Questions for PIs to answer:

A list of general and project area specific questions that will be sent to the PI's 2 weeks before the PI meeting. These questions will be asked during the panel session. In addition the audience can ask questions during the panel session.

Areas:
Security
Distributed Computing
Network Management / Monitoring
Data Transfer
Workflows

General:

1. What is your biggest 'failure', something your research showed it was not promising to pursue?
2. What steps/actions would be required to allow you to leverage/build upon another teams work?
3. What steps/actions would be required to transfer your research results to the DOE facilities and/or science communities?
4. If your project is successful, what will the impact be 5-10 years from now?
5. Research breakthroughs often begin with an unexpected or unanticipated outcome. Can you identify such an outcome from your project and how it has changed your direction?
6. Has your project collaborated or interacted with other DOE (non-NGNS) initiatives. If so, how?
7. From your experience how important is latency to data movement and knowledge discovery on distributed machines?
8. Exascale data may require In-situ processing where data is analyzed, processed and visualized before leaving memory. On the other hand, exascale data processing may utilizes a large distributed computing environment to process data in parallel. what is your view about in-situ v.s. data movement? How can they be integrated into the same computing paradigm?
9. How do we strike a balance between research publications and technology transfer to science communities?

Workflows:

1. What is the role of workflow technologies in exascale systems?
2. How do you see the interplay between resource management and workflow management?
3. What infrastructure services would you like to see developed to support workflow technologies?

Distributed Computing

1. Distributed Resource Utilization: Given the ability to federate and interoperate across heterogeneous resources, how does your project solve/struggle/plan to address the ability to utilize multiple resources for general purpose workloads?
2. "To Distribute or not to Distribute" is the question. When does it make sense to federate resources? When not?
3. How does your distributed computing project fit within (or advance) the DOE exascale landscape?

Security:

1. What are the challenges in integrating Federated IAM with on-demand sharing of Science DMZ resources (e.g., OpenFlow switches, Data transfer nodes) and their orchestration? Do these challenges have similarities with compute and storage orchestration with frameworks such as OpenStack?
2. What are the emerging protocols that a domain can adopt for integration with external/distributed services (non-browser based) provided by other domains in a federated environment? What are the barriers of adoption of these new protocols?
3. Network system resiliency? Intrusion/system failure can affect network and user experience? Do we have mechanism to do early detection and fault mitigation

Data transfer:

1. What makes you think the tools and services now available will work when networks reach terabit/sec speeds?
2. What, if anything, do you need from the network monitoring community?

Network Management / Monitoring:

1. What is preventing a broader adoption of perfSONAR?
2. What are the 3 top tools you use to manage and/or debug the network?
3. What steps/actions need to take place to better integrate security research results into management tools?