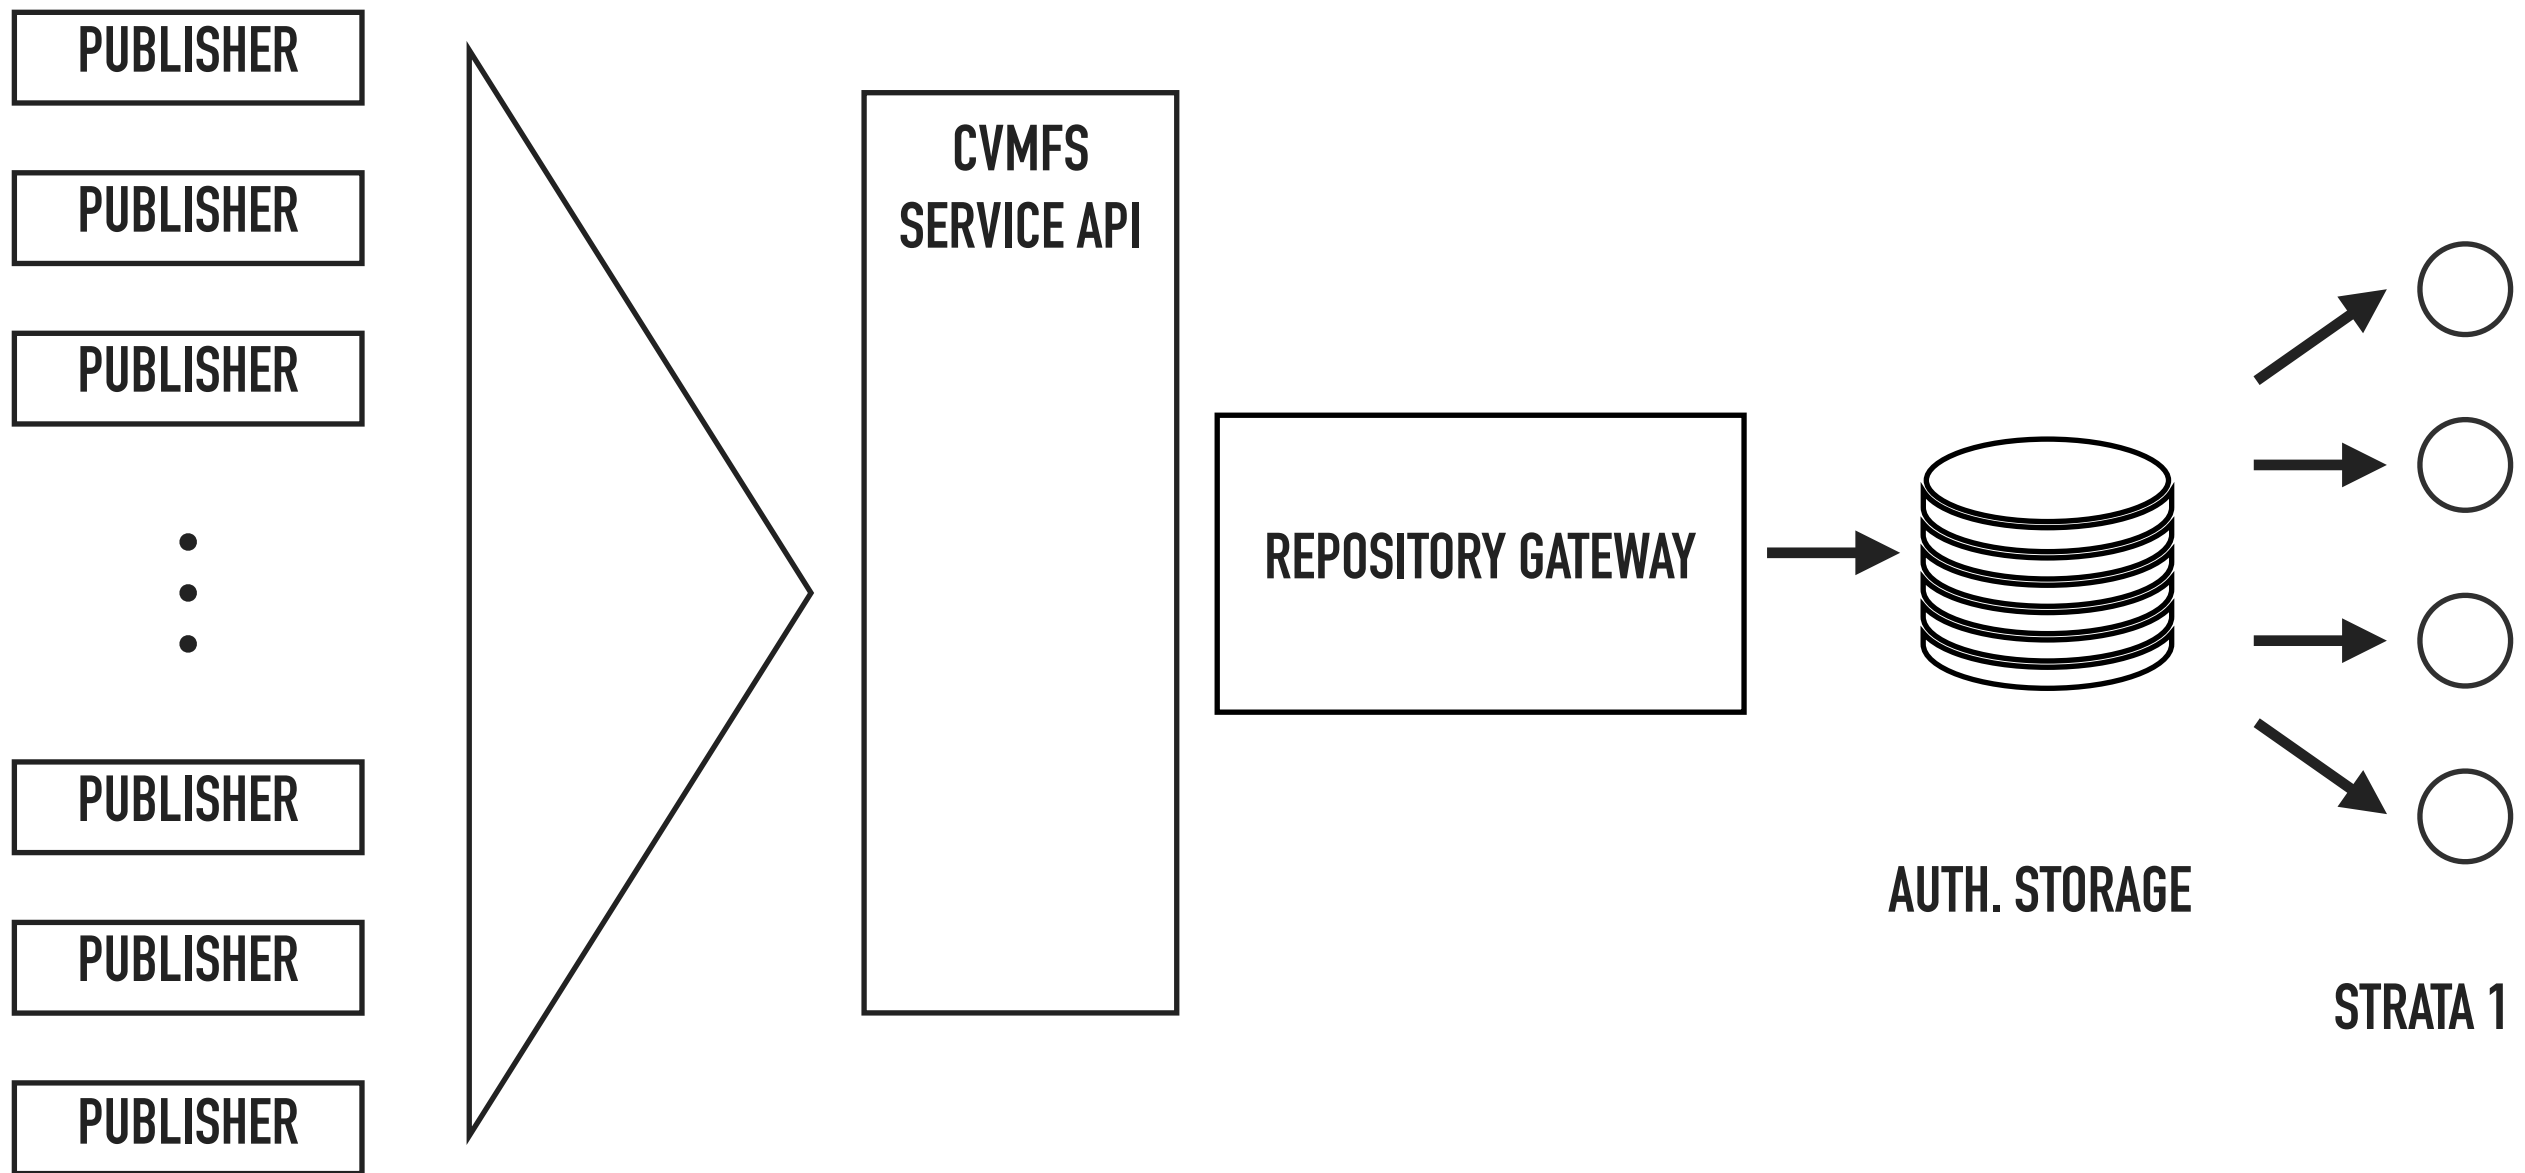
SINGLE WRITER PER REPOSITORY



SINGLE WRITER PER REPOSITORY (CONT.)

| PROS | CONS |
|----------------------|--------------------------------|
| Robust | No concurrency |
| Simple | Shell access |
| Vertical scalability | No fine grained access control |
| | Horizontal scalability |

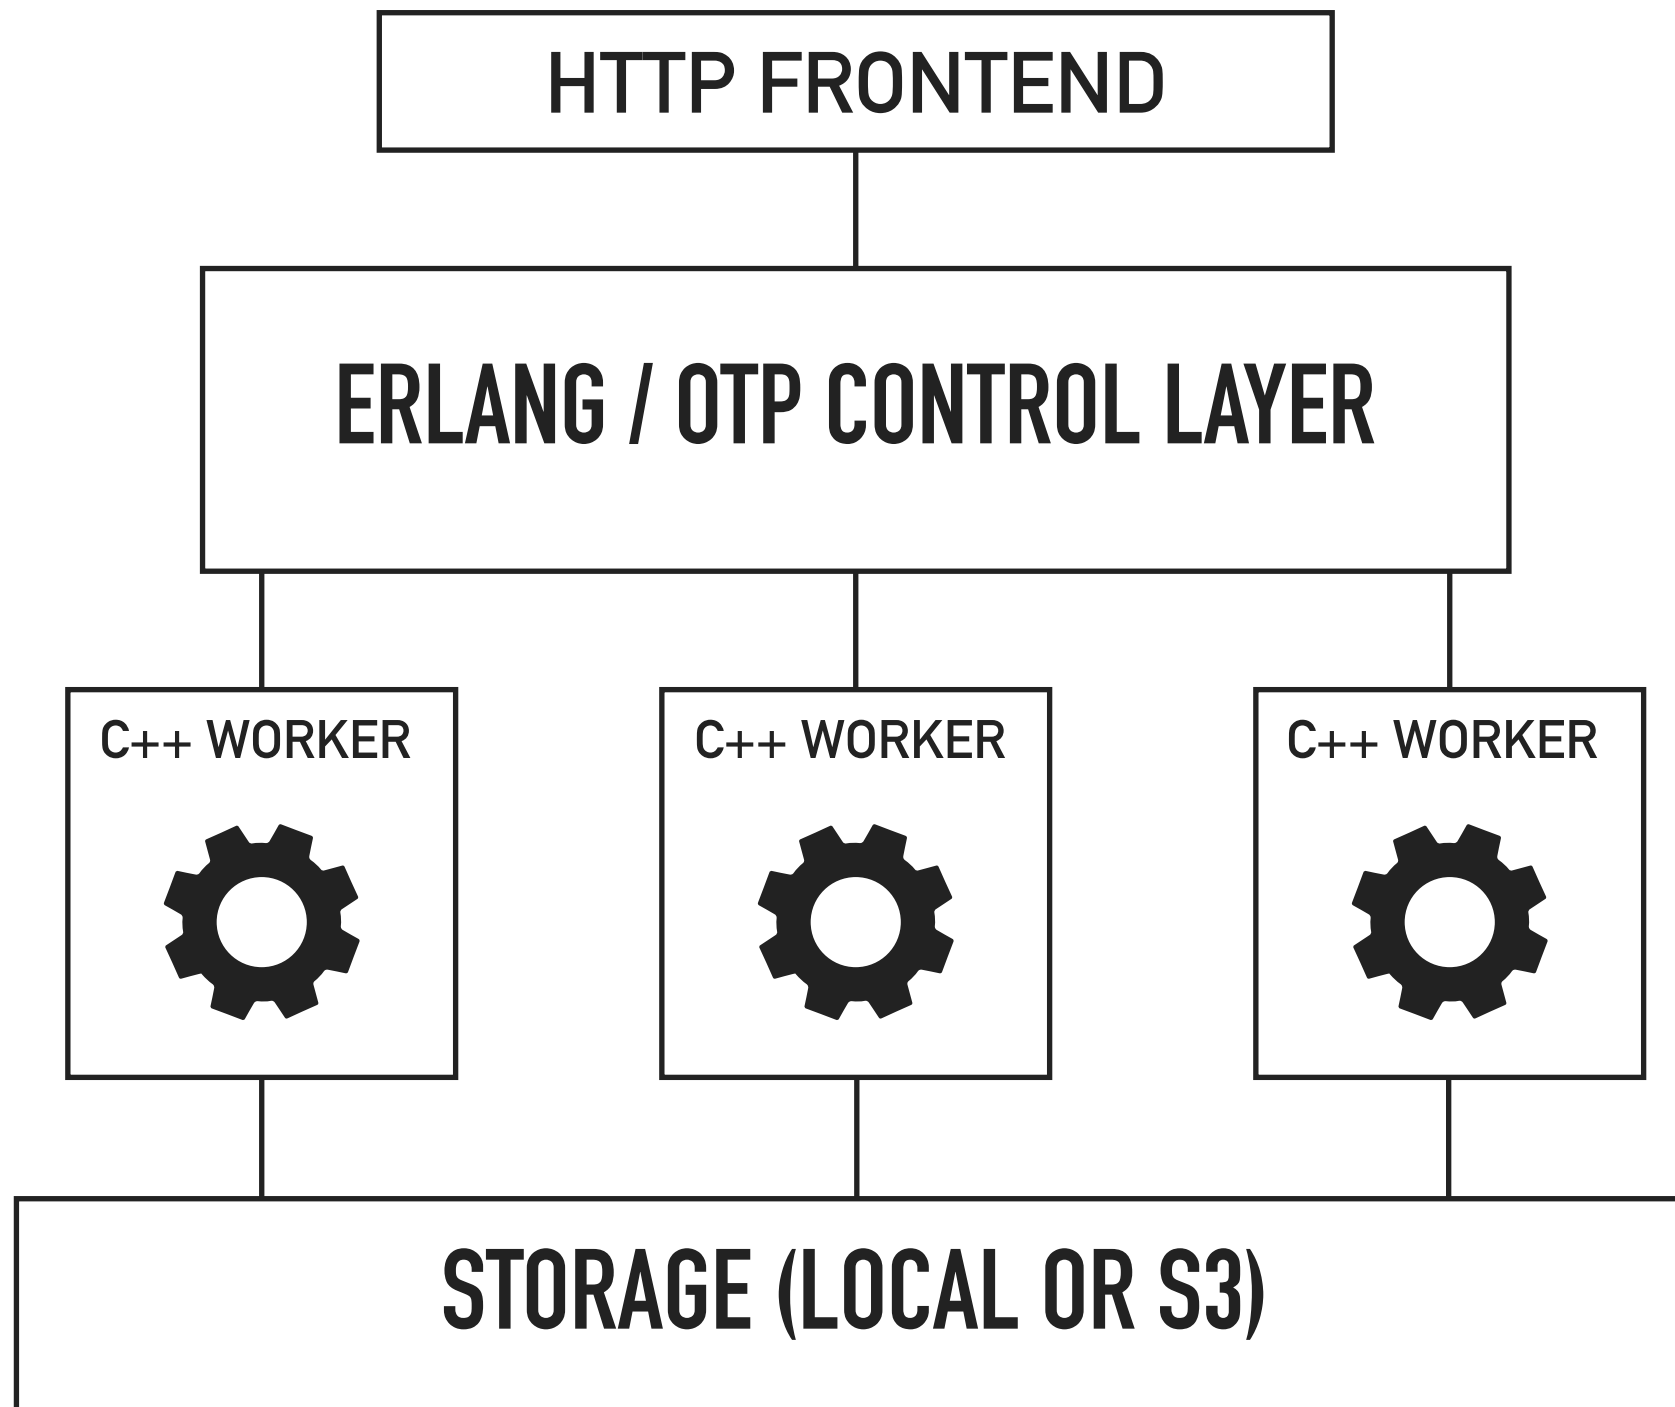
PUBLISHING CONCURRENTLY WITH THE REPOSITORY GATEWAY



THE REPOSITORY GATEWAY

- Mediates access to repository from multiple publishers
- Allows horizontal scalability
- Access control
- Enforces consistency
- HTTP API
- First made available in CVMFS 2.5

ARCHITECTURE



DEVELOPMENT AND OPERATIONAL IMPRESSIONS

- Erlang (language) is very well suited for the domain
- Excellent concurrency
- Simple to interact with worker processes

DEVELOPMENT AND OPERATIONAL IMPRESSIONS

- Erlang (language) is very well suited for the domain
- Excellent concurrency
- Simple to interact with worker processes

But:

DEVELOPMENT AND OPERATIONAL IMPRESSIONS

- Erlang (language) is very well suited for the domain
- Excellent concurrency
- Simple to interact with worker processes

But:

- Ecosystem is not growing
- Tooling is less developed
- Deployment is more complicated
- Predates containers

ERLANG COMPONENTS WERE REWRITTEN IN GO

- Go community is very active, rich ecosystem
- Go verbose, but is easy to pick up, maintainers should have a simpler task
- Good performance and concurrency story
- Good tools (testing, benchmarking, profiling)
- Simple to deploy, container-friendly

RELEASE PLAN

- The repository gateway reached version 1.0 with CVMFS 2.6.0
- Rewrite will be released as version 1.1 (full backward compatibility, upgrade is transparent)
- Bug fixes, performance and scalability improvements, reduced memory footprint
- New features!

ADMINISTRATION OPERATIONS


New API endpoints allow the administrator to:

- Disable repositories
- Inspect and cancel leases
- Trigger garbage collection from publisher machines
- ?

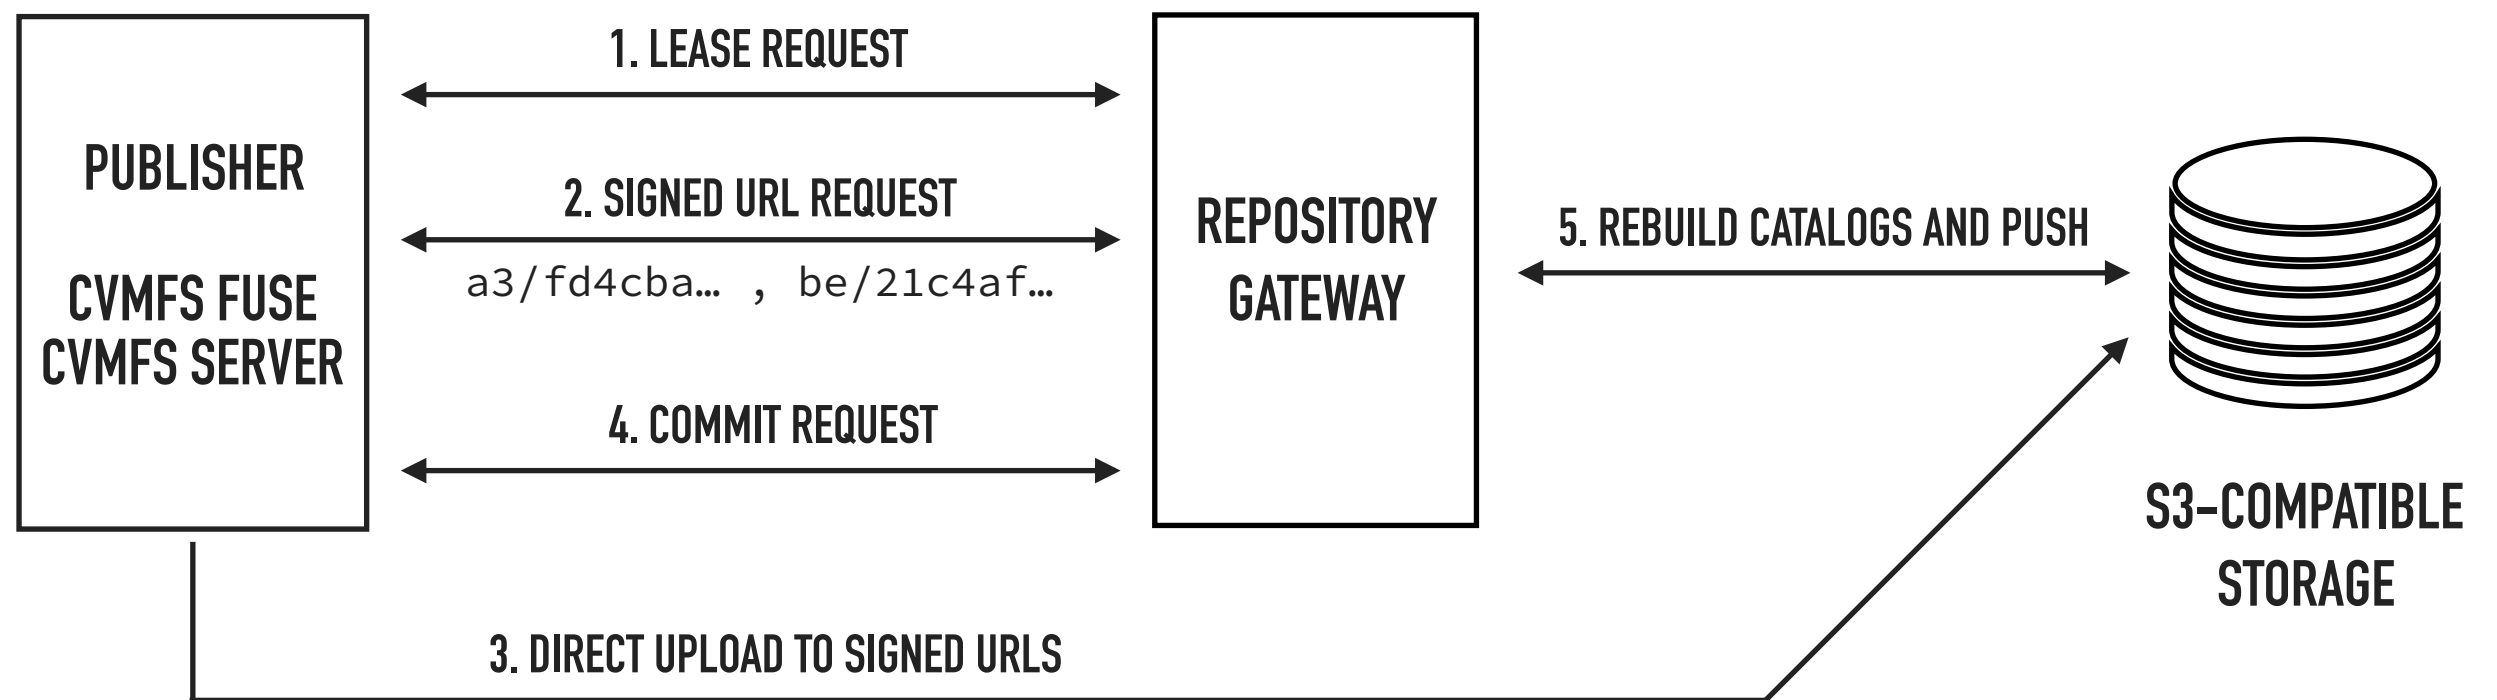
CONSOLIDATION OF SERVICES

- The repository notification server is merged into the repository gateway
- The "Conveyor" job server could also be merged
- Single "cvmfs-gateway" package to install

“STANDARD” PAYLOAD SUBMISSION



DIRECT UPLOAD TO S3 BACKENDS (EXPERIMENTAL)



Requires support in cvmfs-server - 2.6.1 or 2.7

SUMMARY

- A full rewrite of the repository gateway should reduce maintenance workload, more future proof
- Goal is to offer a complete repository API
- New functionality: admin operations, GC
- A collection of improvements in preparation for the “cvmfs enter” world
- Scheduled for CernVM-FS 2.6.1 or 2.7!

THANK YOU!