

# Usability Issues in WLCG Security

WLCG Technology Evolution Group on Security

Editor: Von Welch (vwelch@indiana.edu)

April 12, 2012

Version 0.1

## 1 Introduction

Usability is a key factor in the security of a system, as a lack of security usability will lead to improper deployment, misuse or avoidance of security mechanisms. In this document we describe a number of usability issues identified by the WLCG Technology Evolution Group (TEG) on Security and make some recommendations on addressing those issues. We divide the list of issues into problems typically seen by scientific researchers who comprise the WLCG's users, and the WLCG administrators.

### 1.1 From the Researcher Perspective

- **Credential management:** Some WLCG users have trouble with credential management throughout the lifecycle, including obtaining a credential (including the time required) and maintaining it securely without losing it or forgetting the encryption pass phrase. This problem is increased when a user attempts to undertake multiple activities, which require different VOMs<sup>1</sup> attributes; This requires juggling multiple proxies, with different VOMS attributes in each, and brings the risk of getting them mixed up, accidentally overwriting them, etc.
- **Incoherent proxy storage on complex systems:** Proxy credentials are by default stored in */tmp*, which for some systems (e.g. clusters) is not shared across the whole system, leading to incoherencies in availability of the proxy to activities.
- **Lack of web authentication:** X.509 credentials are difficult to import into or export from browsers, making it very difficult to use a given credential for both web and command-line activities. Additionally, support for RFC 3820 proxy certificates by web services is weak.

---

<sup>1</sup> <https://twiki.cern.ch/twiki/bin/view/LCG/VOMS>

- **Lack of internationalization:** Support for internationalization of names in X.509 distinguished names is not complete and use of non-ASCII characters can cause failures.

## 1.2 From the Administrator's Perspective

- **Managing revocation:** Revocation of certificates is accomplished through certificate revocation lists (CRLs), which have finite lifetimes (typically in the range of a week to a month). When a CRL expires, software still using that CRL will fail safe and fail all authentication attempts. This means that timely updates of CRLs are an operational necessity and any failure of a certificate authority to produce a new CRL or a relying party to obtain and properly install it results in a service outage.
- **Expired host and service certificates:** Certificates used by hosts and services have finite lifetimes, typically a year or small number of years. After they expire, clients attempting to use those hosts and services will experience failures. Currently administrators are responsible for "manually" monitoring these certificates and renewing them before expiration to provide continuous service. The case of a cluster with tens or hundreds of certificates may pose issues of scale.
- **Managing authorization policies:** Authorization is controlled by policies that are encoded in a number of different places: grid-mapfiles, CA signing policies, VOMS, GUMS, etc. Authoring and maintaining all of these to coherently represent the mission of the projects and VOs is a complicated, error-prone task.
- **Client authorization of hosts and services:** *placeholder, to be completed*  
Clients expect services to identify themselves with host certs by default and therefore may need special configuration to accept anything else, viz. a service certificate.
- **Inconsistent user banning mechanisms:** It is occasionally necessary to "ban" a user, or suspend their ability to access resources, due to a suspected or confirmed security incident, misbehavior of the user or their software, or other administrative reasons. This differs from certificate revocation, which is meant to globally revoke a user's authentication credentials, and may not be an appropriate vehicle due to the scope or longer-term authentication effects. Similar to more general problems of managing authorization policies, software services differ in how they allow for banning or whether they even do, leading to a inconsistent, error-prone process.
- **Mixing of authentication and authorization:** Proxy certificates, through their use of impersonation (allowing one entity to appear as another entity)

to support delegation, are inherently a mix of two security concepts - authentication (identification of an entity) and authorization (what is that entity allowed to do). This creates some usability problems such as masking of a delegatee's identity in logging (they get logged as the delegator) and confusion between the use of entity names and VOMS attributes in authorization policies.

- **Lacking debugging and forensics:** Debugging of failures is a challenge caused by the multiple software stacks with different error propagation characteristics and capabilities, and the distributed nature of the system (and the resulting logs) and its users community. This causes challenges in accessing needed log and configuration data for debugging (assuming it exists), and interpreting it as logging across different software components is not standardized in terms of content, syntax or semantics.
- **Inconsistent proxy certificate implementations:** Inconsistent implementations of support for proxy certificates lead to incompatibilities and difficulty in other administrative tasks, such as infrastructure upgrades. An example is the lack of RFC 3820 proxy certificate support in some software, leading to difficulties in removing dependencies on weak cryptographic algorithms.
- **X.509 validation overhead:** Currently the validation of X.509 credentials is done for each network connection within the WLCG. Each of these validations is equivalent to the validation done for a TLS/SSL handshake as part of a web HTTPS connection. And while Google has done significant work in minimizing this overhead [1], the aggregate load across the WLCG may be significant.

## 2 Recommendations

We now turn from problem identification to recommending improvements to address those problems. We sort the recommendations into short-term, defined as something that could be put into place in less than 2 years, and long-term, or something that would require more than 2 years to put into place.

### 2.1 Short-Term

In general, effort should be made to "hide" PKI/X.509 from the end-users of the system as much as possible. Options here include:

- **Minimizing the enrollment process by leveraging existing sources of identity and authentication.** For example, leveraging existing member

authentication services either directly, e.g., the IGTF MICS profile [2], or through federated identity, e.g. the CILogon service<sup>2</sup>.

- **Minimizing the credential management process through the use of short-lived credentials.** Short-lived credentials (i.e., those with lifetimes of less than  $10^6$  seconds or roughly 11 days) as defined by the IGTF SLCS profile [3] in combination with the leveraging of existing authentication mechanisms offers users easy to obtain credentials that are in effect "disposable," removing burdens for long-term curation, migration between computers, renewal, etc.

Other improvements for researchers include:

- **Tools for multiple credentials.** Some users need to juggle multiple credentials with different attributes or from different CAs for connecting to different services with different trust policies, or for acting in different roles. A set of tools or improvements to existing client software to allow users to maintain multiple "roles" each with a different credential and switch easily between those roles would be beneficial.

From the administrator perspective, short-term improvements include:

- **Tools for service credential lifecycle management.** Develop a set of supporting tools to help monitor and renew host and service credentials to prevent expiration failure and reduce workload during normal operation.
  - Automatic renewal may not be possible in many cases, but a cron job could at least alert the admin on time.
- **Improved revocation.** In addition to short-term tools for improved CRL management, an analysis should be undertaken with regards to the trade-off between availability and security with regards to the behavior of software in "failing safe" when an expired CRL is encountered. Allowing configuration of behavior such that software will continue to function with an expired CRL but produce strong warnings may be desirable in some situations and reduce availability failures.
- **Standards for logging.** Having a set of defined use cases, derived requirements and ultimately standards for logging would aid in a variety of activities from debugging to cybersecurity incident response.
- **Usability Evaluation.** Having an evaluation of the usability of grid security and relevant tools performed by expert in the area of usability and human

---

<sup>2</sup> <http://www.cilogon.org/service>

computer interaction could yield improvements or new tools that could be made to improve ease-of-use.

## 2.2 Long-Term

Some of the long-term improvements include:

- **A coherent grid security implementation library.** Having a coherent implementation of grid security features (e.g., RFC 3820 Proxy Certificates [4]) in appropriate languages and frameworks would address many incompatibility challenges being faced today (e.g., difficulties in migrating to modern hash algorithms because of lack of support in some software stacks). A challenge here is the lack of a generic security library across multiple languages - each language has its own de facto standard(s) (e.g., OpenSSL, Bouncy Castle) that are not coordinated. Having a single grid security library, with implementations in the most popular languages and a standard test suite, would address this incoherency and make development of new grid-based applications both easier and more secure.
- **Consider re-implementation leveraging experiences.** Over the past fifteen years, there has been much learned about how global-scale grids operate in practice. Rather than incremental improvements described to this point, a next generation may be considered. Aspects to consider would be to re-approach impersonation and alternatives that separate authentication and authorization more cleanly in the delegation process; push public key cryptography to the outer edges and use some sort of session token (e.g., as in the Condor CEDAR implementation [5]) for transaction authentication; implementation of role and group authorization, particularly in the context of the growth of data.

## 3 References

1. Imperial Violet Blog (authored by Adam Langley, Google Security Engineer). "Overclocking SSL." 25 June 2010.  
<http://www.imperialviolet.org/2010/06/25/overclocking-ssl.html>
2. The Americas Grid Policy Management Authority. Member Integrated X.509 PKI Credential Services (MICS), version 1.1. 2 May 2009.  
<http://www.tagpma.org/node/38>
3. The Americas Grid Policy Management Authority. Short Lived Credential Services X.509 Public Key Certification Authorities, version 2.1b. 3 February 2009. [http://www.tagpma.org/authn\\_profiles/slcs](http://www.tagpma.org/authn_profiles/slcs)
4. Steven Tuecke, Von Welch, Doug Engert, Laura Perlman and Mary Thompson. Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile. RFC3820, 2004.
5. Zach Miller, Dan Bradley, Todd Tannenbaum and Igor Sfiligoi. Flexible session management in a distributed environment. 17th International

Conference on Computing in High Energy and Nuclear Physics (CHEP09).  
<http://dx.doi.org/10.1088/1742-6596/219/4/042017>

