Making GENI REAL: The Response Engineering Application Laboratory

Rick McGeer, Glenn Ricart, Andy Bavier, Joe Kochan, Joe Mambretti

> *Dream no little dreams; they have no power to move the hearts of men*
> *-- Goethe*

Abstract

In this whitepaper, we argue that GENI is not simply a testbed for networking and distributed systems. Rather, it is the prototype for a critical infrastructure of the near future – the Distributed, interoperable, ubiquitous Cloud, which reacts to users and events instantly: the *InstaCloud*. This is not simply a bigger distributed infrastructure; it is an infrastructure with points-of-presence everywhere, which enables a unique class of applications and services. Only this pervasive infrastructure is as fast and facile enough to match the real world around us. We argue that the National Science Foundation should adapt the existing GENI base into a platform to investigate the implementation and properties of the InstaCloud, and the applications and services uniquely enabled by it. We believe that this should be done in the context of a Major Research Infrastructure (MRI) project, which we call the Response Engineering Application Laboratory, or REAL.

Introduction: The Triumph of GENI

We argue that GENI should become a **major operational and research infrastructure for American academics**. We believe that the characteristics of the GENI architecture and deployment make it uniquely suited to form the backbone of a new form of Cloud required to support the services of the era of ubiquitous Big Data and the Internet of Things. These characteristics include: Points-of-Presence across the United States; Layer-2 Programmable connectivity between sites; a demonstrated ability to federate with similar infrastructures across the US; an overlying and standardized API on multiple Cloud infrastructures; inherent federated identity management, demonstrated across every University in the United States and the Canadian SAVI infrastructure; a demonstrated ability to support overlying and embedded infrastructures; inherent resilience against local failures of storage, network, and computation; ability to embed/support/host resources of varying types; demonstrated ability to federate with other, similar infrastructures worldwide, specifically the European Union's Federation for Future Internet Research (FED4FIRE) and Canada's Smart Applications on Virtual Infrastructure (SAVI)

Taken together, these advantages offer a capability unique in computational infrastructure – the ability to construct virtual intranets across the wide area, with admission control, quality of service, and a point of presence near a user, data source, or actuator anywhere – today within the United States, and at the end of this program around the world. Of particular note is the ability to move computation to users and to data, and the ability to rapidly react to users, signals from data sensors, or changes in network connectivity.

Big Data, Small Devices, and Smart Cities

We have entered the era of Big Data. A combination of dramatically falling storage costs and the proliferation of high-bandwidth sensors such as cheap, high-capacity cameras has brought us to the Zettascale era: there are now over 10^{21} bytes on the world's disks, and the number doubles every 24 months. There has been a great deal of attention paid to the requirements for processing this data, but transmission and storage deserve at least as much attention. The doubling rate implies a worldwide data capture rate of approximately 10^{14} bits/second. Most of that data must remain at or near the device which captured it: ingress/egress bandwidth at most sites is orders of magnitude too low to capture this. The National Science Foundation just announced the Pacific Research Partnership, an initiative to spur big-data science at a number of University of California and other California campuses: a 100 Gigabit/second network. This network can only transmit about 0.1% of the data added to the world's disks every second, a ratio which will get smaller over time. A yottabyte will inhabit the world's disks by 2020, and the capture rate will be 10^{17} bits/second. This will far outstrip the capacity of any conceivable wide area network: inevitably, most computation on most data will be in-situ or near in-situ.

A second trend is towards ever-smaller, low-powered, and more mobile devices for personal computation and consumption of information. These devices force storage and computing into the Cloud; but high latency and low bandwidth across the wide area means either a Cloud with a point of presence near the user, or non-interactive applications with greatly-reduced data going to the device. By "near the user" we mean latencies in the range of 5 ms with user/Cloud bandwidth on the order of 1-10 Gb/s. A tight constraint on latency implies many POPs. 5 ms in fiber is 1,000 km, the distance between Seattle and San Francisco; in practice, assuring these latencies requires a much smaller distance. This means that to support truly interactive Big Data on small devices, a Cloud with at least 25 POPs across the United States is required.

A second area of significant national importance is the emergence of the Internet of Things, and in particular Smart Cities. Fundamentally these phrases mean the use of networked information and computer technology to make devices and human systems more efficient. In practice it means networks of sensors controlling actuators on a scale from the individual device to the city. Some of these sensors are quite high-bandwidth (consider traffic cameras or the weather radars in the Cooperative Atmospheric Sensing Apparatus program) and so the observations about the requirement for in-situ storage and processing given above apply here as well. Near-real-time reaction is a requirement for many Smart Cities/IoT applications, which again imply tight limits on latency between sensor, computing agent, and actuator. Bandwidth is as important as latency. The falloff in bandwidth between the local and the wide area is dramatic. Even low-end switches on a local area network have bisection bandwidth of at least 50 Gb/s, and per-node capacity of 1- 10 Gb/s. Standard border routers are generally sized for 1-10 Mb/s/person. Moving IT services to the Cloud therefore represents a bandwidth reduction between service and user of about 1000x. Simple calculations demonstrate that almost all network traffic originating in a city must remain in that city: Chattanooga, TN, has 80,000 homes and offers gigabit connectivity. This gives total bandwidth of about 80 Tb/s; the egress links from any city in the US can handle only a small fraction of that. We must have a distributed Cloud, simply because aggregate edge bandwidth far outstrips the bandwidth available in the core, and because the applications and services of the future cannot tolerate the delays of long-distance network transit. *In the future nearly all network communication will be local.*

In sum, the trend to small devices and massive data is driving computation to the Cloud, as is the computing needs of Smart Cities/Internet of Things. Both real-time control and interactive, rich applications enforce tight constraints on both latency and bandwidth. To get tight response from a Cloud we need a Cloud with a new architecture: the *Instantaneous Cloud*, or *InstaCloud.*
The InstaCloud, the architecture for research for the next decade, is not just a bigger distributed infrastructure. It is an Applications Research Infrastructure as fast and facile as the world around us, which can keep up with the demands of SmartCities and offer rich network applications with desktop speed. We have already shown an exemplar of such an application: the Ignite Distributed Collaborative Scientific System. The Ignite Visualization System provides seamless interaction and immediate updates even under heavy load for widely-separated users. It can fetch a data set consisting of 30,000 points from a server and render it within 150 milliseconds, for a user anywhere in the world, and reflect changes made by a user in one location to all other users within a bound

provided by network latency. The system was demonstrated successfully on a wide variety of clients, including laptop, tablet, and smartphone, with a session featuring six users in five countries on three continents.This example application is the first representative of the class enabled by the InstaCloud. The InstaCloud is characterized by:

- A Cloud that digests and interprets vast amounts of both historical and real-time data and which is designed to provide apparently instantaneous reactions with the accuracy of being big-data-driven, and renders applications on any device from workstation to tablet.

- Applications that can anticipate human and smart thing needs and predict the consequences in both the real and the --sometimes-indistinguishable --cyber world. Imagine a SmartCity emulator which ingests the take from an array of sensors and plugs it into a continuously running emulation of the city, updating scenarios and offering warnings and guidance for possible scenarios. The Cloud that manages this application must combine low latencies, high bandwidth, and significant computational and storage power.

- Applications that can help manage and advise us as we humans navigate a complex information-cyber-real-world. The defining characteristic of the applications of the 2010's and 2020's will be that the distinction between these worlds is blurred, and in that merging of the worlds lie the critical research challenges facing our community.

This InstaCloud best suited to support applications research is the same infrastructure needed for Research and Development on the industrial internet; it is the same infrastructure needed for R&D on Smart Cities, urban sciences and cyberphysical systems. It is an applications research infrastructure that adopts the scalability of PlanetLab, the sliceability and deep programmability of GENI, and adds low, deterministic response times to run in sync with the real world. We need it to manage intersections with autonomous vehicles traveling at full speed; to offer real-time visualizations on a worldwide basis of high-bandwidth sensor takes, and to integrate on-the-ground sensors with offline data; to do in-situ discovery from advanced scientific instruments; to quickly respond to national emergencies; to host rich, immersive online education apps such as the Mars Game. This InstaCloud is the infrastructure for the 2020's, and the mission of the research community is to explore its implications today. We call the prototype of the InstaCloud that we will build now the Response Engineered Application Laboratory – or REAL.
The characteristics of REAL are the defining characteristics of the instantaneous, responsive InstaCloud. Instantaneous response and high bandwidth from the cloud to the user is enabled by points-of-presence at the city level. Moreover, networking becomes simply another allocatable computing resource, as customized to the needs of the application as storage or compute, with per-application quality-of-service, routing, and admission control enforced at the networking level.

GENI As The Infrastructure

In the preceding two sections, we reviewed the characteristics of the GENI infrastructure and of the Cloud required to support the REAL program. GENI has exactly the characteristics needed for REAL. It has 50 POPs, and with the addition of the GENI-architecture-based US IGNITE city network this will grow to 70; the software base is stable and well-tested, and offers the ability to instantiate embedded cloud infrastructures and to include cyberphysical devices; it has sliceable, deeply programmable networking; and it has federated identity.

The GENI infrastructure has already been used to support a number of services that are characteristic of those we expect to be developed under the REAL program. These include many of the demonstrations shown at the Smart Future Summit in Washington, DC in March 2015:

● Mars Rover Game in the Mathematics Classroom, which showed how the InstaCloud could support rich applications on arbitrary devices in support of STEM education in the classroom

● Connected Vehicles Assisting First Responders. This demonstrator showed how connected automobiles and UAVs could rapidly pinpoint and survey accident scenes and transmit information to first responders en route. Again, Cloud connectivity and and programmable networking was vital to rapidly route the information from the vehicle to the first responder, and both bandwidth and latency considerations required that the routing, computation, and processing all be done within the same city as the first responders and the vehicles

● Picturing Pollution Globally and Locally. The Ignite Visualization system permits real-time interaction and visual collaboration around large data sets, and to permit users worldwide to seamlessly interact in real time with the data set and with each other, even when separated across the wide area, using any device from smartphone to workstation. The use of any device implies that the data must be resident in the Cloud; interactivity requires that a Cloud node must be within a few milliseconds to an end user; supporting an end user anywhere means that the Cloud must be *everywhere*. This was shown with nodes (and users) in Tokyo, Victoria, Canada, San Francisco, CA, Washington, DC, Brussels, Belgium and Potsdam, Germany, utilizing nodes from JGN-X, InstaGENI, ExoGENI, FED4FIRE, and SAVI: a true InstaCloud prototype.

These proof points demonstrate that GENI has grown far beyond its initial conception as a testbed for networking and distributed systems researchers, to support the critical applications and services that characterize the InstaCloud. REAL will take these applications and services to the next level, and provide the services phase of infrastructure development.

Phases of Infrastructure Development

Computational infrastructure development typically follows a four-phase path, and we argue that the GENI infrastructure is entering the second phase now. Here, we detail the four phases.

Phase 1: Experimentally-supported research on Basic Infrastructure

This is the earliest phase of infrastructure development. In the development of the Internet, the ARPANet represents this phase. At the end of this phase the basic architecture of the infrastructure and reference implementations are complete. We are just completing this phase of development of the InstaCloud, and the immediate first phase was GENI. The equivalent of the TCP/IP stack is programmable networking and sliceability; the equivalent of Berkeley Unix running on Vaxen is ProtoGENI running on InstaGENI.

Phase 2: Experimentally-supported research on services and applications

This phase takes the basic architecture as a given, and focuses on the services and applications enabled by the new infrastructure. In the case of the Internet, this phase is represented by the original ARPANet, and, later, by NSFNet. The value of the infrastructure is not yet sufficiently established that it can become self-supporting: the value of the services and applications are not yet quantified, and the long-term costs of maintaining the infrastructure on a per-site basis are not yet certain. In this phase, therefore, external support is still required, and the focus of inquiry is the establishment of those services and applications that establish the value and importance of the infrastructure to the community and the nation.

The ARPANet/NSFNet period brought us Telnet, IRC, FTP, sendmail, ARCHIE, Gopher, and the Web. This period is characterized by putting the tools developed in phase (1) in the hands of the wider community. In the case of the Internet, this was represented by the HEPNet, the Magnetic Fusion Energy Net, NASA’s SPAN, and CSNet for the Computer Science community.

In this case, this phase is represented by REAL. We propose to put the tools GENI developed in the hands of the communities that can write services and applications, and simultaneously develop a strategy for a self-sustaining InstaCloud infrastructure, much as Steve Wolff developed a strategy to maintain the Internet infrastructure on a basis independent of direct federal funding. This sustainable model involves establishing: applications and services of value; a cost model for sites going forward, both for support for the global institution and local costs for support and maintenance; and a model that supports institutional practices.

Phase 3: Membership-supported operations

In this phase, members support the infrastructure directly and development of applications and services, and organic growth as institutions access the services. This phase is represented by the Internet Society period immediately following the end of NSF Net. The mechanism of support depends on the details of footprint, deployment, services, and institutional payment preferences.

Phase 4: Commercialization

There will be commercial infrastructures supporting these standards. This can happen much faster than it did for the Internet, since the necessary legal structures are already in place.

From GENI to REAL

REAL is an MRI on the services of the instantaneous Cloud. A specific recommended focus is Smart Cities and Campuses. These themes draw our attention because they illuminate many of the infrastructure issues the InstaCloud will face; because they encapsulate many of the potential applications we see of great social importance; because these draw the attention of researchers from many disciplines across the ambit of NSF; and because these services have the greatest prospect for showing value to the campuses.

REAL will facilitate research guided by applications drawn from the domain sciences. The multi-domain needs of Smart Cities forms an excellent framework to explore the IT and interdisciplinary research enabled by the InstaCloud. A Smart City's environment involves researchers from geosciences and civil engineering; citizen needs for environmental health and safety involves medical, biological, and environmental researchers. Citizen education and training involves education, sociology and economics. Transportation systems involve engineering; social connections in smart cities involve the social and behavioral sciences, modeling and predicting city behavior involves the mathematical and physical sciences as well as economics. All of these, of course, involve computer science and the mathematical sciences.

A university campus is a small city, and therefore provides a wealth of opportunities for collaborative research between cities and campuses. Smart Campus research, utilizing REAL, offers services of significant intrinsic value to the host campus.

The testbeds for SmartCampus/SmartCities are already in place for the REAL MRI – the 50 campuses with GENI Racks and the up to 20 cities which will get US Ignite Racks. This infrastructure, already in place and already largely federated, forms an excellent initial testbed for REAL.

We propose a three-phase transition plan to go from GENI to REAL. In phase (a), REAL adopts the GENI and US Ignite testbeds and transitions them to REAL under the partnership of the GPO and US Ignite as the change agents. Separately, ad-hoc groups from the communities work to propose and adopt grand challenge problems in the Smart Cities area. The grand challenges will be arrived at in a series of regional workshops in the first half of 2016, modeled on the NSF Big Data Charrettes. Since an integral part of this is adoption by the campuses and cities, CIO representatives from GENI campuses and US Ignite communities, as well as leading researchers from the grand challenge areas, will be recruited to lead the drafting of the grand challenges. In phase (b), we move to community governance. The primary roles for the governance organization should be to work with the NSF to devise funding mechanisms for the Smart Cities/Campuses Grand Challenges; to devise and enforce an AUP for the REAL InstaCloud; and to oversee the REAL Central Office. In phase (c), an emphasis should be on the transition to a community-supported InstaCloud, including commercial applications and infrastructure. For the InstaCloud to become what it can be, a self-sustaining model is a must; and sustainable services require the support of commercial enterprises. US Ignite is planning the incorporation of commercial applications and infrastructure, and we believe that REAL should do the same.

REAL Organization: REAL Central

In the medium term, REAL will be too large for a centralized Operations Center (OC). However, GENI today is about 50 sites; adding in the 20 anticipated under the US Ignite buildout gives us a projected scale of 60-70 sites by mid-2016 for REAL. Even anticipating an annual growth rate of 100%, at the end of three years we will only be at 500 sites in a research infrastructure. This will be a research, not an operational infrastructure: we require 3 rather 5 9's of reliability. This has been achieved by PlanetLab.

PlanetLab is an excellent model for REAL over the next 3-5 years: the scales and requirements are quite similar. PlanetLab is 1300 nodes at 600 sites; REAL will have about as many sites and 3-5x PlanetLab's node count. PlanetLab was able to maintain the PlanetLab NOC and develop the code base with a team of 2-3 staff and a handful of grad students, post-docs, and researchers, all doing other things as well. The REAL OC can be managed by an academic group or a small team at a nonprofit. The OC should be modeled on PlanetLab Central rather than the GENI Meta-Operations Center. The OC operators should understand instinctively the needs of the community. This does not preclude a campus IT organization.

Appendix: The SmartCity as a Set of Services Deployed Across the Instantaneous Cloud

There are many challenges and research questions entailed in SmartCities/Internet of Things: this is precisely why, in this white paper, we've argued that this should be the focus of the proposed REAL MRI. It's an area that offers interesting problems for most areas of Computer Science, and puts Computer Science at the heart of a vast number of interdisciplinary projects spanning much of the reach of NSF. Moreover, it's an area of enormous social and political importance. Calit2 Director Larry Smarr is very outspoken on the dangers of global warming, and his take on a solution is very interesting: "Only software can save us": only the efficiencies and responsiveness of the Internet of Things and Smart Campuses and Smart Cities can yield the GHG reductions required without great damage to human society.

One challenge that we can consider now, because it forms a foundational component, is this: what is the software architecture of a SmartCity? This shapes our deployment and emulation options going forward, and thus offers a framework for the easy and reliable emulation, simulation, and deployment of SmartCities technologies.

Consider the fundamental requirements. In any real SmartCities service deployment, we'll need to first simulate, and then emulate, the deployment on a Cloud-based model, which must be as realistic as possible; we'll need to shape the network and more generally the cyber environment, both in the simulation/emulation and in the actual service deployment, to protect the sensors, actuators, and software of the SmartCities service from malfeasance and the operation of other services, and to protect other services from the operation of this service. This means that the operation of the service must be very tightly controlled, on a per-service basis. Moreover, as indicated above, an obvious optimization is to run the model service slightly ahead of the deployed service, examining real-time what-if scenarios and catching problems in the model before they happen in the field. SmartCities deployments must be composable: the take from traffic cameras must be composed with the take from parking sensors, real-time data from public transit, the number of Uber drivers on the road and available, and projections from anticipated demand for various events to come up with a complete picture of current and projected traffic in the city. SmartCities services must scale easily and reliably in the field, without surprise. The deployment will involve VMs and containers on servers located in datacenters, in PoPs in the network core, and at the network edge.

The deployment must fully control placement of computing agents and virtual network topology, use the full network address space, and deploy the network OS of choice to control packet forwarding between computing agents. The SmartCities service entails network-aware services that can position functionality at the optimal point in the network. Finally, deployment must be reliable, resistant to human error, and ideally doable by non-experts.

These requirements imply many things, but one thing they do require is a simple organizing principle: services must be reliable, modeled, composable, and easily and reliably deployed. There do exist computer systems areas with these properties, notably telco Network Functions Virtualization (NFV). NFV entails telco network operators migrating from purpose-built hardware appliances to virtualized commodity servers and software-controlled switches to offer telco services and perform internal network functions. Though in a very different arena, it has the same fundamental properties of SmartCities deployments – network-aware hardware/software, tightly-controlled services deployments.

Appendix: Applications of REAL

For NSF interest, we will focus on application areas of critical interest to CISE and ACI. We identified a number of these areas, as follows:

1. Resilience. Distributed Applications are robust against local failures and provide failover, particularly for disaster relief. This is an area of extreme interest to NSF/OSTP.

2. As an IT infrastructure for academic/city collaboration. This is an area of interest to NSF and OSTP: the big idea is that local universities and academics should collaborate around research areas. this implies a Cloud, so that per-project virtual infrastructures can be rapidly spun up. The advantage of GENI is that the management and software infrastructure of this network of local clouds is externally-maintained, relieving the administration burden on local city and campus CIO's. This is one of the major advantages a commercial Cloud provider offers, but it does so at the expense of bandwidth charges and limitations and unnecessarily high latency. GENI offers the advantage of offloaded administration while reducing latency and bandwidth costs and burden. Telemedicine is an excellent example of this.

3. As a way to offload specialized ACI infrastructure. Currently, ACI maintains a network of specialized supercomputing centers tailored for advanced scientific simulations; however, these are often used for more pedestrian tasks, simply because they are the available IT infrastructure. For example, the supercomputer at LBL (need to look up the name) was used for a large map/reduce computation on the CMIP project, after a three-month data marshalling effort over ESNET and other high-speed national and international networks. this could have been done as efficiently on a far-less specialized resource, or collection of distributed resources.

4. As an alternative to data movement. Much scientific computation combines a large data set search with an analysis job on the results of the search. In (3) we discussed the CMIP experiment run by LBL and ESNET. Astronomers hunt through databases of star and galaxy images for particular phenomena (e.g., standard candles such as Supernovae or Cepheid Variables, or galaxies of specific classifications); physicists for collision events with particular characteristics; geneticists through genome databases. The ability to conduct these searches where the data is collected or lives would be of significant benefit to all of these communities, and is a complement to (3).

5. As a platform for distributed collaboration. This was demonstrated by the Ignite Distributed Collaborative Visualization System at the Future Internet Summit in March 2015. It was demonstrated there that this required a GENI-like infrastructure (in other words, a server within 20 ms of any participant), in order for the participant to be able to have a true interactive experience. This was demonstrated as a Pollution Visualizer, but we believe that virtually every scientific community could use such a tool, as demonstrated by the broad usage of the OptIPortal and OptIPuter. Indeed, the Visualizer demonstrates the possibility of "a handheld OptIPortal", at 2-3 orders of magnitude cost reduction over the OptIPortal. As evidence, the Ignite Visualizer is being explored for use in genomics at the University of Victoria

6. As a platform for the creation of wide-area project-specific Virtual Intranets with tight admission control guaranteed by slicing. This was proposed to the US Government as a key use case for SDN

7. As a framework for the interoperation of NSF's next-generation Cloud projects, CloudLab and Chameleon. Example: we can knit together Chameleon, CloudLab, and GENI today using the GENI Experiment Engine. Other embeddable infrastructures such as OpenCloud can do the same thing.

8. As a platform and framework for international collaboration, both in the domain and computer sciences. We have already begun integration of SAVI in Canada with GENI, and Fed4FIRE in Europe with GENI. We can extend this to Japan's VNode project, and others as the AM API becomes more widespread

9. As a platform for the continuous modeling and monitoring of Smart Cities and Internet Of Things deployment. Smart Cities and IoT have been described as a "hacker's paradise", since malware and cyberattacks -- and simple bugs -- can have significant consequences. Both pre-deployment emulation and continuous during-operation simulation aided by continual feedback from monitoring of the deployment will have significant impact.

10. As a rapid, secure, resilient distribution platform for IoT software updates. Every company that has to do massive software distribution and updates does this over a CDN, generally a purpose-built one. Steam, a video game cloud distributor, has 243 sites worldwide.

11. As a growable infrastructure for second-life ACI infrastructures. See example of FOCUS in New Mexico (more details from Rob or Brian)

12. As an efficiency platform/shared cloud for operational cyberinfrastructure across campusses -- Jim Bottum's "condo of condos" (http://condo-of-condos.org/).

13. As a new platform for scientific publication of reproducible results -- not simply the results, but the VM/container with the operating instructions that anyone can use to play with the data, modify or experiment with it. This has been anticipated by Jay Lepreau's last NSF proposal, LabWiki at U Mass, and APTLab at Utah.

14. As a platform for real-time online education. Note particularly the Mars Game demonstrated at the Future Internet Summit. Note, again, that an interactive application on a small device requires a Server Near You.

15. As a platform for network and distributed systems research.

Appendix: REAL Cost Model

In the transition, we believe that the REAL Central Office will not only perform its long run-duties but may also perform the functions of equipment purchaser, and will perform the functions of network operator.

We anticipate that the standing costs of operating the REAL/GENI infrastructure will be a centralized cost of two FTEs for monitoring and front-line support, and an administrative staff of four FTEs: a Director, a Program Manager responsible for purchasing and contracting, an outreach coordinator, and one administrative support person. The figure of two FTEs is derived from the experience of maintaining Planetlab for a decade. While the PlanetLab Central staff was much larger than that, PlanetLab Central also maintained and developed the PlanetLab code base. We anticipate that the InterCloud Code Base(s) will be developed and maintained under contract, by groups such as Flux Lab at the University of Utah, OpenCloud, and ExoGENI, and the InterCloud Central monitors will function as just that – a small operations center focussed on discovering and reporting problems early. We should note that PlanetLab Central was able to maintain a network of over 1500 nodes at 300 sites for over a decade, with a staff of about that size.

We anticipate that the on-campus (that is, above the costs of REAL Central) labor costs of running a REAL site will be, all in, 20% of an FTE, or 10 FTE-weeks/annum. This assumes three one-week conferences for training, standardization, and clearinghouse activities, and roughly 6 hours/week in the remaining part of the year to respond to node-reboot requests, network and hardware troubleshooting, and installation of software updates.

Capital costs, and in particular minimizing capital expenditure, is one reason to use a central purchasing organization. The workhorse of the GENI network is the InstaGENI project.
InstaGENI was purchased for about $24,000/rack in the initial build, a price about 60% off standard academic pricing for those units. In part this was because of the desire of the InstaGENI vendor to be seen as participating in the leading edge of computer systems research, and the strenuous efforts of the InstaGENI team; but this pricing is unlikely to be replicated in the future. The vendor has told us that replacing an InstaGENI rack today would cost about $60,000.

Volume discounts are available, which would reduce the per-rack cost by 1/3, depending upon the volume being purchased. But this would require a single purchaser, who would distribute the racks to the various sites. For InstaGENI, StarLight/Northwestern played this role; for PlanetLab, Princeton/PlanetLab central did. In fact this is a well-used model. Candidates to be the purchasing agent are:

- REAL Central
- StarLight/Northwestern
- The Quilt
- Internet2
- US Ignite

We anticipate a total hardware budget, accounting for one full refresh cycle over the course of this phase, of roughly $2.8 million. Assuming a three-year phase this has a hardware budget of $1 million/annum.

The networking would be provided where possible by the state and regional networks, or where more appropriate, Internet2. These costs are difficult to estimate, as they are dependent on existing network connections, availability of RENs, etc. However, a worst-case assumption of a new Internet2 Layer-2 port at each site yields a figure of well under $3 million/annum

In sum, the fixed costs of this future are approximately $1.5 million/annum for the REAL Cental Office; labor cost over all sites of roughly $3 million/annum; and hardware costs, all sites, of $1 million/annum, for a total of $5.5 million/annum. Networking costs take this to *at most* $8.3 million/annum. These are very rough order-of-magnitude costs, and are merely given to show the rough total costs over all participants.

It should be noted that this rough calculation includes *all* costs of REAL. In practice, many of these would be absorbed as continuing expenses by the campuses. We expect that the actual direct costs to NSF, even in the very early, heavily-subsidized stages, would include at most REAL Central and some portion of the hardware refresh costs, for a maximum of $3 million/annum.