Topic

- ☐ Aggregate Provider Agreement
  - ☐ Update System
    - ☐ Who is monitoring CVEs, etc? What part of the software stack do they take responsibility for (base OS included)?
    - ☐ What about the base system images (for VMs & bare metal) and not just the management systems/software?
    - ☐ Is there secure update system? Could you handle a zero-day?
  - ☐ How do you maintain integrity of logs?
    - ☐ Is there a default length to keep them?
    - ☐ What sort of system/security logs do you keep?
    - ☐ Do you use a remote syslog server or something similar?
  - ☐ Admin/Management interface(s)
    - ☐ Is there a remote access to an internet-enabled KVM for console access?
      Sounds like it is just physical access once I realized KVM is an overloaded term in the document.
      - ☐ If so how is that configured?
    - ☐ Does it use two factor authentication? Is a root account shared? Do you use root?
      - ☐ Is root or admin password same for each node within a rack? How about between racks?
    - ☐ LDAP is used for AuthN, what about AuthZ?
      - ☐ Is it two-factor?
      - ☐ Is root login allowed (hopefully would have to sudo or su)
    - ☐ If I VPN in what do I see?
      - ☐ All the hypervisor admin interfaces?
      - ☐ Any other management interfaces?
      - ☐ If I compromise one rack, does that get me access through the VPN to the private network of every other?
    - ☐ What can I access from the head node?
      - ☐ Sounds like everything, but this isn't true of the commodity nodes, right?
    - ☐ Are there multiple admin interfaces, say one that is only available to you and not the hosting org?
      Looks like just the VPN and head node, and what is available to hosting org is decided by an undetermined authZ policy implemented in LDAP
  - ☐ Are any user credentials stored anywhere, even temporarily? If so how are they protected and how long do they live?
    Sounds like you argue it is not applicable.
  - ☐ High-level: How are slices isolated? Are there weak points in that isolation?
    - ☐ Sounds like hypervisor isolation for hosts, but what about on the LAN? Is it VLAN isolation between nodes on other slices?
- ☐ LLR Agreement
  - ☐ Does the GENI rack bridge to opt-in users, like WiMAX handsets?
    Sounds N/A, that would only happen on a per experiment basis.

Topic

- ☐ If so, can you still map an incident report about a handset to a slice if the opt-in user does something "bad"?
- ☐ You are not, like MNG, installing software on opt-in user systems, right?

  Nope
- ☐ Can you tell which slices are running on each rack? How about each node on a rack?
- ☐ Can you map timestamp & IP uniquely to a slice? How quickly?

  Sounds like uniqueness maybe hard depending on option used for networking (A,B, or C)
- ☐ Would a university deploying a rack be able to give a clear set of IP addresses that are used exclusively for their rack?

  Sounds feasible

- ☐ Clearinghouse Policy
  - ☐ Is there a way to provide information on resource allocations granted to a slice back to the CH for policy verification?
- ☐ Generic
  - ☐ What is the vetting process for bare metal issues?
    - ☐ Is the bare metal host diskless? Curious about isolation between allocations or experiments in time, not space here.
  - ☐ What sort of connectivity is needed for call back? It is best if admin interfaces can be on tightly controlled networks and not exposed at the sites.

    Looks like head node SSH access and VPN on the back?
  - ☐ Is this Exchange Aggregate at Starlight a special device, or just another GENI rack that is just connected to Starlight?
  - ☐ What is meant by " Since ExoGENI slices have management network access via the commodity Internet, this is the default behavior." on pg 13.