

Quality of proposed research	Qualifications of research team	Comments for proposal PI:	Value of the ANI testbed	Impact of the research	
Good	Good		High	Medium	Accept
Fair	Excellent		Medium	Medium	Accept
		<p>This is an extremely good topic to explore as it is central to the network drivers for DoE: Large data distribution. But it is important also well beyond the ESnet or DoE community. As the larger research and e-science community acquire 100G networks, the topics wrapped around Lustre, HDFS, dCache, and other similar services and systems will be one of the key ROI measures for investing in the new network technology. Do this work well! Be critical and learn not just how these packages work with a 100G core, but what are their architectural deficiencies and limitations that could be corrected/re-designed to make data flow largely transparent to these high end users.</p> <p>As for the proposal – This reviewer sees little objective justification in the proposal to say whether this team is adequately qualified or will be sufficiently engaged to accomplish this research. One grad student for a summer is not adequate staffing level, and not likely adequate FTE's to accomplish the stated work. Fix this. This glaring administrative mistake aside, this is a good proposal. This reviewer has confidence just from the source organization of the proposal that the team is probably suitably qualified. But not all reviewers of this project or the testbed will have that insight or opinion. You need to provide real background information on the team. “they are leaders” is just your opinion...provide specific objective information about what these folks have done over the last five or ten years that makes them qualified to lead or even participate in this project. Let the reviewers decide if “they are leaders” or not. The US citizen is investing a huge sum of money into this testbed at a time when that money could arguably be better spent elsewhere...be sure your team is above reproach in terms of the level of talent DoE is investing to leverage that ANI testbed investment.</p>			
Excellent	Don't know		High	High	Accept
Good	Good		High	Medium	Accept

Good	Good	<p>Much effort has been devoted in the last 15 years to validate the “grid” model on distributed storage and computing. The proposal should bring new technical results to assess the feasibility and “business models” of the various scenarios of interest to the users (from fully distributed to a set of interacting clouds) at 100Gbps for the different layer.</p> <p>It is suggested that it also consider the possibility that the circuits are shared with traffic</p>	High	Medium	Accept
Good	Excellent	<p>While 100 G is not currently available in ANI testbed, high-energy physics applications in ANI testbed are very important. I recommend more concrete plans for research prior to the availability of 100 G.</p>	High	High	Accept
Excellent	Excellent		High	High	Accept
Excellent	Good		High	High	Accept
Good	Fair	<p>What can be done with the phase 1 ANI-testbed is not clear. The proposed experiments seem to require at least the phase 2 configuration. Using VMs on application hosts to run application may result in unstable performance.</p>	Medium	Medium	Accept
Fair	Don't know		High	Medium	Accept