[paccman] Late night thoughts on authorization

43 messages

Keith Hazelton <hazelton@doit.wisc.edu> Thu, Sep 6, 2012 at 11:50 PM
To: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

To paraphrase Roland Hedberg, it is high time to seriously address authorization as we work on our (inter-)federation identity and access management (IAM) infrastructures. Two patterns are commonly found today, depending on whether the locus of authorization evaluation is at the IdP or RP and I would argue that there is a third alternative that is worthy of consideration.

1) If authorization decisions are to be made at the IdP then they are typically expressed as entitlement attribute/values back to the RP. The RP needs to share the business logic for computing the entitlement with the IdP (out of band), and the IdP needs to be willing to do the work to compute authorization using the RP business logic. This scales well for a small set of common entitlements across a(n) (inter-)federation.

2) If authorization decisions are to be made at the RP then the IdP needs to release relevant attributes about the authenticating subject for use by the RP access control system. Group/Role memberships of the subject are among those relevant attributes. In this case, the RP can conceal its business logic, but it may need to obtain consent from the user and/or the user's IdP as a condition for obtaining those attributes.

3) Warning, some concepts that follow are borrowed from the XACML conceptual model. Consider that in simple situations the "business logic" referred to in 1) and 2) above might be expressed in a computable policy rule of the general form "subjects (S) carrying role (or group membership) G may perform actions within the set A on resources in class R". There would seem to be cases in which the desired process would be: Subject S shows up at the RP and requests to do action A1 on resource R1. The RP (somehow, possibly out of band) specifies or references a policy rule as above and asks a "Policy Decision Point (PDP)" for a boolean-valued response whose semantics is T=>Allow, F=>Deny. VERY few of our existing infrastructures include anything like a PDP.

Note that the Role/Group memberships of subjects might be carried in a VO attribute authority (AA) to which the RP belongs (think Policy Information Point, PIP). Those Role/Group memberships might also be MANAGED by VO members with suitable delegated admin rights. One of the challenging use case assumptions here is that those VO delegated admins will have available to them (possibly pseudonymous) identifiers for the subjects whose Role/Group memberships they are managing. If those subject identifiers can be obtained from an IdP in a SAML assertion, then an "undecorated user identifier" would be all that the SP would need from the subject's IdP to enforce the authorization decision. Note that this model locates the Policy Enforcement Point PEP at the RP. One further extension would be useful: The VO could offer a service by which an RP admin could define and persist their RP-specific computable access policy rules (think Policy Administration Point, PAP).

We have not genrally thought of our (inter-)federation IAM infrastructures as containing PIPs (well, we have started to talk about AAs) or PAPs in addition to IdPs and RPs, but if model 3) is interesting to us, we will need to think those thoughts.

--Keith Hazelton

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 2:56 AM
To: David Chadwick <d.w.chadwick@kent.ac.uk>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

> So whilst this type of infrastructure is technically well within
> the grasp of most RPs, there are still many usability problems to
> consider, for example, how to create policies in an intuitive and
> easy way, how to determine which resources are available to which
> users (a simple grant/deny interface cannot tell you this), how to
> coordinate attributes from multiple PIPs/AAs, how to give the user
> control, how to tell the user (or IDP) what the RP's policy is. I
> am afraid that the current metadata distribution mechanism is not
> sufficient for an Internet wide federation. A sophisticated
> DNS-like service is needed. And then you have the issue of semantic
> interoperability, knowing that attribute A from IDP1 is equivalent
> to attribute B from IDP2.
>

Pretty much the same reason SPOCP had trouble taking off. Policy engines exist and are often powerful and relatively easy to deploy.

Application integration is a killer though.

The problem is that there is nothing that corresponds to the simplicity of the REMOTE_USER model in authz land that can be used as a spring-board for integration.

Cheers Leif

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBJqKQACgkQ8Jx8FtbMZnf+IgCeKhg7H/TBufLmPWn4spkTU+d/
uN8An05g2uGXydxBndTo6+M2d5c0tB+B
=zdoL
-----END PGP SIGNATURE-----

Keith Hazelton <hazelton@doit.wisc.edu> Fri, Sep 7, 2012 at 4:15 AM
To: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

On Sep 7, 2012, at 09:56:20, Leif Johansson wrote:

> -----BEGIN PGP SIGNED MESSAGE-----
> Hash: SHA1
>
>
>> So whilst this type of infrastructure is technically well within
>> the grasp of most RPs, there are still many usability problems to
>> consider, for example, how to create policies in an intuitive and
>> easy way, how to determine which resources are available to which
>> users (a simple grant/deny interface cannot tell you this), how to
>> coordinate attributes from multiple PIPs/AAs, how to give the user
>> control, how to tell the user (or IDP) what the RP's policy is. I
>> am afraid that the current metadata distribution mechanism is not
>> sufficient for an Internet wide federation. A sophisticated
>> DNS-like service is needed. And then you have the issue of semantic
>> interoperability, knowing that attribute A from IDP1 is equivalent
>> to attribute B from IDP2.
>>
>
>
> Pretty much the same reason SPOCP had trouble taking off. Policy engines

> exist and are often powerful and relatively easy to deploy.
>
> Application integration is a killer though.
>
> The problem is that there is nothing that corresponds to the simplicity
> of the REMOTE_USER model in authz land that can be used as a spring-
> board for integration.
>
> Cheers Leif

David, Leif,

Yes, lots of prior art, lots of challenges.

On the other hand, if what we collectively desire to do is to take small incremental steps to address some of the more straightforward use cases we want to support under alternative 3), then let's start by asking ourselves "What MUST be in place that isn't there yet?"

1) David, your "standalone authz server which talks the SAML/XACML protocol allowing RPs to make remote call outs for authz decisions" or something else fitting the same description sounds like a missing piece that would be needed in even the most obvious use cases under alternative 3)

2) I believe that the latest version of Grouper, or anything else with comparable capabilities, provides all the functionality a VO would need to offer distributed admin of PIPs and PAPs for a relatively simple class of computable policy rules. So this missing piece would be relatively easy to provide.

3) Given 1) and 2), we're good to go as long as the VO is willing to offer PIP/PAP as a service by which RPs manage their access policies and persist their policy information (Role/Group memberships).

4) There is an interesting subset of use cases for which simple attribute aggregation would suffice. All the essential attributes relevant to the RPs access control policy are available from either the IdP or the VO-hosted PIP as postulated here. Shib offers good-enuf simple aggregation facilities to support this. Do other open source SAML implementations support it? If not, it shouldn't be too hard to add, right?

5) There is a key constraint on the subset of solvable use cases: The subject's chosen IdP must assert a (possibly pseudonymous) user identifier that the PAP admin can use (directly or via a mapping) as a relatively persistent identifier with which to associate groups and roles.

6) Leif, Let's postulate that a particular RP needs more than a SAML assertion from the IdP containing an entitlement or a group membership. In such a case the application integration challenge then boils down to "At each point in your application logic where you need to enforce one of your RP policies, have the application code make a call out to the PDP passing in Subject, Action, Resource and then proceeding according to whether the response was 'allow' or 'deny'".

 --KeithH

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 5:00 AM
To: Keith Hazelton <hazelton@doit.wisc.edu>
Cc: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

> 6) Leif, Let's postulate that a particular RP needs more than a
> SAML assertion from the IdP containing an entitlement or a group
> membership. In such a case the application integration challenge
> then boils down to "At each point in your application logic where
> you need to enforce one of your RP policies, have the application
> code make a call out to the PDP passing in Subject, Action,

> Resource and then proceeding according to whether the response was
> 'allow' or 'deny'".
>

Right. I was trying to make the point that authn was relatively
easy (!) because applications typically have a single entrypoint
for login.

Changing that point or adding a new entrypoint for "the shibboleth
thing" is not that difficult.

Now we're asking applications to _at every point_ they do an
authz decision, add a callout to a PDP.

That is a pretty tall order.

Back in the day when java looked good to most people the J2EE
framework had this nice authz abstraction that made it pretty
easy to stick the PDP callout in the application server. We
used that at my old job with SPOCP and in the enterprise this
stuff is still relatively easy to motivate.

I'd say the best chance we have is to build authz servers
for the post-J2EE-enterprise is to base them on OAuth (perhaps
using UMA) that are *easy* to hook applications up to.

This is sortof what the VOOT community is going for, except
we got lost in portal-land in the form of OpenSocial for a
couple of years :-)

MVH leifj

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBJxcIACgkQ8Jx8FtbMZneNqwCdFEsDrGzblzk+Wl9Q+JNqVHD+
jjwAn2u2bzZwFsuMJ2uupuFz0X1R+21+
=pePD
-----END PGP SIGNATURE-----

Keith Hazelton <hazelton@doit.wisc.edu> Fri, Sep 7, 2012 at 5:12 AM
To: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>, PB IAM <iam@lists.projectbamboo.org>

On Sep 7, 2012, at 12:00:35, Leif Johansson wrote:

> -----BEGIN PGP SIGNED MESSAGE-----
> Hash: SHA1
>
>
>> 6) Leif, Let's postulate that a particular RP needs more than a
>> SAML assertion from the IdP containing an entitlement or a group
>> membership. In such a case the application integration challenge
>> then boils down to "At each point in your application logic where
>> you need to enforce one of your RP policies, have the application
>> code make a call out to the PDP passing in Subject, Action,
>> Resource and then proceeding according to whether the response was

>> 'allow' or 'deny'".
>>
>
> Right. I was trying to make the point that authn was relatively
> easy (!) because applications typically have a single entrypoint
> for login.
>
> Changing that point or adding a new entrypoint for "the shibboleth
> thing" is not that difficult.
>
> Now we're asking applications to _at every point_ they do an
> authz decision, add a callout to a PDP.
>
> That is a pretty tall order.

I guess the question is, if you are developing a new RP app and you need to evaluate authZ questions at multiple points in the app flow, you are in the thicket already and having a place to call for allow/deny decisions would seem like a blessing. If you don't have such a place, what will you do? Seems like it will be some form of hard coding of the access policy rule in application logic based on a set of attributes in the app environment somewhere. Is that easier? Would it still be easier if the RP's VO offered all the essential service bits to support the external AuthZ model?

>
> Back in the day when java looked good to most people the J2EE
> framework had this nice authz abstraction that made it pretty
> easy to stick the PDP callout in the application server. We
> used that at my old job with SPOCP and in the enterprise this
> stuff is still relatively easy to motivate.
>
> I'd say the best chance we have is to build authz servers
> for the post-J2EE-enterprise is to base them on OAuth (perhaps
> using UMA) that are *easy* to hook applications up to.
>
> This is sortof what the VOOT community is going for, except
> we got lost in portal-land in the form of OpenSocial for a
> couple of years :-)

I had already concluded from VAMP side-conversations that I needed to get serious about VOOT. Glad you have finished your desert wanderings =) --Keith

[Quoted text hidden]

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 5:31 AM
To: mace-paccman@internet2.edu

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

On 09/07/2012 12:12 PM, Keith Hazelton wrote:
> On Sep 7, 2012, at 12:00:35, Leif Johansson wrote:
>
>

[Quoted text hidden]

I think most apps think in terms of groups for authz - only they don't agree on what a group is.

For some of the stuff we're doing at NDN we're using django and there (this is just an example) the group concept is very rudimentary but it does what it needs to do in order to support the authz model of that framework.

Slapping VOOT onto the django groups framework turned out to be
relatively easy and that way we were able to avoid having to
change the code in the application.

Just an example but it illustrates my point: if you can change
the container instead of the application, life becomes much
easier.

We should be looking for those integration-points if we want
authz to succeed.

>
> Back in the day when java looked good to most people the J2EE
> framework had this nice authz abstraction that made it pretty easy
> to stick the PDP callout in the application server. We used that at
> my old job with SPOCP and in the enterprise this stuff is still
> relatively easy to motivate.
>
> I'd say the best chance we have is to build authz servers for the
> post-J2EE-enterprise is to base them on OAuth (perhaps using UMA)
> that are *easy* to hook applications up to.
>
> This is sortof what the VOOT community is going for, except we got
> lost in portal-land in the form of OpenSocial for a couple of years
> :-)
>
>> I had already concluded from VAMP side-conversations that I
>> needed to get serious about VOOT. Glad you have finished your
>> desert wanderings =)   --Keith
>
perhaps not finished but we can see the oasis (no pun
intended) from where we stand...
-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBJzQgACgkQ8Jx8FtbMZnfODgCfaWhPoSwKzlDyjedNDOx79pEx
Bg4AmgN08nhJit/2hRBU+p/81B2GHPqy
=WACl
-----END PGP SIGNATURE-----

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 5:52 AM
To: Josh Howlett <Josh.Howlett@ja.net>
Cc: David Chadwick <d.w.chadwick@kent.ac.uk>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

On 09/07/2012 12:45 PM, Josh Howlett wrote:
>
>> The problem is that there is nothing that corresponds to the
>> simplicity of the REMOTE_USER model in authz land that can be
>> used as a spring- board for integration.
>
> Isn't this the whole point of RFC6680?

No, 6680 just gives you another way of shipping attributes into the
applications.

We're talking about using those attributes to make authz decisions
and how _that_ process is externalized to a PDP.

Cheers Leif

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBJ0gIACgkQ8Jx8FtbMZndGtACeJiubboBPtl+bErNTuxaYtmNB
W64AnRonMw40VHCZHKR6KWwUiPdqFMES
=Or8a
-----END PGP SIGNATURE-----

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 7:32 AM
To: Tom Scavo <trscavo@internet2.edu>
Cc: David Chadwick <d.w.chadwick@kent.ac.uk>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>, Keith Hazelton <hazelton@doit.wisc.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

On 09/07/2012 02:21 PM, Tom Scavo wrote:
>
>
>> I am afraid that the current metadata distribution mechanism is
>> not sufficient for an Internet wide federation. A sophisticated
>> DNS-like service is needed.
>
> +1
>

Actually I think the mental model we want is BGP not DNS. The success
of BGP is 100% due to the fact that the resource owner (the AS) is in
control of how traffic flows to and trough the domain. DNS has quite
a few problems that we don't want to repeat.

The problem with the BGP analogy is that most people that have heard me
and Klaas talk about it assumes that in this analogy IP traffic becomes
application access.

That is not the case!

In the Internet of Trust all things are reachable but the BGP-analoge
ensures control over the degree of trust you receive from how the
resource is reached.

In the IoT most things are reachable but some things are reachable
with better level-of-confidence, in the same way that on the regular
Internet most things on the public net are reachable but if you want
working HD video from yutube, you need a private peering with Goog.

We just need to find that BGP analogue that works with deployed
federation technology.

Cheers Leif

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBJ6W0ACgkQ8Jx8FtbMZnd1IQCdF7aChtZpW9jGvW6oKVUZ2y2u
528AnAI+vHFSqEb3aUqBt7t7hwuIkaHx
=D9l0
-----END PGP SIGNATURE-----

Bill Thompson <wgthom@unicon.net> Fri, Sep 7, 2012 at 9:09 AM
To: Leif Johansson <leifj@sunet.se>, mace-paccman@internet2.edu

> Just an example but it illustrates my point: if you can change
> the container instead of the application, life becomes much
> easier.
>
> We should be looking for those integration-points if we want
> authz to succeed.

Chris Hyzer and I are scheduled to give a presentation at I2MM that should be of interested. In the model we've been thinking about the app/framework is the PEP, and the PDP is in the form of Grouper's effective permissions sets, which can be pre-calculated via 3 hierarchies (Roles, Actions, Resources). In this model the app receives the effective permissions via the "integration-points" with the various "containers" (i.e. authZ APIs of various frameworks):
* Apache Shiro
* Spring Security
* .NET AuthZ APIs

The idea is that developers using these authZ frameworks (containers) should be able to leverage the power of Grouper without having to know too much Grouper.

Best,
Bill
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Fri, Sep 7, 2012 at 12:53 PM
To: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I don't think it is reasonable to expect external IdPs to do anything solely for the benefit of a RP. I think there are two main reasons: lack of vested interest in the RP and simple scaling issues.

Going forward, my opinion is that IdPs should be relied on to provide only the most basic authentication. In other words from the RP perspective the various trusted IdPs should be viewed as authentication services and nothing more. IdPs should be in the business of identifying and credentialing subjects (for a fee, or for some other internal benefit).

Regarding authorization, this leaves the RP two options: 1) handle authorization internally based on the authentication credential received from the IdP and/or 2) request authorization information/attributes from a third party based on the authentication credential received from the IdP.

Note the definition of 'authentication credential' I work with - an electronic authentication credential identifies the credential issuer, a level of assurance for the authentication, and the identifier issued to the subject by the credential provider.

So, though patterns 1 and 2 below work, I think they will become increasingly unworkable as the number of RPs explode and more RPs become external as opposed to most RPs being controlled by the same entity as the IdP or where the RP is content with very course authorization attributes such as eduPersonAffiliation.

This thinking will also apply to attribute providers. These third parties, in my opinion, must be motivated to maintain their data for some reason other than the benefit of a specific RP. For instance a university is likely to maintain (be authoritative) if a person is a student or faculty for their own purposes and be willing to provide that data to external parties as an attribute provider.

So looking at this as the RP and the user being the primary actors, and the IdP and one or more attribute providers being third parties, it all depends on these third parties being trusted and being able to discuss the user with the RP in terms of identifiers without ignoring privacy issues.
[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Fri, Sep 7, 2012 at 1:20 PM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I prefer the IdP to provide authorization services, that way there is no reason to release an identifier or any other attributes to applications or services that the person is not authorized to use. The practice of releasing identifiers based only on authentication leads to the use of authorization and attribute silos and can lead to requirements placed on the identifier such as name-based.

Regards,

Brendan Bellina
USC
[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Fri, Sep 7, 2012 at 1:52 PM
To: Brendan Bellina <bbellina@usc.edu>, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

For federated apps (which aren't based on affiliation) authz probably can't be in each user's IdP...

Penn has a federated collaboration application which is kind of like box.net. The components are the user's IdP, Penn's Grouper, and the application SP. The user must first register at the shib'ed Grouper external registration application based on an email-based invite from an admin in one of the "space's". Grouper's SP isn't necessarily the same as the application SP, so a non-opaque ID is used. It would be nice if uApprove worked with the user's IdP, but in half of the cases the user must either ask their IdP operator to release EPPN to our two SP's, or create a ProtectNetwork ID (most common solution). A lot of the users using the system need some help getting started. I look forward to when this is smoother in the future...

The thought was that there is some economy of scale... the protectnetwork ID could be used in other places, and once the user registers in Penn's Grouper, other Penn federated apps could assign authorizations to that user. However, this hasn't really happened yet. BTW, Penn has a support contract with ProtectNetwork...

Lastly, the application just talks to Penn's Grouper directly, it doesn't use a saml attribute resolver which talks to Penn's Grouper, though it could...

Thanks,
Chris
[Quoted text hidden]

Leif Johansson <leifj@sunet.se> Fri, Sep 7, 2012 at 3:41 PM
To: Josh Howlett <Josh.Howlett@ja.net>
Cc: David Chadwick <d.w.chadwick@kent.ac.uk>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

> I agree that's a hard problem if you want a model where the
> attributes flow is AA -> RP -> PDP because it imposes complexity on
> the RP. I think the ideal general flow is AA -> PDP -> RP; but the

> model should also allow other compositions (e.g., if you have a
> particularly picky RP).

Right but that still leaves the question open on how the RP - PDP integration happens.

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBKXAoACgkQ8Jx8FtbMZnfbawCgoRasmCIBDcJqWPRTTrsUXjGE
OBAAoISRbrXO/KVnbI2jJ+d/kGD1xWB4
=o+Ec
-----END PGP SIGNATURE-----

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Fri, Sep 7, 2012 at 7:14 PM
To: Brendan Bellina <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

InCommon has 865 SPs and 263 IdPs. Are you saying that you want your IdP to be responsible for knowing which of your users access each of the 865 SPs and what authorizations within each of the 865 SPs each of your users have been assigned? And by extension are you saying that each of the 263 IdPs should do the same?

Why do you say release of identifiers leads to attribute silos and use of name-based identifiers. I have not observed this to be the case or I don't understand what you are referring to.

-----Original Message-----
From: Brendan Bellina [mailto:bbellina@usc.edu]
Sent: Friday, September 07, 2012 1:21 PM
To: Jones, Mark B
Cc: Keith Hazelton; REFEDS; paccman
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Fri, Sep 7, 2012 at 7:28 PM
To: Chris Hyzer <mchyzer@isc.upenn.edu>, Brendan Bellina <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

This seems like a good use case to discuss.

I assume that your shib'ed Grouper and application SP consume the same identifier and then use that identifier to talk to each other. Which attribute/identifier are you using?

-----Original Message-----
From: Chris Hyzer [mailto:mchyzer@isc.upenn.edu]
Sent: Friday, September 07, 2012 1:53 PM
To: Brendan Bellina; Jones, Mark B
Cc: Keith Hazelton; REFEDS; paccman
[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Fri, Sep 7, 2012 at 9:08 PM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, Brendan Bellina <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

The eppn is used for several purposes:

1. It correlates the user at the SP and Grouper External User Registration page
2. Grouper uses it as the Grouper Subject identifier, so the application behind the SP can look up authorizations based on SP
3. (This is important and can be overlooked) It is used in the Grouper subject description (and the Grouper UI) so when admins of groups or permissions are assigning them, they see something like this (customizable at your institution of course):

[unverifiedInfo] John Smith - Some University [externalUserID] jsmith@someuniversity.edu

The "John Smith" and "Some University" are self entered on the Grouper External User Registration page (could be defaulted to saml data). But the external user id is the EPPN. Granted there are issues with changing EPPNs etc. We didnt go the route of assuming that name is verified if delivered by SAML since do we really know if its the case for all IdPs? (ProtectNetwork is self entered unvetted, and maybe their institution takes it from their directory which could be self entered unvetted...)

4. Its nice to be able to look in the database based on an EPPN and troubleshoot user problems (could be more difficult with an opaque ID). There could be an email address where the invite was sent, but it might not be correct. We have cases where people send the invite to themselves, then get the real invitee who has trouble with technology to share their computer, login, and continue. So the email address where the invite was sent might not be owned by the opaque ID...

Thanks,
Chris

_____
From: Jones, Mark B [Mark.B.Jones@uth.tmc.edu]
Sent: Friday, September 07, 2012 8:28 PM
To: Chris Hyzer; Brendan Bellina
[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Sat, Sep 8, 2012 at 12:04 AM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I think Chris's example answered these questions. They release eppn because it is useful to have a name based identifier when you are assigning privileges and troubleshooting permissions and access issues.

Of course people change names and so identifiers that are name based often need to change as well. If identifiers are released only based in authentication and authorization is managed locally then there will be pressure to use name based identifiers to simplify the authorization. So, as I said, I prefer to not release anything except for authorized individuals. In our case we do not release any attributes to any SPs without review and approval and managing the authorization using groups.

I realize it is more common practice to release attributes and identifiers without this kind of governance at the IdP, but we have found the governance to be more beneficial than inconvenient.

Regards,

Brendan Bellina
USC ITS
[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Sat, Sep 8, 2012 at 7:02 AM
To: Brendan Bellina <bbellina@usc.edu>, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Right, understood. When you say "They release eppn", you mean, "the SP requires EPPN", right? In our case the user probably doesnt know who to ask or how to ask to get their IdP's EEPN released, and might not have

time to deal with it, so a USC user might just make a ProtectNetwork ID to be able to continue working and not use their USC netId for our application... either way it is tedious for the user. If they did change their name, I dont think it would be a burden for the IdP operator, the user would either need to have Penn's Grouper be changed with the new EPPN or go through the intake process again (get invited via email, sign in with new credential, get the application admin to assign necessary authorizations). I could see not releasing EPPN for privacy reasons, but it is very painful for users. uApprove addresses this, or if there were a way to not have SSO, then if the user is typing in their user/pass, I think they expect their username to be sent to the SP...

Thanks,
Chris

_____
From: bbellina@usc.edu [bbellina@usc.edu]
Sent: Saturday, September 08, 2012 1:04 AM
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Sat, Sep 8, 2012 at 9:01 AM
To: "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I still think that managing authorization at the IdP will become increasingly unworkable. This is especially true if the IdP is Google or Verizon. Chris says that many of his users elect to use ProtectNetwork for example. Would you be comfortable allowing ProtectNetwork manage authorization to your SPs and would ProtectNetwork even be willing to do such a thing?
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Sat, Sep 8, 2012 at 9:07 AM
To: Chris Hyzer <mchyzer@isc.upenn.edu>, Brendan Bellina <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Did you customize Grouper to use Shib or can you use Shib 'out-of-the-box'?
Is Grouper using ePPN as the internal identifier or just to map the user to an internal identifier?
Could something other than ePPN be used by Grouper?

-----Original Message-----
From: Chris Hyzer [mailto:mchyzer@isc.upenn.edu]
[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Sat, Sep 8, 2012 at 10:28 AM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

We have SPs that elected to trust ProtectNetwork logins. The user has to self-register at our guest page and is redirected to their IdP of choice (off an approved list we have agreements with, currently includes ProtectNetwork but will be expanding to oAuth shortly). This creates an entry for them in our directory and allows their authorizations to be managed by our groups processing. So while in this case the IdP of their choice is controlling only authentication, our IdP is providing the authorization services. This still prevents the need for authorizations to be managed at the SP or application.

Russ Beall will be presenting on our solution at the I2MM, focusing on the oAuth integration.

Regards,

Brendan Bellina
USC

[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Sat, Sep 8, 2012 at 7:15 PM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, Brendan Bellina <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

No, Grouper is not customized, works with shib out of the box. Actually, at Penn we have cosign on most of grouper, and shib only on the external registration part (for extra security).

Grouper uses a uuid as the internal id, so the eppn (or whatever) could be changed in one place easily.

Doesnt have to be eppn, could be something else.

Here is the doc:

https://spaces.internet2.edu/display/Grouper/Grouper+external+subjects

Thanks,
Chris

_____
From: Jones, Mark B [Mark.B.Jones@uth.tmc.edu]
Sent: Saturday, September 08, 2012 10:07 AM
[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Sat, Sep 8, 2012 at 7:20 PM
To: Brendan Bellina <bbellina@usc.edu>, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

If your IdP is doing authorizations for external subjects (resolving attributes) for your applications then I think we are all saying the same thing... that the external end-user's idp isnt doing all authorizations for federated applications, its something related to the SP... (which could be the local user's idp)

Thanks,
Chris

_____
From: bbellina@usc.edu [bbellina@usc.edu]
Sent: Saturday, September 08, 2012 11:28 AM
[Quoted text hidden]

Leif Johansson <leifj@sunet.se> Sun, Sep 9, 2012 at 7:54 AM
To: David Chadwick <d.w.chadwick@kent.ac.uk>
Cc: Josh Howlett <Josh.Howlett@ja.net>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

On 09/08/2012 07:40 AM, David Chadwick wrote:
>
>
> On 07/09/2012 21:41, Leif Johansson wrote:
>
>>>> I agree that's a hard problem if you want a model where the
>>>> attributes flow is AA -> RP -> PDP because it imposes
>>>> complexity on the RP. I think the ideal general flow is AA ->
>>>> PDP -> RP; but the model should also allow other compositions
>>>> (e.g., if you have a particularly picky RP).

> Right but that still leaves the question open on how the RP - PDP
> integration happens.
>
>> Several years ago the OGF did a lot of work on this and published
>> specs to show all these message flows (both push and pull), and
>> built working demonstrators for the grid using a variety of
>> opensource products.
>
>> See http://www.ogf.org/gf/docs/?final
>
>> and pick up docs
>
>> 156,157,158,159
>
>> and the earlier 66,67 for the first attempts
>
>
OK I guess I was not being clear. The xacml stuff is fine etc but
I was talking about how to make it easy for applications sitting
behind RPs to talk to PDPs, and I don't think "make an xacml query"
is a good answer.

Cheers Leif
-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBMkYgACgkQ8Jx8FtbMZndUlQCgjwmxg331X+hA9ZHypGB4I2Nf
EvoAn0TmrOdWtqzhjfHLt3RH+sYFsNlm
=Qj2Z
-----END PGP SIGNATURE-----

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Sun, Sep 9, 2012 at 10:25 PM
To: "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

This is the type of use case I expect to be more and more common. Authorization is not handled by the user's IdP or the SP.

One thing that is not clear... are you saying that your "groups processing" is the same thing as your IdP? If the user is not authenticating to your IdP how is it providing authorization services?

-----Original Message-----
From: bbellina@usc.edu [mailto:bbellina@usc.edu]
[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Sun, Sep 9, 2012 at 10:56 PM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Our groups processing allows entitlements and other attributes to be set based on memberships in groups. The groups are managed in our enterprise directory. Because the guest registers he/she gets an entry in our directory and that can be placed in groups so that entitlements are associated with it. When the guest tries to access one of our SPs his/her IdP provides authentication and an identifier and our IdP allows the release of those attributes and any that are in our directory based on the entitlements. So our IdP provides authorization and attribute enrichment based on the information given at registration and group memberships.

Regards,

Brendan
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Sun, Sep 9, 2012 at 11:01 PM
To: "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Is your IdP something other than Shibboleth? Wouldn't the user have to login at your IdP as well in order for it to release attributes?

-----Original Message-----
From: bbellina@usc.edu [mailto:bbellina@usc.edu]
[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Mon, Sep 10, 2012 at 9:13 AM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

We use standard Shibboleth. The Shibboleth SP can be configured to chain multiple IdPs so that attributes can be consolidated from multiple sources. In this case the user's IdP provides eppn and our IdP provides group memberships, name, email, our internal id, and entitlements.

I think we gave a presentation on this at the I2MM Shibboleth session a year ago or so on this. I'll see if I can find a link to it.

Regards,

Brendan Bellina
USC ITS

[Quoted text hidden]

Chris Hyzer <mchyzer@isc.upenn.edu> Mon, Sep 10, 2012 at 9:51 AM
To: "bbellina@usc.edu" <bbellina@usc.edu>, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I wonder if there is some terminology that we could use that would be more clear. i.e. we started this discussion out saying "I prefer the IdP to provide authorization", but really the IdP is providing the identity and the attribute resolver is providing the authorization (and the attribute resolver just happens to be an IdP for some users as well). Just curious, how are these two things differentiated in the talk at I2? Can we refer to the non-identity IdP something else besides IdP? :)

Thanks,
Chris

-----Original Message-----
From: bbellina@usc.edu [mailto:bbellina@usc.edu]
[Quoted text hidden]

Cantor, Scott <cantor.2@osu.edu> Mon, Sep 10, 2012 at 10:13 AM
To: Chris Hyzer <mchyzer@isc.upenn.edu>

Cc: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

> On 9/10/12 10:51 AM, "Chris Hyzer" <mchyzer@isc.upenn.edu> wrote:
>
> >I wonder if there is some terminology that we could use that would be
> >more clear. i.e. we started this discussion out saying "I prefer the IdP
> >to provide authorization", but really the IdP is providing the identity
> >and the attribute resolver is providing the authorization (and the
> >attribute resolver just happens to be an IdP for some users as well).
> >Just curious, how are these two things differentiated in the talk at I2?
> >Can we refer to the non-identity IdP something else besides IdP? :)
>
> Well, in SAML a system is called an Attribute Authority if it responds to
> stand alone queries for information about users.
>
> -- Scott

Cantor, Scott <cantor.2@osu.edu> Mon, Sep 10, 2012 at 10:45 AM
To: Tom Scavo <trscavo@internet2.edu>
Cc: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

> On 9/10/12 11:33 AM, "Tom Scavo" <trscavo@internet2.edu> wrote:
> >
> >Yes, and the US government calls this "Backend Attribute Exchange" (you
> >can google that term for more info :)
>
> Well, they're referring to a deployment and a profile overall, not just
> the concept of an attribute authority, but yes, they use them. I'm told by
> commercial implementers I won't name that they've managed to violate the
> SAML spec in their profile.
>
> -- Scott

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Mon, Sep 10, 2012 at 10:53 AM
To: Chris Hyzer <mchyzer@isc.upenn.edu>, "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

> I agree. I do feel as if I am tripping over terminology. 'IdP' does not seem specific enough for this conversation.
> At a minimum we need to specify 'authentication credential provider' or 'attribute provider'.
>
> -----Original Message-----
> From: Chris Hyzer [mailto:mchyzer@isc.upenn.edu]
> Sent: Monday, September 10, 2012 9:52 AM
> To: bbellina@usc.edu; Jones, Mark B
> Cc: Keith Hazelton; REFEDS; paccman
> [Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Mon, Sep 10, 2012 at 10:56 AM
To: "Cantor, Scott" <cantor.2@osu.edu>, Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

> What is the SAML term for a system that provides authentication?

-----Original Message-----
From: mace-paccman-request@internet2.edu [mailto:mace-paccman-request@internet2.edu] On Behalf Of Cantor, Scott
Sent: Monday, September 10, 2012 10:13 AM
To: Chris Hyzer
Cc: REFEDS; paccman
Subject: Re: [paccman] Late night thoughts on authorization

[Quoted text hidden]

Brendan Bellina <bbellina@usc.edu> Mon, Sep 10, 2012 at 10:59 AM
To: Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

The start of the discussion was Mark's statements that indicated that IdP's should provide only basic authentication and identifier release and that authorization should be handled by the applications. I don't agree that local authorization is the way to go and I also think that it is best not to release identifiers or any attributes about an individual if you can determine ahead of time that this individual is not authorized. So I am a fan of external authorization and I think we should be encouraging application developers to externalize permissions rather than continue internalizing them. Tough battle I know.

Regards,

Brendan
[Quoted text hidden]

Cantor, Scott <cantor.2@osu.edu> Mon, Sep 10, 2012 at 11:04 AM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>
Cc: REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

On 9/10/12 11:56 AM, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu> wrote:

>What is the SAML term for a system that provides authentication?

Well, if by "authentication" one means "supports the SAML Authentication Request protocol", that's what IdP is used to mean. Formally anyway.

-- Scott

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Mon, Sep 10, 2012 at 11:06 AM
To: Brendan Bellina <bbellina@usc.edu>, Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Well we have some level of miscommunication. I was not recommending that authorization be handled locally by applications.

I'm still confused about your not wanting to release identifiers. Your own example has the user's IdP releasing ePPN and then authorization attributes being released by your IdP. Isn't this a release of an identifier by the authenticating IdP and then a request to a third party for authorization attributes?

-----Original Message-----
From: Brendan Bellina [mailto:bbellina@usc.edu]
Sent: Monday, September 10, 2012 10:59 AM
To: Chris Hyzer

[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Mon, Sep 10, 2012 at 11:13 AM
To: "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

So in this chaining of multiple IdPs, are the 'IdPs' subsequent to the first treated as 'SAML Attribute Authorities'.

-----Original Message-----
From: bbellina@usc.edu [mailto:bbellina@usc.edu]
[Quoted text hidden]

Cantor, Scott <cantor.2@osu.edu> Mon, Sep 10, 2012 at 11:29 AM
To: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, "bbellina@usc.edu" <bbellina@usc.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

On 9/10/12 12:13 PM, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu> wrote:

>So in this chaining of multiple IdPs, are the 'IdPs' subsequent to the
>first treated as 'SAML Attribute Authorities'.

Yes, it's not chaining IdPs, it's adding queries to additional AAs and if you gave it a plugin to do something else it would do that also. As far as the SP is concerned, the only IdP is the one that initiated the login.

-- Scott

Andy Dale <dalea@oclc.org> Mon, Sep 10, 2012 at 11:34 AM
To: Brendan Bellina <bbellina@usc.edu>, Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I think we need to embrace the fact that different policies can and should be evaluated in different places.

There are policies that can reasonably be delegated to a trusted IDP to evaluate such as "Over 18". (Globally Abstracted Policies)

There are policies that can reasonably be delegated to an Authorization Service. Generally these are policies that are non resource specific but more complex than one would want to have to move out to each IDP. For example: Is a member of Group A or (Group B and Group C). (Enterprise Abstraction)

Often SPs will have custom authorization policy that is for example, resource specific. The level of expert knowledge needed to evaluate these policies is often unreasonable to externalize. (Implemented in place)

I think we need semantics to describe and make use of all of these patterns. Not suggesting that any of these are easy, just that I think we need all three rather than thinking we can pick one.

In my recent implementations we use all three patterns and then use a 'Voter Pattern' with multiple Access Decision Managers to bring the three together in a simple configuration/implementation.

Andy Dale
[Quoted text hidden]

Jones, Mark B <Mark.B.Jones@uth.tmc.edu> Mon, Sep 10, 2012 at 11:43 AM
To: Andy Dale <dalea@oclc.org>, Brendan Bellina <bbellina@usc.edu>, Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

Good point that there is no 'one size fits all' authorization strategy.
[Quoted text hidden]

Leif Johansson <leifj@sunet.se> Mon, Sep 10, 2012 at 3:18 PM
To: David Chadwick <d.w.chadwick@kent.ac.uk>
Cc: Josh Howlett <Josh.Howlett@ja.net>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

On 09/10/2012 04:23 PM, David Chadwick wrote:
> Agreed
>
> You need an API in the programming language of the application. We
> are currently working on this in Python for OpenStack. We produced
> a set of APIs in PHP last year and published them here
>

Before we need APIs we need concepts and paradigms.

Are we talking about a group or an RBAC abstraction for instance? What
are applications doing today?

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.11 (GNU/Linux)
Comment: Using GnuPG with Mozilla - http://www.enigmail.net/

iEYEARECAAYFAlBOSy4ACgkQ8Jx8FtbMZnfkQQCfY0tD8wPiPiccoUEFs7gV9S2B
0r8AnixYXFfJYiW/rNph99DxT8IHBOtT
=f1dH
-----END PGP SIGNATURE-----

Tom Dopirak <tgd@cmu.edu> Wed, Sep 12, 2012 at 10:18 AM
To: Chris Hyzer <mchyzer@isc.upenn.edu>
Cc: "bbellina@usc.edu" <bbellina@usc.edu>, "Jones, Mark B" <Mark.B.Jones@uth.tmc.edu>, Keith Hazelton <hazelton@doit.wisc.edu>, REFEDS <refeds@terena.org>, paccman <mace-paccman@internet2.edu>

I was wondering if we wanted to set aside some time at the Paccman Working Group at I2 to talk about this further and who might want to summarize this thread for the greater group and lead the conversation. We do expect to touch o this topic on this Thursday's 1PM eastern call.

Tom Dopirak
tgd@andrew.cmu.edu
Senior Consulting Architect
Carnegie Mellon University
412-268-8691

"I'm just glad to be feeling better. I really thought I'd be seeing Elvis soon." - Dylan

[Quoted text hidden]

Keith Hazelton <hazelton@doit.wisc.edu> Wed, Sep 12, 2012 at 10:28 AM
To: paccman <mace-paccman@internet2.edu>
Cc: Chris Hyzer <mchyzer@isc.upenn.edu>, Brendan Bellina <bbellina@usc.edu>, Mark Jones <Mark.B.Jones@uth.tmc.edu>

On Sep 12, 2012, at 10:18:26, Tom Dopirak wrote:

> I was wondering if we wanted to set aside some time at the Paccman Working Group at I2 to talk about this further and who might want to summarize this thread for the greater group and lead the conversation. We do expect to touch o this topic on this Thursday's 1PM eastern call.

I'll offer a summary of the thread in advance of the Paccman WG call. --Keith
[Quoted text hidden]