Discussion of the Issues (Feb 10 and later)

From: David Walker <DHWalker@ucdavis.edu>
To: Prakashan Korambath <ppk@ats.ucla.edu>
Cc: Arlene Allen <arlene.allen@isc.ucsb.edu>, Labate, Bill <labate@ats.ucla.edu>, Jin, Kejian <kjin@ats.ucla.edu>, Wu, Albert <albertwu@ucla.edu>, Russ Hobby <rdhobby@ucdavis.edu>
Subject: Re: UC-wide authentication for UC Grid
Date: Tue, 10 Feb 2009 15:00:33 -0800

OK, I've done some reading...

First a proposed vision: UC Grid will become a the infrastructure of services for UC, and it is used by all members of the community. Very few of these people will be programming, though; they'll be using canned applications.

This means that we need an authentication strategy that can scale to hundreds of thousands of people, which I why I keep hammering on UCTrust integration. UCTrust has already solved the scaling problem.

For UC Grid, I see two issues:

- technical integration of Shibboleth into UC Grid's authentication, and
- an extremely efficient registration process for nearly all grid users.

Technical Integration
It appears to me that the GridSphere account used by the UC Grid Portal is the key to everything. It is mapped by MyProxy servers to a certificate, and by clusters to local cluster logins. If we can federate the GridSphere account, we should be home free.

I did a little reading of the GridSphere documentation, and it looks like it can integrate any JAAS-compliant authentication method. My memory is that the AYSO people integrated Shib via JAAS, but even if they did not, I'm sure we can find someone who has. The potential gotcha would be if GridSphere has limitations on the size or composition of its user identifiers, since federated identifiers tend to be long and have an at-sign in them.

Registration
The current registration process's workflow starts with a cluster administrator (or a campus grid administrator for "pool" users). UC-wide GridSphere accounts are only allocated after approvals are granted. Given that UC-wide GridSphere accounts must be unique, and they are (currently) not federated, I understand that they are a precious resource, so giving them out without need should be avoided.

In my proposed vision of the future, though, I believe that a very small percentage of all users will have explicit access to specific clusters. In other words, 99%+ will be pool users.

This argues for reversing the registration process: Everyone gets a (federated) GridSphere account without human approval. A small percentage of users ask for specific cluster access, after they are fully registered and have their digitally-signed certificate.

Does any of this make sense?

Daivd

From: Kejian Jin <kjin@ats.ucla.edu>
To: David H Walker <dhwalker@ucdavis.edu>
Cc: Prakashan Korambath <ppk@ats.ucla.edu>, Arlene Allen <arlene.allen@isc.ucsb.edu>, Labate, Bill <labate@ats.ucla.edu>, Wu, Albert <albertwu@ucla.edu>, Russ Hobby <rdhobby@ucdavis.edu>
Subject: Re: UC-wide authentication for UC Grid
Date: Tue, 10 Feb 2009 16:40:41 -0800

Hi,

Shibbolizing Gridsphere already exists. People had changed the
gridsphere so that it allows the shibboleth as a PAM for gridsphere.

That project is from Australia: https://mams.melcoe.mq.edu.au/zope/mams/kb/all/GridSphere%20Wink%20demo.zip/view

I have their source code.

Basically, the shibboleth IdP has to pass userName, surName, givenName,
Organization, email, IDP, Role to gridsphere which internally create a
User Object for gridsphere at real time. It modified the login of
gridsphere by doing that.

I am more insterested to discuss how we could create a short live
credential for user that will authorize the user to use certain resources.
(authorization) I will really like to tell the meaning of
"single-Sign-On". Most people uses that shibbolized gridsphere for
portal authentication of portal sign-in, but it is never used for
authorization of resources.

I have worked on a project before. That project will
create the Unix virtual workspace and create a certificate (once) and
generate proxy at real time and submit job and application
to the cluster as pool user.
please see https://research.ucgrid.org
anyone with UCLA ID will be able to have a virtual desktop and submit
job, or run some grid applications...

I hope the discuss we planed will help us to generate some ideas about
how to do that in a secure way and easy-to-deploy way. There are many ways
of doing the same thing, but we really like to have your input:

The following are my thought for doing that:

Method 1: IdP gets the username and password, it used the username and
password to retrieve a proxy from myproxy.ucgrid.org. It is just a command
something like this: get-myproxy -h myproxy.ucgrid.org username password.
it will include that delegated and short live credential in Assertion.

Method 2: IdP passes a unique string (like the one in openldap) to
SP, SP
understand
that and lookup some sort database to figure out the username and password
and generate user proxy and use that for job submission.

....

more input and discussion is needed!

Thank you very much for your time...

Regards,

Kejian Jin
UCLA Grid Team
UCLA Web: http://grid.ucla.edu
UC Web: http://portal.ucgrid.org
University Of California, Los Angeles

From: David Walker <DHWalker@ucdavis.edu>
To: Kejian Jin <kjin@ats.ucla.edu>
Cc: Prakashan Korambath <ppk@ats.ucla.edu>, Arlene Allen <arlene.allen@isc.ucsb.edu>, Labate, Bill <labate@ats.ucla.edu>, Wu, Albert <albertwu@ucla.edu>, Russ Hobby <rdhobby@ucdavis.edu>, ...
Subject: Re: UC-wide authentication for UC Grid
Date: Tue, 17 Feb 2009 11:50:59 -0800

Keijian,

Good news about the software from Australia. Looking over their demo, it seems to me that their "guest" user would correspond to our "pool" users. Does that sound right to you? I'm not sure that all of our campuses would have all of the attributes that the Australians plan to use, but that may not be a big issue. The user could also be prompted for the missing information during registration.

I also like your ideas about using UCTrust to pass authorization, as well as authentication, although that's probably a longer-term (phase 2?) issue. It strikes me, also, that it will probably apply only to clusters that have given authorization controls to the grid; many will not.

FYI, I've created an area within the UCTrust wiki space for us at:

https://spaces.ais.ucla.edu//x/SA43AQ

See you tomorrow.

David