

OSG Technology - Years 3-5

Brian Bockelman

“Big Picture” Goals

- Finish HTCondor-CE transition. Years 3, 4
- Tackle “common input data” problem. Years 3, 4, 5
- More flexible software delivery (OASIS). Year 3, 4
- Integrate BOSCO / HTCondor-G+SSH into the production grid. Year 3
- Modernize and solidify authz support. Year 3, 4, 5
- Blueprint meetings - reboot and reorganize by M. Livny. Year 3

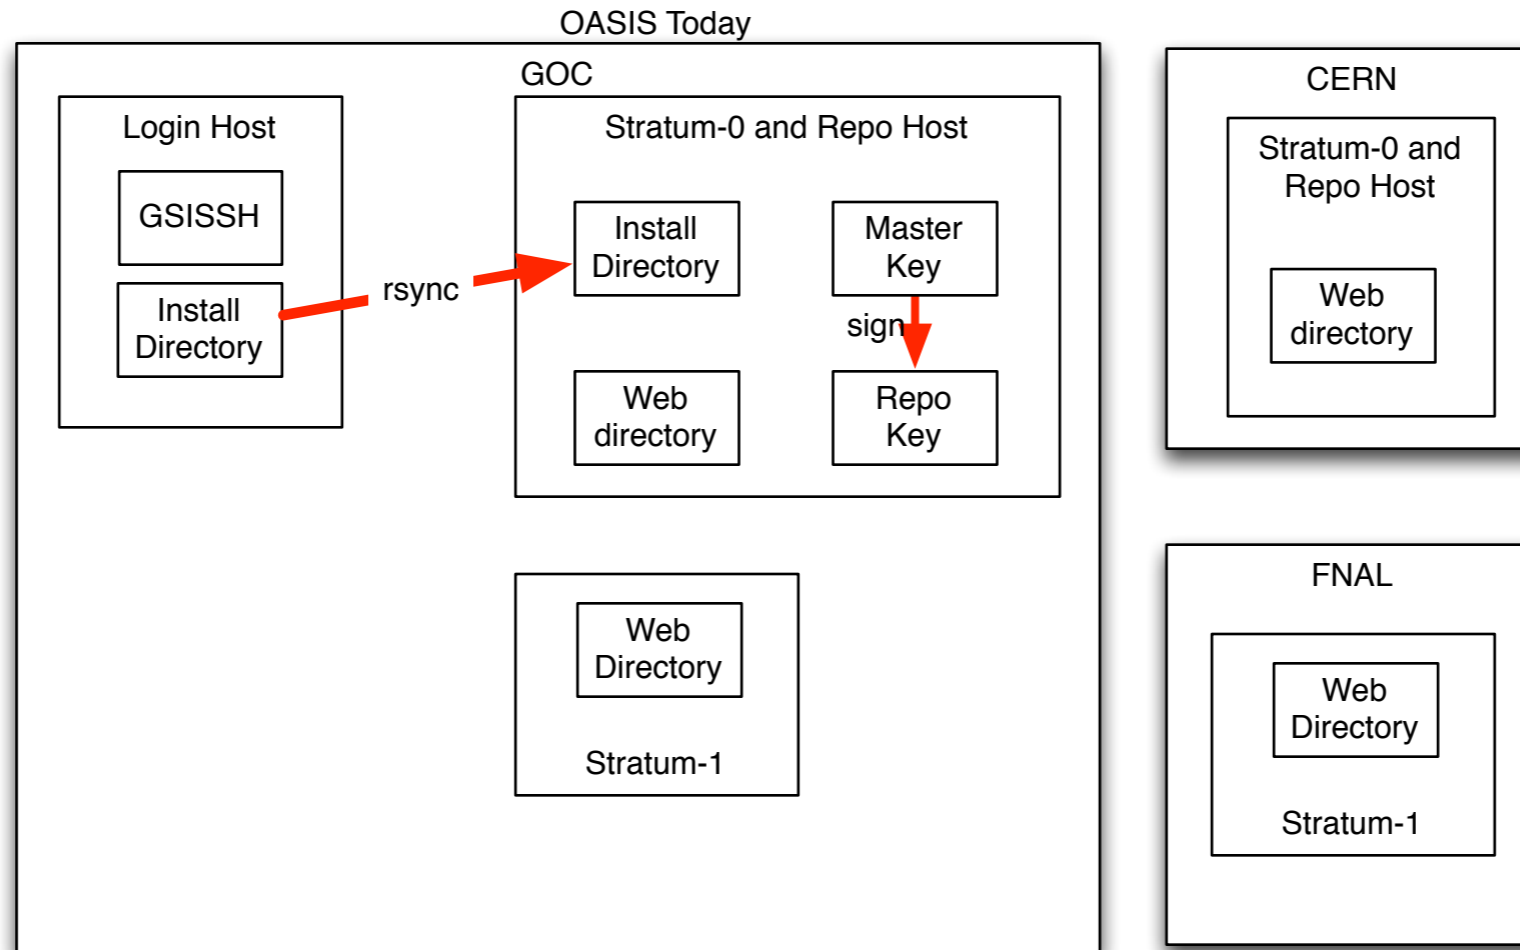
Finish HTCondor-CE Transition

- **Decisions needed ~ now:** cutover default documentation and recommendations to HTCondor-CE.
 - In Year 3, I would like to start removing GRAM from our default osg-ce package.
- **Blockers for some WLCG sites:** SAM tests only use glideWMS. No progress since 12 months; switchover was supposed to happen this month. GRAM and HTCondor-CE can run side-by-side.
- **Blocker for some sites:** Testing on LSF and SGE has been minimal.
- Other than those two items, sites should be able to switch.
- **Info services:** I'd like to have all HTCondor-CE installs report to a central OSG collector; use built-in mechanisms instead of separate software like GIP or osg-info-services. **Year 4.**
- **Effort for Year 3:** Marian Zvada (UNL) has been doing hands-on site support.

Improve software distribution

- We now allow VOs access to a central server where they can publish updates into CVMFS.
 - GOC hosts Stratum-0, repo server, login server, and Stratum-1 server.
 - Stratum-1 servers are also located at BNL and FNAL.
 - Stratum-0 and repo servers are also located at CERN.
- Vocabulary:
 - **Stratum-0**: Master key that signs the repo server. Done once a week.
 - **Repo server**: host which takes files installed by user and generates appropriate CVMFS metadata and web directory structure.
 - **Stratum-1**: host(s) which synchronize from multiple repo servers and exports data to the sites.
 - No site should contact the stratum-0 directly!

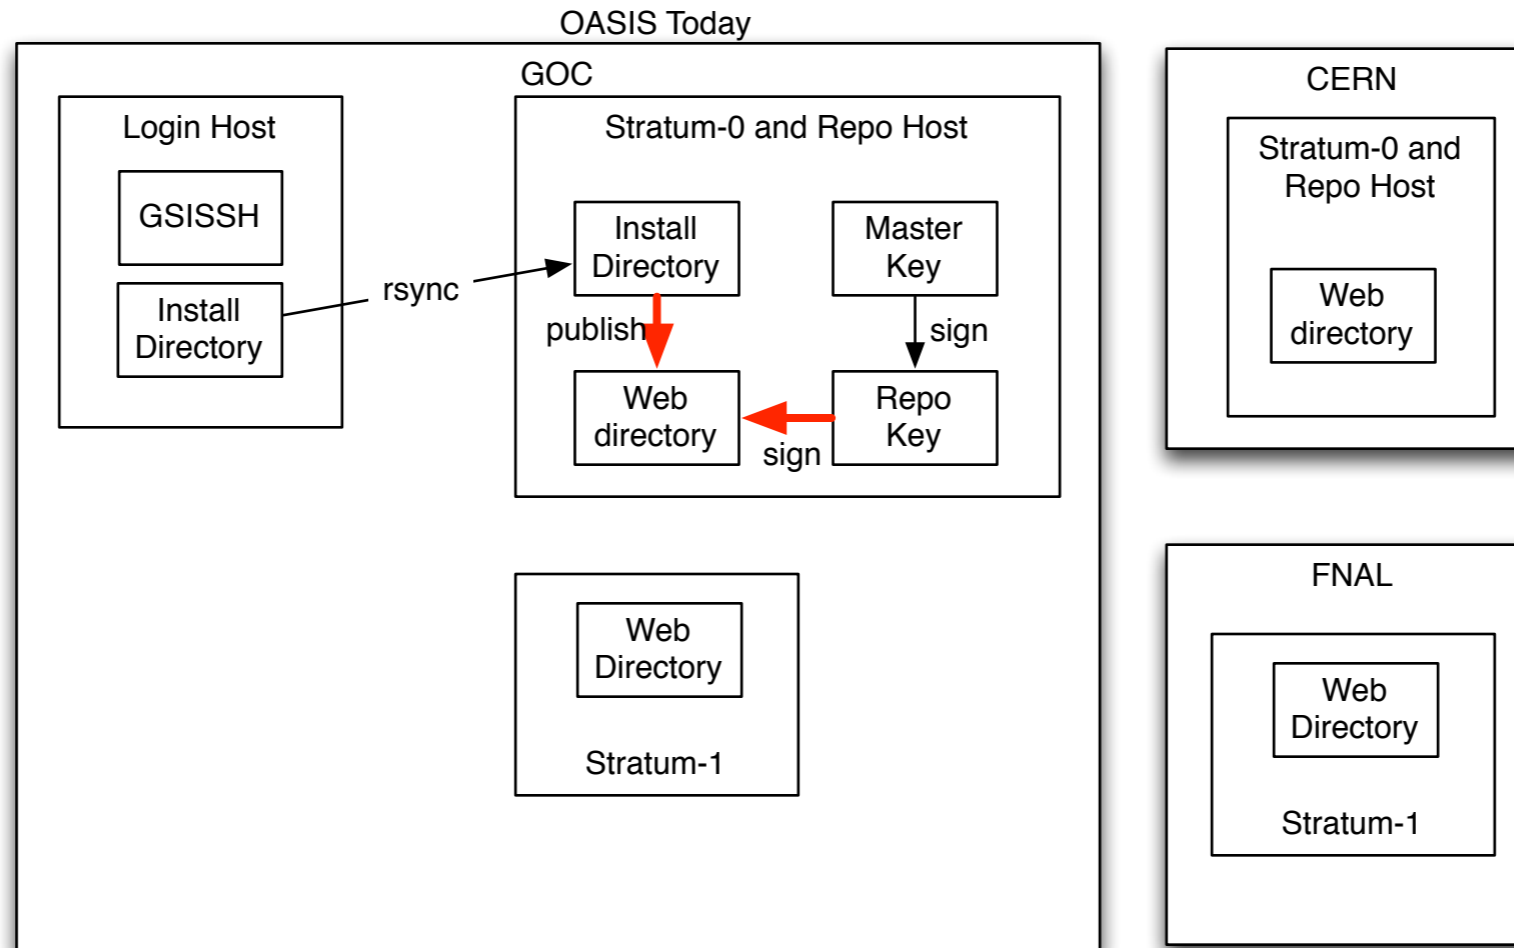
OASIS Today



Notes:

- BNL Stratum-1 not shown here.
- Rsync happens on user request.
- Key signing happens once a week.

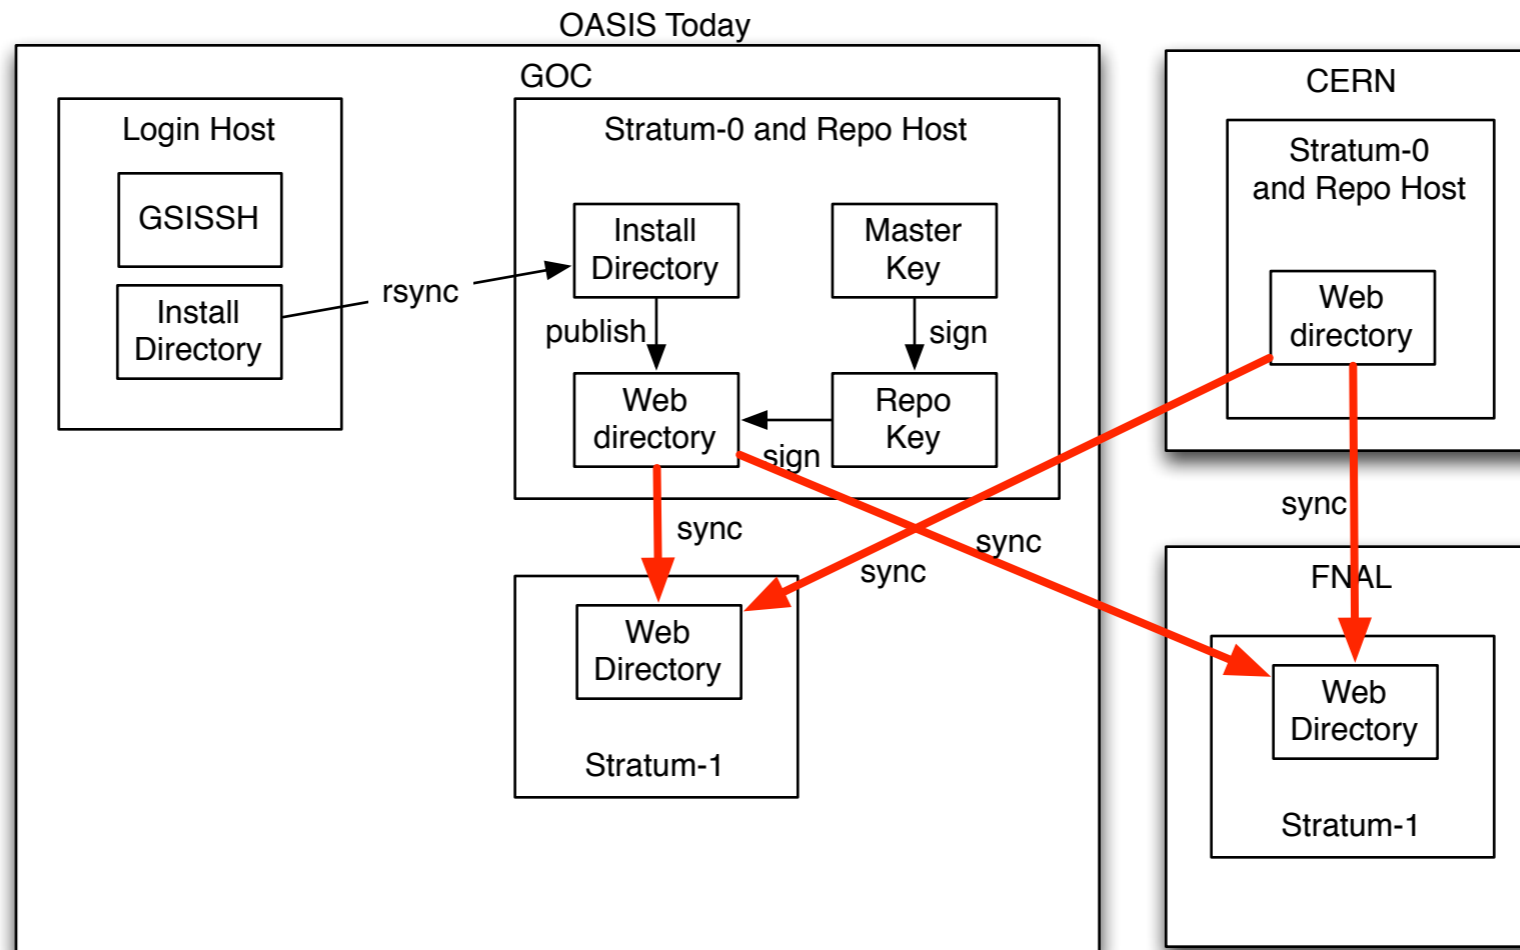
OASIS Today



Notes:

- Publish and sign happens after rsync completes.

OASIS Today



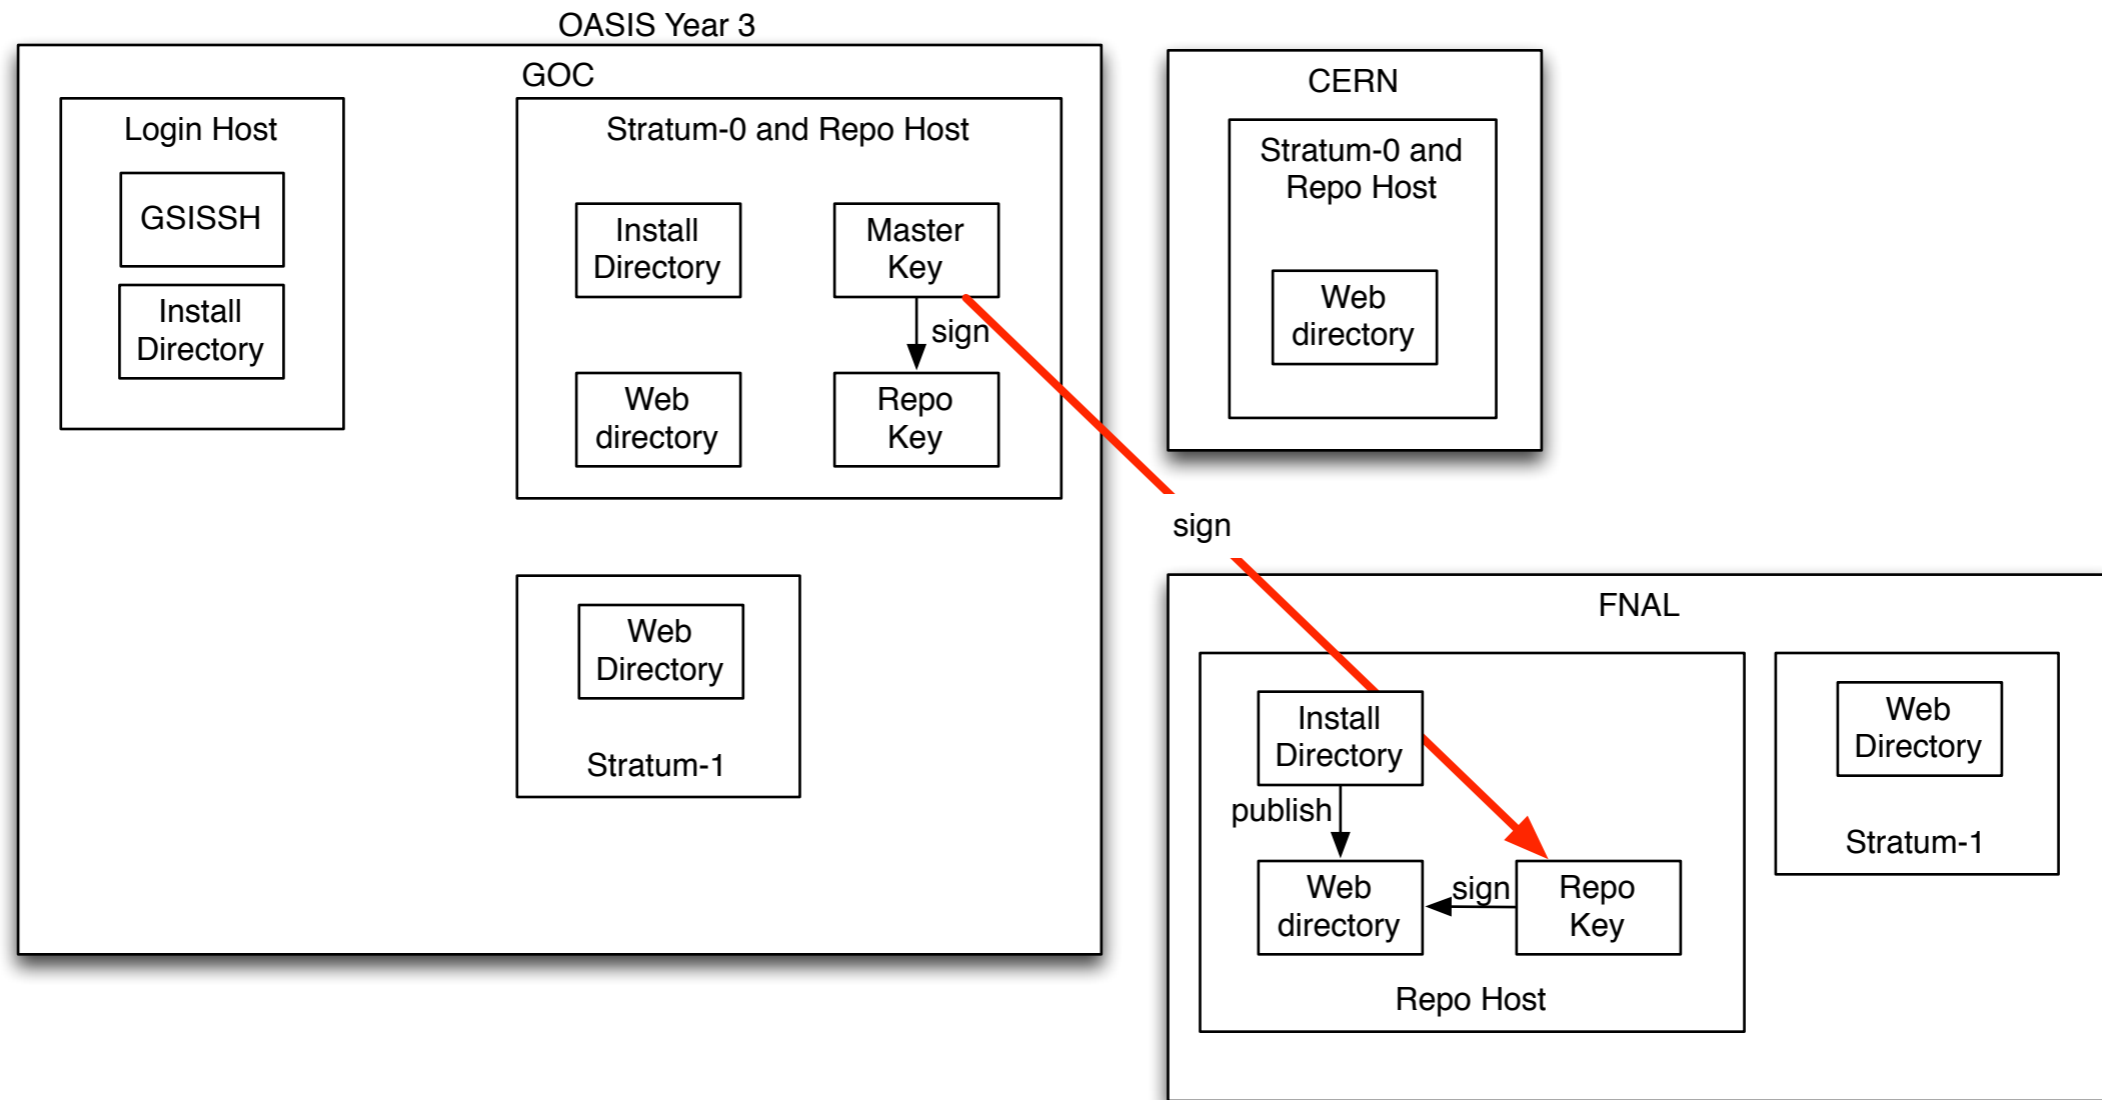
Notes:

- sync happens at a set interval (about 15 minutes).
- Each Stratum-1 can sync multiple repo servers.
 - For example, I show CERN, which hosts the WLCG repos.

OASIS - Year 3

- Remove references to \$OSG_APP in all documentation and RSV.
- Add support for external repo servers that are signed by OSG.
- Year 3 should target 1-2 external repos (**proposal**: start with IF and OSG-Connect).
- “*Install locally; run globally*”.

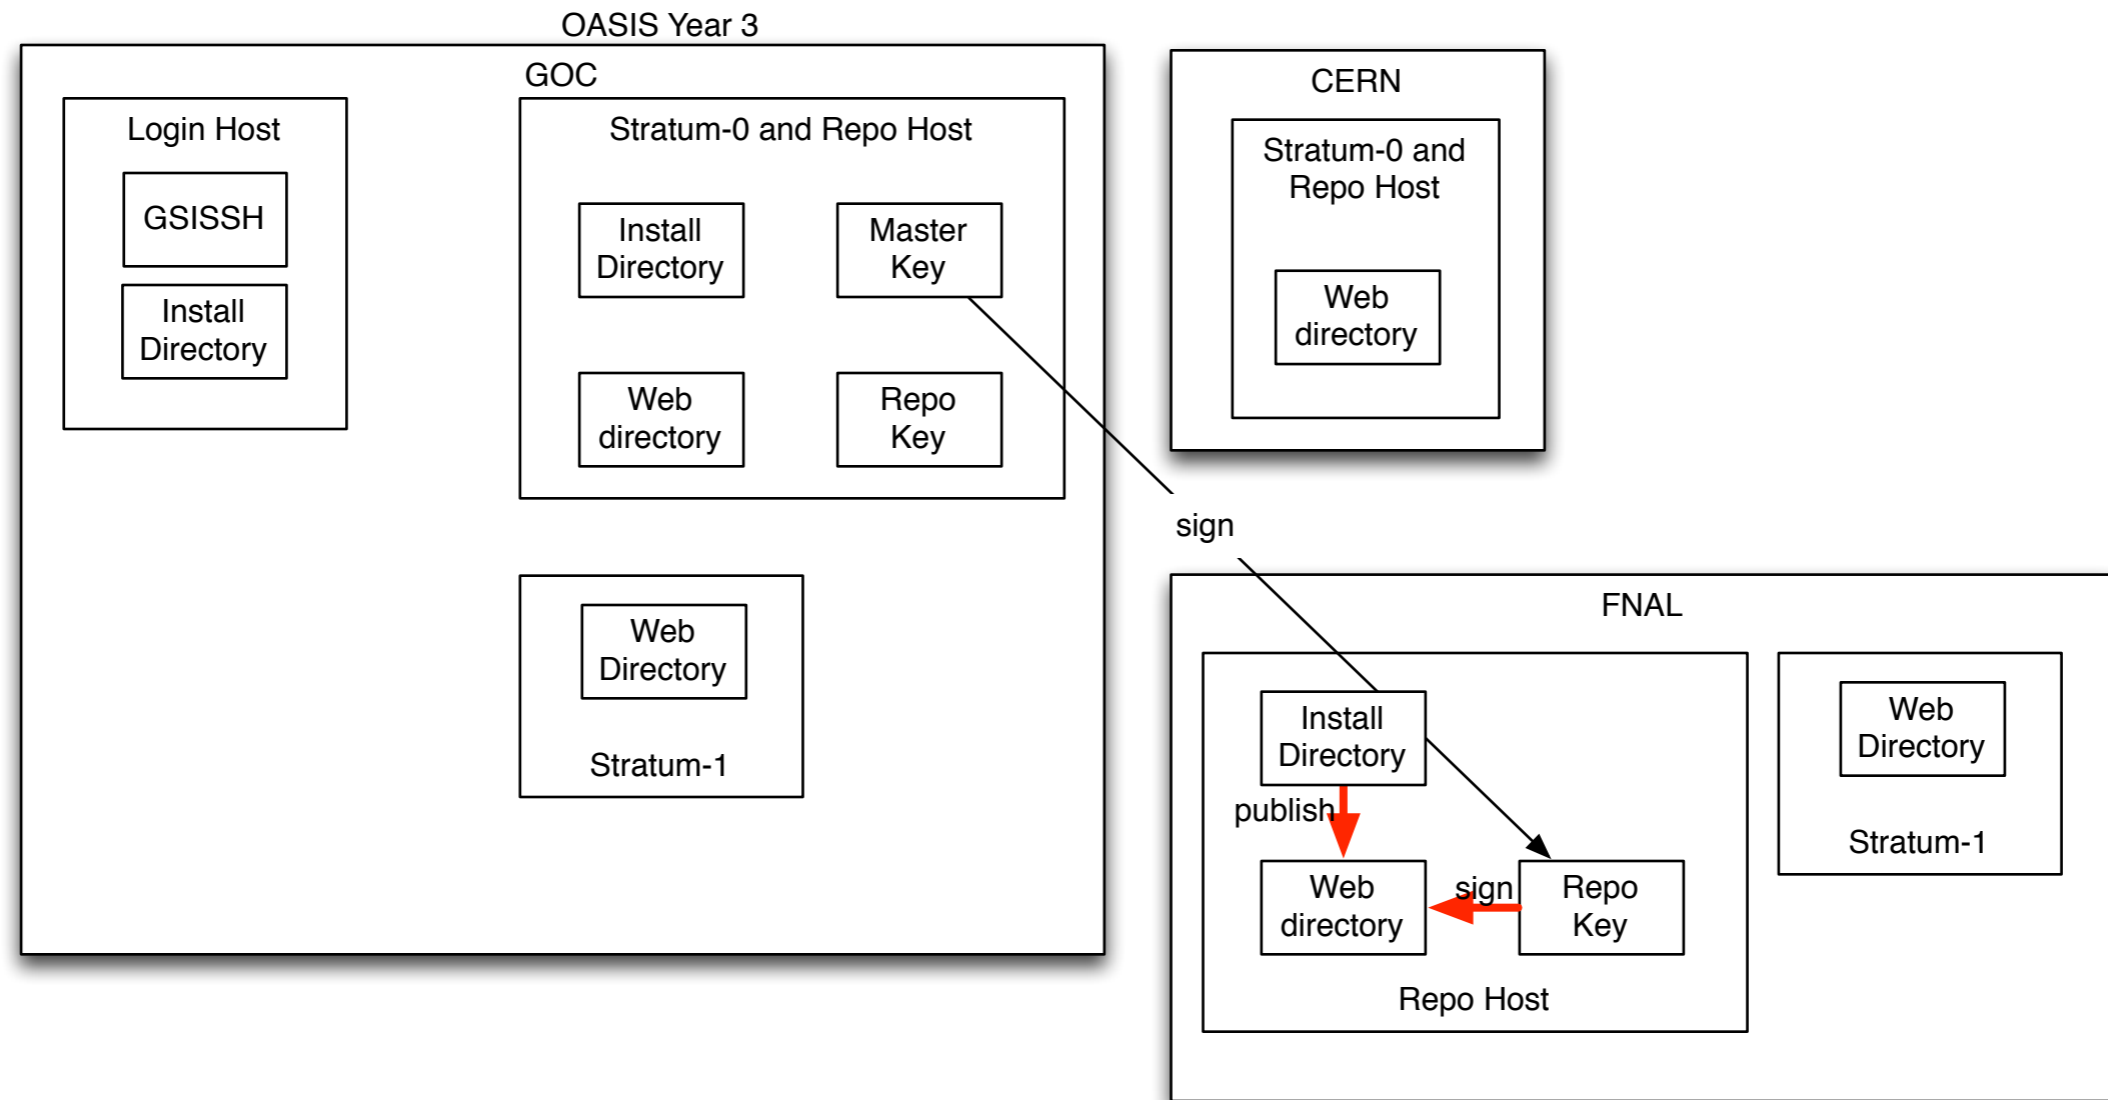
OASIS - Year 3



Notes:

- GOC signs repo key once a week.

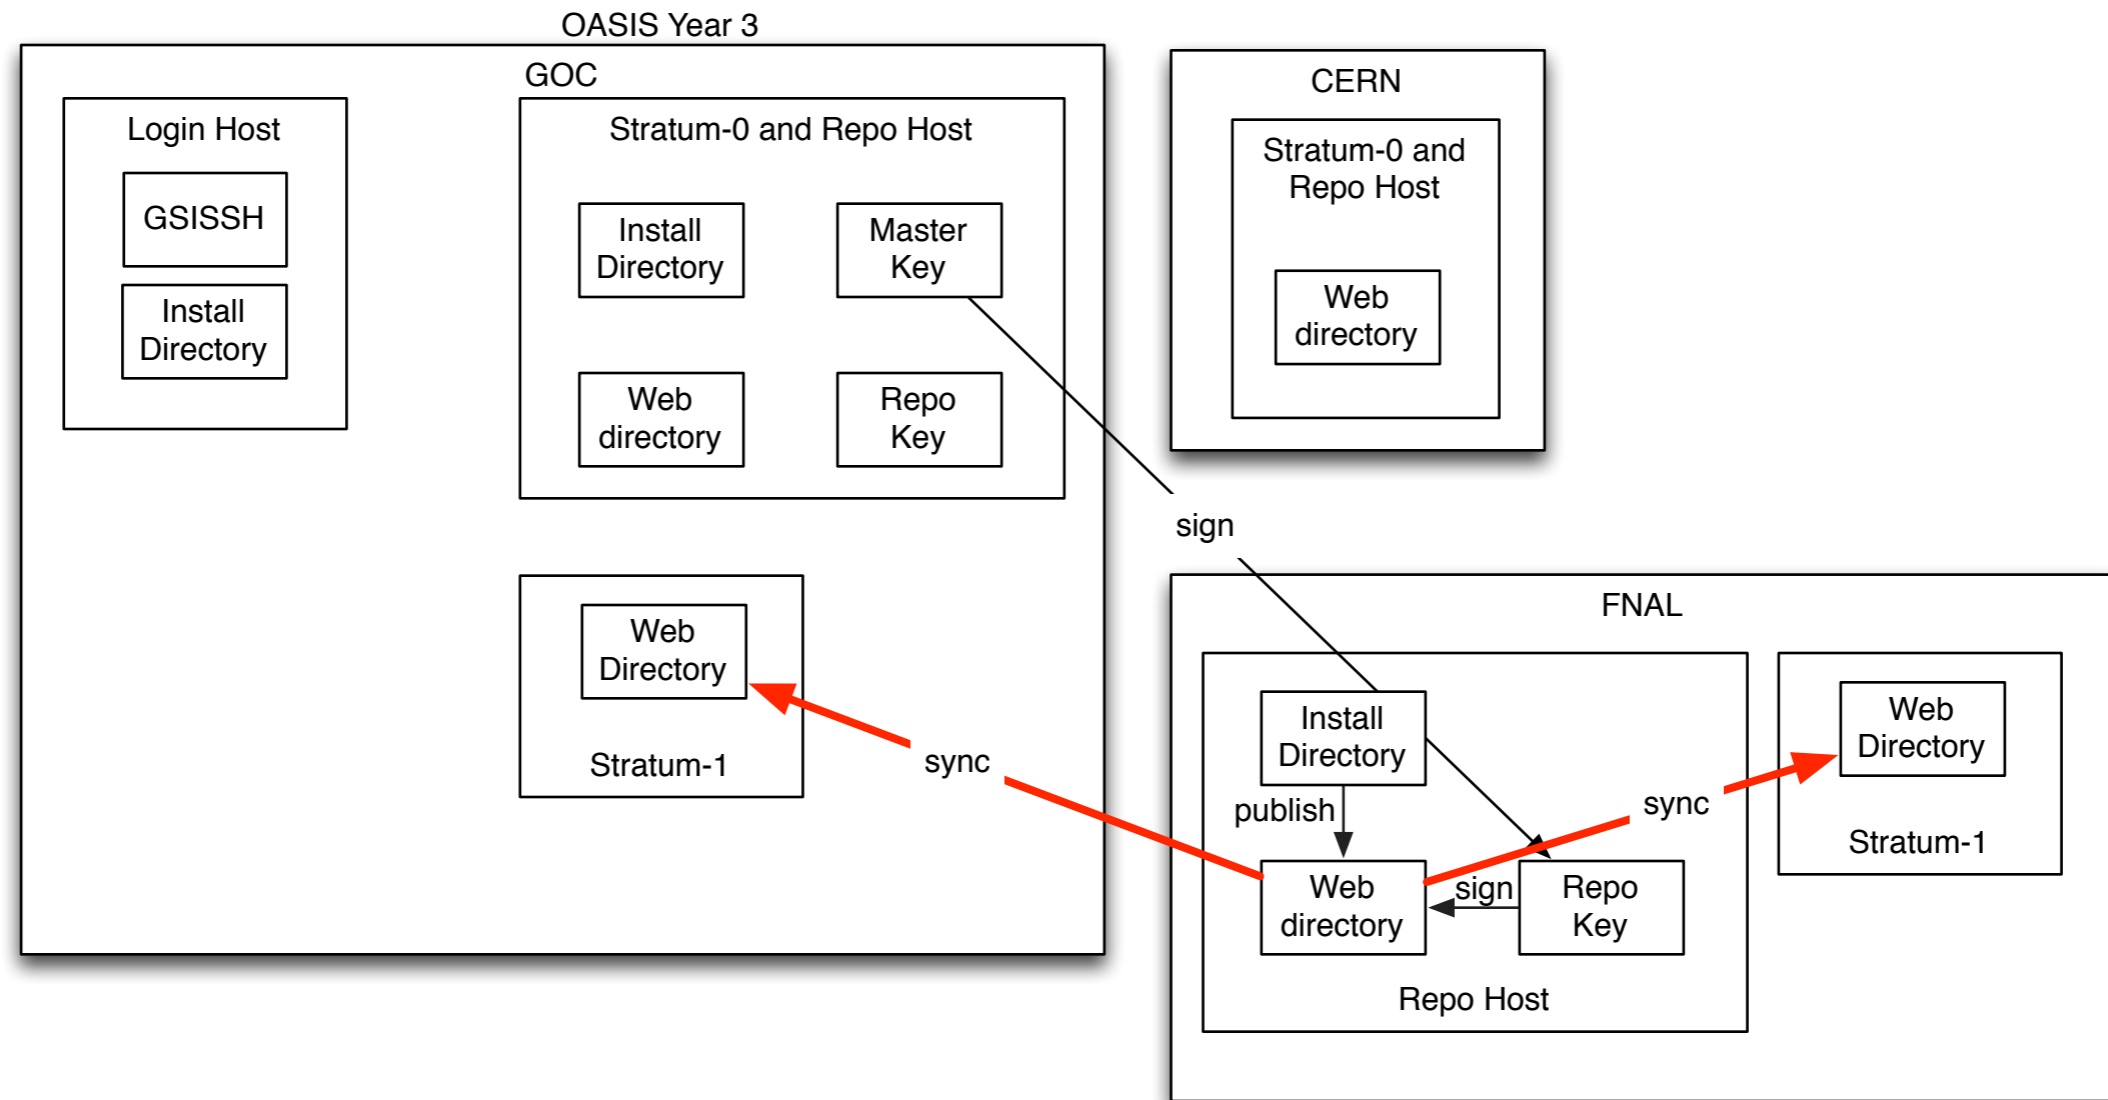
OASIS - Year 3



Notes:

- OSG-approved repos can publish and sign according to their own schedule.
- The installs can be done with whatever fits the VO.

OASIS - Year 3



Notes:

- OSG retains the ability to “blank” repos for security issues; this forces all clients to show empty directories for that repo.

Common Input Data

- **Problem statement:** How do we support a workflow with 1,000 jobs, each of which need the same 1-10GB of input data?
- In Year 2, we worked with the HTCondor team to design a new daemon for managing caches.
 - See <https://docs.google.com/document/d/1shVt1q6qiR-h83kZaCJ3sxV737JhTltCbk9QTuxcvUg/edit>
 - Unfortunately, staff transitions derailed the implementation.
 - I still think this is the best way forward; working to find external staff to restart the effort.
 - Alternate to this approach is to “beef up” squid infrastructure and use OASIS (once CVMFS 2.1 transition is done).
 - Downside is we have far less ability to “manage” squid caches.
 - IF experiments are investigating this on their dedicated resources - in Year 3, we can observe their progress.

New Resource Type Integration

- We failed to get BOSCO / HTCondor-G+SSH resources into the Production Grid during Year 2 using glideinWMS.
 - In parallel, it appears the Campus Grids was able to do this for OSG Connect.
- Cloud resources are in better shape, but it's not obvious we have a sustainable approach in glideinWMS.
- **Decision needed:** Do we want to proceed with the glideinWMS approach or focus on using OSG-Connect to extend the Open Facility to these sites?
- Goal for Year 3 is to use the OSG staff allocation to demonstrate integration with XSEDE.
 - **Decision needed:** Can we ask XSEDE to deploy CVMFS on appropriate resources?

Site authz modernization

- GUMS has had little maintenance in the last 3+ years; I worry this will become critical soon!
- Immediate **decision needed**: IF experiments at FNAL have requested we add support for Unix groups to GUMS.
 - The request is reasonable and implementation should be straightforward (2 weeks).
 - I recommend we accept the request.
- Year 3 **Proposal**: Software group works with BNL to transition GUMS to OSG. Technology group will implement the new feature; aim will be to have this feature be part of the first release done by OSG.
- In year 3, request that Miron organize a blueprint on authz to plan long-term needs.
- In year 4, execute plan put together by blueprint.
 - At the minimum, perform extensive cleanup of dependencies. During Year 4, make a release that contains no “binary blob” dependencies.

Risks to program of work

- Technology evolution necessarily depends on external software projects.
 - Sometimes these projects don't have the effort or don't prioritize OSG (glideinWMS).
 - We will heavily rely on the HTCondor team for improvements to the common input data.
 - We rely on the CVMFS team for continued updates and improvements to that software.
- Building the right team: often, we need a specific set of skills for a project; it can be hard to find some in OSG.
 - In particular, we lack depth in Java (needed for bestman2, jglobus2, and GUMS).
- “Outrunning” our ability to get new technology into production.
 - There are cases where software was released and documentation or osg-configure wasn't ready.

Focus areas

- Provide value to VOs and resource owners:
 - HTCondor-CE, BOSCO integration, HTCondor scalability work, authz modernization.
- Provide opportunities for the long tail of science:
 - Common input data project, OASIS
- Promote use of DHTC for science:
 - Blueprints activity.