

CILogon Fall 2009 Workshop Report

<http://www.cilogon.org/fall-2009-workshop>

Jim Basney, Von Welch, and Randy Butler
info@cilogon.org

The [CILogon Project](#) held a [one-day requirements-gathering workshop](#), on Tuesday, September 29, 2009 in the NCSA Building at the University of Illinois in Urbana, Illinois, to discuss challenges related to securely accessing cyberinfrastructure in the science and engineering community and solicit community input on the priorities for the project to address those challenges. We invited representatives from NSF domain science cyberinfrastructure (CI) projects to attend the meeting. The CI representatives in attendance were:

- [Dark Energy Survey](#) - Tony Darnell and Michelle Gower
- [Institute for Computing in Humanities, Arts, and Social Science](#) - Alex Yahja
- [Laser Interferometer Gravitational Wave Observatory](#) - Scott Koranda
- [Large Synoptic Survey Telescope](#) - Mike Freemon
- [Long Term Ecological Research Network](#) - Mark Servilla
- [NSF Office of CyberInfrastructure](#) - Jennifer Schopf
- [Ocean Observatories Initiative](#) - Von Welch
- [US National Virtual Observatory](#) - Bill Baker and Ray Plante
- [Plant Science Cyberinfrastructure \(iPlant\) Collaborative](#) - Stephen Mock
- [nanoHUB, HUBzero](#) - Steven Clark and Steven Snyder

The CILogon team summarize their conclusions from the workshop presentations and discussions as follows:

- Different NSF CI projects are in different phases of their lifecycle and hence have different types of requirements.
- We received helpful feedback on our proposed development tasks and their priorities.
- Several attendees expressed strong interest in best practice and "How To" documentation.
- Several attendees identified OpenId as a topic for collaborative exploration.
- Multiple project representatives identified certificate expiration as a problem.
- Attendees agreed to the value of usage metrics collection.
- Attendees expressed both interest in InCommon and ProtectNetwork, as well as a need to better understand those services.
- We will facilitate follow-on collaboration with the participating projects through the [cilogon.org website](#), site visits, and phone calls.

We discuss our conclusions in more detail in the following sections.

Projects in different lifecycle phases

The CILogon team concluded from project presentations and project discussions that the represented projects were in different phases of their lifecycle, which in turn meant the attendees presented different types of requirements for CILogon. We categorize the different project phases as:

- *Early Inception*: Projects still in an early requirements gathering or CI design phase who did not have concrete requirements for identity management established yet. E.g. iPlant, LTER, LSST, and DES.
- *Exploratory*: Projects that had a general identity management conceptual architecture (possibly to replace an existing implementation) and that are undertaking research and/or prototyping to establish a specific technology and implementation plan. E.g. NVO.
- *Completed Design with Identified Gaps or Desires*: Projects that have a completed design, are ready to begin on implementation, and have identified needs and/or desired functionality that is unmet by existing implementations. E.g. LIGO.
- *Operational*: Projects that are in operation, not looking to make big changes until obvious benefit or user demand emerges, but which have minor issues they wanted help with or have interest in staying educated on emerging technology. E.g., nanoHUB, Hubzero.

Input on software priorities

We reviewed CILogon's proposed software development and support activities with attendees and received valuable input from them that will help us re-prioritize tasks in our project plan. One clear message we received was that the CILogon project should give high priority to documentation of best practices, facilitating information sharing between CI projects, and consulting with CI projects regarding their specific identity management needs. While attendees expressed interest in many of the proposed software development tasks, we conclude that support efforts for existing capabilities will likely be more important than major new software development tasks. When we demonstrated the capabilities of the [GridShib CA](#), as used in the go.teragrid.org portal, multiple attendees remarked that the capabilities appeared useful for their projects and they were not aware that the capabilities were available. Thus, we see the need for renewed effort to document, advertise, and support/maintain the existing capabilities of CILogon software so CI projects can take full advantage of what is already available.

Best practice and "How To" documentation

As discussed in the summary, workshop attendees requested guidance on how to use existing functionality to accomplish their project tasks. We believe such needs could be satisfied with "How To" documentation augmenting the existing software documentation, explaining how to use the software to accomplish these specific tasks. Attendees also requested documentation that could be used to demonstrate the appropriateness of a particular approach in order to convince others (typically operational security personnel) of its acceptability. We envision such documentation taking the form of Best Practices collected across the broader set of CI communities. Having such practices documented as being in broad use would enable a community to adopt it themselves with greater assurance.

Interest in OpenID

[OpenID](#) is a protocol for web-based authentication and access control that has been adopted by many commercial service providers. The CILogon Project plan includes adding OpenID support to the GridShib CA, so the GridShib CA can issue certificates based on OpenID authentication in addition to Shibboleth/SAML authentication. Attendees expressed interest in this capability along with a general interest in experiences and recommendations for integrating OpenID into CI web applications and portals. Bill Baker presented plans in NVO to adopt OpenID. We also discussed the use of OpenID in [Earth System Grid](#) and HUBzero. We concluded that, in addition to the specific development task of adding OpenID support to the

GridShib CA, the CILogon Project could make a valuable impact in the community by gathering and documenting experiences and recommendations for the use of OpenID in CI projects.

Certificate expiration

The CILogon team presented their plans to develop a credential renewal service to address the problem of certificate expiration for long-lived workflows using short-lived certificates. Attendees representing DES, iPlant, and LIGO confirmed that certificate expiration is a problem that requires additional attention. We discussed the specific use case of automated data processing jobs (for example, the daily processing of images in DES), which are overseen by system administrators or CI operators, in contrast to long-running jobs initiated on demand by end-users. We discussed the use of long-lived "robot" certificates for these automated jobs, and we discovered that attendees had different understandings of best practices for managing these certificates. We discussed methods for addressing certificate expiration using existing capabilities of Condor-G and MyProxy. We conclude that certificate expiration is an important focus area for the CILogon project, but that education about existing software capabilities and current best practices may be more valuable than a new software effort to develop a credential renewal service.

Usage metrics

The CILogon team presented their plans to incorporate [dev.globus usage metrics](#) into the MyProxy, GridShib, and GSI-OpenSSH software to enable CI projects to gather usage statistics and allow the CILogon project to report usage information to NSF. We discussed that the usage collection would be opt-out, so that by default, the software would send usage information (without user-identifying data) to the CILogon usage collector. We also discussed how the software could be configured to send usage information to a CI project's own collector, including user-identifying data if desired. Attendees generally agreed to the value of the collection of usage information, understanding its importance in justifying further funding and support for the software projects. Attendees did not raise any privacy concerns. The CILogon team will proceed with their plans in this regard, which will include announcements to the full user community before any usage data is collected.

Interest in InCommon and ProtectNetwork education

While most of the project participants were aware of the [InCommon Federation](#), there was a general interest in learning more about it, including basic information such as "What does joining InCommon really entail?", "What does joining InCommon really provide?" and "What impact does an Organization meeting the InCommon Silver level of assurance have on that organization's user community?" There was similarly interest in the [ProtectNetwork](#) identity provider. A specific question that was raised was to what degree InCommon could be "blessed" by NSF such that leveraging it would not present the project with problems justifying it at their annual review. We conclude that it would be valuable for the CILogon project to develop educational material addressing these topics.

Follow-on collaboration

The participants indicated a desire to continue to stay up-to-date on both the activities of the CILogon project and the activities of the other projects. For example, hubZERO and NVO had experiences with ProtectNetwork and OpenID of broader interest to the community, and there was also interest in learning more about [Earth System Grid](#) experiences with OpenID. As a first concrete step for facilitating follow-on collaboration, we created cilogon.org accounts for all attendees, providing a [collaborative web space](#) for sharing experiences and technical

documentation. We discussed the need to share announcements of significant advances in identity management across CI projects and will experiment with using the [CILogon News RSS feed](#), which all cilogon.org account-holders can post to, for this purpose.

Conclusion

In conclusion, the workshop organizers found the attendees' input very valuable and are very grateful for their participation. The presentations by the LIGO and NVO representatives were very helpful, and we plan to invite additional community presentations in next year's CILogon workshop. We found that assistance from NSF program officers in advertising the workshop noticeably helped in attracting participants, resulting in representation from a total of ten NSF CI projects.

Additional materials from the workshop (including presentation slides and meeting notes) are available from the [workshop page](#) on the [cilogon.org web site](#).

Acknowledgements

The workshop organizers would like to thank Scott Koranda of LIGO and Bill Baker of NVO for their presentations during the workshop.

This material is based upon work supported by the National Science Foundation under grant numbers [0850557](#) and [0943633](#). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.