GENI Clearinghouse Policy

Name: GENI Clearinghouse Policy
Author: Adam Slagell
Version: 0.2.1 (Working Draft)
Date: July 21, 2011

1. Goals & Scope

This is a policy document describing to the GENI federation actors (e.g., aggregate authorities, identity providers, project leaders, and experimenters) what a clearinghouse is, the services it will provide, and the policies it will follow to operationally meet the needs of the community. This is not an agreement itself, but instead these actors will individually sign agreements specific to their relationships with the clearinghouse, which acts as a trust anchor for the GENI federation. The critical federation agreements and plans are therefore specified in this high-level policy document. **This policy and the referenced agreements do not preclude any bilateral agreements between experimenters, aggregate authorities, identity providers and other GENI service providers, though.**

2. Related Documents

These related documents contributed strongly to our understanding of what the GENI Clearinghouse is: its roles, responsibilities and services provided. A changing and developing understanding of GENI's architecture or the federation's organization may result in changes to this concept of clearinghouse operations. Any such changes will require the definitions in this document to be updated, and their impact on the principles espoused here will have to be determined.

Document ID	Document Title and Issue Date
GeniSysOvrvw	"GENI System Overview", September 29, 2008. http://groups.geni.net/geni/wiki/GeniSysOvrvw
SysReqDoc	"GENI System Requirements Document", July 7, 2009. http://groups.geni.net/geni/wiki/SysReqDoc
ExperimentLifecycleDocument	"Lifecycle of a GENI Experiment", April 30, 2009 http://groups.geni.net/geni/wiki/ExperimentLifecycleDocument
GEC 11 Slides	"Federation in GENI & the GENI Clearinghouse", July 27, 2011

3. What is the GENI Clearinghouse?

Before we discuss the purpose and goals of this agreement, it is crucial to come to a common understanding of what a clearinghouse is. In this section we briefly discuss what a clearinghouse is, its role in the GENI federation, and what sort of services it will/may provide.

Above all else, the GENI Clearinghouse (hereto referred to as "the clearinghouse") is a trust anchor. It acts as a root certificate authority (CA) for any party issuing credentials (including itself) or providing resources (e.g., identity providers, aggregate managers, and slice registries). It also serves as a trust bridge between aggregate authorities (AAs), identity providers (IdPs) and experimenters by each party establishing agreements with the clearinghouse. Alternatively, we could have N^2 pairwise agreements between all parties, but that is clearly not scalable. **Therefore, the clearinghouse becomes (1) a bridge between parties to minimize the number of agreements and (2) associated with a set of policies/agreements that define how the federation operates.**

The clearinghouse will also provide general portal operations for setting up and creating experiments. At a minimum, this means supporting the creation and registration of projects, project leaders, slices and experimenters. Being a clearinghouse for all GENI-approved aggregates (i.e., those that have signed an Aggregate Provider Agreement with the clearinghouse), the clearinghouse becomes essential for resource discovery with a well-maintained component registry. **While some aggregates may choose to accept slices not signed by and registered with the clearinghouse, some will require its endorsement because that endorsement means that the other aggregates and the users have signed off on a common set of policies with certain expectations.**

Most federation-level resource allocation policies could only be verified at the clearinghouse. An aggregate may not have a wide-enough view to implement such a policy without help from the clearinghouse. Whether this is a policy from NSF (e.g., "student created slices can only have X hosts") or part of an agreement with an IdP and AA such as FIRE (e.g., "GENI users can only take Y Mbps of link Z to the EU"), the clearinghouse is the only place that can verify whether such policies are being followed. **However, the AAs remain completely autonomous and can decide whether or not to act upon information from the clearinghouse when deciding whether or not to grant sliver requests to valid slices.** Aggregates are only asked to report to the clearinghouse what resources they gave to registered slices.

In addition to the critical services of resource discovery, project creation and slice registration, the clearinghouse may one day provide many secondary functions: experiment template discovery services, experiment archiving and sunsetting services, and deployment tools. A user-friendly and fully functional portal should provide all the tools (or interfaces to them) necessary throughout an experiment's lifecycle in a single place. For any such additional services, this clearinghouse agreement should include, or point to, service-level agreements telling AAs, project leaders and experimenters what they should expect.

The clearinghouse will also have many accounting responsibilities with registering projects, project leads, experimenters and slices; all activities which will be logged. These activities in turn involve issuing credentials, verifying attributes, presenting agreements for acceptance and checking resource allocation policy. **Several databases and logs will need to be maintained for all these activities, and the privacy policies stated in this document describe how these logs will be used and protected.**

4. GENI Federation Agreements & Actors

For all the actors below, the clearinghouse should maintain the authoritative list. Aggregate providers who have signed an agreement with the clearinghouse will implicitly be listed within the list of all components provided to slice creators for resource discovery. Valid identity providers and slice registries should be listed prominently on the clearinghouse web page. GENI credentials issued by a valid identity provider will be signed by either the clearinghouse or another key whose certificate is signed by the clearinghouse. Valid GENI slices will be signed when registered with the clearinghouse. While not public, a list of all registered project leaders will be held at the clearinghouse. All experimenters with valid GENI credentials will either be listed in the principal registry at the clearinghouse, or in the registries at IdPs who are authorized to mint GENI credentials, such as, Emulab.

Aggregate Providers

The aggregate providers are vital to the community as they are providing all the resources (e.g., virtual machines, routers, network links, etc) for experiments in GENI. Because of their importance, and because the aggregates appeared before even the users, an **Aggregate Provider Agreement** was developed early on. This agreement sets a baseline of expectations for the behavior and responsibilities of GENI aggregates. Experimenters building slices will know that resources listed at the clearinghouse are those of the aggregates that have signed off on this agreement. Aggregates seeing requests for resources from slices signed by the clearinghouse will know that they are dealing with users who have signed off on a GENI Acceptable Use Policy.

While an initial Aggregate Provider Agreement was approved at earlier GECs, it should be updated to reflect new language and concepts that have developed along with this conception of the clearinghouse. Therefore, as this concept and agreement solidifies, the Aggregate Provider Agreement must be reviewed and updated.

Identity Providers

The concept of multiple identity providers for GENI is relatively new, at least in its formal conception. Campuses, through InCommon, and other testbeds will all act as identity providers, vetting users, passing attributes, issuing and revoking credentials. Some (e.g. Emulab) will also serve as full GENI principal registries, while others (e.g., InCommon) may just provide authentication and use the clearinghouse to create credentials and register users . Therefore, agreements between the clearinghouse and different IdPs will have to be customized based upon the services provided. However, each agreement will have to address common issues such as the level of assurance for user registration and authentication, revocation capabilities, log retention, and the types of attributes that are collected and passed on to the clearinghouse.

A slice registry (besides the one at the clearing house) is a special type of identity provider and requires a different sort of agreement. If slices are created not by the clearinghouse and by users who are not registered with the clearinghouse but another identity provider, then there must still be a way to map slices to real people. The clearinghouse, when registering and signing such a slice, is trusting that the slice registry (potentially just another function of an existing identity provider) is (1) logging all transactions of slice creation, (2) maintaining a mapping of slice creators to principals (That might be in the slice credential and verified at registration, but this registry should know everyone that is authorized to act on a slice), (3) creating slices only for users from a valid GENI IdP who has presented to users a GENI AUP (This can probably be checked at registration by CH), and (4) maintaining an ability to revoke slices (Might be enough that the clearinghouse can revoke its signature). All four of those points are important parts of an effective agreement with a slice registry and must be present, even if they are wrapped up as part of a larger IdP agreement instead of a separate slice registry agreement

Experimenters

While there is a Recommended Use Policy for users in GENI, there is *currently* no formal Acceptable Use Policy (AUP). Further, the RUP is not actually presented to anyone. **It is important that a formal AUP be presented to users upon registration and before they are issued any GENI credentials.** This should be based upon the RUP, which states things, such as, "users must follow all applicable laws and regulations". However, it must go further and clarify to experimenters that they are ultimately responsible for the behavior of their experiments. Furthermore, proper treatment towards opt-in users and the need to register any software they install on opt-in user systems must be made clear. Experimenters should respect the privacy of opt-in users and their wishes to no longer be a part of an experiment. Lastly, this AUP should point to the Legal, Law Enforcement and Regulatory (LLR) Plan.

Project Leaders

Project Leaders may be users themselves, but need not be the ones creating slices and actually running experiments. However, they have a responsibilities as the ones who guide projects and delegate the ability to create slices. Therefore, they need to register with the clearinghouse, and any slice that is to be registered with the clearinghouse must be tied back to a registered project. The way this will work is that the clearinghouse will give the project leader a project credential. The project leader can use that credential to delegate the ability to create slices for that project to another user. That user will create a slice with a valid slice registry. To register the slice with the clearinghouse, the slice creator will need to present the delegation credential along with the slice credential to the clearinghouse, perhaps with some additional self-certified information requested at that time.

The project leader therefore has responsibility for all the slices created with their project. Therefore, an agreement should be presented to the project leader, when registering with the clearinghouse, which has all the content of the AUP, but additionally holds them responsible for being truthful in the project description. It must also make them aware of their responsibility for all the project's slices, and hence importance of not carelessly delegating the authority to create slices.

Status of Discussed Federation Agreements and Plans

- ✓ Aggregate Provider Agreement
  - *It will need some updates as the clearinghouse is implemented, but overall it is in good shape.*
- ✗ Identity Provider Agreement
  - *No IdP agreements have been created, but GENI will need ones with InCommon, Emulab, and Planetlab at a minimum. Agreements may also be needed in the future with FIRE, DETER and ORCA. Each will be unique, but I could imagine two or three basic templates.*
- ✗ Slice Registry Agreements
  - *These are likely to be rolled into IdP agreements as one is unlikely to be an authorized slice registry if not and IdP.*
- ✗ Project Leader Agreement
  - *There is none, but this is not needed until there is a functional clearinghouse*
- ✗ Federation Charter
  - *There should be a federation charter that describes the governance structure of GENI and references all the agreements in this subsection. This section will then be less relevant in this document, and some of the policies (e.g., how disputes are resolved between actors) should be moved there.*
- ✗ Acceptable Use Policy
  - *There is none, but one could and should be created soon based off of the RUP. Perhaps more thought is needed on opt-in user issues, though.*
- ✓ Legal, Law Enforcement & Regulatory Plan
  - *It will need to be updated with new terminology such as "Aggregate Authority". Once the definitions of settle, this can be done.*
- ⚠ Clearinghouse Agreement
  - *This is becoming less an agreement than a root policy document.*
  - *There is a lot to be filled in, but questions need to be answered.* ***Blue is for stuff to be filled in, and red is just for comments or notes.***
- ⚠ Incident Response Plan
  - *It is in an early draft form, but it cannot be implemented without some decisions being made, positions created and funding allocated. The last section should inform future projects to be funded, though.*
- ? Certificate Authority Policy
  - *I'm still not sure whether or not this should be in this document. These policies tend to be rather lengthy and dull. So I lean towards pulling it out. However, a complete CH agreement really depends on it. Either way, we need to understand the technical details of how credentials will be signed and by whom before going too far down this path.*

5. Clearinghouse Services

Project Creation

Every experiment and every slice in GENI will be associated with a GENI Project. For a slice to be registered with the clearinghouse, it must be associated with an official project.

The clearinghouse will provide an interface that allows projects to be registered and approved. First, however, the project leader must register an account with the clearinghouse through an approved identity provider, such as InCommon. Attributes (yet to be defined) will be collected and verified by clearinghouse operations staff, and a GENI credential will be issued to the project leader by the clearinghouse (could they use their Emulab credential? You would still want them to register with the clearinghouse if a Project Leader).

Once an account has been created for the project leader, they can create one or more projects for approval. Projects will have a unique ID as well as a common name. When making the request they should describe the project purpose, it's scope and size, and expected duration (Is there anything else needed? Do we need to know how many participants there will be?). Projects will fall in one of 3 sizes: small, medium or large (Need to define sizes). Small projects will be approved in near real-time and will depend only upon the attributes of the project leader (What attributes? I imagine cannot be a student). Medium projects will be approved (or additional information requested) by the clearinghouse operator within 2 business days. Large projects must be approved by the GENI Oversight Group (Is this right, or a separate approval committee) who will work to get responses be to project leaders within 7 business days. (We need some sort of QOS promise here. Is this a decent benchmark?)

After the successful registration of a project, the project leader will receive a signed project credential allowing them, or their delegates, to register slices associated with the project. Of course, to register such slices, they must be created by a principle registered at the clearinghouse or another authorized identity provider. Further, the slice must be created at either the clearinghouse or another authorized slice registry. This is all to ensure that the creator of the slice is tied to a GENI principle at the clearinghouse or one of the trusted identity providers. Without this, a slice presented to the clearinghouse could be owned by someone the clearinghouse does not know and with a registry that keeps no records of slice creation.

Principal Registration & Revocation

Experimenters can create an account and obtain GENI credentials through the clearinghouse. Using an identity provider such as InCommon, they can authenticate to the clearinghouse. Additional attributes not provided through the IdP (we need to define what these are: contact info, home org, student status, etc?) will be supplied by the experimenter and verified manually by the clearinghouse operators. Because of the need to manually verify attributes, at least during early operations, the clearinghouse will work to get accounts approved and credentials issued within two business days (What is the appropriate QOS guarantee here? Hard to tell without knowing attributes collected).

Some IdPs will be able to mint the GENI credentials themselves, but they must support revocation and present new users with the AUP before issuing credentials. Furthermore, an IdP that creates GENI credentials for principals must log that transaction so that it can be determined when the AUP was presented to the user. The clearinghouse will also log each new account creation, presenting the AUP to users before allowing the application process to complete. It will record a pointer to the principal record, timestamp of issuance, and the remote IP address from which the account request came.

For all credentials issued at the clearinghouse, there must be a mechanism for revocation of compromised credentials. Furthermore, users may have their accounts revoked for violating the AUP. (Need process: who does this? GENI Governance Group? Is this or the AUP the place for defining that process) In either case, there must a be a prompt mechanism for revocation. At a minimum, the clearinghouse will maintain a Certificate Revocation List (CRL) that is referenced in any credential issued (More on this in the CA policy). Ideally, the clearinghouse would also support OCSP for real-time revocation information, though that would not be required of all IdPs. Only support of CRLs is required in the agreements with IdPs.

Slice Creation and Registration and Revocation

Any approved slice registry (whether at the clearinghouse or an approved IdP) can mint slice credentials for GENI if they agree to log the transactions, maintain a registry of all associated principles with the slice, and support revocation of the credentials. However, to be registered with the clearinghouse and associated with a project, the following steps must be followed:

1. **Slice registry verified:** The clearinghouse verifies that it is from a source trusted to mint slices. The clearinghouse will not know the list of authorized experimenters for the slice if it is not maintaining the registry that created the slice.
2. **GENI credential verified:** The clearinghouse must either have the principal registered, such as if it issued the credential, or must verify that it was issued by a trusted IdP who is collecting and verifying all the same attributes that the clearinghouse would. Such an IdP would also have agreed to present the user with the GENI AUP as well as to have logged the credential creation.
3. **Project verified:** The slice registration request must ask for registration with a valid project. Furthermore, the slice creator must either be the project leader or someone who can prove that they have been delegated the right to mint slices for that project.
4. **Slice information collected:** The creator describes the purpose of the slice and experiments on it and fills in contact information if they are not using a GENI principal created and registered at the clearinghouse.
5. **Slice credential is signed:** The clearinghouse signs the slice credentials in a way that binds it to the project as well.

This whole process is logged by the clearinghouse, including credential generation if that is done by the clearinghouse. The slice credential, associated identity credentials, project ID and any relevant delegation credentials are logged. Timestamps and IP addresses are also recorded, as well as the results of all the verification steps. Furthermore, the slice registry will keep track of all experimenters authorized to act on the slice (Is this even possible, or can the slice owner delegate without involving the slice registry?) and all resources allocated to the slice (They will depend upon this information coming back from the aggregates. Otherwise, they cannot even tell if the resource allocation policies are being met. This needs to be added to aggregate provider agreement). These logs are retained and protected as specified in the privacy and logging policies.

Whether due to a compromise or as part of the emergency shutdown process, it may be necessary to revoke a slice. The clearinghouse will maintain a CRL that allows it to deregister slices. If the clearinghouse issued the slice, it can also revoke the slice credential itself. Details on revocation mechanisms are deferred to the certificate authority policy.

Resource Discovery

The clearinghouse will provide experimenters with mechanisms to discovery components as a clearinghouse for all the resources available in the federation. This list maintained in the component registry is not tracking real-time availability of resources. The experimenter, perhaps with help from clearinghouse tools, will have to go directly to the aggregate to reserve resources for specific time-periods.

While the clearinghouse does not and cannot act as a scheduler for the aggregates, who may be sharing the same resources with other federations, it does try to keep track of resources that are down or offline. In the aggregate provider agreement, aggregate operators agree to notify the clearinghouse (needs update of language in agreement) of planned and unplanned outages. This information plus regular pings from the clearinghouse will help streamline resource discovery for experimenters.

Because the operation of resources is outside the control of the clearinghouse, they do not make any promises regarding the operational status of aggregates. Resource discovery is therefore a best effort service. All the clearinghouse guarantees is that those aggregates listed have agreed to the Aggregate Provider Agreement and have not violated the agreement in such a way that the GENI Oversight Group has revoked their federation membership. Operational problems with an aggregate should always first be brought up with the aggregate operator. GENI tickets created for problems such as a malfunctioning aggregate will be routed to the aggregate operator.

Federation Resource Allocation Policy Verification

Resource allocation policies may come from either the GENI Oversight Group through NSF or through special agreements with other identity providers, such as, FIRE. The former should be either listed in this document's policy section or referenced in this document as an external policy (We'll need to add such a section later). An example would be something like, "students are only allowed to have 10 hosts per slice that they create". The latter type of policy would be in agreements that should be publicly available on the clearinghouse web site. Examples would be a limit on slice size for FIRE users or a limit on bandwidth usage by GENI users for an international link with FIRE.

Because individual aggregates are autonomous and free to allocate resources to whomever they please, **this is policy verification and not enforcement**. The service provided by the clearinghouse is to give information to any aggregate as to whether or not a resource request for a slice would violate such federation policies. **This is necessarily a best effort service because the clearinghouse will rely on aggregates to inform it of resources allocated for slices registered with the clearinghouse**. Aggregate Authorities have agreed to provide this information as part of the Aggregate Provider Agreement.

Some aggregates, such as the GENI Racks, will only accept resource requests through the clearinghouse. They will then implicitly receive only compliant requests. Other aggregates will accept requests directly from the experimenters. An interface to the clearinghouse will be provided to them to check policy compliance of a request in real-time should they choose to use it. **A separate interface will be provided to them to inform the clearinghouse of allocation requests granted.**

Future Services

This clearinghouse policy is a living document. As new services and responsibilities are added to the role of GENI Clearinghouse, this document must be updated. There are several services that have been discussed, but not deemed critical to implement any time soon. These include experiment sunsetting and archiving, experiment template creation, and several deployment tools to make the creation of new experiments simpler. If the clearinghouse is ever expanded into a more general portal with these services, descriptions of the services and quality of service guarantees will need to be added here. Furthermore, privacy policies will need to account for any new types of data collected.

6. Policies

In this section we describe the policies that guide the operations of the clearinghouse. The service level promises are not in this section, but are instead discussed in section 5 along with the descriptions of each service.

Governance

The GENI Clearinghouse receives its direction from the GENI Oversight Group, established in the federation charter (Don't have one yet). The Clearinghouse Operator manages the Clearinghouse Operations Team to fulfill its mission as handed down from the oversight group and described in this document.

The clearinghouse does not have authority over any other party. Instead, agreements are established with the community it serves, and if agreements are violated, it may stop providing services to those parties. For example, the Aggregate Provider Agreement sets a standard of service that aggregates will provide to the GENI community. If they cannot maintain the security and integrity of their services, then they could be delisted in the clearinghouse. Similarly, an experimenter could use a slice maliciously and violate the AUP. In that case their GENI credential generated by the clearinghouse could be revoked. If another IdP issued the credential, they would need to revoke the credential, at least the GENI one. They could still hand out credentials to that user for their own test bed, for example. Again, this would only be enforceable through agreements: the IdP agreement with the clearinghouse in this case. (Of course it doesn't exist yet)

Responsibilities

The clearinghouse's responsibilities to the community (e.g., the services it will provide and the way it protects personal information) are described in this document. The responsibilities of the other actors (e.g. aggregate authorities, identity providers and experimenters) towards the clearinghouse are to be spelled out in agreements they sign. (or it will be when revised)

The parties that ultimately bear responsibility for the behavior of an experiment are the experimenter and project lead, as spelled-out in the AUP, project lead agreement, and Legal, Law Enforcement and Regulatory Plan (needs updating). The project lead is expected to delegate slice creation capabilities to trusted individuals with due care, and the project lead is expect to be aware of the different experiments and experimenters associated with their project. The experimenter who runs the actual experiment within a slice is most directly responsible for the behavior of a slice, particularly in regards to careless, disruptive or illegal activities of an experiment. Individual aggregates are not in the position to take responsibility for experiments as they are not given the information to do so. All that an aggregate operator can do is shut down slivers if informed by another, such as a law enforcement agency, that a slice is misbehaving. Through all of this, the GENI LLR Representative has the responsibility to route requests and put the different parties in contact to resolve problems but again is not responsible for the behaviors of the experiments.

Conflict Resolution Process

The clearinghouse could take actions against slices, experiments, project leaders and aggregates based upon violations of agreements or under the order of a court in rare circumstances. The clearinghouse may:

- Delist an aggregate's resources from the component registry;

- No longer trust an IdP for authentication and/or credential creation;
- Revoke a registered slice credential;
- Revoke a project credential; and
- Revoke a GENI principal credential for an experimenter or project lead.

These are the only tools available to the clearinghouse for enforcement of agreements. Directions to take such actions could come from the GENI Oversight Group, a law enforcement agent with proper authority and documentation; or the Clearinghouse Operator if they determine an agreement has been violated. However, there is a process that is followed to try and resolve conflicts before any such actions are taken, and there are mechanisms of appeal to any decision. The process will follow the following steps for any violation or complaint of a violation.

1. A complaint is made with the clearinghouse by a GENI actor or an LLR (Legal, Law Enforcement & Regulatory) agent.
2. The clearinghouse operator reviews the complaint and tries to gather as much information regarding the validity of the complaint. She may consult with a lawyer in the case of an LLR request, with the assistance of the GENI LLR Representative.
3. The clearinghouse operator informs the accused actor of the complaint and let's them respond. Unless there is some legal reason that they cannot inform the accused or must act more quickly than the accused can respond, they will wait for a response.
4. The clearinghouse operator, maybe with guidance from the GENI Oversight Group, will make a decision. If the complaint came with a directive to take action from the GENI Oversight Group directly, steps 2 and 3 may be skipped.
5. The response is informed to all parties involved and carried out. This whole process will be as open and transparent as legally possible in all situations.

Those who are accused of breaking an agreement and have had credentials or services revoked can always make an appeal to the GENI Oversight Group who will have the final say on any decisions regarding the above actions that could be taken by the clearinghouse. The following steps are taken in the case of an appeal.

1. The GENI Oversight Group receives an appeal.
2. The oversight group informs the clearinghouse. They may or may not request a stay on an action taken while the appeal is being reviewed.
3. The oversight group gathers as much information as possible from all relevant parties and makes a decision.
4. The decision is communicated to the party making the appeal and the clearinghouse, who will follow their decision.
5. If the appeal is not resolved in favor of the accused, they are given instructions on what they must do to resolve the problem and reverse the action. Except for serious legal issues, it is unlikely that anyone would be permanently removed from the federation. However, to rejoin they would need to address the issues that were considered breaches of their agreements.

Privacy Policy

The clearinghouse respects the privacy and rights of its users and recognizes the responsibility and trust put in it by the community. Therefore, every effort is made to protect that privacy.

The clearinghouse collects information at registration for attribution and generates logs of all transaction (e.g., signing slices, creating credentials, etc) performed while providing its services. These logs will be stored for at least 90 days. Clearinghouse operations uses this information only to provide for the security and operational stability of GENI and to determine the parties responsible for the actions of a given slice. Identifying information is never shared to those not providing GENI operational services, unless legally required, and never for anything beyond providing secure and highly available services. If identifying information is ever shared in such cases, effort is made to inform the identified party of exactly what was shared and why.

The clearinghouse follows the Legal, Law Enforcement and Regulatory (LLR) Plan, and may receive requests from an LLR agent either directly or referred from the GENI LLR Representative. **If the clearinghouse is asked who the owner of a slice is, it will first offer to point the requestor to the email alias created for each slice when it is first registered**. If the clearinghouse does not have full contact information on the user, it may redirect the request to the appropriate IdP who will follow their institutional policies on releasing such information. If ordered by a court or recommended by legal council for the institution operating the clearinghouse, they will reveal full contact info to the LLR agent making the request. However, they will also inform the user that they have done so unless legally prohibited from doing so.

Finally, the clearinghouse operations team will take due diligence and follow industry best practices to secure the registries and logs against attackers. They will follow hardening guidelines if developed by the GENI Computer Security and Incident Response Team (GENI-CSIRT). Furthermore, if these databases ever do get attacked and there is reason to believe identifying information is exposed, the affected users will be promptly notified.

User Attributes

Several attributes are collected during user registration. Some, but not all, of these are passed from InCommon. Trusted IdPs who have signed an agreement with the GENI federation will need to collect the same attributes to put into valid GENI credentials. The following list is the minimum set of attributes that will be collected so as resource allocation policies in GENI can be implemented. (This is problematic to determine without knowing the kinds of polices that we might get from NetSC)

- thing1
- thing2

Aggregate Provided Data

Because some aggregates will not require resource requests to come through the clearinghouse, the clearinghouse will not see all resource allocations to slices. Without knowing what resources are given to a slice, it would have no way to verify that federation level resource allocation policies from either the NetSC Council or as part of agreements with other test beds (e.g., FIRE) are actually being followed. Therefore, aggregates are required per the Aggregate Provider Agreement (needs updating here) to provide the clearinghouse with information regarding requests they have granted for GENI slices, i.e. those registered and signed by the clearinghouse.

In addition to resources allocated, aggregates will be providing the clearinghouse with updates on resources that malfunction or go offline (probably need to update the aggregate provider agreement). While the clearinghouse could regularly ping components, that is costly to do frequently for a large set of components. Therefore, it is more efficient to operate under the assumption that resources are operating normally, unless notified by an aggregate operator. Relatedly, aggregates are expected to inform the clearinghouse of components are permanently retired. (Of course, how this is technically achieved is outside of scope.)

Lastly, it is requested that aggregate operators provide a list of all publicly facing IPs potentially used by its slivers, if possible. By having a list of public IPs that are associated with GENI experiments, the LLR representative can respond to requests more efficiently and promptly. However, this is not required in the Aggregate Provider Agreement (do we want to change this?) and may not be technically feasible for aggregates to separate from other IP addresses at their institution.

Certificate Authority Policy

Since the clearinghouse is issuing certificates, signing others certificates and acting as a trusted root or bridge between entities, it is very much acting like a CA (certificate authority). Therefore, it needs a policy governing its operation (e.g., how keys are protected, how certificates are revoked, etc) and what is required to trust other CAs (e.g., level of assurance for vetting, encryption algorithms used, etc). We recommend that this is a separate policy document, though linked to here. Ideally, the federation would agree on a standard CA policy that would be followed for anyone issuing certificates or credentials (including the clearinghouse), and it would be a requirement to meet the baseline established in that policy to be a CA which is trusted in the federation.

Glossary

- **Aggregate:** is a system containing a collection of resources (i.e. components) under common administration running an aggregate manager service (defined in the GENI Software Framework Architecture).
- **Aggregate Authority (AA):** is responsible for the management of the aggregate, but can delegate selected functions to other actors. The aggregate authority is the one who can enter into agreements for the aggregate.
- **Aggregate Administrator:** is one who has been delegated the responsibility, by the aggregate authority, to set local resource allocation policy for an aggregate and its components.
- **Aggregate Operator:** is appointed by the Aggregate Authority to operate the Aggregate Manager and any components of the aggregate. This may be a the AA itself, or someone from another organization operating on its behalf.
- **Component:** encapsulates a collection of resources, including physical resources (e.g., CPU, memory, disk, bandwidth) logical resources (e.g., file descriptors, port numbers), and synthetic resources (e.g., packet forwarding fast paths).
- **Component Registry:** holds or points to a record for each affiliated substrate component or aggregate. Each record includes: the responsible organization (i.e., aggregate authority); associated principals (e.g., operators) and their individual permissions; the interface(s) to query for available resources (e.g., on the aggregate manager); other contact information; and (optionally) policy to be applied to use of this component or aggregate.
- **Clearinghouse:** is system consisting of software, operations, and policy to enable authoritative resource allocation in a GENI federation.
- **Clearinghouse Operations:** is the operations team responsible for the GENI clearinghouse. This could be a role of the GENI meta-operations group. The **Clearinghouse Operator**, is the one managing this team.
- **Experimenter:** is a user with credentials recognized by the clearinghouse to act upon a slice.
- **GENI-CSIRT:** Computer Security and Incident Response Team for GENI composed of people from several major stakeholder institutions and guided under the authority of the GENI Oversight Group.
- **GENI Project:** is a grouping of experimenters and slices working on a common effort. It may have multiple slices concurrently and over time.
- **GENI Oversight Group:** is the group responsible for ensuring that meta-operations and the clearinghouse operations groups fulfill their responsibilities. It is also the governance body for the GENI federation, responsible for guiding project direction and resolving disputes between other actors.
- **Identity Provider:** is system that issues credentials for principals. At a minimum they are responsible for user user enrollment, vetting and authentication. An identity provider may or may not go further and run a complete principle registry, or it may rely on the clearinghouse to implement a full registry with all the needed attributes.
- **Identity Provider Authority:** is one who can form agreements with other GENI actors on behalf of the identity provider.
- **Identity Provider Operator:** is one who has been delegated the responsibility for maintaining the operational readiness of that identity provider.
- **Meta-operations:** is the operational group responsible for cross-aggregate issues & experimenter support.
- **Project Leader:** is the actor who is ultimately responsible for the behavior of a GENI project.
- **Principle Registry:** holds or points to a record for each GENI-associated experimenter, project leader, operator, etc. Each principal record includes: a global name, contact information, authentication key (or a pointer to it in another trusted registry), roles, and status (active, suspended, etc.). Depending upon the agreements with an IdP, the registry may be split across the IdP and the clearinghouse.
- **Slice Authority:** is a system that issues credentials for slices. There are minimal requirements, (such as supporting revocation), needed of slice authorities, but slices are not trusted by GENI until the clearinghouse registers the slice and signs it. However, agreements with slice authorities could give them extra authority to act as a GENI slice registry and sign slices on behalf of the clearinghouse.
- **Slice Registry:** holds or points to a record for each slice, equivalent to a “bank account” for that slice in that it records transactions and authorized accesses. Each slice record includes: the responsible organization (e.g., the slice creator) and its permissions; the associated GENI project; associated principals (e.g., experimenters,) and their individual permissions; and the slice status (active, suspended, etc.).