

Report on the Workshop on GENI and Security

Matt Bishop
Department of Computer Science
University of California at Davis
Davis, CA 95616-8562
bishop@cs.ucdavis.edu

August 15, 2009

This workshop was funded by the National Science Foundation's award number CNS-0646965 to the University of California at Davis. All opinions expressed in this report, and the presentations, are not necessarily those of the National Science Foundation, the University of California at Davis, or the institutions of the authors and presenters.

Abstract

The Workshop on GENI and Security, held at the University of California at Davis on January 22–23, 2009. Its goal was to engage the security community in GENI’s design and prototyping, to ensure that security issues are properly considered during its development. The workshop was designed to discuss questions such as: What security-related experiments would researchers like to run on GENI, and what benefit would they expect from those experiments? What constraints or requirements would they need to carry out the experiments? How can GENI shield other experiments and work being done using GENI from the effects of those experiments? How can successful attacks against GENI be prevented?

An additional goal was to encourage the submission of security-related proposals in response to a request for GENI analysis and prototyping proposals.

The participants were enthusiastic about the GENI project, and had a myriad of ideas about security and GENI. In addition, many responses to the solicitation for GENI analysis and prototyping proposals were received.

Contents

1	The Workshop on GENI and Security	3
1.1	Key Points	3
1.2	Requirements and Issues	4
1.3	Resource Management	4
1.4	Logging, Recording, and Capturing Events	4
1.5	Privacy	5
1.6	Architecture and Infrastructure	5
1.7	Experiments	7
1.8	GENI Itself	8
1.9	Acknowledgement	9
2	Organization and Agenda	10
2.1	Organization	10
2.1.1	Breakout Groups	10
2.1.1.1	Experimental Issues Relating to Security	11
2.1.1.2	Infrastructure Issues Relating to Security	11
2.1.1.3	Issues Relating to the Security of GENI	11
2.1.2	Organizers and Arrangers	11
2.1.3	Co-Chairs	12
2.1.4	Steering Committee	12
2.1.5	General Assistance	12
2.1.6	Local Arrangements	12
2.2	First Day’s Agenda: January 22, 2009	13
2.3	Second Day’s Agenda: January 23, 2009	14
3	Slides from the Workshop	15
3.1	Workshop on GENI and Security: Matt Bishop	16
3.2	GENI: Chip Elliott	25
3.3	Experimentation with Network-Based Security Mechanisms: George Kesidis	52
3.4	Ingredients of an Early Design for Protecting the GENI Facility: Mike Reiter	59
3.5	Some GENI Thoughts: Nicholas Weaver	67
3.6	GENITor: Nikita Borisov	70
3.7	Adaptive Security Slice Monitoring: João W. Cangussu and Ram Dantu	73
3.8	Exploiting Insecurity to Secure Software Update Systems: Justin Cappos	77
3.9	Campus Testbed for Network Management and Operations: Nick Feamster	83
3.10	GENI as an Infrastructure to Study Malicious Overlay Networks: Wenke Lee	89
3.11	Gacks: Secure Resource Allocation for GENI: John Hartman	92
3.12	Trust, Identity Management and GENI: Ken Klingenstein	96
3.13	An Adversarial Experimental Platform for Privacy and Anonymity: Ben Zhao	109
3.14	Observations on Operations/Security from a (Former) Tier 1 Builder/Operator: Chase Cotton	113

3.15	GENI Ideas: Instrumentation, Experiments and Security: Richard Ford and Ronda Henning .	115
3.16	Establishing and Communicating Trust in GENI: Raquel Hill and Jean Camp	119
3.17	(Integrity Justified) Experimental Provenance: Patrick McDaniel	123
3.18	GENI Security Configuration In a Box: Ehab Al-Shaer	127
3.19	GENI Infrastructure and Proposed GENI Experiment: Brian Hay and Kara Nance	137
3.20	GENI Security Services: Calvin Ko, Alefiya Hussain, Steve Schwab, Jim Horning, and Sandy Murphy	139
3.21	Attribution in GENI: Carrie Gates, Jeffrey Hunker, and Matt Bishop	145
3.22	GENI Trace Collection for Security Studies: Yan Luo	147
3.23	Security Event Standardization: Doug Pearson and Wes Young	155
3.24	ReAssure and SELinux: Jacques Thomas, Pascal Meunier, Patrick Eugster, and Jan Vitek . .	164
3.25	Security for High-End CyberInfrastructure: Lessons Learned: Randy Butler, Roy Campbell, Himanshu Khurana, Adam Slagell, and Von Welch	170
3.26	Supporting Study of High-Confidence Criticality-Aware Distributed CPHS in GENI: Sandeep K. S. Gupta	175
3.27	Privacy in the GENI Project: Robin Wilton	192
3.28	Secure Multi-Party Computation: Manoj Prabhakaran	197

Chapter 1

The Workshop on GENI and Security

The goal of the Workshop on GENI and Security was to engage the security community in GENI's design and prototyping, to ensure that security issues are properly considered during its development. The specific issues of interest were:

1. What classes of experiments should GENI support, and what capabilities will GENI require in order to support them?
2. How can GENI itself be secured and protected from attack? Moreover, how can those networks and cyberphysical mechanisms connected to GENI be protected from attacks originating from GENI, or malfunctioning GENI experiments?

An additional goal of the workshop was to encourage the security community to respond to a solicitation for GENI analysis and prototyping proposals released in mid-December by the GENI Projects Office.

1.1 Key Points

All participants in the workshop felt that GENI must foster a culture of scientific experimentation from the very beginning. To do this:

1. GENI must provide capabilities to enable a science of security that involves the experimental validation of security-related hypotheses that could not be validated in current testbed settings.
2. The construction of formal security experiments with hypotheses, controls, and well-articulated measurements will require substantial care and review to assure reproducibility and scientific and statistical validity.
3. GENI must provide the capabilities to enable experimenters to capture all the data needed to enable others to reproduce the experiment.
4. The deployment of GENI will require the development of mechanisms to reconcile conflicting requirements, constraints, and customs in different parts of the network.
5. The operation of GENI will require careful planning to enable communication among the federated organizations to handle (security and other) problems. The GENI infrastructure should support security testing, to ensure that security breaches can be handled quickly and effectively.

The participants were enthusiastic about the need for security in GENI, and had a myriad of ideas about the subject. Several responses to the solicitation for proposals focused on security, thus achieving the additional goal of the workshop.

1.2 Requirements and Issues

The nature of the GENI network itself raised security problems. As GENI is not controlled centrally, but is composed of autonomous federated networks, different organizations (indeed, different *types* of organizations—academic, governmental, and commercial) must provide resources and access for GENI to succeed. This raises technological, policy, procedural, and legal issues.

This narrative highlights some of those issues supporting the key points, above.

1.3 Resource Management

The question of resource management raises several security issues. First, who has the right to use resources? This requires identification and credentialing of the entities involved, and the ability to track delegation of rights. GENI will require cross-federation agreements and mechanisms to enable such management. The enforcement mechanisms must be able to reconcile disparate organizational practices and researcher identity management systems, and translate capabilities between the federated constituents. In addition, the ability to account for actions—to tie actions to the entities that take them—is normally considered a critical aspect of resource management. Will GENI researchers be held accountable for disruptive experiments? Interestingly, the participants split on this, a substantial number holding that too strict accountability might violate privacy. This raises a key issue that is explored below (see Section 1.5, **Privacy**).

GENI provides virtual networks running on a large number of systems, most of which use virtualization to support the virtual network. (For future reference, a virtual network is called a *slice*, and that part of a virtual network supported by a single system is called a *sliver*.) Managing and securing virtualization to support the virtual networks and machines, and managing and securing the slices, is a question of resource management, and one critical to the success of GENI.

Another key issue in GENI is isolation: how to prevent an experiment in one slice (or a set of slices) from interfering with experiments in other slices. If two slices share the same CPU on a particular system, do the two experiments interfere? Managing resources to both mitigate and make visible such interference is critical—and depends on an equally critical issue, the definition of “interference”. The issue of covert channels is an old one, and still a vexing one; thus questions of interference are likely to involve shades of gray, rather than binary black/white clarity. Furthermore, the degree of possible isolation may vary substantially across the heterogeneous technologies embedded within GENI. The meta-issue of how the environment and very nature of GENI affects experiments run on GENI must be understood in order to determine whether the results of the experiment will hold in other environments.

1.4 Logging, Recording, and Capturing Events

The participants expected that GENI would enable experimental validation of security-related hypotheses on large-scale networks. A key aspect of experimental science is *reproducibility*, not only by the experimenters, but also by others in order to verify the claimed results. This basic requirement implies that GENI must have specific capabilities.

GENI must be able to *record* events that occur during the experiment. This means it must support various types and levels of logging, at the level of “bits on the wire”. The ability to capture packets, for example using a program like *tcpdump*, is not sufficient because we expect GENI to be used to test new protocols, including those not based on IP (and therefore that conventional network analysis programs will not record). But the ability to measure and record everything, including background traffic and timings, leads to privacy issues (see Section 1.5, **Privacy**). The multi-national federation of networks forming GENI exacerbates this conflict.

Second, GENI must be able to *replicate* the environment of an experiment so the experiment can be repeated under the same condition as the original experiment. An experimenter should be able to take the data recorded for an experiment and from that recreate the relevant parts of the background traffic, the

slice on which the experiment was run (including all components—internal slivers, end points, etc.), and any other parts of the environment. Then the experiment can be rerun and the results can be validated. As in other experimental fields, perfect replication may be impossible in many scenarios, which raises the important scientific question of the degree of replication and repeatability required for experimental validation of security-related hypotheses.

1.5 Privacy

Because GENI is a federated testbed, the definition of “privacy” will vary among the federated networks. In particular, the federation is planned to include organizations in Europe and Japan, where privacy laws are very different than those in the United States. This has several consequences.

First is the impact on what can be recorded. Synthesized data (especially synthetic background traffic) should not be a problem anywhere, but such data is often not realistic. For example, intrusion detection systems often use the synthesized 1999 IDS Challenge data set to demonstrate their effectiveness; in the research community, any such results are considered suspect. Various proposals for recording and replaying real network traffic would avoid this problem, but raise many others, both technical and legal. In the context of privacy, one is whether the traffic can be used, or whether it must be anonymized and if so, to what degree.

Two approaches for this were discussed. The first is simply to record data elsewhere, anonymize it, and construct a framework for seeding it with attacks should the experimenter decide to do so. Then one could replay this data for experimental purposes. The second approach is to encourage ordinary users to use GENI, in effect making GENI a network that the public (or segments of the public, such as students or academic institutions) could use. Both raise issues of privacy, but the approach for handling privacy is different. The first can be anonymized before it is used; the second would have to be anonymized on the fly or recorded and subsequently anonymized. Further, the transformed data would have to be shown to have the same characteristics (specifically, those that can affect the experiment) as those of the untransformed data. Finally, whether “perfect” anonymization is in fact possible is an open question; often private data can be reconstructed from anonymized data when the attacker has access to external information.

There was some discussion of requiring the users of GENI to consent to monitoring, but this was felt infeasible unless the set of users is tightly controlled. We could solve this problem by limiting the measurement and recording to those parts of the data relevant to network analysis and that did not violate privacy rights; but this raises other issues, such as the reuse of the data for other experiments.

To put this problem of balancing as starkly as possible: under what conditions can we decide whether an experiment is doing something that violates the rules of usage without compromising the privacy of the experiment? Indeed, who is the “we” that decides this? And how are disputes handled (see Section 1.6, **Architecture and Infrastructure**, below).

A further question of privacy arises in regard to visibility of the measurements themselves. Are an experiment’s measured results visible only to the researcher(s) running the experiment? Or must they be made open and transparent to all researchers? Since some researchers may wish to preserve their own privacy (e.g. until they publish), there may be good social reasons to keep measurements private at least for some time.

Thus, the entire process of data collection, and controlling the data once collected, is key not only to the success of GENI as an experimental testbed but also to the acceptability of GENI under the law, regulations, and policies of its constituent networks.

1.6 Architecture and Infrastructure

Considerable discussion of the infrastructure for GENI revolved around the human and policy aspects, as opposed to the purely technical. As security is primarily a human endeavor, this was not surprising. Several interesting questions emerged.

First, *what is security?* An early discussion brought this out. Consider an experimenter who is designing a new protocol with attribution of its packets as its goal—that is, every packet can be traced back to its host of origin. This enables one to deduce, for example, origins of distributed denial of service attacks—generally considered a good thing.¹ However, if a dissident is emailing anonymous messages to the press identifying corruption, the anonymity may be that dissident’s only protection from trial and imprisonment; here, attribution would be considered a bad thing.² So, is attribution a security requirement? The best answer is that it depends upon policy—and the exact policy will undoubtedly vary among the various federated networks (especially among those in different countries, and therefore subject to different laws).

These considerations suggest that using automated mechanisms to monitor and enforce security is problematic. Several specific mechanisms were discussed in the workshop. One issue is whether these mechanisms could provide a high enough probability of detection at a sufficiently low probability of false alarm. More broadly, it is by no means obvious how these mechanisms can be aligned with the deep policy issues discussed above.

Second, what (security) support services must the organizations with networks federated with GENI supply? To a large degree, this question is poorly understood because, so far, very few federated systems have crossed national and international boundaries. The participants in the workshop knew of no direct experience with such systems within the field of computer science research. There may, however, be lessons that are directly relevant from “big science” federations of recent years, such as the Grid endeavors and large-scale physics experiments (for example, the Large Hadron Collider).

Further, our experience with the Internet is disquieting. As an example, consider incident response. Different incident response groups have tried electronic means for communicating among themselves; these usually are not effective enough. The most effective communication mechanism is personal contact, either because you know your counterpart personally, or you can reach your counterpart with the aid of others who know you both. Whom do you call when one part of GENI is malfunctioning and blocking access to your resources? Further, if there is a dispute, how will it be arbitrated? While legal recourse is available, in the United States at least, this often takes a very long time and is expensive. International litigation is probably even more expensive and time-consuming, even in those cases where it is feasible. An arbitration mechanism would work better.

The infrastructure ideally would supply timely response to questions, and take action when problems are reported. This essentially requires that someone be available at all times. If the model for GENI is to federate academic, government, and commercial institutions’ networks, many of the constituent networks will not be able to provide that level of support—for example, academic computer science experimental networks run by faculty and students. Decreeing that a certain minimum standard must be met in order to federate with GENI was felt to be impractical, based on past experience with other types of voluntary federations. Invariably, some constituent fails to meet the minimum standards; but unless the failure is egregious, it is in practice difficult to expel a volunteer, and much more so if the volunteer is supplying needed or valuable resources. In general, social and peer pressure work better to encourage conformance to a minimal standard than do consequences that are costly for the federation.

This then brings up the question of how the federation works. Each constituent brings resources into the federation. Who decides how to assign these resources, and to whom? This affects availability, a key security service. For example, a policy may require that disruptive or misbehaving experiments have their priority, and hence their access to resources, reduced.³ If this is a centralized decision, then the central controller must have control of all experiments—many in the workshop believed that this was highly unlikely in a federated network or system, and felt it antithetical to the nature of GENI. (GENI Spiral 1 does by contrast posit exactly such a centralized control system, located within the clearinghouse.) A distributed decision must take into account local policies as well as global policies, and there must be a mechanism for reconciling

¹But not always. In war time, if a country were to use a distributed denial of service (DDoS) attack to hinder its adversary, that country would want the DDoS attack to appear to come from an ally of the adversary to sow discord among the country’s enemies. There, attribution is exactly what the country does *not* want.

²Except by the government trying to identify and catch the dissident.

³This also raises the question of what “disruptive or misbehaving” experiments are. See the discussions about defining privacy and security, above.

differences.

This takes us back to the requirements—what support services must GENI provide? It is not clear that a single set of services would meet with universal acceptance because of the tension between privacy and accountability, as discussed above. Thus, ultimately GENI’s stakeholders must set its requirements. The workshop identified three major types of stakeholders: *those who provide the resources* (the constituents), *those who provide the data* (for example, sample background traffic or measurements), and *those who will use the resources* (the experimenters, managers, and other users). There was some discussion as to whether the experimental subjects also represented a set of stakeholders that needed to be represented beyond a human experiment review board (IRB). Additional stakeholders may include governments, regulatory bodies, and other political, legal, and social entities.

Ultimately, the owners of resources must manage their resources, because few will voluntarily give full control of their resources to the distributed system (see Section 1.3, *Resource Management*, above). Some aspects will probably be done locally. Others would require a common clearinghouse. For example, if an experiment needs access to a SCADA testbed connected to GENI, the experimenter can query the clearinghouse asking where she can access a SCADA testbed connected to GENI and having specific properties. The clearinghouse can then suggest other constituents whose SCADA testbeds meet the stated requirements, and the experimenter can then schedule time on them with either the local controller or (better) using the clearinghouse.

Workshop participants also discussed the nature of experiments. Some larger, long-term experiments will take on a provenance of their own, creating a meta-structure within GENI. Participants also raised the issue of ownership of the experiment, and of how to handle the transfer of intellectual property regarding the experiment should it be transferred into production mode, for example as a new security service for which organizations would pay.

Participants pointed out that the GENI must be easy for the constituents to manage. As GENI is a federation of networks, the goal is to get institutions to allow GENI to use their resources. Without funds from GENI, this requires volunteers. Experience shows that if volunteers must spend great amounts of time, effort, and other resources to do their tasks, they quickly become “former volunteers”. For GENI, this would be disastrous. Further, the principle of psychological acceptability says that if management is not easy, configuration and other errors will occur, possibly disrupting experiments, and the GENI testbed itself. Therefore, considerable care must be given to making joining the GENI testbed, and maintaining membership in it, inexpensive in both efforts and funds.

It was also noted that GENI should enable an experimenter to specify and acquire specific classes of resources. For example, an experimenter should be able to acquire a computer to use as a router, rather than being forced to use a slice of the computer as a router.

1.7 Experiments

GENI must provide capabilities to enable a science of security that involves the experimental validation of security-related hypotheses that could not be validated in current testbed settings. The participants viewed GENI as a vehicle for instilling a culture of scientific instrumentation and experimentation into the security research community. With the availability of such a testbed there would be no excuse for failing to experimentally verify security claims that cannot be verified using other means. Further, several participants pointed out that GENI could be used as a teaching tool for how to carry out scientific experimentation in computer science and, especially, computer security. Given the need for emphasis on rigorous scientific testing in computer science curricula, this use may be the most important for the future of computer science and computer security.

Basic scientific and experimental capabilities include mechanisms to collect information and make measurements. This raises privacy issues, as discussed above. As more people want to use GENI, teaching them how to implement experiments correctly, and analyze the results, becomes critically important. The participants all felt that GENI should provide a set of detailed experiments that users could modify to learn how to do experiments on GENI—even simple experiments would be very helpful. Recipes or cookbooks

for constructing and running experiments will prove invaluable, too. A supportive experimental community willing to share its knowledge and tools, combined with a GENI help desk for experimenters, is an important asset

Creating a methodology for experimentation involving security, especially for experimentation on GENI, is important. This methodology should address topics such as the validation of the experiments themselves, validation of the data used by the experiment, and how GENI itself affects experimental results due to instrumentation effects, communication delays, and other attributes not present in the environment being experimented about.

Considerable discussion focused on the type of experiments users of GENI might want to run. Throughout the discussion, the focus was on experiments that are infeasible on current systems and testbeds because they are too small or not isolated; and infeasible on the Internet, again because it is not isolated.

The two experiments with the most immediate impact are the validation of models for distributed denial of service attacks and defenses on a large scale; and for the development of new architectures to inhibit botnets. Validation in this context requires the deployment and running of both types of attacks, because often experimental results show that the models we have developed do not match the reality of what happens in the network, and thus must be tuned. Worse, some phenomena may well be chaotic and so effects cannot be predicted, only described once they occur. Without experimentation, we will not know how good our models are, and whether they can be used to predict results on systems, especially those other than GENI (such as the Internet).

Both these experiments would disrupt the use of the Internet if tried on that. Other examples are cascading failure (where end or infrastructure systems begin failing), or simulations of changing large distributed networks with properties different than that of the Internet (such as the power grid). Thus, more generally, any experiment that would disrupt the Internet if run on the Internet would be appropriate for GENI. Further, GENI has a programmable infrastructure, so the routers and other infrastructure systems can be reprogrammed from other nodes (unlike the current Internet). This allows the edges (end nodes) and the core (infrastructure) to collaborate, for example on security defenses or measurements; this is not possible in the current Internet, in general.

Three other types of experiments were discussed. GENI offers the opportunity to evaluate the security of deployed solutions on a large-scale distributed network and/or system. For example, one can use GENI as the testbed for a large distributed system or application, and then analyze it to determine whether it *really* is secure, robust, and scalable. One can also use GENI to test (or simulate) very high cost, but low probability, events for complex scenarios and novel threats.⁴ The third type was an exercise like CyberStorm to develop plans and procedures to deal with threats against large distributed networks and systems.

Concepts that start as an experiment may develop constituencies of users who depend on the implementation. Thus, the experiment may evolve into a service. As noted above, this raises the problems of handling intellectual property, and transitioning the experiment to a production service. GENI needs to express the rules governing solutions to these problems in its environment, and develop mechanisms to support and promote this growth.

1.8 GENI Itself

The workshop also discussed protecting GENI, and ensuring experiments stayed on GENI. The phrasing here is critical. It is not possible to prevent attacks on GENI, and undoubtedly some will succeed. The issue then is how to minimize the effects that those attacks have on GENI, and on the experiments being run; and how to ensure the experimenters are notified of the attack so they can take that into account when analyzing their results.

A key issue is legal liability. For example, suppose a malware experiment in GENI goes awry because GENI's mechanism for isolating the slice fails. What are the legal ramifications, especially when the network crosses international borders? How do we ensure that the GENI constituents can communicate among

⁴Some participants referred to GENI supporting an "Underwriters Laboratory" for security technology. This raises many issues such as quality control, requirements testing, and such, that the workshop did not pursue.

themselves to deal with terminating the worm, and repairing the damage to GENI and to others, effectively? As another example, suppose an attacker compromises a system belonging to GENI and implants a botnet on GENI. This not only compromises GENI, but it also renders many experiments (for example, those relating to network throughput) suspect.

One approach that many (especially the practitioners) thought would help would be to use “red teams” to compromise GENI to test the ability of the GENI organization, and the federated organizations making up GENI, to respond to attacks. The goal would be for the red team to disrupt some aspects of the GENI testbed (preferably those not being used for experimentation) and see how long it took to detect and restore those parts. This tests not only the technical protections but also the procedures, the availability, and the readiness of the constituents and the managers to act quickly.

Finally, the participants noted that GENI itself is an experiment: a federation and testbed of this complexity has not yet been created. Therefore, we should consider having social scientists study GENI itself and how users, organizations, and others interact with GENI and with one another. The goal here is to improve the usefulness and usability of GENI to make it as effective a testbed as possible.

1.9 Acknowledgement

This workshop was funded by the National Science Foundation’s award number CNS-0646965 to the University of California at Davis. All opinions expressed in this report, and the presentations, are not necessarily those of the National Science Foundation, the University of California at Davis, or the institutions of the authors and presenters.

Chapter 2

Organization and Agenda

2.1 Organization

The goals of the workshop led to its organization as a brainstorming meeting. Each participant was asked to prepare a short¹ statement of ideas addressing the following two issues:

- What classes of security experiments should GENI support? What capabilities will GENI require to allow the conduct of these experiments? The capabilities may be intrinsic to GENI (such as equipment or software of a particular kind) or extrinsic (such as organizational management, or external interfaces and connectivity). Experiments involving malware or vulnerabilities analysis may require that parts of the infrastructure suite be partitioned from other parts. Deploying and testing new protocols may require that the suite be partitioned to prevent errors in the implementation or in the protocol itself from interfering with other uses of the infrastructure.
- How can GENI itself be adequately secured and protected from attack? What forms of authentication, authorization, and accountability would be most appropriate? As access to GENI will be from the Internet, GENI will be exposed to potential attackers. Other types of attack may involve physical compromise of the systems making up GENI, or of the Internet (or other) infrastructure that provides support for GENI. Protocols, management and organizational procedures and processes, and access control mechanisms must be developed to safeguard both the GENI resource and the data and software that researchers deploy on it.

Those who were invited were asked to give a 5 minute presentation on their ideas. These presentations were intermingled with the breakout sessions, again to foster ideas and encourage people to discover and pursue common interests.

Several speakers were also invited to present longer talks on GENI, experimentation, and an early design for protecting GENI.

2.1.1 Breakout Groups

The workshop meeting had three main themes, with one breakout group for each:

1. Experimental issues relating to security (see Section 2.1.1.1)
2. Infrastructure issues relating to security (see Section 2.1.1.2); and
3. Issues relating to the security of GENI (see Section 2.1.1.3).

¹At most 1 page

The breakout groups ran simultaneously in three sessions of 75 minutes each, and participants were encouraged to move between groups. The intent was to enable all members to contribute to all groups, and bring ideas from one group to another.

The results of the breakout groups, and the workshop as a whole, are presented in the previous section. Because many of the issues in the groups overlapped and complimented one another, we felt it inadvisable to present them separately.

2.1.1.1 Experimental Issues Relating to Security

Moderators: Terry Benzel and Karl Levitt

The questions this group was asked to examine were:

- What types of experiments will people want to do on GENI?
- What resources (construed broadly) do they need?
- What types of experiments could be started quickly assuming resources were available?
- What types of these resources could be made available quickly?

2.1.1.2 Infrastructure Issues Relating to Security

Moderator: Deborah Frincke

The questions this group was asked to examine were:

- What technological issues arise from a federated, heterogeneous environment?
- What procedural issues arise from a federated, heterogeneous environment? Can we separate these from technological issues?
- What legal issues arise from this network being spread over multiple jurisdictions? This includes international jurisdictions such as Europe.

2.1.1.3 Issues Relating to the Security of GENI

Moderator: J. F. Mergen

The questions this group was asked to examine were:

- What does “GENI security” mean? What does “protecting (a federated) GENI” mean?
- How do we protect GENI experiments from one another?
- How do we protect GENI itself from both experiments and outside attack? When things go wrong within GENI, how do we restore it?
- How do we protect the outside world from experiments? This includes the Internet, SCADA, cellular systems, and any other type of system or network connected to GENI.

2.1.2 Organizers and Arrangers

The following people helped organize, support, and run the workshop. In addition to the sponsorship of the National Science Foundation under award CNS-0646965, the University of California at Davis provided substantial assistance arranging the local accommodations and facilities that the workshop used, and we are grateful to both NSF and UC Davis for their support.

All opinions expressed in this report, and the presentations, are not necessarily those of the National Science Foundation, the University of California at Davis, or the institutions of the authors and presenters.

2.1.3 Co-Chairs

- Matt Bishop, University of California at Davis
- Chip Elliott, BBN

2.1.4 Steering Committee

- Heidi Picher Dempsey, BBN
- Deborah Frincke, Pacific Northwest National Laboratories
- Suzanne Iacono, National Science Foundation
- Karl Levitt, University of California at Davis and National Science Foundation
- John Mitchell, Stanford University
- Vern Paxson, University of California at Berkeley
- Taieb Znati, National Science Foundation

2.1.5 General Assistance

- Sean Peisert, University of California at Davis

2.1.6 Local Arrangements

- Greg Gibbs, University of California at Davis
- Linda Tsang, University of California at Davis

2.2 First Day's Agenda: January 22, 2009

- Welcome and Orientation, Matt Bishop (page 16)
- “GENI,” Chip Elliott, GENI Project Director (page 25)
- Break
- “Experimentation with Network-Based Security Mechanisms,” George Kesidis (page 52)
- “Ingredients of an Early Design for Protecting the GENI Facility,” Mike Reiter (page 59)
- “National Cyber Range,” Mike Van Putte
- Participant talks
 - “Some GENI Thoughts,” Nicholas Weaver (page 67)
 - “GENITor,” Nikita Borisov (page 70)
 - “Adaptive Security Slice Monitoring,” João Cangussu, Ram Dantu (page 73)
 - “Exploiting Insecurity to Secure Software Update Systems,” Justin Cappos (page 77)
 - “Campus Testbed for Network Management and Operations,” Nick Feamster (page 83)
 - “GENI as an Infrastructure to Study Malicious Overlay Networks,” Wenke Lee (page 89)
 - “Gacks: Secure Resource Allocation for GENI,” John Hartman (92)
 - “Trust, Identity Management and GENI,” Ken Klingenstein (page 96)
- Lunch at Tercero (a dining hall)
- Breakout Sessions #1
- Reports from breakout groups (5 min each max)
- Participant talks
 - “An Adversarial Experimental Platform for Privacy and Anonymity,” Ben Zhao (page 109)
 - “Observations on Operations/Security from a (Former) Tier 1 Builder/Operator,” Chase Cotton (page 113)
 - “GENI Ideas: Instrumentation, Experiments and Security,” Richard Ford, Ronda Henning (page 115)
 - “Establishing and Communicating Trust in GENI,” Raquel Hill, Jean Camp (page 119)
 - “(Integrity Justified) Experimental Provenance,” Patrick McDaniel (page 123)
 - “GENI Security Configuration In a Box,” Ehab Al-Shaer (page 127)
- Break
- Breakout Sessions #2
- Reports from breakout groups (5 min each max)

2.3 Second Day's Agenda: January 23, 2009

- Logistics and Review of Yesterday's Results
- Participant talks
 - “GENI Infrastructure and Proposed GENI Experiment,” Brian Hay, Kara Nance (page 137)
 - “GENI Security Services,” Calvin Ko, Alefiya Hussain, Steve Schwab, Jim Horning, Sandy Murphy (page 139)
 - “Attribution in GENI,” Carrie Gates, Jeffrey Hunker, Matt Bishop (page 145)
 - “GENI Trace Collection for Security Studies.” Yan Luo (page 147)
 - “Security Event Standardization,” Doug Pearson, Wes Young (page 155)
 - “ReAssure and SELinux,” Jacques Thomas, Pascal Meunier, Patrick Eugster, Jan Vitek (page 164)
 - “Security for High-End CyberInfrastructure: Lessons Learned,” Randy Butler, Roy Campbell, Himanshu Khurana, Adam Slagell, Von Welch (page 170)
 - “Supporting Study of High-Confidence Criticality-Aware Distributed CPHS in GENI,” Sandeep Gupta (page 175)
 - “Privacy in the GENI Project,” Robin Wilton (page 192)
 - “Secure Multi-Party Computation,” Manoj Prabhakaran (page 197)
- Breakout Sessions #3
- Break
- Reports from breakout groups (5 min each max)
- Plenary discussion, summing up, etc.
- Concluded

Chapter 3

Slides from the Workshop

The following are the slides from the workshop.

3.1 Workshop on GENI and Security: Matt Bishop

Workshop on GENI and Security

January 22–23, 2009



Workshop Goal

- To engage the security community in GENI's planning, design, prototyping, and early trial experiments (now underway) to ensure that security issues are properly considered during development.
- To give ideas on how to make GENI as useful as possible to the security community
- To answer, or suggest approaches to determine the answers, to several questions
 - See breakout sessions
- To solicit security-related proposals



Agenda: Thursday Morning

- 8:30–9:00 Welcome, schedule, orientation, discussion of goals, *etc*
- 9:00–10:00 Talk: All about GENI, both organizational and technical
- 10:00–10:15 **Break**
- 10:15–10:45 Talk: Experimentation in Security
- 10:45–11:00 Talk: Reflections on GENI Security
- 11:00–11:15 Talk: National Cyber Range
- 11:15–12:00 Participant talks
- 12:00–1:15 **Lunch** (at Tercero)

3



Agenda: Thursday Afternoon

- 1:15–2:25 Breakout #1
- 2:25–2:40 Reports from breakout groups (5 minutes each)
- 2:40–3:20 Participant talks
- 3:20–3:35 **Break**
- 3:35–4:45 Breakout #2
- 4:45–5:00 Reports from breakout groups (5 minutes each)
- 5:00–6:00 **Reception** (in the lobby)

4



Agenda: Friday Morning

- 8:30–9:00 Logistics, review of previous day's results
- 9:00–9:50 Participant talks
- 9:50–11:00 Breakout #3
- 11:00–11:15 **Break**
- 11:15–11:30 Reports from breakout groups (5 minutes each)
- 11:30–12:00 Plenary discussion, summing up, *etc.*
- 12:00 **Workshop ends**

5



Participant Talks: Thursday Morning

- 11:15 Nick Weaver
- 11:20 Nikita Borosov
- 11:25 Joao Cangussu
- 11:30 Justin Cappos
- 11:35 Nick Feamster
- 11:40 Wenke Lee
- 11:45 John Hartman
- 11:50 Ken Klingenstein

6



Participant Talks: Thursday Afternoon

- 2:40 Ben Zhao
- 2:45 Chase Cotton
- 2:50 Richard Ford
- 2:55 Raquel Hill
- 3:00 Patrick McDaniel
- 3:05 Deborah Frincke
- 3:10 Ehab Al-Shaer

7



Participant Talks: Friday Afternoon

- 9:00 Brian Hay
- 9:05 Calvin Ko
- 9:10 Carrie Gates
- 9:15 Yan Luo
- 9:20 Doug Pearson
- 9:25 Jacques Thomas
- 9:30 Von Welch₂₀
- 9:35 Sandeep Gupta
- 9:40 Robin Wilton

8



Breakout Sessions

- Designed to focus on specific questions, areas
- Three sessions of 70 minutes each
 - People assigned to first sessions
 - You can go to any of the second and third sessions—we urge you to go to a different one each time!
- At end of each session, 5 minutes to present answers, thoughts, *etc.*, to series of questions
 - See following slides

9



Breakout #1: Experimental Issues

Moderators: Terry Benzel, Karl Levitt

- What types of experiments will people want to do on GENI?
- What resources (construed broadly) do they need?
- What types of experiments could be started quickly assuming resources were available? What types of these resources could be made available quickly?

10



Breakout #2: Infrastructure Issues

Moderator: Deborah Frincke

- What technological issues arise from a federated, heterogeneous environment?
- What procedural issues arise from a federated, heterogeneous environment? Can we separate these from technological issues?
- What legal issues arise from this network being spread over multiple jurisdictions? This includes international jurisdictions such as Europe.

11



Breakout #3: GENI Security Issues

Moderator: J. F. Mergen

- What does “GENI security” mean? What does “protecting (a federated) GENI” mean?
- How do we protect GENI experiments from one another?
- How do we protect GENI itself from both experiments and outside attack? When things go wrong within GENI, how do we restore it?
- How do we protect the outside world from experiments? This includes the Internet, SCADA, cellular systems, and any other type of system or network connected to GENI.

12



Logistical Information

- Bathrooms
 - Midway down the corridor outside this room
- Wireless network
 - Called “moobilenet” (yes, *two* “o”s—think cows)
 - Bring up web browser
 - Login is bishop@cs.ucdavis.edu
 - Password is “geni-sec”
 - This will work in many places on campus
 - *Warning:* this network may not be accessible on other floors of this building

13



Web Site

- We want to put a list of participants there
 - Plan is: Name, affiliation
 - If you do not want to be listed, please tell us!
- We also would like to put slides, results of breakout sessions there
 - If you do not want your slides put up, please tell us!
- Web site:
 - <http://seclab.cs.ucdavis.edu/meetings/gen-sec>

14



Travel Support

- Form in your packet
 - Also available on the web at
<http://www.cs.ucdavis.edu/department/forms/travelexpense.pdf>
- Fill it out, sign it, send it *along with the original receipts* to:
 - Matt Bishop
 - Dept. of Computer Science
 - University of California at Davis
 - 1 Shields Ave.
 - Davis, CA 95616-8562
- Please do so within 30 days!

15



People

- Co-chairs:
 - Matt Bishop, UC Davis
 - Chip Elliott, BBN
- General Assistance
 - Sean Peisert, UC Davis
- Local Arrangements
 - Greg Gibbs, UC Davis
 - Linda Tsang, UC Davis
- Sponsor
 - **National Science Foundation**
- Steering Committee
 - Heidi Picher Dempsey, BBN
 - Deborah Frincke, PNNL
 - Suzanne Iacono, NSF
 - Karl Levitt, NSF
 - John Mitchell, Stanford
 - Vern Paxson, UC Berkeley
 - Taieb Znati, NSF

24

16



3.2 GENI: Chip Elliott



GENI

Global Environment for Network Innovations

Chip Elliott
GENI Project Director
celliott@bbn.com

www.geni.net

Clearing house for all GENI news and documents

www.geni.net

1



Thank you Matt! and Karl!



There once was a Bishop from Davis . . .

And also introducing . . .

➤ National Science Foundation

- Dr. Suzi Iacono
- Dr. Karl Levitt

➤ DARPA

- Dr. Mike VanPutte

➤ GENI Project Office

- Dr. Harry Mussman
- Dr. Vic Thomas

26



GPO goals for this workshop

- Engage the security community to play an active, central role in GENI's planning, prototyping, and early trial experiments (now rolling out as Spiral 1; first demos in March)
- Very concretely – encourage you to submit proposals for GPO Solicitation #2, due Feb. 20

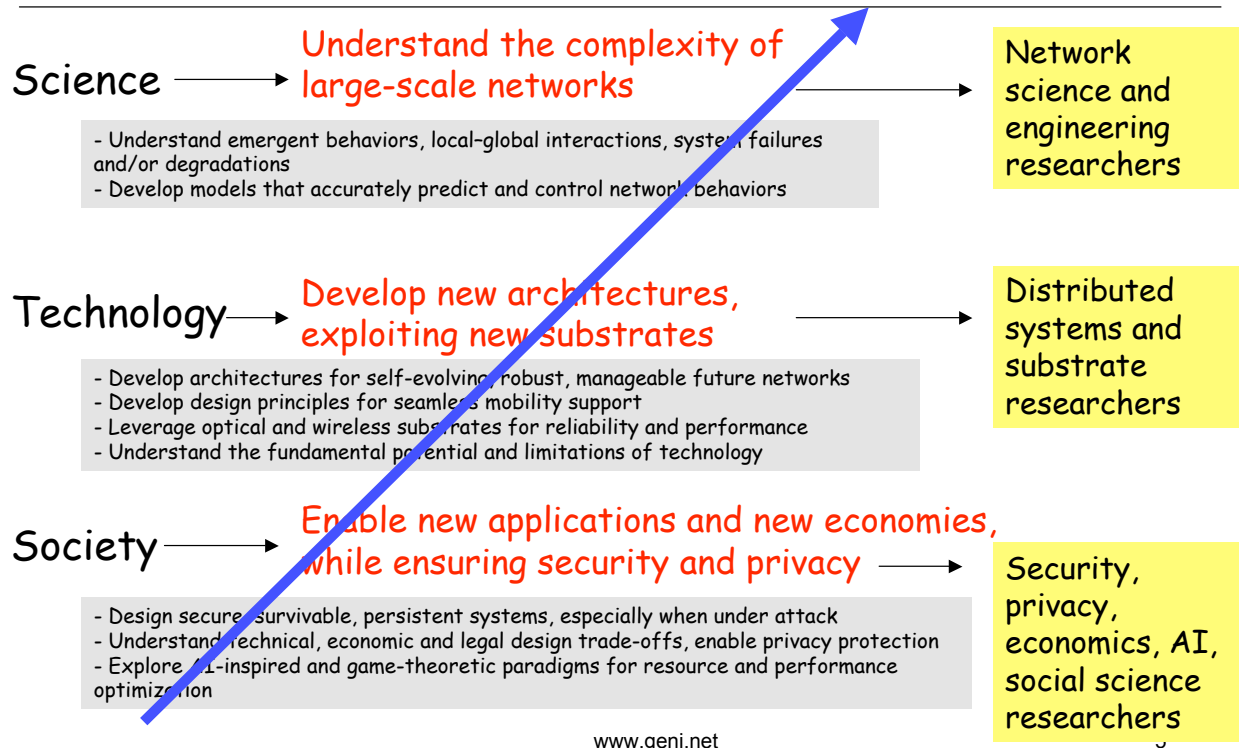


Outline

- What is GENI?
- How we'll build it, how we'll use it (Two Comic Books)
- The GENI system concept
- GENI Spiral 1
- How can you participate?

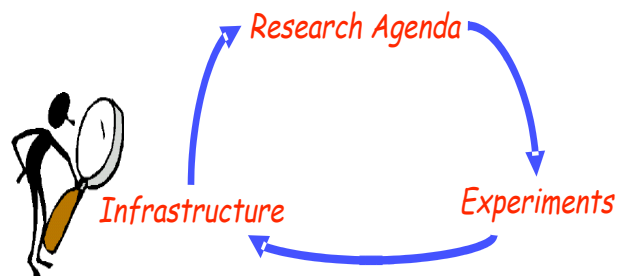


GENI supports Fundamental Challenges Network Science & Engineering (NetSE)



Research Agenda to Experiments to Infrastructure

- Research agenda
 - Identifies fundamental questions
 - Drives a set of experiments to validate theories and models
- Experiments & requirements
 - Drives what infrastructure and facilities are needed
- Infrastructure could range from
 - Existing Internet, existing testbeds, federation of testbeds, something brand new (from small to large), federation of all of the above, to federation with international efforts
 - No pre-ordained outcome



Existing Input

- | | |
|--|---|
| <ul style="list-style-type: none"> • Clark et al. planning document for Global Environment for Network Innovations • Shenker et al. "I Dream of GENI" document • Kearns and Forrest ISAT study • Feigenbaum, Mitzenmacher, and others on Theory of Networked Computation | <ul style="list-style-type: none"> • Hendler and others in Web Science • 28. Ruzena Bajcsy, Fran Berman, and others on CS-plus-Social Sciences • NSF/OECD Workshop "Social and Economic Factors Shaping the Future of the Internet" • NSF "networking" programs <ul style="list-style-type: none"> - FIND, SING, NGNI |
|--|---|



“Our founders”

The GENI Planning Group and Many, Many Working Group Volunteers

Larry Peterson, Princeton (Chair)	Nick McKeown, Stanford
Tom Anderson, Washington	Dipankar Raychaudhuri, Rutgers
Dan Blumenthal, UCSB	Mike Reiter, CMU
Dean Casey, NGENET Research	Jennifer Rexford, Princeton
David Clark, MIT	Scott Shenker, Berkeley
Deborah Estrin, UCLA	Amin Vahdat, UCSD
Joe Evans, Kansas	John Wroclawski, USC/ISI
Terry Benzel, USC/ISI	CK Ong, Princeton

And Within NSF

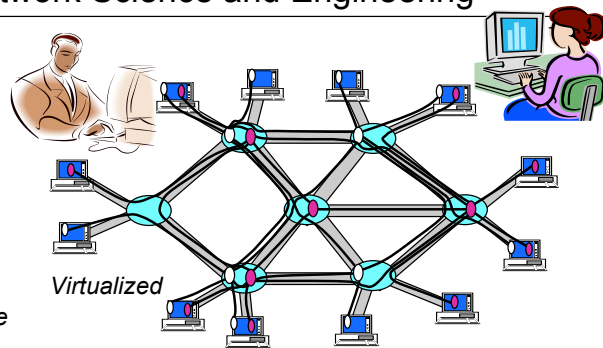
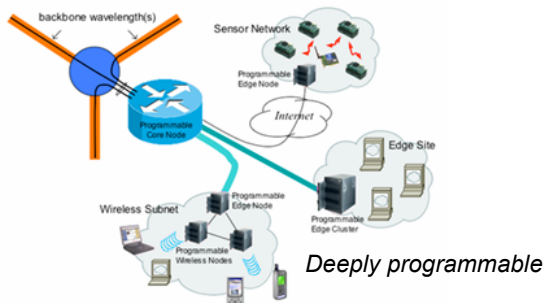
Peter Freeman	Guru Parulkar	Ty Znati
Debbie Crawford	Darleen Fisher	Gracie Narcho
Larry Landweber	Cheryl Albus	Paul Morton
Suzi Iacono	Allison Mankin	

Their hard work has created GENI’s Conceptual Design, the starting point for all our work going forward.

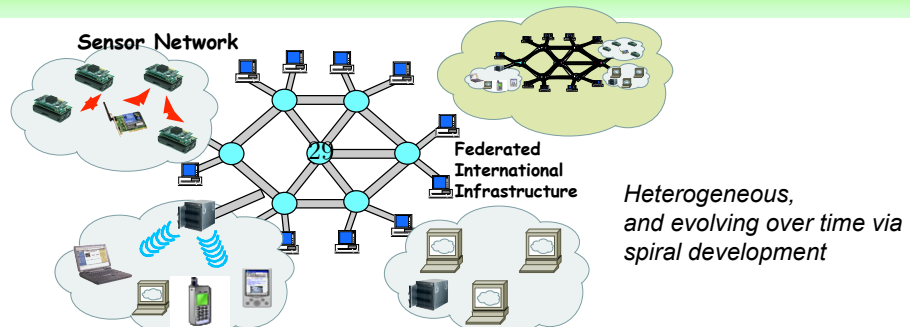


The GENI Vision

A national-scale suite of infrastructure for long-running, realistic experiments in Network Science and Engineering



Programmable & federated, with end-to-end virtualized “slices”





Outline

- What is GENI?
- How we'll build it, how we'll use it
(Two Comic Books)
- The GENI system concept
- GENI Spiral 1
- How can you participate?

www.geni.net

9



How We'll Use GENI

Note that this is the “classics illustrated” version – a comic book!

Please read the Network Science and Engineering Research Agenda to learn all about the community's vision for the research it will enable.

Your suggestions are very much appreciated!

www.geni.net

10



A bright idea




I have a great idea! The original Internet architecture was designed to connect one computer to another – but a better architecture would be fundamentally based on PEOPLE and CONTENT!

That will never work! It won't scale! What about security? It's impossible to implement or operate! Show me!

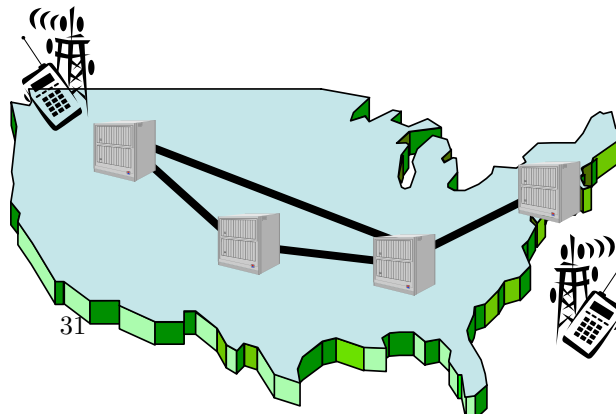


Trying it out



My new architecture worked great in the lab, so now I'm going to try a larger experiment for a few months.


And so he poured his experimental software into clusters of CPUs and disks, bulk data transfer devices ('routers'), and wireless access devices throughout the GENI suite, and started taking measurements . . .



He uses a modest slice of GENI, sharing its infrastructure with many other concurrent experiments.

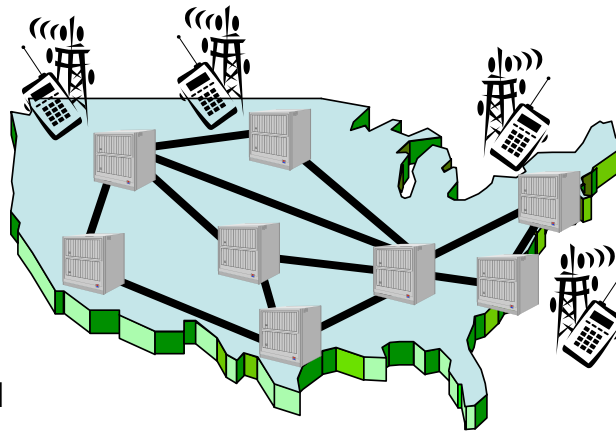


It turns into a really good idea



Boy did I learn a lot! I've published papers, the architecture has evolved in major ways, and I'm even attracting real users!

Location-based social networks are really cool!



His experiment grew larger and continued to evolve as more and more real users opted in . . .

His slice of GENI keeps growing, but GENI is still running many other concurrent experiments.



Experiment turns into reality



My experiment was a real success, and my architecture turned out to be mostly compatible with today's Internet after all – so I'm taking it off GENI and spinning it out as a real company.




I always said it was a good idea, but way too conservative.





Meanwhile . . .



I have a great idea! If the Internet were augmented with a scalable control plane and realtime measurement tools, it could be 100x as reliable as it is today . . . !

And I have a great concept for incorporating live sensor feeds into our daily lives !



If **you** have a great idea, check out the **NSF CISE Network Science and Engineering** program.

www.geni.net

15



Moral of this story

- GENI is meant to enable . . .
 - Trials of new architectures, which may or may not be compatible with today's Internet
 - Long-running, realistic experiments with enough instrumentation to provide real insights and data
 - 'Opt in' for real users into long-running experiments
 - Large-scale growth for successful experiments, so good ideas can be shaken down at scale
- A reminder . . .
 - GENI itself is not an experiment !
 - GENI is a suite of infrastructure on which experiments run

GENI creates a huge opportunity for ambitious research!

www.geni.net

16



How We'll Build GENI

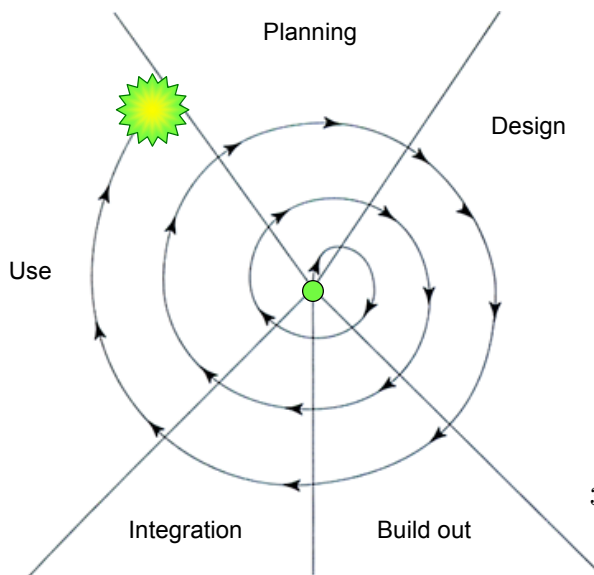
Note that this is the “classics illustrated” version – a comic book!

Please read the GENI System Overview and GENI Spiral 1 Overview for detailed planning information.



Spiral Development

GENI grows through a well-structured, adaptive process



GENI Prototyping Plan

- An achievable **Spiral 1**
Rev 1 control frameworks, federation of multiple substrates (clusters, wireless, regional / national optical net with early GENI 'routers', some existing testbeds), Rev 1 user interface and instrumentation.

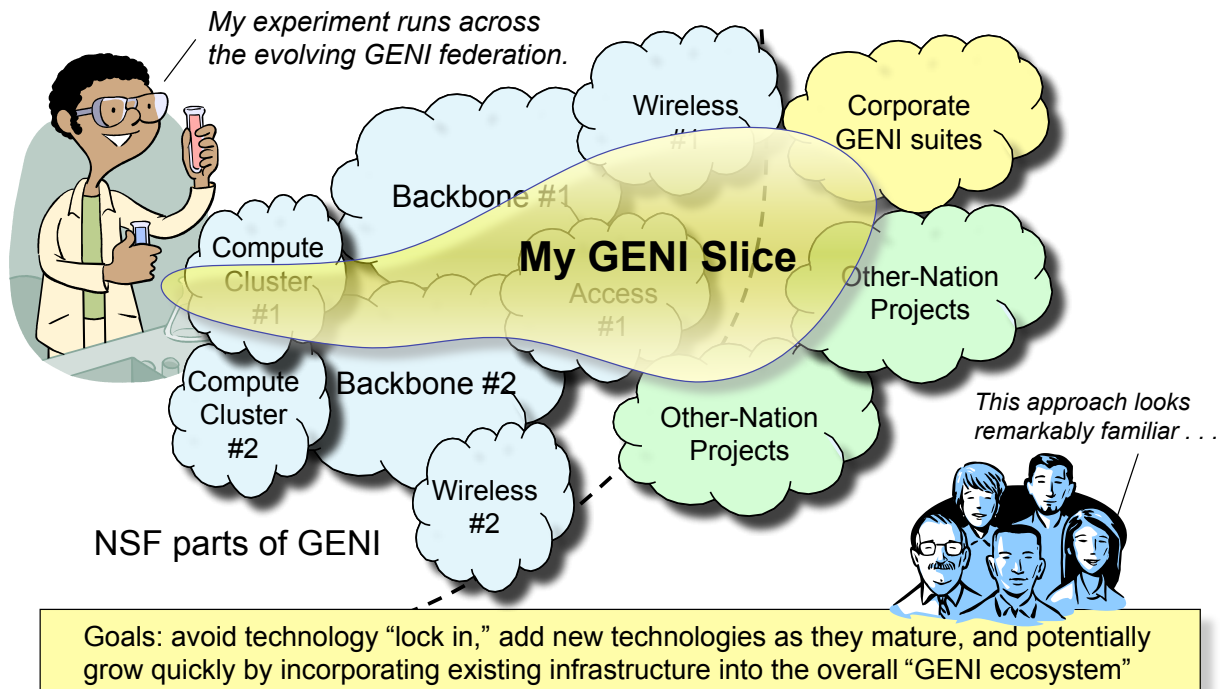
- **Envisioned ultimate goal**
Example: Planning Group's desired GENI suite, probably trimmed some ways and expanded others. Incorporates large-scale distributed computing resources, high-speed backbone nodes, nationwide optical networks, wireless & sensor nets, etc.

- **Spiral Development Process**
Re-evaluate goals and technologies yearly by a systematic process, decide what to prototype and build next.



Federation

GENI grows by “gluing together” heterogeneous infrastructure



www.geni.net

19



Outline

- What is GENI?
- How we'll build it, how we'll use it (Two Comic Books)
- The GENI system concept
- GENI Spiral 1
- How can you participate?

35

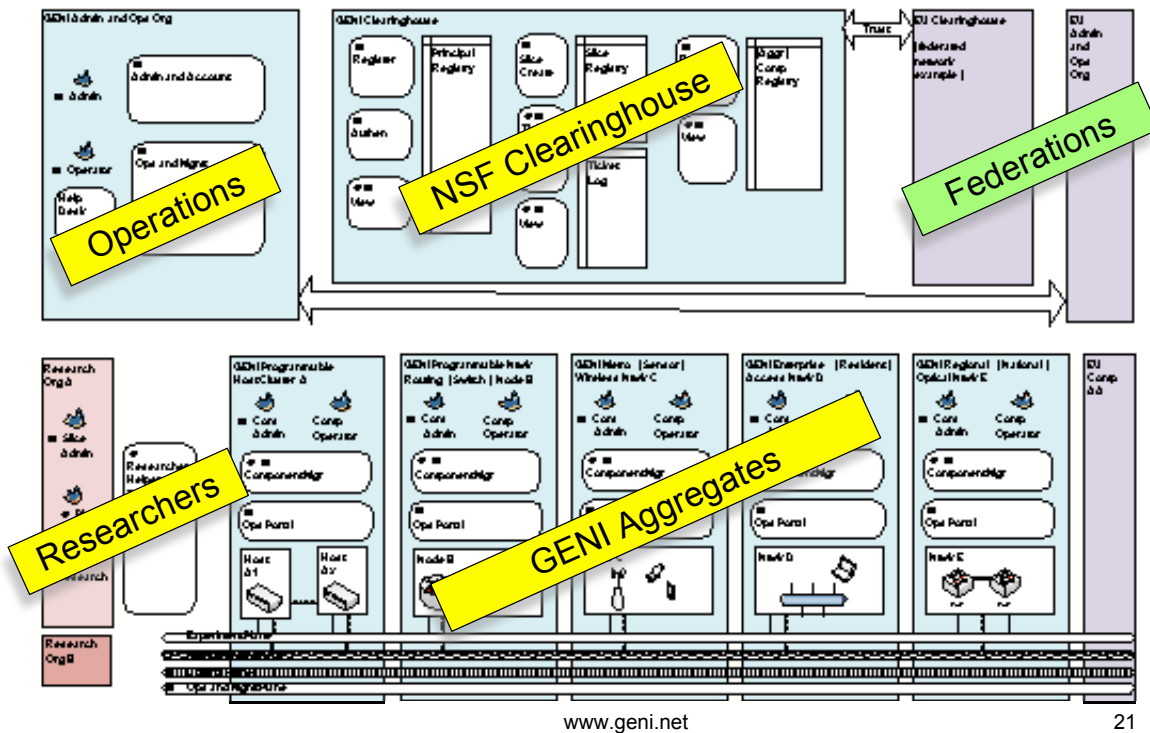
www.geni.net

20



GENI System Decomposition (simplified)

Engineering analysis drives Spiral 1 integration



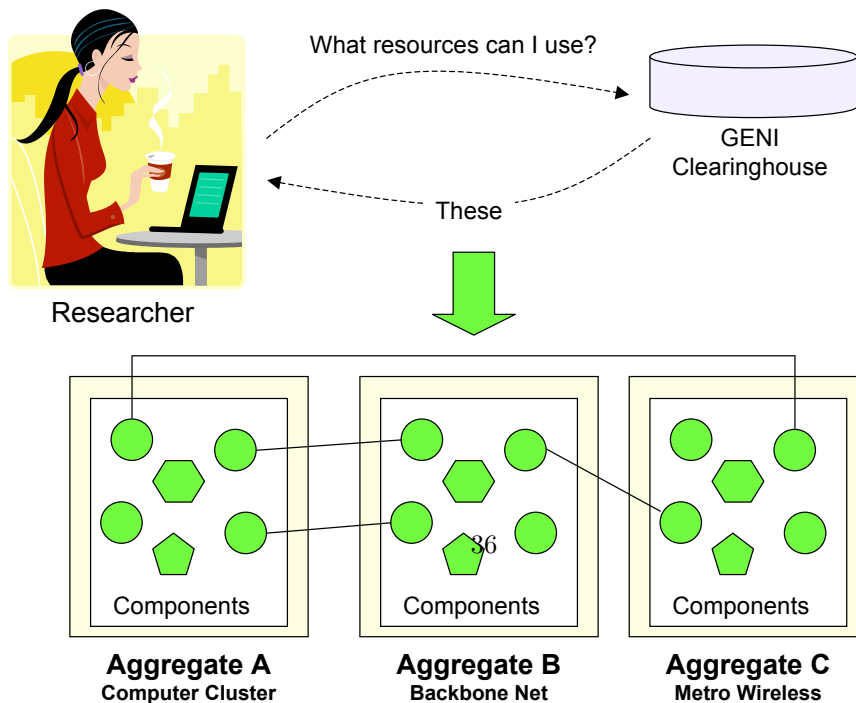
www.geni.net

21



Resource discovery

Aggregates publish resources, schedules, etc., via clearinghouses



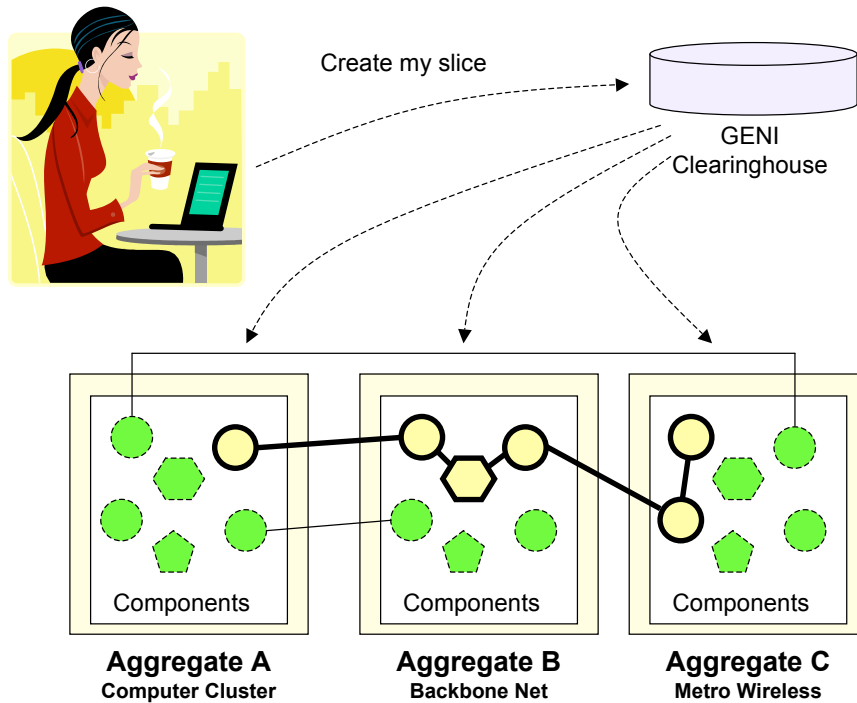
www.geni.net

22



Slice creation

Clearinghouse checks credentials & enforces policy
Aggregates allocate resources & create topologies

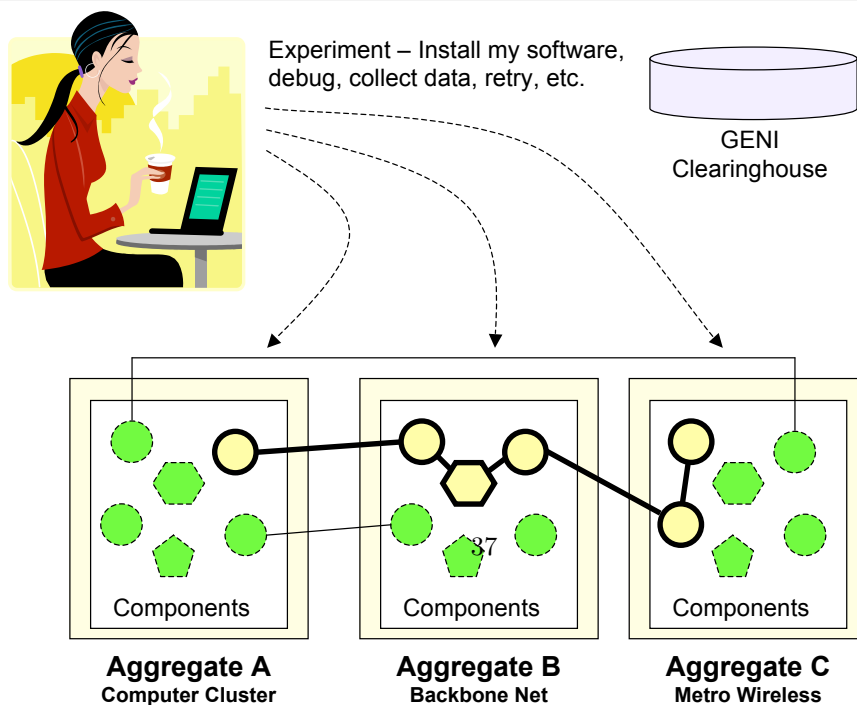


23



Experimentation

Researcher loads software, debugs, collects measurements

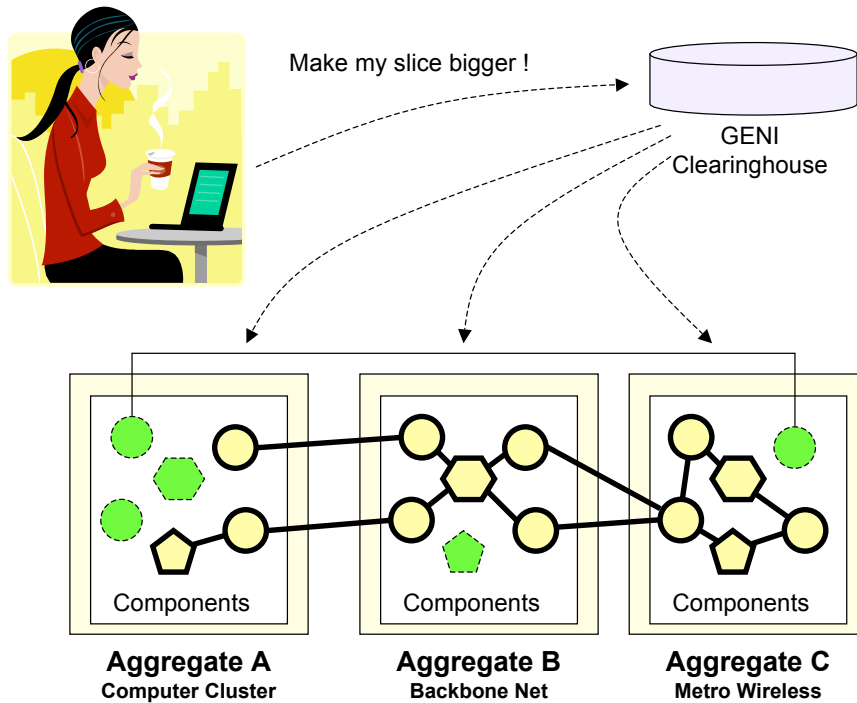


24



Slice growth & revision

Allows successful, long-running experiments to grow larger



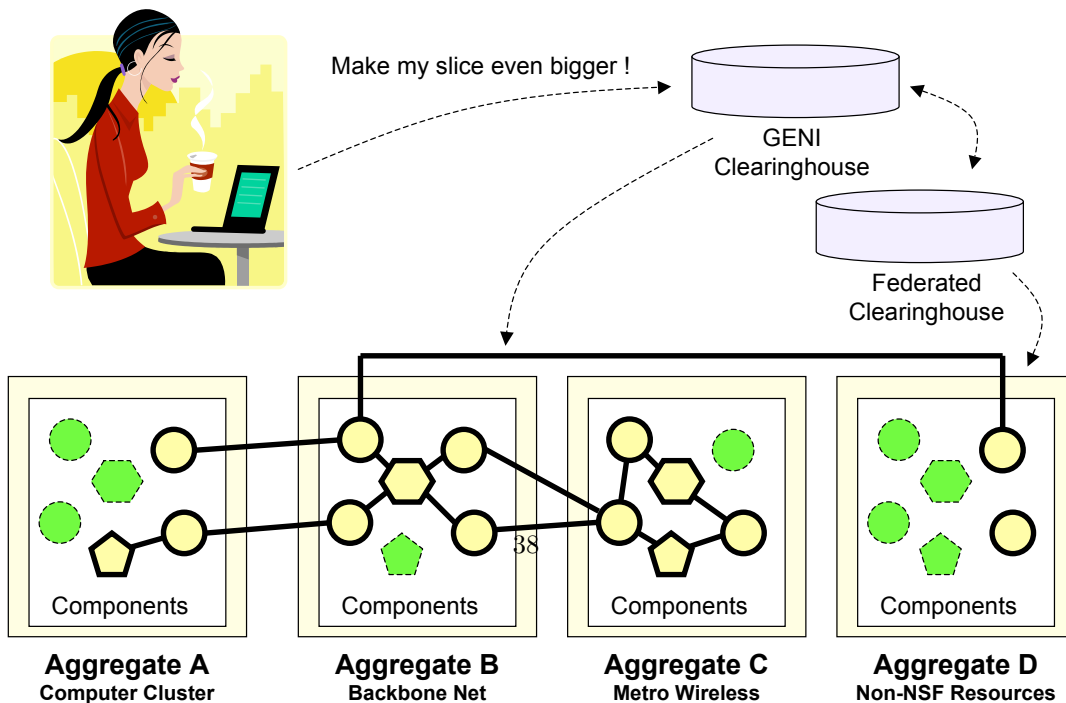
www.geni.net

25



Federation of Clearinghouses

Growth path to international, semi-private, and commercial GENIs



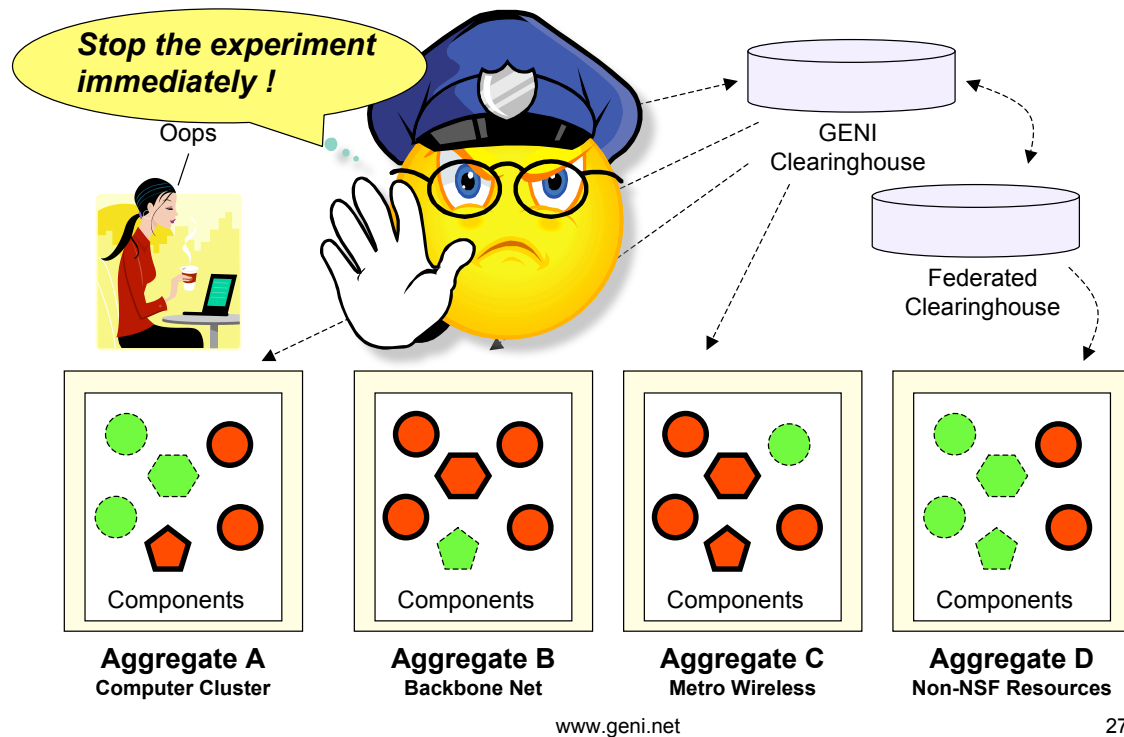
www.geni.net

26



Operations & Management

Always present in background for usual reasons
Will need an 'emergency shutdown' mechanism



Outline

- What is GENI?
- How we'll build it, how we'll use it (Two Comic Books)
- The GENI system concept
- GENI Spiral 1
- How can you participate?



GENI Spiral 1 has now begun!

First results expected in 6-12 months

GENI Project Office Announces \$12M for Community-Based GENI Prototype Development

July 22, 2008

The GENI Project Office, operated by BBN Technologies, an advanced technologies solutions firm, announced today that it has been awarded a **three year grant worth approximately \$4M a year** from the US National Science Foundation to perform GENI design and risk-reduction prototyping.

The funds will be used to contract with **29 university-industrial teams** selected through an open, peer-reviewed process. The first year funding will be used to **construct GENI Spiral 1, a set of early, functional prototypes** of key elements of the GENI system.

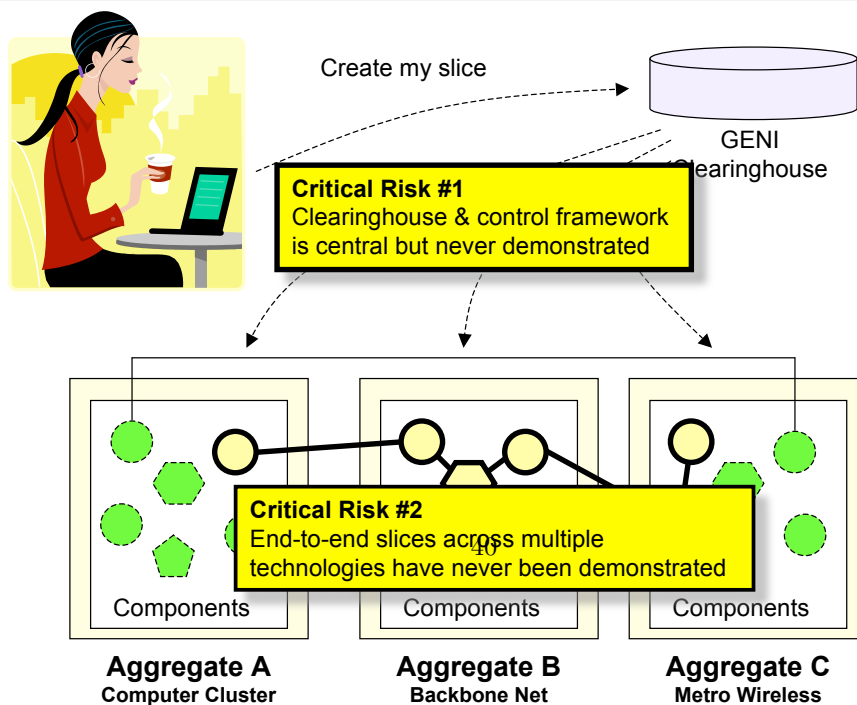
www.geni.net

29



GENI's Critical Technical Risks

These risks drive the Prototyping Goals for GENI Spiral 1



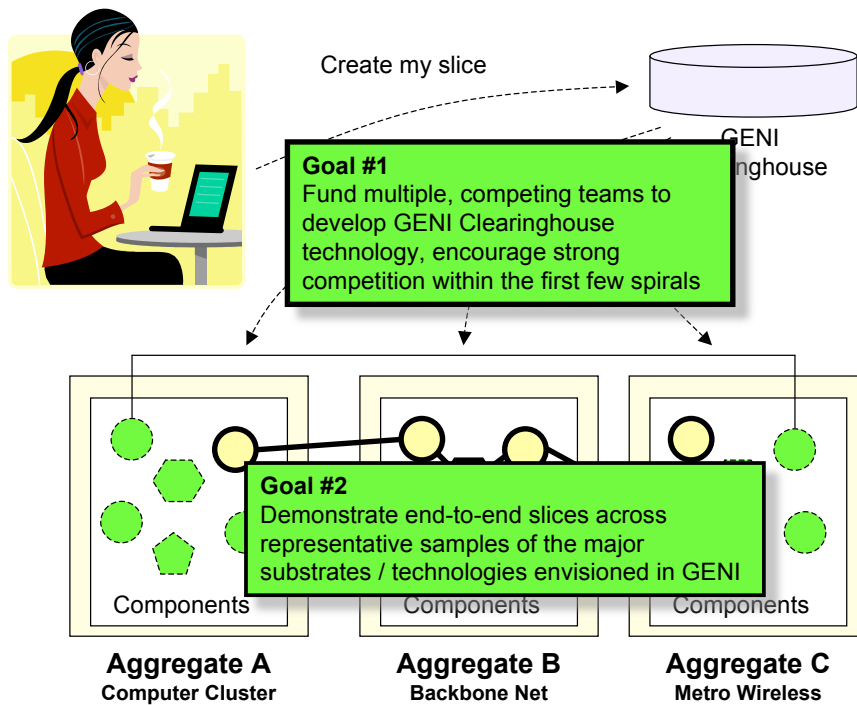
www.geni.net

30



Key Goals for GENI Spiral 1

Drive down the critical technical risks in GENI's concept

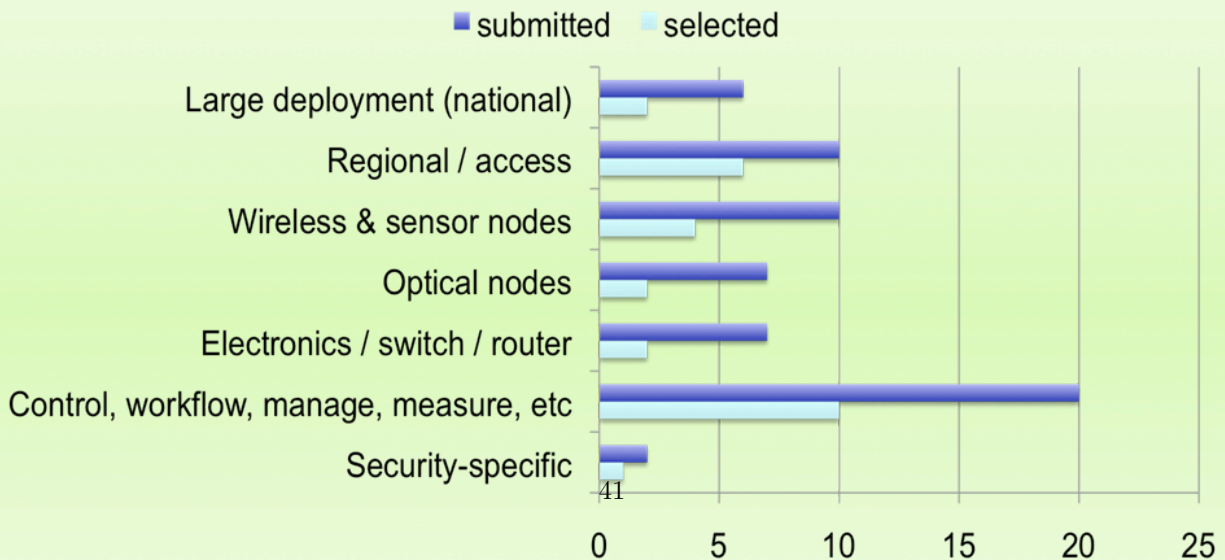


www.geni.net

31



1st GENI Solicitation – proposal areas



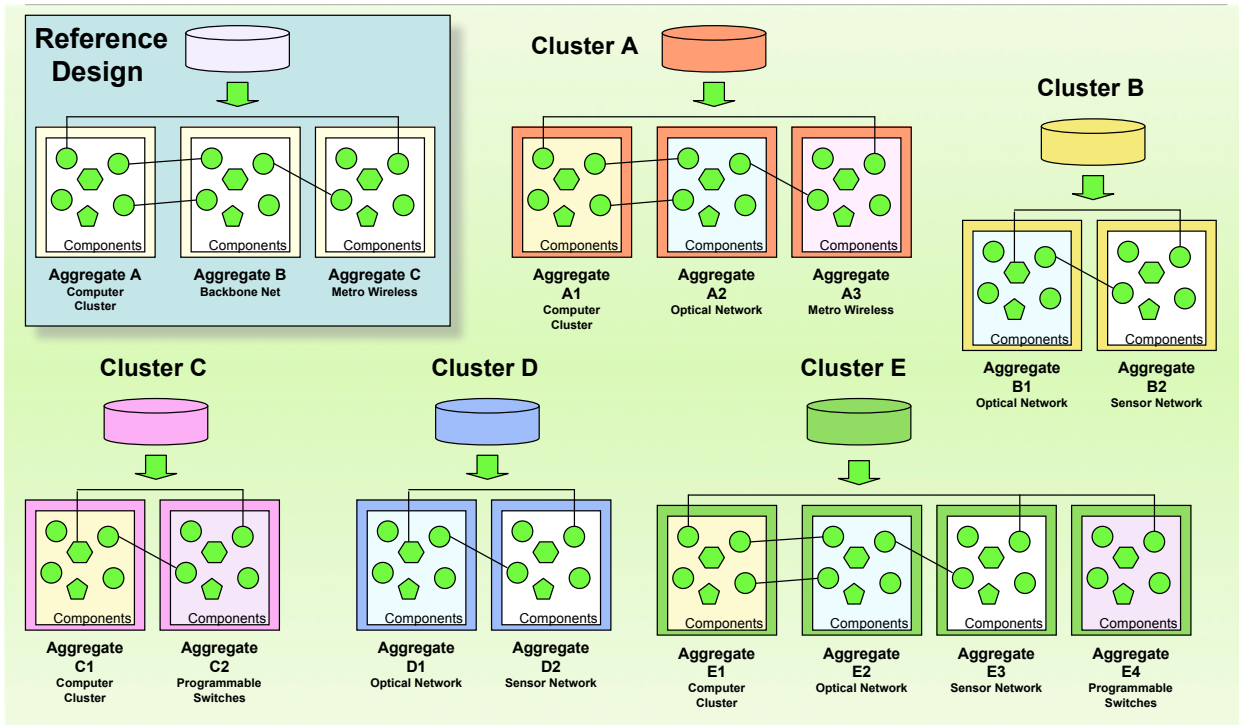
www.geni.net

32



Spiral 1 integration and trial operations

Five competing control frameworks, wide variety of substrates



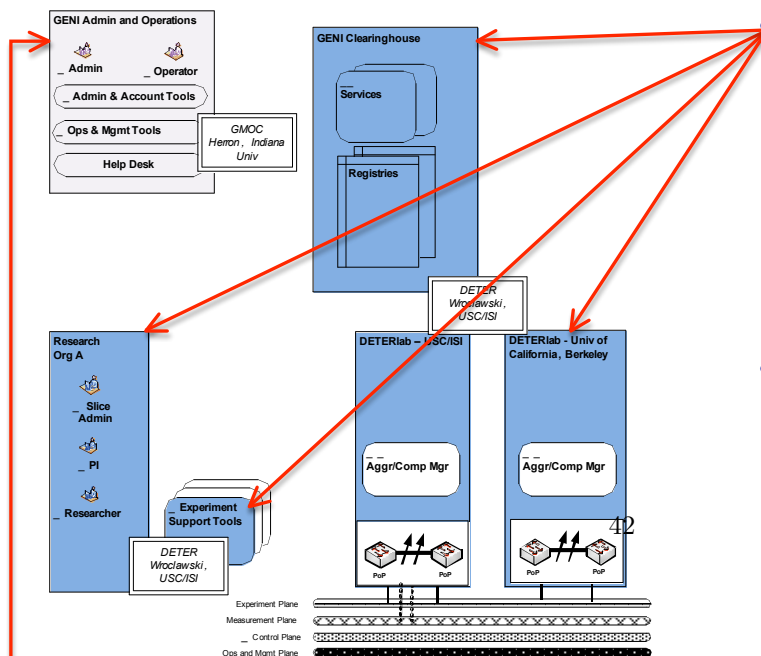
www.geni.net

33



Cluster A Integration

(uses TIED/DETER control framework)



DETER Trial Integration

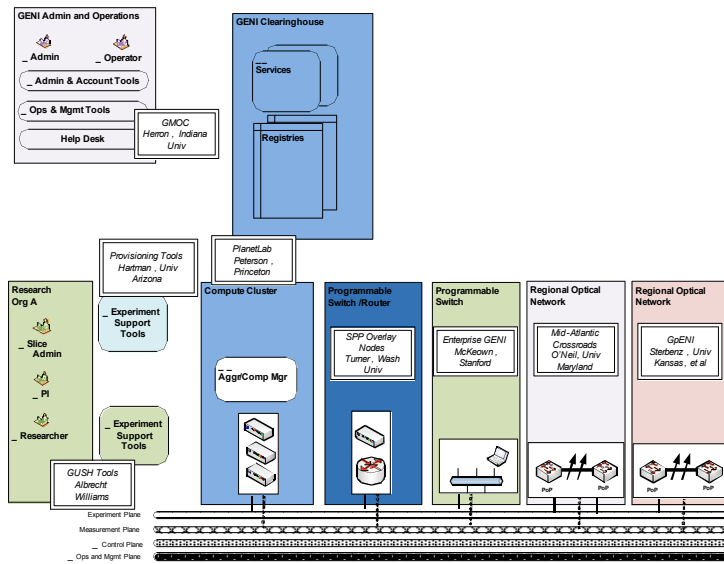
- DE TER security testbed
- Emphasis on federation
- Clearinghouse, CM
- 100+ nodes at ISI, UC Berkeley

GMOC

- Global Research NOC (Indiana)



Cluster B Integration (uses PlanetLab control framework)



- PlanetLab
 - Clearinghouse, CM
 - 800+ nodes
 - VINI (virtual topologies)
- Enterprise GENI
 - GENI VLANs on enterprise nets
- SPP Overlay Nodes
 - Programmable routers
- GUSH Tools
 - Experiment design tools
- Provisioning Service
 - Slice & experiment management tools
- Mid-Atlantic Crossroads
 - Regional network with VLAN control plane
- GpENI
 - Regional network with sliceable optics & routers
- GMOC

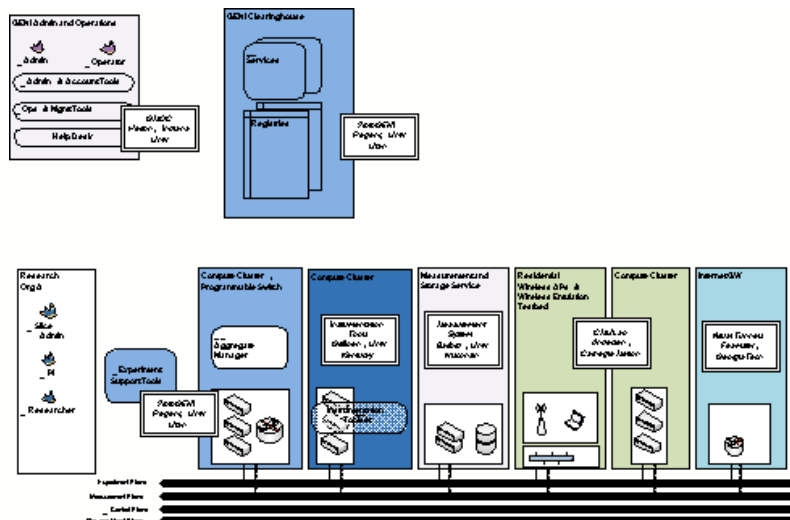
GEC3

www.geni.net

35



Cluster C Integration (uses ProtoGENI/Emulab Control Framework)



- ProtoGENI
 - Clearinghouse, CM
 - Emulab resources
 - (370+ nodes)
- CMULab
 - Home Wireless APs
 - Emulab cluster
 - Wireless emulation testbed
- Instrumentation Tools
 - UK Edulab (compute/store)
- Measurement System
 - GIMS prototype
- Virtual Tunnels
 - Dynamic tunnel tools
 - BGP distribution tools
- GMOC

43

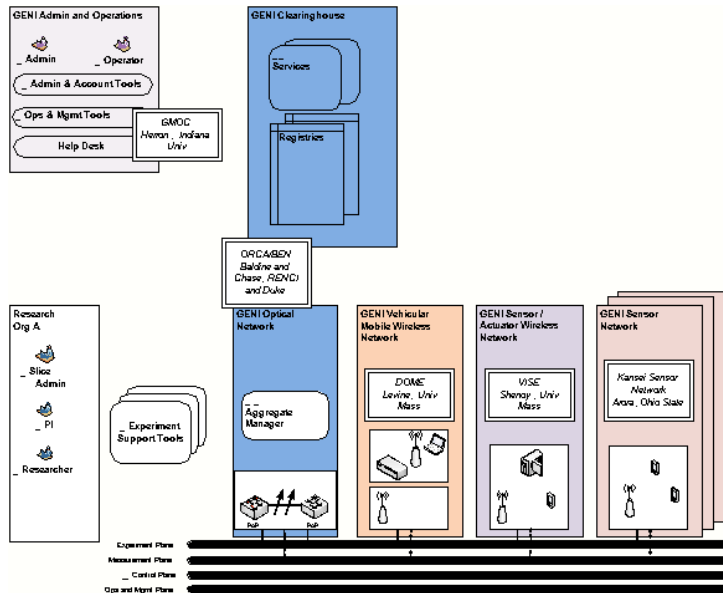
GEC3

www.geni.net

36



Cluster D Integration (uses ORCA Control Framework)



- ORCA/BEN
 - ORCA resource leasing software
 - Metro-Scale Optical Testbed (BEN)
- VISE
 - CASA (radar, video, weather sensors)
- Kansei Sensor Network
 - Wireless sensor network arrays
 - 3 federated sites each w/~100 sensor nodes
- Diverse Outdoor Mobile Environment (DOME)
 - Programmable nodes with radios on city busses
- GMOC

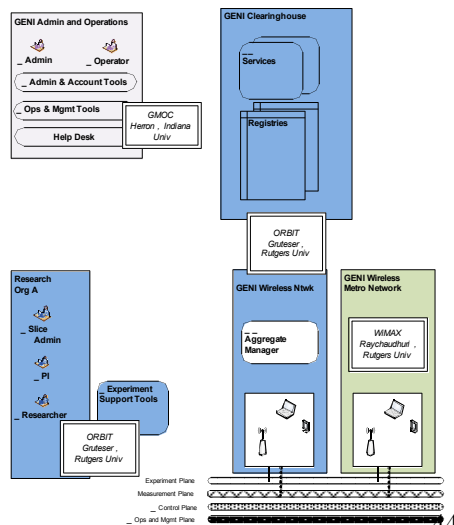
GEC3

www.geni.net

37



Cluster E Integration (uses ORBIT control framework)



- ORBIT
 - Heterogeneous testbed control, management, & measurement software
 - WINLAB wireless testbeds resources (400+ sensor nodes)
 - NICTA (Australia) wireless outdoor traffic testbed
- WiMAX
 - Open, programmable WiMAX base station
- GMOC

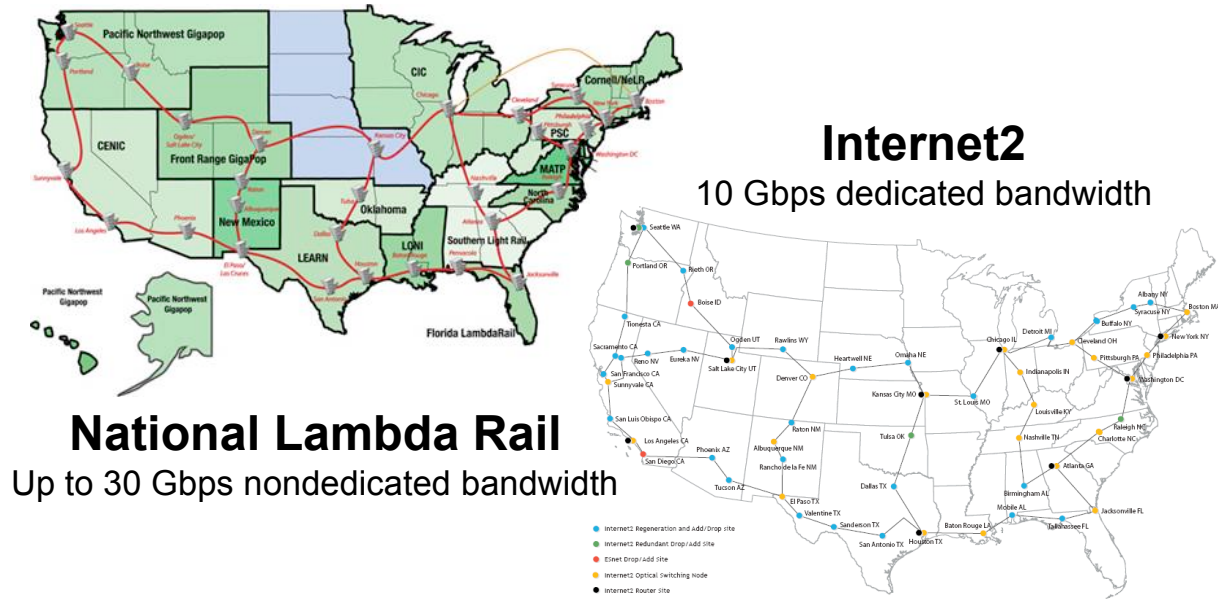
GEC3

www.geni.net

38



Generous Donations to GENI Prototyping Internet2 and National Lambda Rail



40 Gbps capacity for GENI prototyping on two national footprints to provide Layer 2 Ethernet VLANs as slices (IP or non-IP)

www.geni.net

39



Currently in the works Prototyping GENI through campuses

- August Meeting at O'Hare
 - Thanks to EduCause (Mark Luker, Garret Sern)
 - Stimulated by Larry Landweber
- CIOs from 11 major research universities
 - Berkeley, Clemson, GA Tech, Indiana, MIT, Penn State, Rice, U. Alaska, UIUC, UT Austin, U. Wisconsin
- Discussions of representative GENI prototypes
 - Nick McKeown, Stanford (OpenFlow)
 - Arvind Krishnamurthy, UW (Million Node GENI)
 - GPO Staff
- Near-term GENI / CIO activities
 - How to “GENI-enable” campus IT infrastructure
 - Coordinated policy for handling side-effects of network research (Larry Peterson, Helen Nissenbaum)



GENI Spiral 1

- Provides the very first, national-scale prototype of an interoperable infrastructure suite for Network Science and Engineering experiments
- Creates an end-to-end GENI prototype in 6-12 months with broad academic and industrial participation, while encouraging strong competition in the design and implementation of GENI's control framework and clearinghouse
- Includes multiple national backbones and regional optical networks, campuses, compute and storage clusters, metropolitan wireless and sensor networks, instrumentation and measurement, and user opt-in
- Because the GENI control framework software presents very high technical and programmatic risk, the GPO has funded multiple, competing teams to integrate and demonstrate competing versions of the control software in Spiral 1

Nothing like GENI has ever existed; the integrated, end-to-end, virtualized, and sliceable infrastructure suite created in Spiral 1 will be entirely novel.



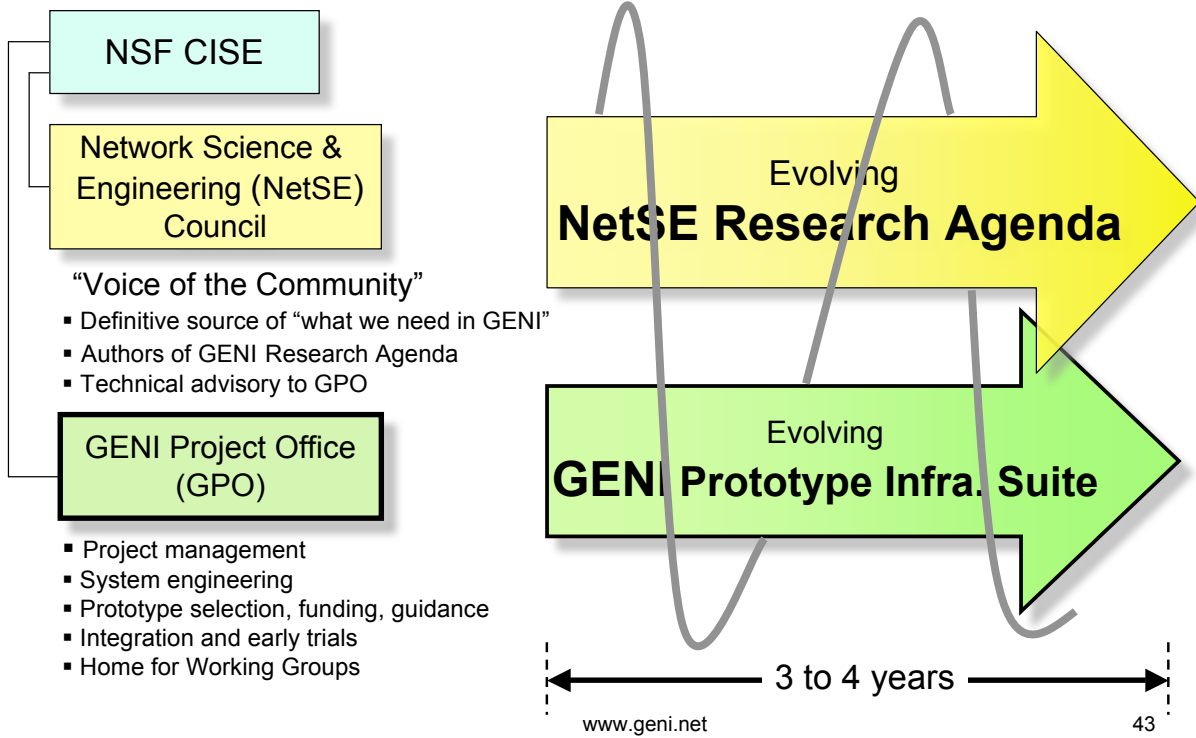
Outline

- What is GENI?
- How we'll build it, how we'll use it (Two Comic Books)
- The GENI system concept
- GENI Spiral 1
- How can you participate?



GENI in Context

Supports the Evolving NetSE Research Agenda



NetSE Council



Ellen Zegura (Chair) Tom Anderson (UW) Joe Berthold (Ciena) Charlie Catlett (Argonne) Mike Dahlin (UT Austin) Chip Elliott (GPO)



Joan Feigenbarum (Yale) Stephanie Forrest (UNM) Jim Hendler (RPI) Michael Kearns (U.Penn) Ed Lazowska (UW) Peter Lee (CMU)



And not shown . . .

Roscoe Giles
Helen Nissenbaum



GENI is being Designed & Built by the Community Via an Open, Transparent, & Fair GPO Process

- All design, prototyping, & construction will be performed by the research community (academia & industry)
- Openness is emphasized
 - Design process is open, transparent, and broadly inclusive
 - Open-source solutions are strongly preferred
 - Intellectual property is OK, under no-fee license for GENI use
- GPO will be fair and even-handed
 - BBN brings no technology to the table
 - BBN does not intend to write any GENI software, nor does it envision bidding on any prototyping or construction activities (but “never say never”)
 - If BBN does create any GENI technology, it will be made public at no cost

www.geni.net

45



Working Groups drive GENI's Technical Design Meet every 4 Months to Review Progress Together

- **Working Groups**, open to all
 - The locus for all GENI technical design
 - Patterned on the early IETF
 - Discuss by email, create documents, meet 3x per year in person
 - Each led by Chair(s), plus a professional System Engineer
- **GENI Engineering Conferences**, open to all who fit in the room
 - Held at regular 4-month periods
 - Held on / near university campuses (volunteers?)
 - All GPO-funded teams required to participate
 - Systematic, open review of each Working Group status (all documents and prototypes / trials / etc.)
 - Also time for Working Groups to meet face-to-face
 - Results in prioritized list for next round of prototype funding areas (priorities decided by NetSE and GPO)

www.geni.net

46



GENI Working Groups (WGs)

Open to all, participate via **geni.net** email lists

Shaded areas pose major security / privacy challenges

- **Substrates**
All hardware, real-estate, facilities, etc., required for the GENI infrastructure suite (including optical networks, wireless, computers, etc.)
- **Control Framework with Federation**
Written definitions of the core GENI mechanisms for providing experimental control of a node or collection of nodes. The very earliest version must incorporate federation.
- **Experiment Workflow**
Tools and mechanisms by which a researcher designs and performs experiments using GENI. Includes all user interfaces for researchers, as well as data collection, archiving, etc.
- **User Opt-In**
How do “real users” (not researchers) participate in GENI experiments. Includes both mechanisms and considerations such as privacy, etc.
- **Operations, Management, Integration, and Security**
How do operators provision, operate, manage, and trouble-shoot GENI? Includes all mechanisms for integrating and securely operating the GENI infrastructure suite.



GENI Engineering Conferences

Meet every 4 months to review progress together

- **4th meeting March 31-April 2, 2009, Miami, open to all**
 - Team meetings, integrated demos, Working Group meetings
 - Also discuss GPO solicitation, how to submit a proposal, evaluation process & criteria, how much money, etc.
 - **Travel grants** to US academics for participant diversity
- **Subsequent Meetings, open to all who fit in the room**
 - Held at regular 4-month periods
 - Held on / near university campuses (volunteers?)
 - All GPO-funded teams required to participate
 - Systematic, open review of each Working Group status (all documents and prototypes / trials / etc.)
 - Also time for Working Groups to meet face-to-face
 - Discussion will provide input to subsequent spiral goals



GPO Solicitations

Academic-industrial teams favored but not required

- **Second solicitation active – proposals due Feb. 20 !**

- What kinds of proposals do we solicit?
 - Analyses & idea papers
 - Prototypes of high-risk GENI technology
 - Integrations and trials of prototypes
- How are proposals judged?
 - Merit review
 - Joint academic / industrial teams are favored but not required
 - Open source will be favored but not required
(IP licenses on www.geni.net)

www.geni.net

49



GENI Solicitation 2 – Proposals due Feb. 20

- Overview
 - Solicitation issued December 2008
 - Proposals due February 20, 2009
 - Total funds ~ \$3.5 M / yr for 3 years, as always subject to availability of funds
 - Existing / new GENI participants both welcome
- Strong preference given to . . .
 - Joint Academic / Industrial teams
 - Active participation of campus / regional infrastructure providers (e.g., letter from campus CIO)
- Main solicitation interests
 - Security design and analysis for GENI
 - Experimental workflow prototypes
 - Instrumentation and measurement prototypes
 - Early tries at international federation
 - Other good ideas

www.geni.net
Solicitation and background information

www.geni.net

50



GENI is a Huge Opportunity

- **GENI is an unbelievably exciting project for the community**
 - Our research community has changed the world profoundly. GENI opens up a space to do it again.
- **We believe the whole community will build GENI together**
 - Our vision is for a very lean, fast-moving GPO, with substantially all design and prototyping performed by academic and industry research teams.
- **GENI Spiral 1 is now underway !**
 - within a GENI project framework that is open, transparent, and broadly inclusive.

www.geni.net

Clearing house for all GENI news and documents

3.3 Experimentation with Network-Based Security Mechanisms: George Kesidis

Experimentation with network-based security mechanisms

GENI Security Workshop

January 22-23, 2009
UC Davis

G. Kesidis

EE and CSE Depts

Penn State

kesidis@enr.psu.edu

George Kesidis, Penn State University
GENI Security Workshop 01/22/09

1

Outline

- GENI experimental context.
- Experimental progression as part of a generic engineering design cycle.
- Theoretical/formal phase.
- Simulation/emulation phase.
- Prototypical deployment.
- Specific examples and component problems, *e.g.*,
 - Experimental scale-down, and
 - Traffic generation.

53

GENI experimental context

- In the following, we are considering:
 - a network-based security mechanism under test in the context of
 - a “clean slate” network architecture.
- A security experiment would therefore need to specify:
 - a network topology (open or closed) including peripheral end-systems,
 - background and attack traffic,
 - a network architecture spanning:
 - addressing/packet-format,
 - name resolution,
 - routing/forwarding,
 - and possibly layer-3 protocols for connection establishment, authentication, *etc.*,
 - and the security mechanism under test (possibly implicitly part of the network architecture).

Generic engineering design cycle

1. Device conception/design.
2. Formal/theoretical evaluation based on models of the designed device/system and its operating conditions.
3. Testing with increasingly greater realism and cost:
 1. Simulation
 2. Emulation
 3. Prototypical deployments
 - Redesigns possible after each test phase.
 - Each test phase should be conducted and documented so as to be “repeatable” by a third party, to within assessed statistical confidences in the performance and complexity metrics.
 - Each test phase may involve consideration of:
 - Presence of and interoperation with competitive devices/systems,
 - Incremental deployment strategies to improve rate of adoption,
 - Assessment of a “control” device/system and comparison against the competition.

Different perspectives on theoretical/formal study

- “Unfortunately, understanding network performance is more of an art than a science. There is little underlying theory that is actually of any use in practice. The best we can do is give rules of thumb gained from hard experience and present examples taken from the real world.” from A.S. Tanenbaum. Computer Networks, 3rd Ed. Prentice Hall, 1996, p. 555,556.
- V. Paxson and S. Floyd, “Wide-Area Traffic: The Failure of Poisson Modeling”, *ACM/IEEE ToN*, 1995.
- J. Cao, W.S. Cleveland, D. Lin and D.X. Sun, “Internet traffic tends toward Poisson and independent as the load increases”, in *Nonlinear Estimation and Classification*, Springer, 2002.
- My own experience is that theoretical/formal performance evaluation, consciously conducted in highly idealized and simplified network settings, are pursued and valued by industry.

Trace-based traffic replay,

- Attack traffic recreation.
 - In a network setting, might not need to recreate the host exploit and thereby avoid containment issues.
 - How to develop “variations” of known attacks to test the robustness of a defense, while avoiding the perception of developing new attacks.
- Background traffic recreation is obviously important for assessment of false positives.
 - Far more activity in the research community on forensics than on employment of network trace traffic traces for testing purposes.
 - Dissemination of anonymized traces greatly improved through the activities of, e.g., CAIDA and PREDICT.
- Methods of realistically “light” salting of traces by, e.g., cover low-intensity attack activity (b/g traffic rarely captured with interesting attacks *in situ*).
- Need to characterize session-level “demand” from traces motivated by the need to experiment with: 55
 - high volume attacks, and
 - new network architectures (even just different layer 4).

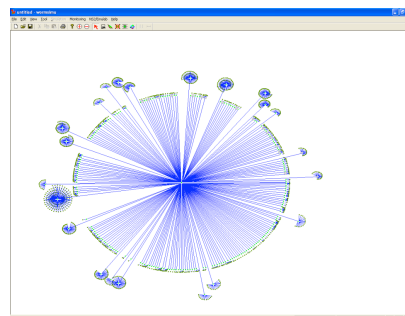
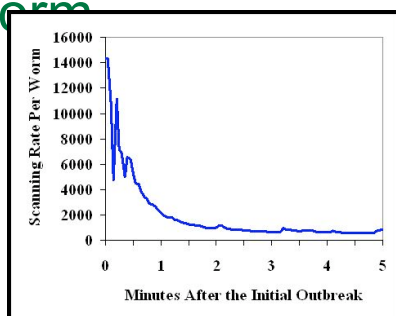
Scale-down for theoretical, simulation or emulation testing phases

- Given limited resources for simulation and emulation, may need to reduce the scale of the experiment by using a much smaller “open” or “closed” network.
- Scale-down of an open topology as in Thevenin-Norton equivalent circuits.
- Clearly, also need metrics to assess fidelity of scaled-down model to the original.
- Some preliminary results by the DHS/NSF EMIST team for scale-down of
 - attack traffic (scanning worms), and
 - inter-AS network topologies (attacks targeting BGP).

George Kesidis, Penn State University
GENI Security Workshop 01/22/09

7

Example: 128:1 scaled-down SQL Slammer Worm



- Characteristic outbound scan-rate saturation, but total attack traffic negligible compared to core background traffic.
- A SIR model of worm spread recreated this characteristic saturation [TOMACS'08].
- 128:1 scale-down experiment on DETER [WORM'04]:
 - Idealized Internet core connecting access (stub) links
 - 600 susceptible SQLs residing behind <1000 stubs
 - Stub links to core are distinct access (bandwidth) bottlenecks
- Need not emulate actual method of host infection, but can vary scanning strategy, see EMIST's scanning worm tool on DETER experimenter's dashboard.
- EMIST's development of tools for experiment specification, visualization, scale-down, traffic generation, etc., was potent outreach activity.

George Kesidis, Penn State University
GENI Security Workshop 01/22/09

8

Incentives and security

- Network trace data can inform “utilities” modeling user behavior – important for games studying effects of economic incentives.
- As a result of the network neutrality ruling by the FCC, renewed interest in incentives to deter excessive consumption (BitTorrent) under flat rate plans.
- Comcast residential broadband access recently migrating from flat rate F toward pricing formula involving usage-based charges above a threshold (*i.e.*, overages, as commonly used at NNI): $F + R(U-\Theta)^+$
- Such a pricing formula naturally leads to an authentication problem, renewed interest in differentiated services, *etc.*
- Examples many such mechanisms have already been standardized and are already deployed in the commodity Internet, *e.g.*, CMTS DOCSIS, AT&T’s DSL U-verse, MIDCOM, and a lot of “traffic engineering” (TE) technology including diffserv, scheduling, Ethernet (802.1p).

Prototypical deployment - GENI

- I understood GENI’s original story as a testbed intended to be able to:
 - simultaneously mount different network architectures under test,
 - somehow assess them through engaging
 - actual data/service providers (& p2p networks) which could shop data or services they want to mount among the architectures under test, and
 - actual end-users that would access the “GENI” data/services through interfaces with the existing Internet.
- Security experimentation:
 - Need for attack isolation may preclude “deliberate” attack experimentation.
 - Need to facilitate defense deployment and assessment tools.
 - Need to facilitate deployment of other types of security mechanisms to manage different types of authentications, reputation/referral systems, *etc.*
- Generally, the testbed thus conceived faces problems, *e.g.*,
 - reconfigurable routers, associated experimental artifacts, and
 - Management of experimental resource allocation and associated fairness issues.
- Again, interesting research problem of *incremental deployment strategies* to
 - maximize performance in a “hybrid” environment and, thereby,
 - maximize likelihood of growth in deployment/adoption, *i.e.*, survival of the fittest.

GENI defaults

- Availability of a “default” network architecture, or one chosen from a library, which can be modified by the experimenter.
- In particular, availability of default:
 - connection-oriented network architectures,
 - incentive systems, and
 - associated billing/book-keeping and authentication mechanisms.
- Note that modifications may need oversight so that commodity Internet does not experience unexpected problems through the GENI interface.

Acknowledgements

- EMIST(-DETER) ‘03-’07 project team, particularly
 - S.F. Wu, J. Rowe of UC Davis
 - V. Paxson and N. Weaver of ICSI/UC Berkeley
 - P. Liu and D.J. Miller of Penn State
 - S. Fahmy of Purdue
 - P. Porras of SRI
 - S. Schwab of SPARTA
 - J. Evans (NSF) and D. Maughan (DHS)
 - Tools: DETER’s experimenter’s dashboard at www.isi.edu/deter
- Talks on security experimentation at the NSF ‘08 “Science of Security” workshop by
 - Roy Maxion
 - John Mitchell

3.4 Ingredients of an Early Design for Protecting the GENI Facility: Mike Reiter

Ingredients of an Early Design for Protecting the GENI Facility

GENI Distributed Services Working Group

Tom Anderson, David Andersen, Mic Bowman, Frans Kaaskhoek, Rick McGeer, Vivek Pai, Mike Reiter, Mothy Roscoe, Ion Stoica, Amin Vahdat

Disclaimer

- This talk summarizes the early design of security mechanisms to protect against abuse of the GENI facility
 - Prior to establishment of BBN as GPO
- I have no knowledge of how this relates to the security facilities envisioned today for GENI
- In particular, I in no way⁶⁰ speak for BBN or the current state of GENI on this matter

Some Topics We Considered

- Threat model
- Goals/requirements

- Access control
- Authentication and key management
- Auditing
- Intrusion detection

Threat model

Exploitation of a slice

- Runaway experiments
 - Unwanted Internet traffic
 - Exhausting disk space
- Misuse of experimental service by end users
 - E.g., to traffic in illegal content
- Corruption of a slice
 - Via theft of experimenter's credentials or compromise of slice software

61

Exploitation of GENI itself

- Compromise of host O/S
- DoS or compromise of GENI management plane

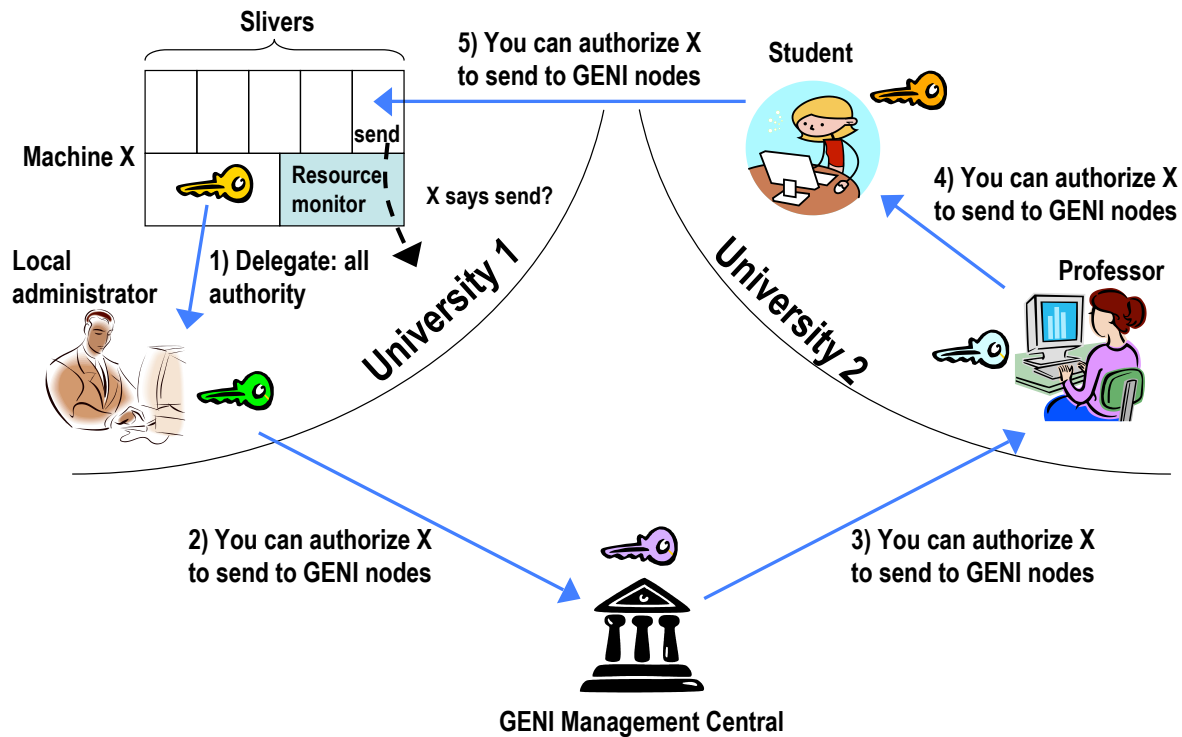
Requirements: Do no harm

- Explicit delegations of authority
 - Node owner → GMC → Researcher → students → ...
- Least privilege
 - Goes a long way toward confining rogue activities
- Revocation
 - Keys and systems will be compromised
- Auditability
- Scalability/Performance
- Autonomy/Federation/Policy Neutrality
 - Control ultimately rests with node owners, can delegate selected rights to GMC

Access Control Requirements

- Arbitrarily flexible
 - Did not want to “hard code” policy into the system
- Dynamically extensible
- Verifiably sound and principled
 - Avoid ad hoc approaches
- Auditable
 - Must be able to determine why an access was granted, and who was responsible

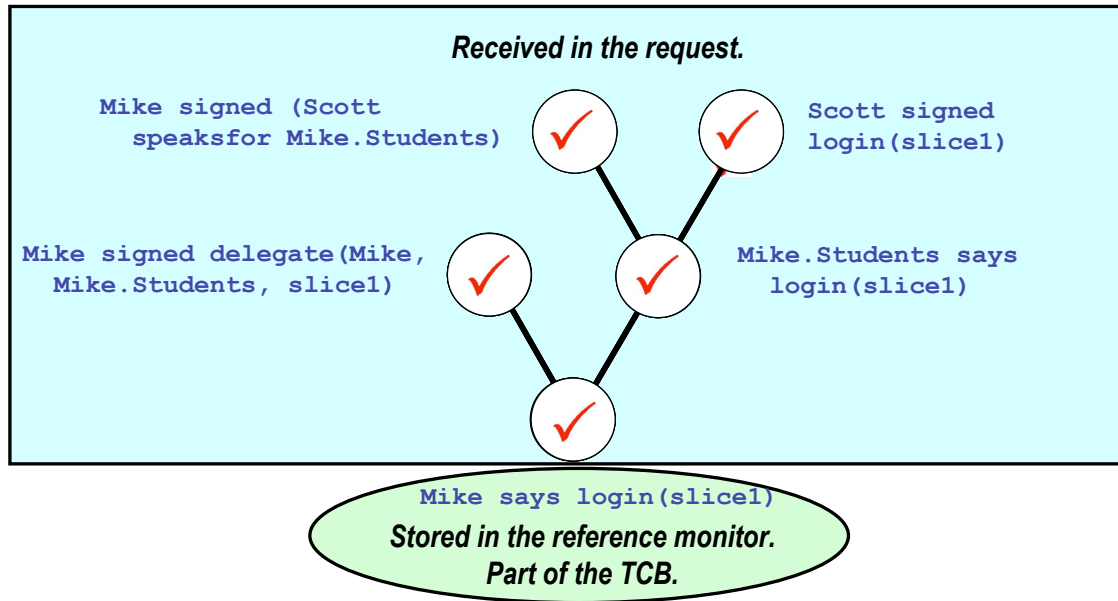
Authorization Example



A Proof-Carrying Approach

- Encode access control decision procedure in a formal logic
 - Can be used to express groups, roles, delegations, and new constructs
 - Can encode other, specific access-control mechanisms
- Digitally signed statements (e.g., certificates) used to instantiate logical statements
- Client submits a proof that its request complies with access-control policy
- Reference monitor checks that the proof is a valid proof of required policy

A Tiny Example



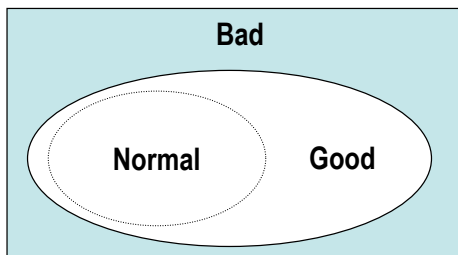
Authentication and Key Management

- GENI would have a PKI (as a corollary of the authorization framework)
 - Every principal would have a public/private key
 - ▶ E.g., users, administrators, nodes
 - Certified by local administrator
 - Keys sign certificates to make statements in the authorization logic (identity, groups, authorization, delegation, ...)
- Private key compromise an issue
 - Encrypted with user's password? Off-line attacks
 - Smart card/dongle? Most secure, but less usable
 - Capture-resilient protocols: A middle ground
 - ▶ An (untrusted) capture-protection server can disable use of a key, e.g., when observing a password-guessing attack

Intrusion Detection

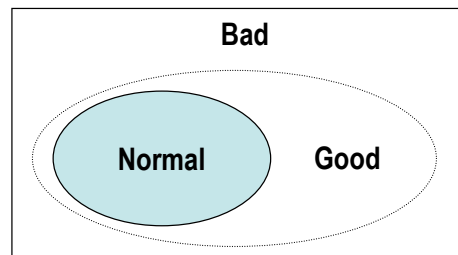
- Traditional intrusion detection methods may not suffice *for monitoring experiments*

Misuse detection
Specify bad behavior and watch for it



Problem: Experiments do lots of things that look "bad"

(Learning-based) Anomaly detection
Learn "normal" behavior and watch for exceptions

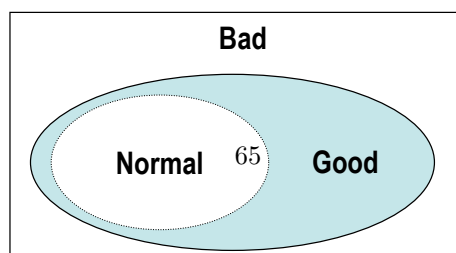


Problem: Experiments may be too short-lived or ill-behaved to establish "normal" baseline

Intrusion Detection

- Specification-based intrusion detection is more appropriate for monitoring experiments
 - Fits in naturally with authorization framework, as well

Specification-based intrusion detection
Specify good behavior and watch for violations



Audit Log Prototype: PlanetFlow

[Huang et al.]

- PlanetFlow: logs packet headers sent and received from each node to Internet
 - Enables operations staff to trace complaints back to originating slice
 - Notify experimenter; in an emergency, suspend slice
- All access control decisions can be logged and analyzed post-hoc
 - To understand why a request was granted (e.g., to give attacker permission to create a sliver)

Issues Left Open

- DoS-resistant GENI control plane
 - Initial control plane would employ IP and inherit the DoS vulnerabilities thereof
 - GENI experimentation may demonstrate a control plane that is more resistant
- Privacy of operational data in GENI
 - Could be a great source of research data
- Operational procedures and practices
 - Central to security of the facility

3.5 Some GENI Thoughts: Nicholas Weaver

Some GENI thoughts

Nicholas Weaver

International Computer Science Institute

All opinions are my own



GENI's Vision Appears to be Separate Testbeds

- GENI appears to be a group of separate facilities:
- The *Emulab*-style component:
 - Centralized cluster
 - Interconnect is as critical as processing
 - Commercial equivalent:
EC2 with VLANs
- The *Planetlab*-style component:
 - Many distributed endpoints over commodity Internet
 - Diversity of locations is the greatest asset⁶⁸
 - Commercial equivalent:
Akamai



The Emulab Component: Federation Considered Dangerous

- For small experiments, federation doesn't help
 - It will fit in a single sub-testbed
- For large experiments, federation is **dangerous**
 - Unless the experimental topology's bottlenecks match those in the federation's arrangement, the experiment **can't be mapped**
- Focus on federation distracts from the proper structure: All emulab-style components should be in a **single, centralized facility**
 - Bandwidth bottlenecks can not be **economically** removed
 - Compare the price of 10 Gbps national links to the price of 10-Gig Ethernet cables
 - Latency bottlenecks **can never be removed**



The PlanetLab Component: Needs *Many* More Endpoints

- Goal should be 2000+ end-points, worldwide, as close to the end-users as possible.
 - This requires that ISPs **want** to deploy Geni-PlanetNodes
- Develop an **application** which saves **everybody** money
 - P2P without caches is **very bad** for ISPs and customers
 - P2P **with caches** is **very good** for ISPs, customers, and content providers
- EG. A BitTorrent cache architecture
 - Troll for torrents
 - Connect only with **isp local** clients and trade cache data
 - Handle the DMCA complaints automatically
 - If the infrastructure is missing, applications still work unchanged
 - Calculate bandwidth savings on the nodes
 - A fraction of the saved bandwidth is now available for Geni-Planetlab style experiments



3.6 GENITor: Nikita Borisov

GENITor

Nikita Borisov

University of Illinois at Urbana-
Champaign



- A network for anonymous communication
- Many current users (~200M), most legitimate, some not
- Two objectives:
 - Anonymity
 - Usability -> *Performance*
- Research requires test-bed evaluation
 - Best experiments will involve real, “Opt-in” users

Tor as an Opt-In application in GENI

- Is it even legal?
- How do we protect the rest of the world?
but also...
- How do we protect the users from GENI
 - Tor experimenters (i.e., me)
 - Other experimenters (may not even be malicious)
 - GENI infrastructure

Privacy Guarantees

- Federation
 - Distribute Tor across multiple sub-units
 - Interconnect GeniTor with real Tor
- Privacy vs. measurement
 - Put some restrictions on what measurements I can collect
 - And offer credible guarantees!
 - Provide restrictions on measurements of co-hosted experiments

3.7 Adaptive Security Slice Monitoring: João W. Cangussu and Ram Dantu

Adaptive Security Slice Monitoring

João W. Cangussu
University of Texas at Dallas
Department of Computer Science
cangussu@utdallas.edu

Ram Dantu
University of North Texas
Department of Computer Science
rdantu@unt.edu

Workshop on
GENI and
Security

GLOBAL ENVIRONMENT FOR NETWORK INNOVATIONS

The Erik Jonsson School of Engineering and Computer Science

- ◆ GENI will be hosting the testing of innovative networking techniques which can bring with them a series of new flaws that could be malicious or accidentally exploited.
- ◆ Security attacks can happen intentionally when running an experiment or it can be the consequence of a series of unexpected events.
- ◆ GENI should be prepared to identify the existing attacks and it should be able to learn how to identify any new attack. It should also be able to point out the causes of the attacks.

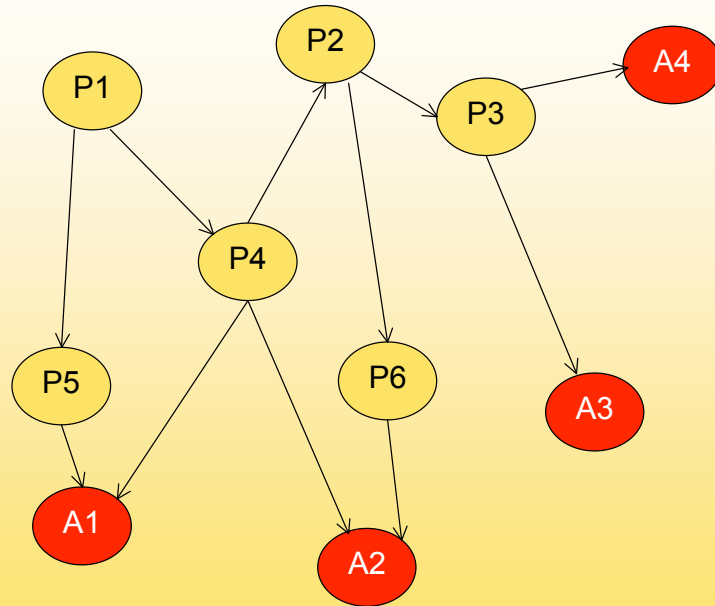
- ◆ Here we propose the use of Bayesian Belief Networks (BBNs) to model the cause-effect relationship between elements of a network experiment and associated attacks/flaws.
- ◆ Data is collected as experiments are running and used to create the structure of a BBN as well as to train it. In this way whenever an attack is identified the events leading to the attack are accounted for.
- ◆ The created BBN can be later used not only to identify the attack when other experiments are running but also to allow the computation of the probability of the attack under the occurrence of a series of events; it can sound an alarm when an attack or specific flaw is about to happen.
- ◆ The BBNs can be dynamically adapted/updated for any number of attacks and flaws, including newly discovered problems.

The Erik Jonsson School of Engineering and Computer Science

- ◆ Observations extracted from system are stored in the source nodes (system configuration based on parameters values)
- ◆ Hypotheses regarding the state of the system are stored at the terminal nodes (security state of the system)
- ◆ We can run queries to determine the probability of a given attack as well as the source of the attack.
- ◆ The BBN is automatically constructed using algorithms to create the structure of the network based on collected data.
- ◆ Training the BBN is also based on data.
- ◆ New attacks can be dynamically incorporated to the BBN.

Pi: Slice Monitored Parameters
Ai: Attacks/Security Issues

Experiments data is used to create the structure and the CPT of the BBN.



3.8 Exploiting Insecurity to Secure Software Update Systems: Justin Cappos

Exploiting Insecurity to Secure Software Update Systems

Justin Cappos

Department of Computer Science and Engineering
University of Washington

Introduction

software update system -- a piece of software that installs, updates, removes, or patches software or firmware on a device by retrieving information (**software updates**) from a trusted, external source (**repository**)

Software update systems are widely insecure [**Bellissimo HotSec 06, Cappos CCS 08**]

Software update systems are ubiquitous

But security is simple, right?

Just use HTTPS

Common errors in how certificates are handled

Online data becomes single point of weakness

... and add signatures to the software updates

Attackers can perform a replay attack

... and add version numbers to the software updates

Attackers can launch freeze attacks

But security is simple, right? (cont.)

..... and add a quorum of keys signature system for the root of trust, add signing by different compartmentalized key types, use online keys only to provide freeze attack protection and bound their trust window, etc. [\[Thandy software updater for Tor\]](#)

We still found 8 design or implementation flaws

Having each developer build their own "secure" software update system will fail

Is there a practical risk?

PlanetLab uses YUM -- updates come both from Fedora 9 and PLC

Lease a server and have it listed as an official Fedora mirror

Ensure that PlanetLab nodes contact only your mirror

Find existing exploit code for an old version of a package that isn't installed

Change the package metadata so the old version of the package is installed with any update

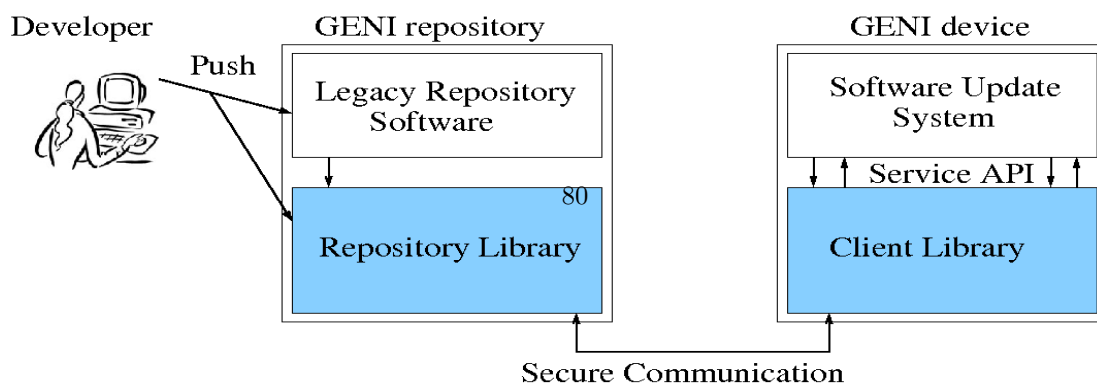
After the PlanetLab node does an update, remotely exploit it

A knowledgeable attacker can root any system on PlanetLab today!

Our approach for new systems

Build a client library that provides security for software update systems

Build a repository library that correctly signs developer updates

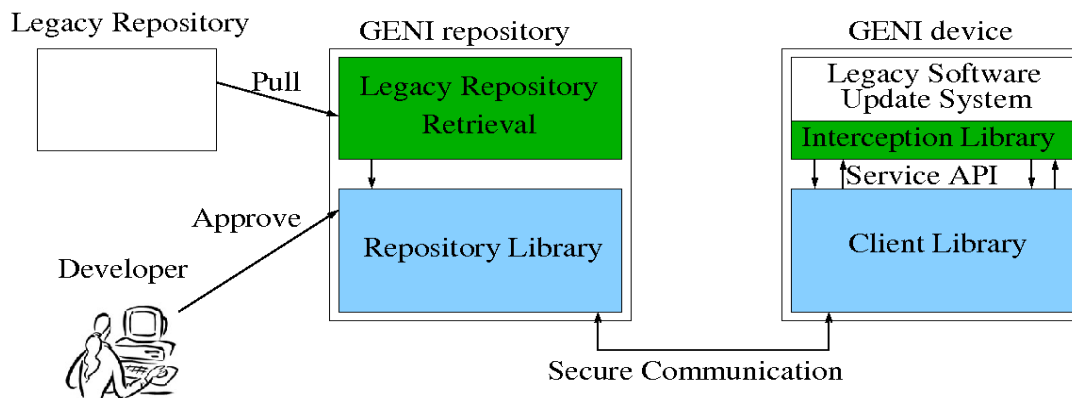


Our approach for legacy systems

Must retain functionality of existing system

Intercept traffic from insecure software update systems to transparently force it through the client library

Provide feedback to the user / system administrator



Proposal Overview

Work with the Tor project

Many pairs of eyes uncover bugs more easily

Build an artifact early, add security mechanisms gradually

Portability of the client library is key

Focus on supporting the developer / repository interface(s) used by GENI and Tor

Conclusion

Software update systems are extremely vulnerable

Subtle issues in building a secure software update system

We propose to:

- Build a library for securing software update systems

- Secure legacy systems by exploiting their insecurity

- Work with the open source community to ensure quality

Why focus on this threat?

Existing implementations are insecure [[Bellissimo 06, Cappos 08](#)]

Software update systems run as root

Traditional defenses don't protect against attacks

Ubiquitous

An attack often appears benign

Attack code can be easily reused [[EvilGrade](#)]

Trends show server attacks are on the rise [[CERT](#)]

3.9 Campus Testbed for Network Management and Operations: Nick Feamster

Campus Testbed for Network Management and Operations

Nick Feamster
Georgia Tech

Joint with Ankur Nayak, Russ Clark, Ron Hutchins, Campus OIT
Also input from Wenke Lee

Summary

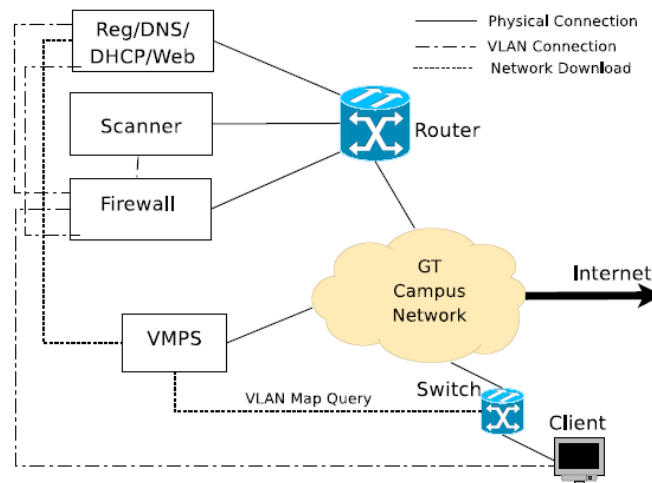
- We are building an experimental network at Georgia Tech
 - Programmable network switches (OpenFlow)
 - Multiple on-campus sites
 - Dedicated fiber between these sites
 - Upstream connectivity and IP address space (“own AS”)
- Initial testing platform for network solutions deployed on-campus
- We are building this to test our own ideas in network management and operations

Network Management Tasks

- Security-related network management tasks
 - Authentication and access control
 - Resource allocation
- Today: Many solutions require operator vigilance, hacks, magic, etc.
- We are exploring how to make these tasks easier with programmable networking

3

Access Control and Monitoring



85

- New hosts
 - Assigned to private VLAN
 - Given private IP address space
 - Authenticated and scanned

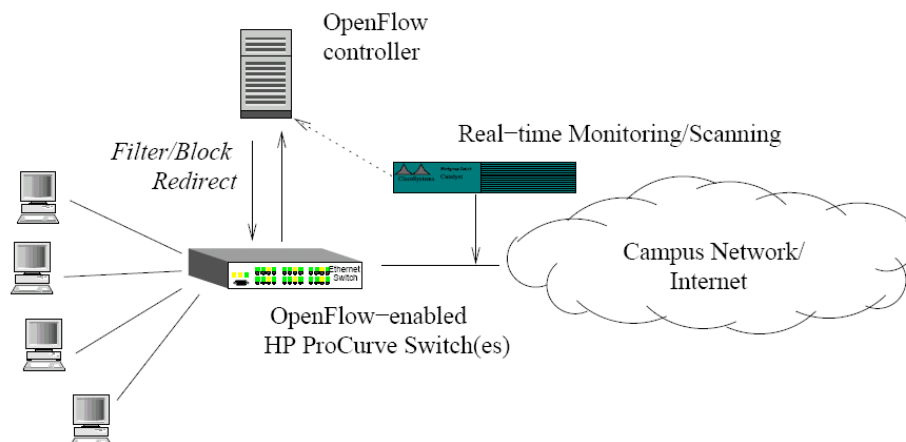
4

Problems with Current Architecture

- Access control is too coarse-grained
 - All unauthenticated/unscanned hosts are on the same subnet
 - Hosts with access are all on the same VLAN
- Lack of dynamism
 - Hosts cannot be dynamically remapped
- Monitoring is not continuous
 - Reaction to alarms is manual

5

Simplify/Enhance: Programmable Networks



86

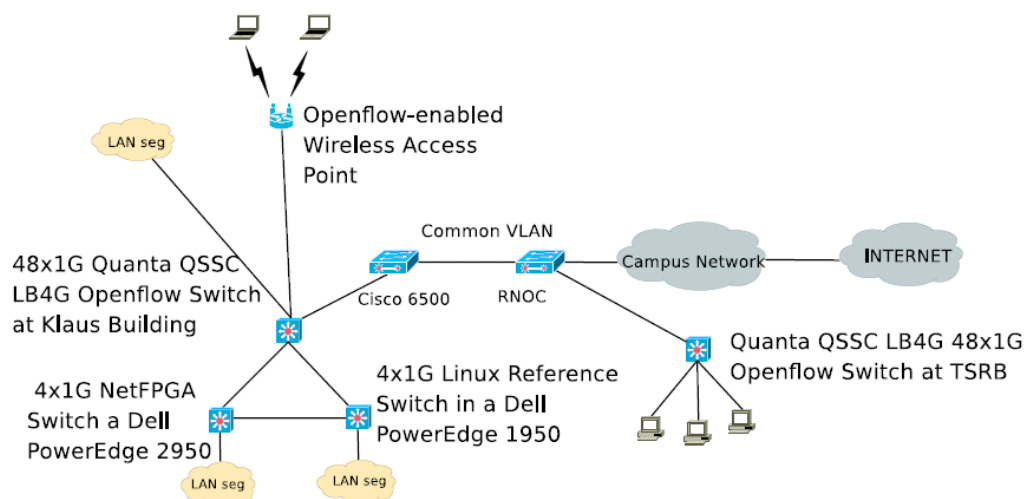
- Flow-table entries in switches redirect hosts to gardenwall
- Traffic is remapped with flow table entries per-host
- Continuous, real-time monitoring integrated with controller₆

“Outsourcing” Network Management

- Lots of independently operated networks
 - Each with view of network traffic
 - Including home networks (a known large source of unwanted traffic)
- Lots of distributed inference algorithms
 - SpamTracker
 - BotMiner
- **What if these networks had programmable switches?**
 - Use output from distributed inference to control network elements across many networks

7

Current Campus Testbed



87

- Space for running real-world projects and applications
- Need: Ability to “re-enact” network events

8

Looking Forward

- Campus-wide deployment
 - Network has 275 switches for access control that *can* run OpenFlow today
 - Firmware upgrade scheduled for Spring 2010
- Big questions
 - Sharing between production network and research
 - Connectivity to other campuses
 - Integration with measurement?

3.10 GENI as an Infrastructure to Study Malicious Overlay Networks: Wenke Lee

GENI as an Infrastructure to Study Malicious Overlay Networks

Wenke Lee

Georgia Institute of Technology

Goals

- Use GENI as a large-scale distributed test-bed for security research
 - The best we can get if we can't experiment on the real Internet
- Leapfrog our ability to understand large-scale malicious networks (botnets) and predict their future trends
 - Essential properties of botnets, how botnets must rely on core network services, trade-offs of botnet design considerations, etc.
- Evaluate botnet detection and removal technologies

A New Look at Botnets

- Analyze essential properties of botnet lifecycle
 - E.g., botnets are valuable, long-term resources
- Derive *axioms* that directly follow from the properties
 - E.g., botnets need to have *agility* to evade detection and removal
- Derive *theories* from the axioms
 - E.g., a particular kind of botnet structure has better *network agility* than the others
 - E.g., by detecting and neutralizing the sources of *network agility*, we can limit botnets' evasion capabilities and thus make botnets easier to detect and remove
- Apply the theories to *practice*
 - E.g., what are the ways that network agility can be realized?
 - E.g., an on-line detection of naming (DNS) based agility.

An Experimental Approach

- Experiment with design and deployment, as well as detection and removal of botnets on GENI, e.g.,
 - design various types of botnets – topology structures, characteristics/values of essential properties, etc.
 - deploy these botnets – measure their propagation speed, size, aggregate⁹¹attack power, etc.
 - evaluate detection and removal techniques

3.11 Gacks: Secure Resource Allocation for GENI: John Hartman

Gacks

Secure Resource Allocation for GENI

John H. Hartman
University of Arizona

Scott Baker
SB Software

Justin Cappos
University of Washington

Larry Peterson
Princeton University



Overview

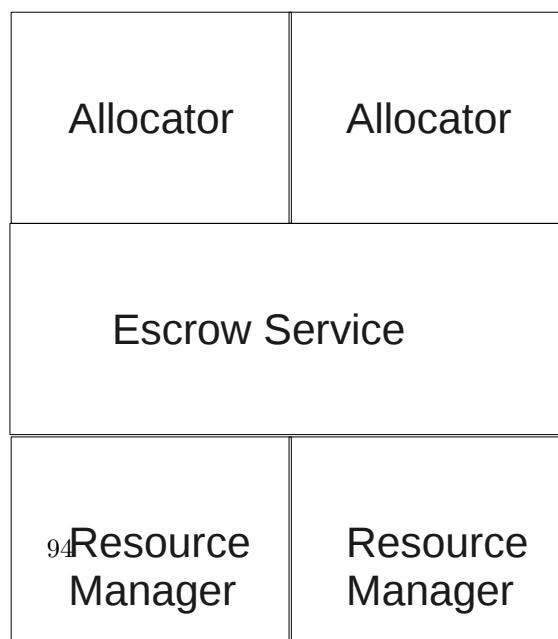
- Secure binding of resources to slices
- Infrastructure to support a variety of resource allocation policies (e.g. auctions)
- Allow distrustful entities to exchange resources
- Different types of resource bindings:
 - Permanent (owned)
 - Temporary (borrowed)

Players

- Slices – consume resources
- Resources – consumed by slices
- Components – provide resources
 - Resource manager – enforces resource bindings
- Experimenters – bind resources to slices
- Allocators – allocate resources to experimenters

Gacks Architecture

- Escrow service
 - Secure resource exchange
 - Audit trail
 - Waist of the hourglass



Resource Manager

- Resources as first-class objects
 - Named
 - Owned
 - Borrowed
- *Receipts* enable auditing
- Authentication based on GENI credentials
 - (PlanetLab/geniwrapper implementation)

3.12 Trust, Identity Management and GENI: Ken Klingenstein

Trust, Identity Management and GENI

Dr. Ken Klingenstein,
Senior Director, Middleware and Security, Internet2
Technologist, University of Colorado at Boulder



Topics

- Internet identity update
 - Technology updates
 - ISOC, IETF "Identity, Trust and the Internet"
- R&E identity federations
- Some thoughts on federation and trust



Internet identity

- Federated identity
 - Enterprise centric, exponentially growing, privacy preserving, rich attribute mechanisms
 - Requires lawyers, infrastructure, etc
- User centric identity
 - P2P, rapidly growing, light-weight
 - Marketplace is fractured; products are getting heavier to deal with privacy, attributes, etc.
- Unifying layers emerging – Cardspace, Higgins

Federated identity

- Convergence around SAML 2.0 – even MS; increasing use of Shibboleth as the interoperability standard.
- Exponential growth in national and international R&E sectors
- Emerging verticals in the automobile industry, real-estate, government, medical
- Policy convergence for LOA, basic attributes (eduPerson), but all else, including interfederation, remains to be developed
- Application use growing steadily
- Visibility is about to increase significantly through end-user interactions with identity selectors and privacy managers

User-centric identity

- Driven by social networking {Facebook, MySpace, etc} and {Google, AOL, MSN}, growing rapidly
- Relatively lightweight to implement for both application developers and identity providers
- Separates unique identifier and trust (reputation systems, etc.)
- Fractured by lack of standards, vying corporate interests, lack of relying parties, etc.
- OpenId, Facebook Connect, Google Connect, AOL

Unifying the user experience

- Among various identity providers, including P2P, self-issued, federated
- Need to manage discovery, authentication, and attribute release
- Cardspace, Higgins, uApprove, etc.
- Consistent metaphors, somewhat different technical approaches
- Starting to deploy
- Integrating enterprise and social identity

Trust, Identity and the Internet

- Acknowledges the assumptions of the original protocols about the fine nature of our friends on the Internet and the subsequent realities
- <http://www.isoc.org/isoc/mission/initiative/trust.shtml>
- ISOC initiative to introduce trust and identity-leveraged capabilities to many RFC's and protocols
- First target area is DKIM; subsequent targets include SIP and firewall traversal (trust-mediated transparency)

Privacy

- A broad and complex term, like security, encompassing many different themes
- In the GENI case, at least several instances
 - Protection of research data and collaborative materials
 - Consent for personal data release for access controls, particularly in international collaborations
 - Likely others
- International federations have already explored some of the privacy issues.

Federation Update

- R&E federations sprouting at national, state, regional, university system, library alliance, and elsewhere
- Federated identity growing in business
 - Many bilateral outsourced relationships
 - Hub and spoke
 - Multilateral relationships growing in some verticals

R&E Federation Killer Apps

- Content access – Elsevier, OCLC, JSTOR, iTunes
- Government access – NIH, NSF and research.gov
- Access to collaboration tools – wikis, moodle, drupal, foodle
- Roaming network access
- Outsourced services – National Student Clearing House, student travel, plagiarism testing, travel accounting
- MS Dreamspark
- Google Apps for Education

International R&E federations

- More than 25 national federations
- Several countries at 100% coverage, including Norway, Switzerland, Finland; communities served varies somewhat by country, but all are multi-application and include HE
- UK intends a single federation for HE and Further Education ~ tens of millions of users
- EU-wide identity effort now rolling out - IDABC and the Stork Project (www.eid-stork.eu)
- Key issues around EU Privacy and the EPTID
- Some early interfederation – Kalmar Union and US-UK

kjk@internet2.edu



InCommon

InCommon®

- **Over 123 members now**
- **More than two million “users”**
- **Most of the major research institutions**
- **Other types of members**
 - Non usual suspects – Lafayette, NITEL, Univ of Mary Washington, etc.
 - National Institute of Health, NSF and research.gov
 - Energy Labs, ESnet, TeraGrid
 - MS, Apple, Elsevier, etc.
 - Student service providers
- **Steering Committee chaired by Lois Brooks of Stanford;**
Technical Committee chaired by Renee Shuey of Penn State

kjk@internet2.edu



InCommon Update

- Growth is quite strong; doubled in size for the fifth year straight...
- Potential size estimates (pre-interfederation) could grow > 5,000 enterprises; revenue stream....
- Overarching MoU for federal agencies to join may happen
- Silver profile approved
- Major planning effort on the future of InCommon now underway, including governance, community served, pricing and packaging principles, business models

NIH

- Driving agency for much of our government activity
- Several types of applications, spanning two levels of LOA and a number of attributes
 - Wikis, access to genome databases, etc
 - CTSA
 - Electronic grants administration
- “Why should external users have internal NIH accounts?”
- Easier stuff – technology, clue at NIH
- Harder stuff – attributes (e.g. “organization”), dynamically supplied versus statically-supplied info

Federation Soup

- Within the US, federations happening in many ways – state, university system, library, regional, etc
- Until we do interfederation, and probably afterwards, federations will form among enterprises that need to collaborate, regardless of their sector
- Common issues include business models, legal models, LOA and attributes, sustainability of soup
- Overlapping memberships and policy differences creates lots of complexity in user experience, membership models, business models, etc.
- One workshop in, so far...
- <https://spaces.internet2.edu/display/FederationSoup/Home>



kjk@internet2.edu

Examples of federation soup

- Texas: UT, Texas TACC/Digital library, LEARN
- North Carolina – the MCNC federation
- California – UCOP, Cal State, State of Cal, etc...
- New Jersey - NJEdge

104



kjk@internet2.edu

A point in time

- We're about ten years into federated identity
 - Much has been accomplished – strong use cases, SAML 2.0, national level R&E federations, redirection of government efforts, corporate deployments, etc.
 - Many positive if unexpected outcomes (secrecy, revenue)
- There are significant gaps to fill in
 - Building a real global Internet identity layer
 - Nothing looks technically intractable; policies are harder
 - Integration of enterprise and social identity



kjk@internet2.edu

Federated what...

- Not all things federated fit together well
 - E.g. federated search meets federated identity is an uneven fit.
 - Federated resources may not overlap with federated users and identities
- The hardest part of federation is the policy space.
- What parts of the existing policy space should/must GENI use?¹⁰⁵



kjk@internet2.edu

Even in identity federation...

- Which federation(s) to be in
- The alignment of resource owners to federations
- Levels of LOA
- Common schema
 - For people
 - For almost everything else – devices, measurements, etc

Virtual Organizations and Federations

- VO's can leverage peered federations
 - Use local authentication, integrate local and external privileges, etc.
 - Improve end-user experience, create a layer of privacy, better security
- A VO, or a cluster of VO's sharing an IdM or a CA, can be considered a federation
- COmanage might be a useful tool.

Access control

- Web versus web services vs other protocols
 - Shib is web right now, with some web services extensions and a few non-web buried instances
 - SAML can be bound to almost any protocol, but hasn't been yet
- Sources of authority for privileges on all sorts of things...
- Using groups
- Using privileges

Externalizing identity management from the management apps

- <http://groups.geni.net/geni/wiki/GeniServices> is not federated...
- The collaboration apps
- The domain apps
- The admin users

Trust-mediated transparency

- Security is not just threats; it is also opportunities
- The biggest problem, for the R&E community, is the TDA's (traffic disruption appliance) – firewalls, NAT's , packetshapers, etc
- A deeply layered problem, with vicious feedback loops
- Dave Clark talked (~2003) about trust-mediated transparency as an essential aspect of the next-gen Internet...

3.13 An Adversarial Experimental Platform for Privacy and Anonymity: Ben Zhao



An Adversarial Experimental Platform for Privacy and Anonymity

Ben Zhao, U. C. Santa Barbara
NSF GENI Security Workshop, January 2009

Disclaimer

- This sounds a lot like Nikita's ideas on testing Tor on GENI
 - Independent idea, similar in some ways, different in others
- Key difference
 - Focus on general privacy/anonymization techniques
 - Focus on fine grain data collection and data measurement, tracing, and replay
- Clearly, he and I will talk ☺

A Need for Experimental Privacy

- Internet privacy an increasingly important topic
 - Anonymity relevant to new popular applications
 - E.g. VoIP, content sharing, remote machine control, secure data access, social networks
- Experimental evaluation critical, but challenging
 - Real world often different from analysis of idealized protocols
 - Assumptions often unrealistic
 - Real world factors key to breaking secure protocols
 - E.g. network/node dynamics, resource heterogeneity
 - Challenging to setup and deploy
 - Thorny legal issues w/ deployed services



Dream for Privacy Experimentalists

- What would we really like to have?
 - Experiments on popular privacy protocols with real users
 - What are real traffic patterns and user behavior patterns?
 - How do users react to attacks/DoS in real time?
 - Publicly available traces for repeatable, realistic experimentation
 - Adversarial evaluation of anonymity protocols and attacks



An Adversarial Measurement Platform

- Outcomes
 - Real users, waived legal rights (naïve?)
 - Re-evaluation of commonly accepted assumptions
 - Real-time anonymity attacks and defenses
 - Detailed, anonymized traces for public consumption



What Do We Need / Questions

- Possible requirements from GENI
 - Detailed traffic capture/logging at routers
 - Well-instrumented VMs for user-controlled network dynamics
 - IP- and DNS-level firewalls to enforce AUPs
 - Central directory for privacy-enhanced/anonymous applications and services
- Questions and issues
 - Timers, time synchronization, accuracy
 - Access or access control to external sites
 - Preventing pollution by legally questionable content
 - Isolating/identifying anonymous traffic



3.14 Observations on Operations/Security from a (Former) Tier 1 Builder/Operator: Chase Cotton

observations on operations/security from a (former) tier 1 builder/operator



- tool (the network) vs. experiments (the customers using the “service”)
 - prior requirements work is inspiring – but need “hard” strawman use cases to guide “tool” design/build phase (think multicast effect re: IP); this effort - security examples
 - given current cost constrains – use old tech in the tool where possible / new tech where required (e.g. virtualization/partitioning)
 - as a service, think super VPN – may ease some of the security / virtualization issues
- excluding forensics, operations and security are typically a mated pair (design if not org)
 - distributed ops (organizationally) of a single network problematic; global Internet special case
 - out-of-band (OOB) constructed from the controlled network is problematic
 - as element provisioning/surveillance hard/closed and not homogeneous – partitioning and virtualization extra difficult
 - actual and virtual – ticket systems / fix agents / “remote hands”
- given state of element programmability – recognize performance realities
 - functional/logical (adequate) speed experiments --- focus here 1st (APIs)
 - higher speed experiments once (if) at-speed programmable elements can be built/acquired
- all above apply to reduced security experiment space
- old axiom: “better/faster/cheaper – pick 2”
 - GENI version? “experiment flexibility”/ “i/f simplicity” / “security/stability” / other?

**3.15 GENI Ideas: Instrumentation, Experiments and Security:
Richard Ford and Ronda Henning**



GENI Ideas: Instrumentation, Experiments and Security

Richard Ford (rford@fit.edu)
Ronda Henning (rhenning@harris.com)

1

The Harris Institute for Assured Information 1/29/09

Three ideas, One slide...

- › GENI Ideas: Instrumentation, Experiments and Security
 - Richard Ford (rford@fit.edu)
 - Ronda Henning (rhenning@harris.com)
- › Three Ideas: Monitoring
 - › Develop a unified, modular monitoring protocol for GENI nodes
 - › Single set of APIs implemented on each platform at the virtualization layer
 - › Backplane logging channel required
 - › Modular logging allows for maximum reuse of code
 - › Logging should not change the results... but how will we know?
 - › No real "opt in" for external users (those running outside GENI slices) whose data we will be snarfing
 - › BTW, this is going to generate a LOT of data...
 - › GENI enablement of campus environments: how to adhere to campus policies (for example, RIAA-related issues)
 - › Privacy, privacy, privacy... oh, and privacy
 - As AOL release taught us, pseudonymity is of little help
- › Experiments
 - › Malware...
 - › Per Nick: write a viable worm and he will mutilate you in interesting novel way!
 - › Do need to ensure containment of effect (spread too obviously, but there's no excuse)
 - See my comment on monitoring previously
 - › Desperate need for background traffic – experimentation without this is meaningless
 - Furthermore, should follow the type of extremes we see in reality
 - Don't require experimenters to be experts in this!
 - Replay of stored traffic is okay, but it's unclean and doesn't reflect some very interesting environments (like MANETs)
 - › How will we get users to "opt in" to these experiments?
 - And opt in to the monitoring we'll need
- › Security
 - › Statefulness is (often) the enemy of security
 - › Reducing saved state of GENI between and during runs narrows the window for an attacker
 - › What stops a cluster owner stealing IP from experimenters?
 - › Where cluster owner could be, for example, a hostile government...
 - › What happens when GENI gets used for evil (be a great target for a botherder, for example...)
 - › Should be rate limits and heuristics at the GENI/Internet border that can shutdown a slice... but this is HUGELY double-edged
 - › Need a federated, distributed framework for detection
 - › Outliers are really the interesting parts in many experiments we shouldn't shut these down "accidently"
 - › What stops an experimenter (or someone posing as an experimenter) deploying hostile code to user nodes?
- › Contact
 - › Richard: rford@fit.edu
 - › Ronda: rhenning@harris.com

116



Monitoring

- ▶ **Must develop a unified, modular monitoring protocol for GENI nodes**
 - ▶ Single set of APIs implemented on each platform at the virtualization layer
 - ▶ For example, system API logging... solve generic problem and configure
 - ▶ Backplane logging channel required
 - ▶ Modular logging allows for maximum reuse of code
 - ▶ ... rolled up per slice
 - ▶ Logging should not change the results... but how will we know?
 - ▶ No real “opt in” for external users (those running outside GENI slices) whose data we will be snarfing
 - ▶ BTW, this is going to generate a LOT of data...
 - ▶ GENI enablement of campus environments: how to adhere to campus policies (for example, RIAA-related issues)
 - ▶ Flexibility of demarq points?
 - ▶ Privacy, privacy, privacy, privacy... oh, and privacy
 - ▶ As AOL release taught us, pseudonymity is of little help



▶ 3

The Harris Institute for Assured Information 1/29/09

Experiments

- ▶ **Malware...**
 - ▶ Per Nick: write a viable worm and he will mutilate you in interesting novel ways! (Must check with IRB)
 - ▶ Do need to ensure containment of effect (spread too obviously, but there's no excuse)
 - ▶ See my comment on monitoring previously
 - ▶ Desperate need for *good* background traffic – experimentation without this is meaningless
 - ▶ Furthermore, should follow the type of extremes we see in reality
 - ▶ Don't require experimenters to be experts in this (allow as bolt on)
 - ▶ Replay of stored traffic is okay, but it's unclean and doesn't reflect some very interesting environments (like MANETs)
 - ▶ How will we get users to “opt in” to these experiments?



▶ 4

▶ And to opt in to the monitoring we'll need The Harris Institute for Assured Information 1/29/09

Security

- ▶ Statefulness is (often) the enemy of security
 - ▶ Reducing saved state of GENI between and during runs narrows the window for an attacker
- ▶ What stops a cluster owner stealing IP from experimenters?
 - ▶ Where cluster owner could be, for example, a hostile government...
- ▶ What happens when GENI gets used for evil (be a great target for a botherder, for example...)
 - ▶ Should be rate limits and heuristics at the GENI/Internet border that can shutdown a slice... but this is HUGELY double-edged
 - ▶ Need a federated, distributed framework for detection (ties back to monitoring)
 - ▶ Outliers are really the interesting parts in many experiments we shouldn't shut these down "accidentally"
 - ▶ What stops an experimenter (or someone posing as an experimenter) deploying hostile code to user nodes?



▶ 5

The Harris Institute for Assured Information 1/29/09

Contact

- ▶ Richard: rford@fit.edu
- ▶ Ronda: rhenning@harris.com



▶ 6

The Harris Institute for Assured Information 1/29/09

3.16 Establishing and Communicating Trust in GENI: Raquel Hill and Jean Camp

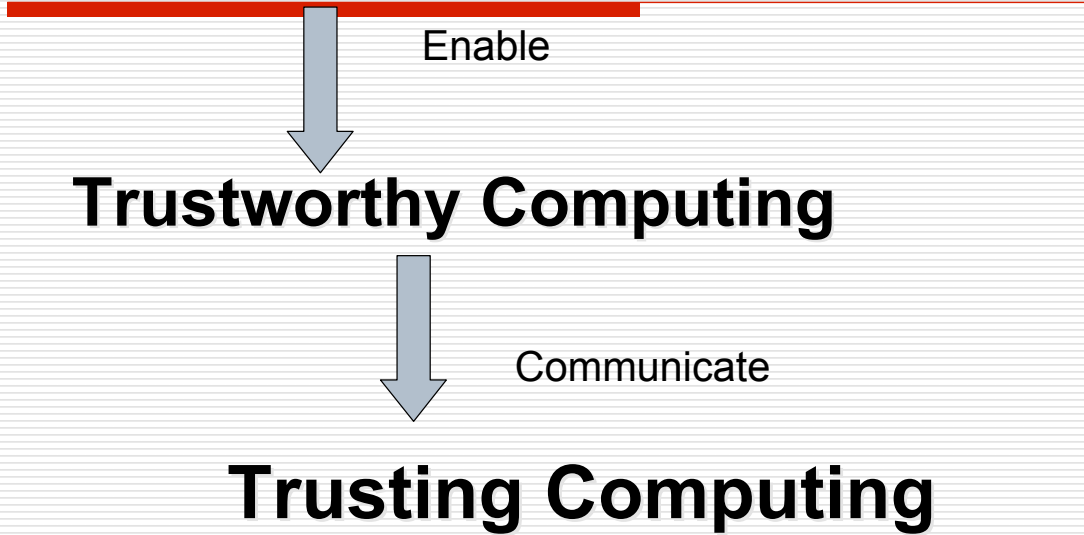
Establishing and Communicating Trust in GENI

Raquel Hill and Jean Camp
Indiana University

Trust Issues

- One goal of GENIE O&M is to detect malicious behavior and return infrastructure to trusted state
 - What is this trusted state?
 - Does it include rolling back software
- Clearinghouses
 - Specify criteria for trust (i.e. identity, behavior)
 - Extend trust model beyond authentication and access control
 - Include set of conditions that will enable trustors to evaluate an entity
- Users

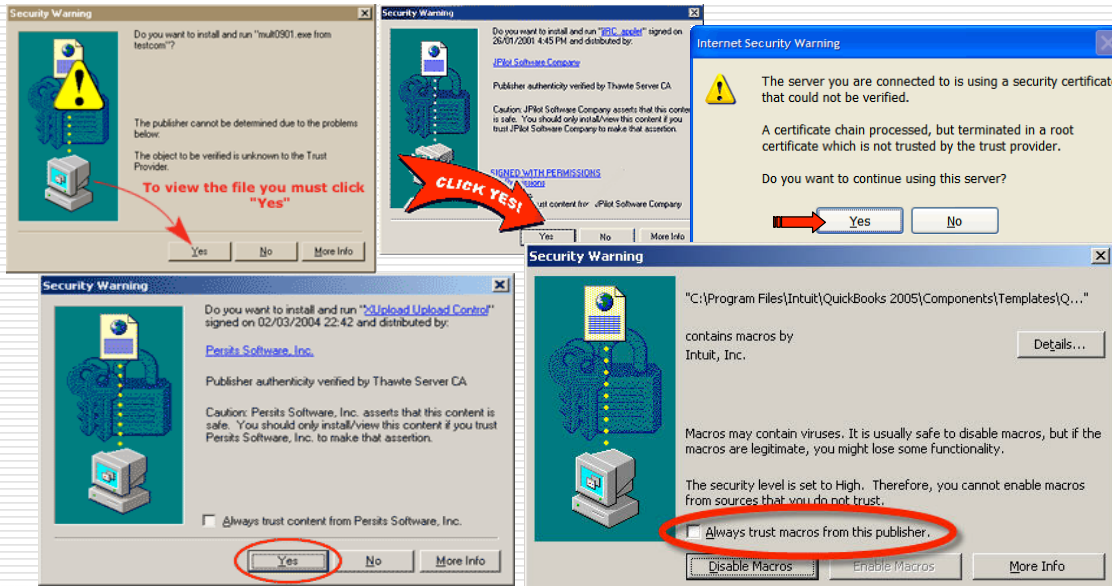
Trusted Computing



Security Communication as Risk Communication



The Current Response to Security Risks...



Communications Challenge

- ❑ Map changes in trust levels to risk
- ❑ Understand how individuals differentiate between changes in risk levels
- ❑ Assist the network, and GENI, in effectively responding to the change in risk

3.17 (Integrity Justified) Experimental Provenance: Patrick Mc-Daniel

Presented by Trent Jaeger.



(Integrity Justified) Experimental Provenance

Patrick McDaniel, Pennsylvania State University
Workshop on GENI and Security
Davis, CA -- January 22, 2009

Provenance

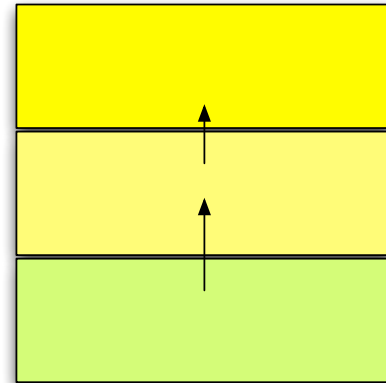


- A human scale problem:
 - ▶ Data often comes from many sources ...
 - ▶ ... is synthesized/influenced by complex/hidden processes ...
 - ▶ ... thus, how do you really know what the data means?
- *Data provenance* immutably identifies how data came to be in the state it is.
 - ▶ **Who/what** contributed to it?
 - ▶ **What** was it based on?
 - ▶ **When** was it generated? 124
 - ▶ **Why** was it generated?
 - ▶ **How** was it generated?



Why GENI provenance?

- Error handling
 - Detection, isolation, and recovery
- Source attribution
 - Forensics, consistency, believability
- Experimental Reproducibility
 - Extension, instrumentation
- Data revision
 - Updates, correction, extension, refinement
- Evidentiary
 - Evidence that data is legitimate/legal (certification, verification)
- *Experimental data can only be judged in light of how, when and where it comes from*

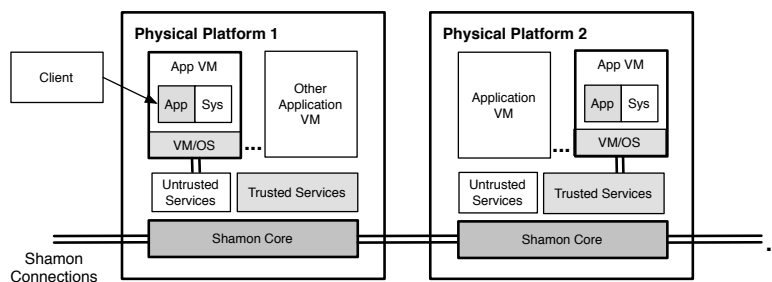


GENI System Provenance

- Assessing system provenance is key to understanding achieving the goals of GENI
 - What software was a component (slice/aggregate) running?
 - What inputs and configuration were used?
 - What security policy was being enforced?
 - e.g., isolation, data protection, privacy
- Stated as experimental *criteria* during the setup/acceptance
 - Think about sensitive experiments: *NCR*-esque, proprietary algorithms, opt-in with personal information
 - Determines apparatus acceptability of validation

GENI adoption requires answers to these questions

- Integrity measurement techniques provide information about the instantaneous state of a system, but *not* its data, or over time, or for other computational elements (VMs)
- What if you could build an aggregate of mutually attesting components that uses that apparatus to attest to the system state, protection state, data, and environment.
 - ... and tie a proof of that aggregate to experimental results.
- Building on the shared reference monitor (Shamon)



3.18 GENI Security Configuration In a Box: Ehab Al-Shaer

GENI Security Configuration In a Box

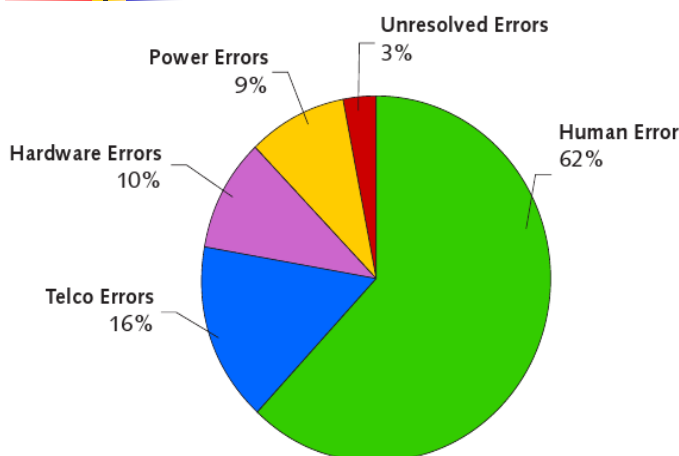
Ehab Al-Shaer

*Assurable Networking Recent Center (ARC)
School of Computing, DePaul University,
Chicago, IL, USA*

NSF GENI Workshop

**University of California Davis,
January 22, 2009**

State of Security Configuration Management



“Eighty percent of IT budgets is used to maintain the status quo.”, Kerravala, Zeus. “As the Value of Enterprise Networks Escalates, So Does the Need for Configuration Management.” The Yankee Group January 2004 [2].
“Most of network outages are caused by operators errors rather than equipment failure.”, Z. Kerravala. Configuration Management Delivers Business Resiliency. The Yankee Group, November 2002.

- “It is estimated that configuration errors enable 65% of cyber attacks and cause 62% of infrastructure downtime”, Network World, July 2006.
- *Recent surveys show Configuration errors are a large portion of operator errors which are in turn the largest contributor to failures and repair time [1].*
- *“Management of ACLs was the most critical missing or limited feature, Arbor Networks’ Worldwide Infrastructure Security Report, Sept 2007.*

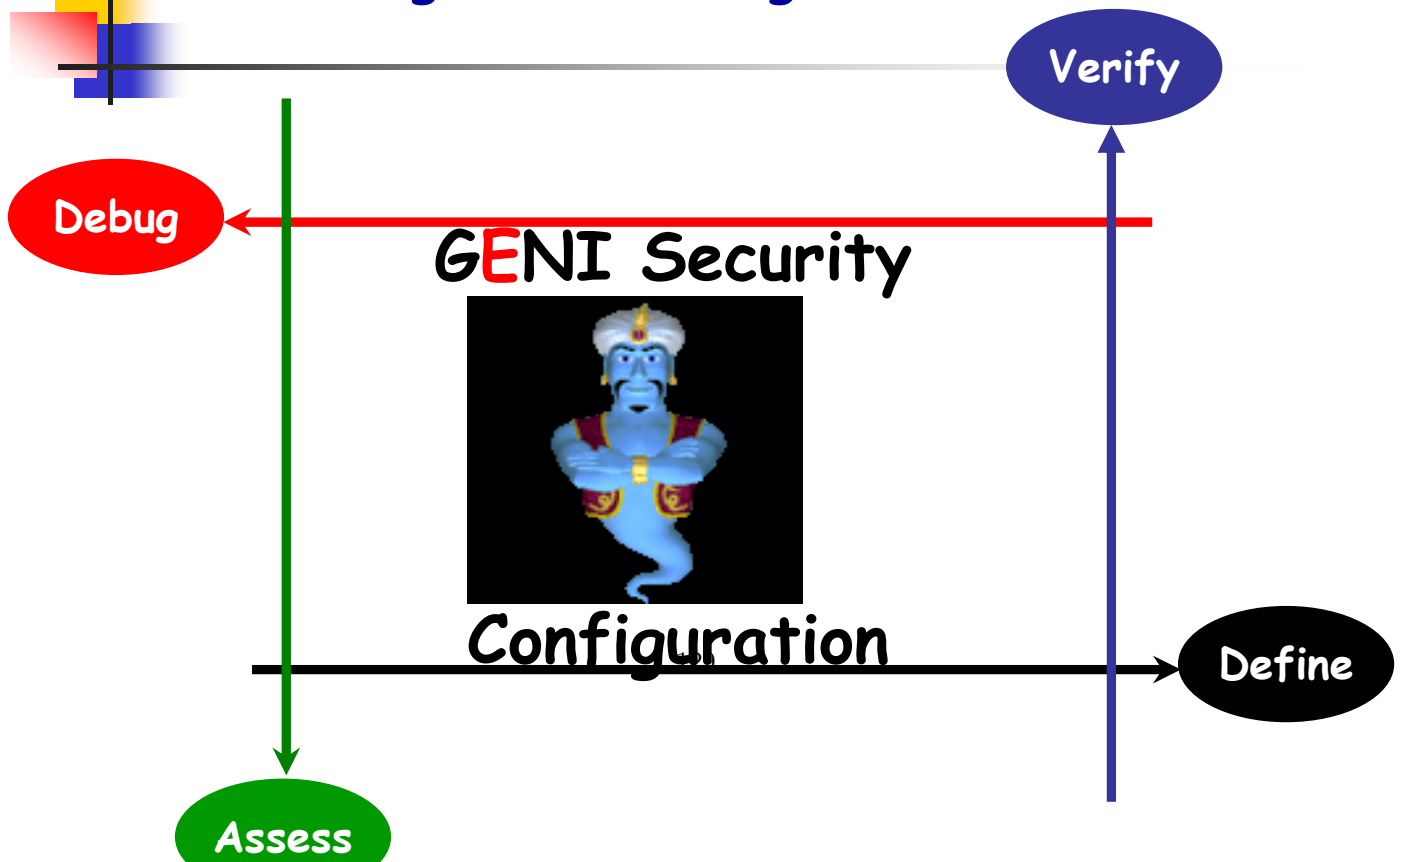
[1] D. Oppenheimer, A. Ganapathi, and D. A. Patterson. *Why Internet*

GENI Challenges

- Distributed resources
- Distributed control
- Dynamic policy coordination, interaction/federation, adaptation
- But still the goal is to keep it programmable, usable, assurable, and consistent → complex configuration

- How to provide end-to-end security configuration assurability/provability?
- How to make security systems configuration usable: high-level, distribution transparency?
- How to measure and assess configuration in term of risk, privacy, flexibility and cost?

Putting GENI Configuration in a Box



Idea#1: ConfigChecker & ConfigLego— Automated Security Configuration Verification

Goals

- Global end-to-end unified verification across heterogeneous devices: unifying the representation of the security configurations of all network devices.
- Integrating network and host security configuration checking: having a single model that can analyze both network and application level devices and services is the main focus.
- Abstraction and Composability
- Scalability (10,000 of nodes)

Approaches

- Bottom-up
- Modeling configuration semantic using Binary Decision Diagrams (BDD) gives canonical representation regardless of the syntax
- ConfigChecker**: models the network as a giant state machines and used model checker and CTL to query and verify security configuration
 - Modeling packet transformations is an increasingly hard task.
 - Problems on a network-wide scale are impossible to detect manually, and automated tools focus on a single device or devices of a single type.
- ConfigLego** allows for abstracting and composing portions of the network under-
investigation

© Ehab Al-Shaer

5

Modeling Access Control Policies

- Single-trigger policy** is an access policy where only one action is triggered for a given packet. C_i is the **1st** match leads to action a

$$P_a = \bigvee_{i \in \text{index}(a)} (\neg C_1 \wedge \neg C_2 \dots \neg C_{i-1} \wedge C_i)$$

$$P_a = \bigvee_{i \in \text{index}(a)} \bigwedge_{j=1}^{i-1} \neg C_j \wedge C_i$$

- Multiple-trigger policy** is an access policy where multiple different actions may be triggered for the same packet. C_i is **any** match leads to action a

$$P_a = \bigvee_{i \in \text{index}(a)} C_i$$

where

$$\text{index}(a) = \{i \mid R_i = C_i \rightsquigarrow a\}$$

Intra-Policy Conflicts Formalization : Crypto-access List

- Policy expression S_a represents a policy that incorporates rule R_i , and S'_a is the policy with R_i excluded. R_i may be involved in the following conflicts:

- Shadowing:**

$$[(S'_{a_i} \Leftrightarrow S_{a_i}) = true] \text{ and } [(C_i \Rightarrow S'_{a_i}) = false]$$

- Redundancy:**

$$[(S'_{a_i} \Leftrightarrow S_{a_i}) = true] \text{ and } [(C_i \Rightarrow S'_{a_i}) \neq false]$$

- Exception:**

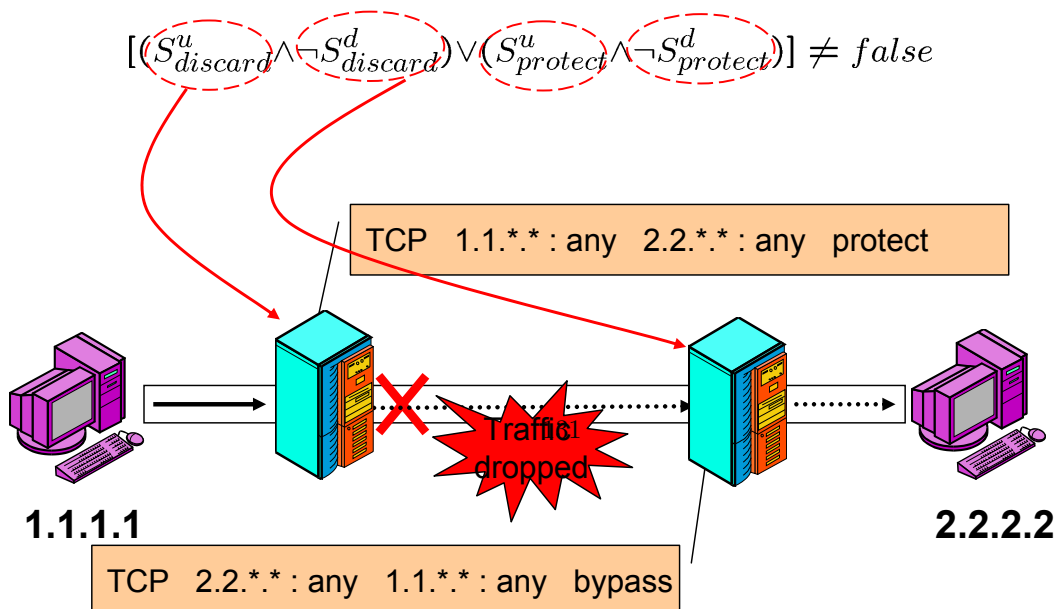
$$[(S'_{a_i} \Leftrightarrow S_{a_i}) \neq true] \text{ and } [(C_i \Rightarrow S'_{a_i}) = false]$$

- Correlation:**

$$[(S'_{a_i} \Leftrightarrow S_{a_i}) \neq true] \text{ and } [(C_i \Rightarrow S'_{a_i}) \neq false]$$

IPSec Inter-Policy Conflicts Formalization: Crypto-access Lists

- Shadowing: upstream policy blocks traffic



Diagnosing Unreachability Problems between Routers and Firewalls

- **Flow-level Analysis:** Is flow C_k forwarded by routers in L (each of routing tables BDD T_j^i for router i and port j) but **Blocked** due to conflict between *Routing* and *FW Filtering*.

$$[(C_k \Rightarrow \bigwedge_{(i,j) \in L} T_j^i) \wedge (C_k \Rightarrow \neg S_A^n)] \neq false$$

- This shows that a traffic C_j is forwarded by the routing policy, T_j^i , from node i to n but yet blocked by the filtering policy, $S_{discard}^n$, of the destination domain.

- **Path-level Analysis:** Discovering Any Unreachability Conflicts between *Routing* and *Filtering*.

$$\phi_k \leftarrow [SAT(\bigwedge_{(i,j) \in path(x)} T_j^i \wedge \neg S_A^n \wedge \neg(\bigwedge_{i=1, k-1} \phi_i))] \neq false$$

- For $\phi_i=1$, n misconfiguration examples, and $\phi_i(0) = true$

- **Network or Federated-level Analysis:** Spurious conflict between downstream d and upstream u ISP domains:

$$[(S_{bypass}^u \wedge \neg S_{bypass}^d) \vee (S_{limit}^u \wedge S_{discard}^d)] \neq false$$

- Notice that $S_{discard}$, S_{bypass} and S_{limit} are filtering policies representations related to the filtering actions as described in [ICNP05, CommMag06].

ConfigChecker Queries (Model Checker approach)

- **Q1: Reachability Soundness:**

- From any source node $ip1$ if there is a next-hop to destination $ip2$, then there must be a way that eventually leads to $ip2$ from $ip1$.

$$Q = (loc(ip1) \wedge EX(dest = ip2)) \rightarrow loc(ip1) \wedge EF(dest(ip2) \Leftrightarrow loc(ip1))$$

- **Q2: Discovering Broken End-to-end IPSec Tunnel:**

- Given a specific flow, will it stay in a tunnel until the final destination? (assuming the IPSec gateways are a hop away from the source and destination)

$$Q = (src = a1 \wedge dest = a2 \wedge loc(a1) \wedge IPSec(encT)) \rightarrow AU((IPSec(encT) \vee loc \rightarrow \mathcal{G}), loc(a2))$$

- **Q3: What nodes have access to the plain-text packet:**

- Given a specific flow, which nodes will eventually have access to the packet without encryption?

$$Q = AF_{-}(flow(ip1, ip2) \wedge loc(ip1)) \wedge \neg IPSec(encrypt)$$

ConfigChecker Queries

- **Q4: Back-door access after route changes:**
 - What is difference in the new configuration as compared with the ordinary original one. Is there any backdoor?

$$C_{org} \triangleq [\neg multiroute \wedge src = a1 \wedge dest = a2 \wedge loc(a1) \rightarrow AF(loc(a2) \wedge src = a1 \wedge dest = a2)]$$

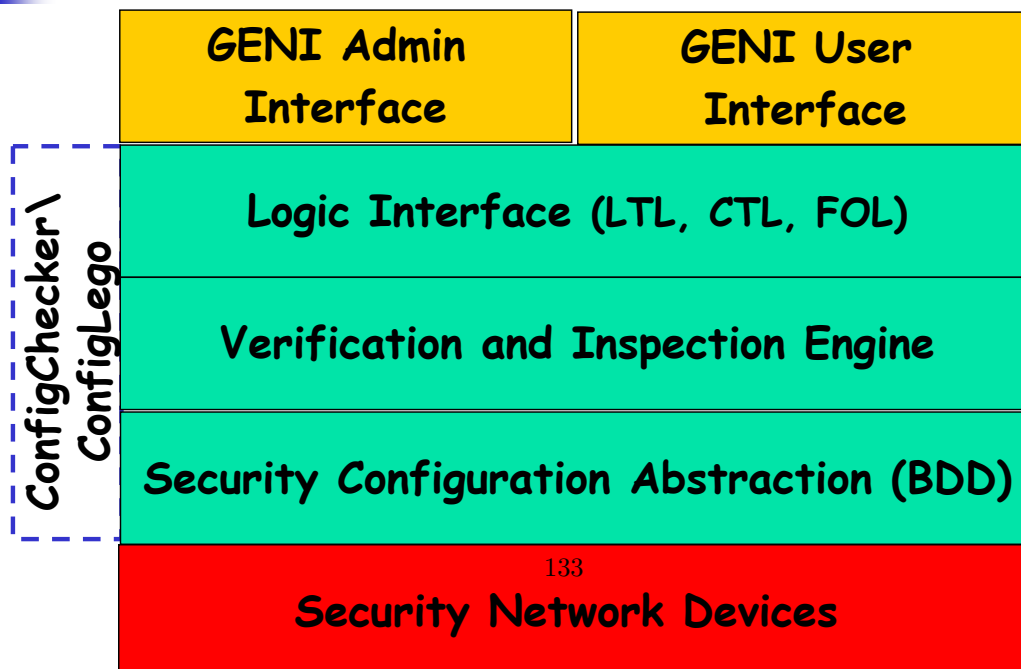
$$C_{new} \triangleq [multiroute \wedge src = a1 \wedge dest = a2 \wedge loc(a1) \rightarrow AF(loc(a2) \wedge src = a1 \wedge dest = a2)]$$

$$Backdoors: \neg C_{org} \wedge C_{new}$$

$$Broken\ flows: \neg C_{new} \wedge C_{org}$$

More information on ConfigChecker: www.arc.depaul.edu

Idea#1: GENI ConfigChecker / ConfigLego



Policy Advisor Tool for Distributed Policy (Firewall & IPSec) Management

IPSec2 Access Rules

Rule	Protocol	Source	Destination	Action
A1	tcp	10.0.0.0/24:0	10.0.2.2/32:0	Accept
A2	tcp	10.0.0.0/24:0	10.0.2.3/32:0	Protect
A3	tcp	10.0.0.0/24:0	10.0.3.2/32:0	Accept
A4	tcp	10.0.0.0/24:0	10.0.3.3/32:0	Protect
A5	tcp	10.0.0.0/24:0	10.0.3.0/24:0	Accept
A6	tcp	0.0.0.0/0:0	0.0.0.0/0:0	Deny

IPSec2 Transform Rules

Rule	Protocol	Source	Destination	Transform	Tunnel
T1	tcp	10.0.0.0/24:0	10.0.3.0/24:0	ESP-Transport	
T2	tcp	10.0.0.0/24:0	10.0.2.0/24:0	AH-Transport	

Inter-Policy Conflict Analysis Report

Device	Rule	Conflict description
IPSec1	A3	Access is totally spurious
	A5	Access is partially spurious
IPSec2	A3	Access is totally spurious
	A5	Access is partially spurious
IPSec1	A1	Access is totally shadowed
	A2	Access is totally shadowed
	T2	Transform is stronger than rule IPSec2/T2

© Ehab Al-Shaar

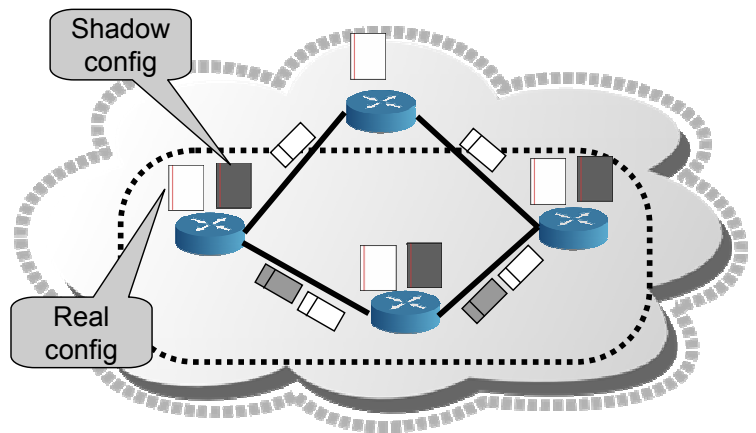
13

Intra-Policy Advisor Tool is used by the following 43 companies and institutions as of November, 2006

- Lisle Technology Partners, USA;
- Phontech, Norway;
- Naval Surface Warfare Center, Panama City, USA;
- Cisco Systems, USA;
- AT&T, USA;
- Gateshead Council, UK;
- ISRC, Queensland University of Technology, Australia;
- Imperial College and UCL, London, UK;
- Danet Group, Germany;
- TNT Express Worldwide, UK Ltd, United Kingdom;
- Checkpoint, USA;
- FireWall-1, The Netherlands;
- UFRGS, Brasil;
- DataConsult, Lebanon;
- Rosebank Consulting, GB;
- Columbia University, USA;
- Mayer Consulting, USA;
- Panduit Corp, USA;
- UPMC Paris 5 University, France;
- Royal institute of Science, Sweden;
- GE, US;
- Aligo, USA.
- Others not listed

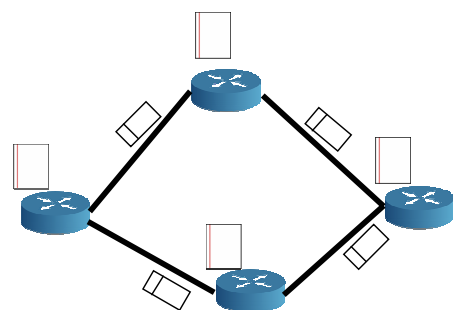
Idea#2: Shadow Configurations for On-line Configuration Debugging

- Use Deployed Network
- Allow an additional shadow configuration on each router
 - Routing, ACLs, interface addresses, etc.
- Scalable and realistic (no modeling)
- Two key capabilities
 - Pre-deployment testing/debugging
 - Does not affect real traffic



Scenario: Config Changes

- Scenario: Change configuration parameters
 - Address performance/security issues
 - Deploy new services (e.g., filters, IDS probes and QoS)
- Operation
 - 1) Copy real traffic to shadow plan
 - 2) Change shadow and test
 - 3) Store and aggregate traces
 - 4) Debug, compare and isolate
 - 5) Commit real and shadow
- Prototype for Routing only (with Richard Wang, Yale) – see SIGCOMM 2008





Summary & Future Work

- **GENI success will be greatly dependant on assurability and usability of security configuration: define, verify, evaluate/metrics and optimize**
- **Other Issues**
 - How integrate application level and network level access control
 - How to build API and high-level user interfaces to help using the underlying configuration engines
 - Measuring security
 - Top-down approach: Balancing security, usability, privacy and cost

Thank You!!



3.19 GENI Infrastructure and Proposed GENI Experiment: Brian Hay and Kara Nance

GENI Infrastructure

- ▶ **Need to have ability to monitor network traffic**
 - ▶ Results of experiments
 - ▶ Debugging an experiment or the environment
 - ▶ Forensic analysis of intrusions
 - ▶ Replay
 - ▶ Defensive capabilities
- ▶ **Need to easily manage the traffic monitoring**
 - ▶ Efficient identification and tagging of relevant packets
 - ▶ Instruct all subset of slivers to record traffic that meets some condition(s).
 - ▶ Ability for slivers to validate packets
 - ▶ Out-of-band control and packet collection?



Proposed GENI Experiment

- ▶ **Dynamics of Large Scale Networks**
 - ▶ Worked on power systems model, including agent actors (done)
 - ▶ Building and coupling communications network model to power systems model (in progress)
- ▶ **Need GENI to facilitate this work**
 1. Deploy such networks in GENI
 2. Determine characteristics (topologies, protocols, ...) that inhibit or exacerbate network failures
 3. Develop and validate models
- ▶ **GENI gets a real use case that we cannot perform today.**



3.20 GENI Security Services: Calvin Ko, Alefiya Hussain, Steve Schwab, Jim Horning, and Sandy Murphy

GENI Security Services

*Calvin Ko, Alefiya Hussain, Steve Schwab,
Jim Horning and Sandy Murphy*
Sparta, Inc.

The Threat Model

- External attacks on the GENI infrastructure, a DoS attack
- Contain and prevent the impact of accidental/malicious misbehaving experiments on the outside world
- Isolation between experiments, so that one experimenter cannot disrupt another

Security Requirements

- **Explicit Trust:** All principal privileges should be managed explicitly
- **Least Privilege:** each principal is given the authority needed to perform a particular task
- **Revocation:** compromised components can be recalled from an experiment
- **Auditability:** compromise incident traceable
-**Scalability, Usability, Autonomy, Performance**

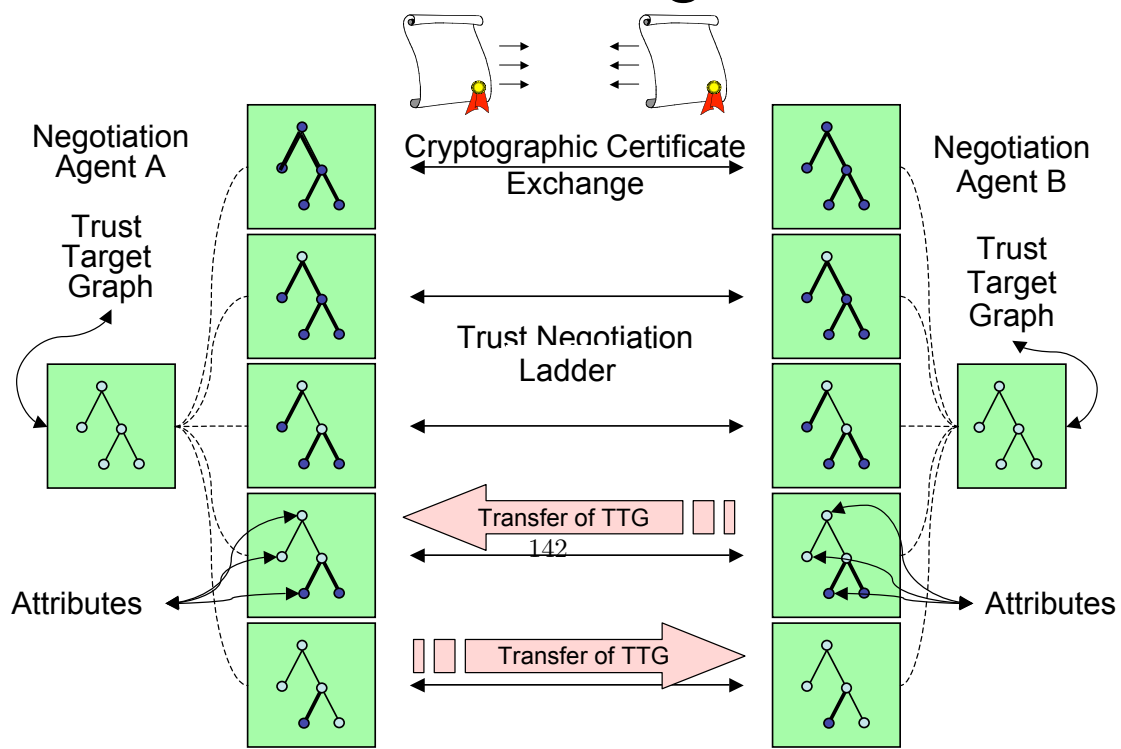
Identities, Authorization, and Accountability Requirements

- Identities and Identity records
 - Every principal will have an identify for accountability
 - Principal may have multiple identities
 - GENI identity will map to a real world identity
- Authorization
 - Before any GENI resource is accessed, but should also allow for support of anonymous use
- Accountability
 - Activity should be traceable to a principal, mainly to identify sources of bad traffic

Attribute Based Access Control

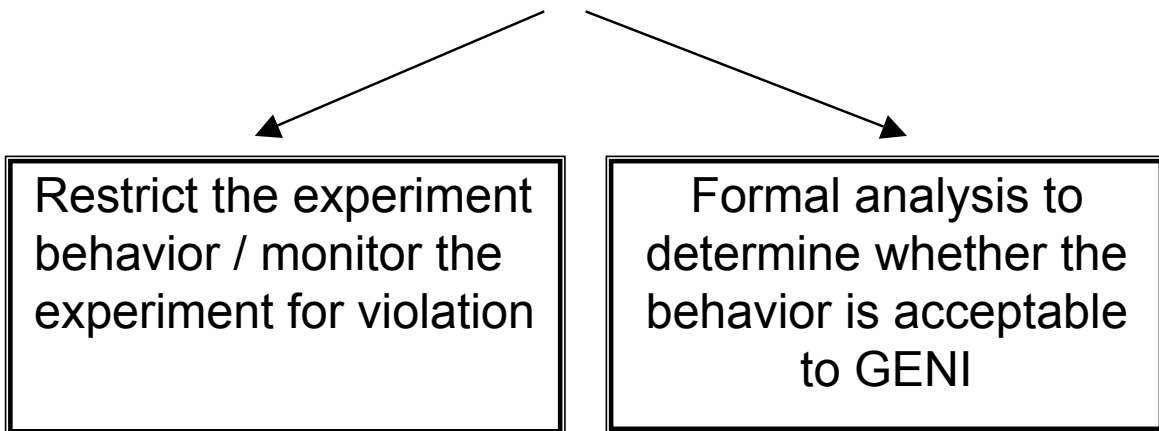
- Authorization decision is based on attributes of the principal
 - Credentials carry cryptographically signed claims about an principal's attributes
 - Requestor and provider may be strangers
 - Automated **Trust Negotiation** protects sensitive attributes while enabling unanticipated users to gain access in accordance with policy, and principal's authorization attributes

ABAC Trust Negotiation



Experiment Monitoring

- Specification-based detection
 - Specify valid experiment behavior of a slice
 - interaction with the outside



GENI Security Services for Experimental Slices

- Provide security services in the form of a security library and/or toolkit interface
- Reusable security components
- Allow researchers to plug in GENI facility security mechanisms without having to re-invent the wheel.

Issues and Questions

- Should each slice have a standard way of accessing security services exported from the GENI substrate
- What services are relevant?
- How should access be provided?
- Should a GENI slice be able to access services in another GENI slice?

3.21 Attribution in GENI: Carrie Gates, Jeffrey Hunker, and Matt Bishop

separate infrastructures for attribution

① infrastructure support: **attribution**

→ trust hardware vendors, software vendors
 → trust configuration, installation, etc, those who do it

needed to experiment with different levels, types of attribution?

1) What is attribution?
 a) Why do we want attribution?
 b) What do we want to know?
 c) When do we NOT want attribution?

2) What is the required infrastructure?
 a) How can GENI provide attribution?
 3) A whole bunch of policy stuff...

US w/ machine
 infrastructure
 at 1st level
 at second level

infrastructure can't be perfect
 some agents mess with proxy & Transf

no requirement that messages be delivered or

Technical coverage behavioral
 did not want to know ID.

clery

3.22 GENI Trace Collection for Security Studies: Yan Luo



GENI Trace Collection for Security Studies

Yan Luo
Department of Electrical and Computer Engineering
University of Massachusetts Lowell



Questions

1. What assets are you trying to protect?
2. What are the risks to these assets?
3. What are the security solutions?
4. How well does the security solution mitigate those risks?
5. What other risks does the security solution cause?
6. What cost and trade-offs does the security solution impose?

148

Bruce Schneier, *Beyond Fear*

What assets are you trying to protect?

- Computing nodes
- Programmable routers/switches
- Radar, sensors, ...
- Network bandwidth
- Application data
 - E-commerce data?
 - Healthcare app data?
- Experiments
- Users

What are the risks to the assets?

- Shutdown/disable GENI hardware
- Breach of privacy data
- Misuse of allocated GENI resources
- Unauthorized usage of GENI resources
- Interrupting user experiments
- Users losing interest due unavailability

What are the security solutions?

- Our proposed solution:
 - providing a mechanism of capturing and analyzing packet traces on GENI.

Trace Collection and Analysis in Network Research

- Long history
 - 1994: DAG card developed by University of Waikato networking research group
 - 1995: NLANR established the NLANR/Fix-West real time flow data web site
- Popular trace archives
 - NLANR
 - Internet Measurement Data Catalog
<http://imdc.datcat.org/Home>
 - WITS: Waikato Internet Traffic Storage
<http://www.wand.net.nz/wits/>
- Proved to be beneficial
 - Hundreds of papers

Challenges in Network Trace Studies

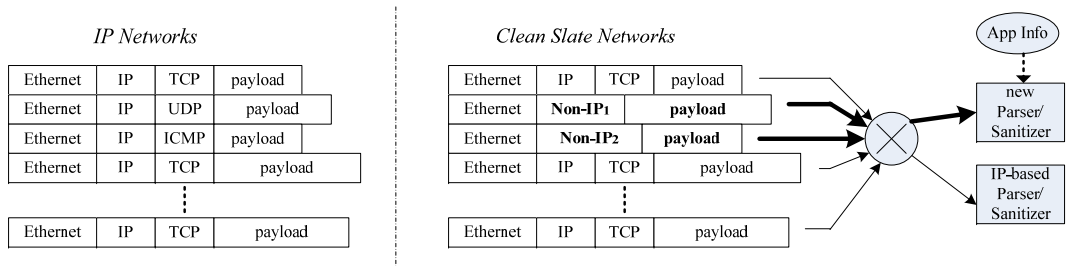
- High speed trace capture and archive
 - Specialized hardware
 - Enormous storage space
- Trace anonymization
 - Protect privacy
 - Facilitate trace sharing
- New challenges for GENI
 - Packet formats
 - Experimental applications

1/20/2009

Yan Luo, GENI Security Workshop
Davis, CA

8

Mixed Network Traffic in Clean Slate Networks



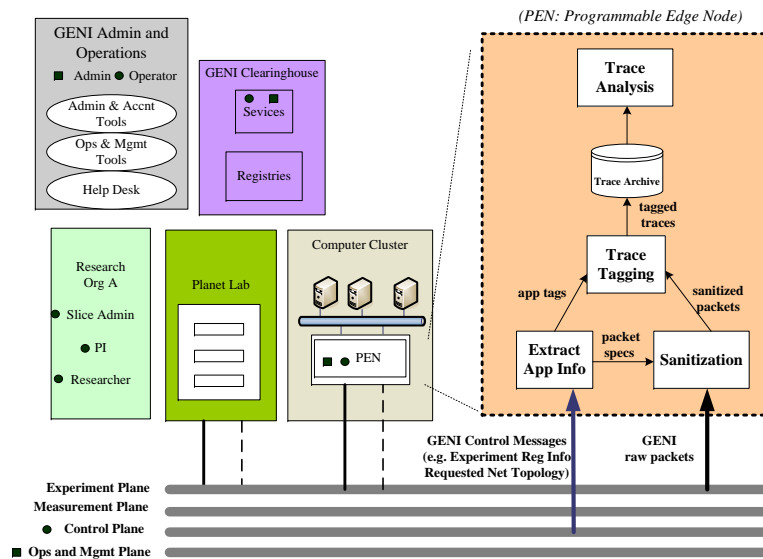
151

1/20/2009

Yan Luo, GENI Security Workshop
Davis, CA

9

Proposed Trace Collection Architecture

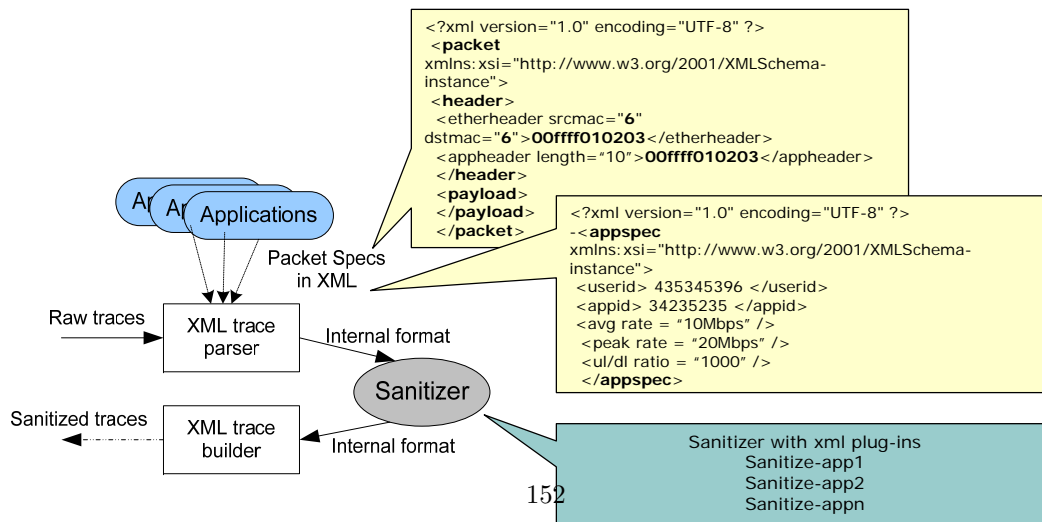


1/20/2009

Yan Luo, GENI Security Workshop
Davis, CA

10

Trace Specification and Anonymization



1/20/2009

Yan Luo, GENI Security Workshop
Davis, CA

11

How well does the security solution mitigate those risks?

- Online capture and anonymization of packet traces for post-analysis
- Facilitate trace sharing and publication
- Audit experiments and their packets
- Detect abnormal behavior
 - Invalid packet formats
 - Misuse of GENI resource
 - unauthorized usage
 - Unexpected experiment behavior (duration, burst, ul/dl, etc)

What other risks does the security solution cause?

- Additional design complexity of GENI infra.
- Users lose interest because of cumbersome application specs
- Weakest link targeted by phony/malicious application/packet specs
- Performance distortion/degradation of experiments
 - Additional packet processing

What cost and trade-offs does the security solution impose?

- Packet and application specs expected from GENI user
 - Detailed specs
 - better understanding of the apps and network activities
 - Cumbersome to user

vs.

- Simplified specs
- minimal knowledge of apps and activities
- simple to user

3.23 Security Event Standardization: Doug Pearson and Wes Young

Security Event Standardization

“SES”, Moving security messages throughout the ether.

Workshop on GENI and Security
UC-Davis, January 2009
Doug Pearson / Wes Young

Addressing

- Addressing the Workshop question:
 - How can GENI itself be adequately secured and protected from attack?
- Operationally protecting GENI, experiments, and connected infrastructures

Idea

- ◉ Share, in real-time, security event information within a trusted federation, and among federations; and
- ◉ Apply the shared information to local protection and response.

Partial Solution

- ◉ The Idea is just one small part of a necessary total security solution
- ◉ Is designed to augment and enhance other components of a total solution; and
- ◉ Is designed to integrate with other operational processes

At its roots, not a new Idea

- Lots of security event information is being shared right now
 - Private communities
 - Semi-private communities
 - Public sources

Issues

- Current methods cumbersome
 - Many rely on e-mail
 - Not easily automated
 - Requires the “human interrupt” signal
 - Not structured for correlation
- Multiple data representations
 - Non-standard
 - Not easily parsed
 - Not easily acted on
 - Hard to measure confidence

Issues

- Long-term Intelligence
 - Hostage to our inboxes
 - Difficulty of correlation
 - Difficulty of coordinated or cooperative analysis
- Multiple Federations
 - Trust relationships
 - Political and organizational boundaries

Building a Solution

- Based on work started at Argonne National Laboratory – “Federated Model”
- Development in progress
 - REN-ISAC
 - In cooperation with Internet2/CSI2
 - Funded by DoJ grant to Internet2 for a number of security projects and activities
 - Cooperating with parallel work at Argonne, funded by DoE.

Building a Solution

- Standardization
 - IDMEF - Security standard for representing mid-level security messages in XML
 - Developed in early 2000's
- Extensions
 - Understanding "Sites" (via ASN, CIDR)
 - Understanding "Federations"

Building a Solution

- Interoperation with EDDY (End-to-end Diagnostic Discovery)
 - Transport option
 - Local option for advanced event management
- Request Tracker (RT) – Solves the "UI", "ACL" and "Workflow" problem. Allows us to build on existing, rich, open-source technology.

Phase I Solution

- ◉ Local log (IDS, firewall, sshd, DNS, darknet sensor, etc.) parsing to yield “mid-level events”.
- ◉ Normalized data description in IDMEF
- ◉ Transport, storage, and retrieval
- ◉ Trusted federation
- ◉ Real-time security event information sharing → protection and response.

Phase I Solution

- ◉ Pilot Deployment
 - Sharing of data within REN-ISAC and Department of Energy federations
 - Sharing between REN-ISAC and DOE federations
 - Sharing real-time event and analysis (e.g. top-offending) data
- ◉ Production deployments in REN-ISAC and DOE

Building a Framework

- ◉ Framework for the incorporation of additional correlation and analysis tools
- ◉ Interface with systems that notify abuse contacts regarding infected systems, e.g. the REN-ISAC notification system
- ◉ Interface with systems that treat higher-level incident information in a federated context

Extending the Framework

- ◉ Long term intelligence storage
- ◉ Feed of security intelligence to other federations and mitigation communities
- ◉ Threat analysis platform
- ◉ The Future
 - Rapid application development
 - “Super Crunching” of data


The Result

- A better understanding of:
 - Who our attackers are
 - What they're doing
 - How they're doing it
- Rapid and comprehensive protection

Contacts

- Doug Pearson
 - dodpears@ren-isac.net
- Wes Young
 - wes@barely3am.com

3.24 ReAssure and SELinux: Jacques Thomas, Pascal Meunier, Patrick Eugster, and Jan Vitek

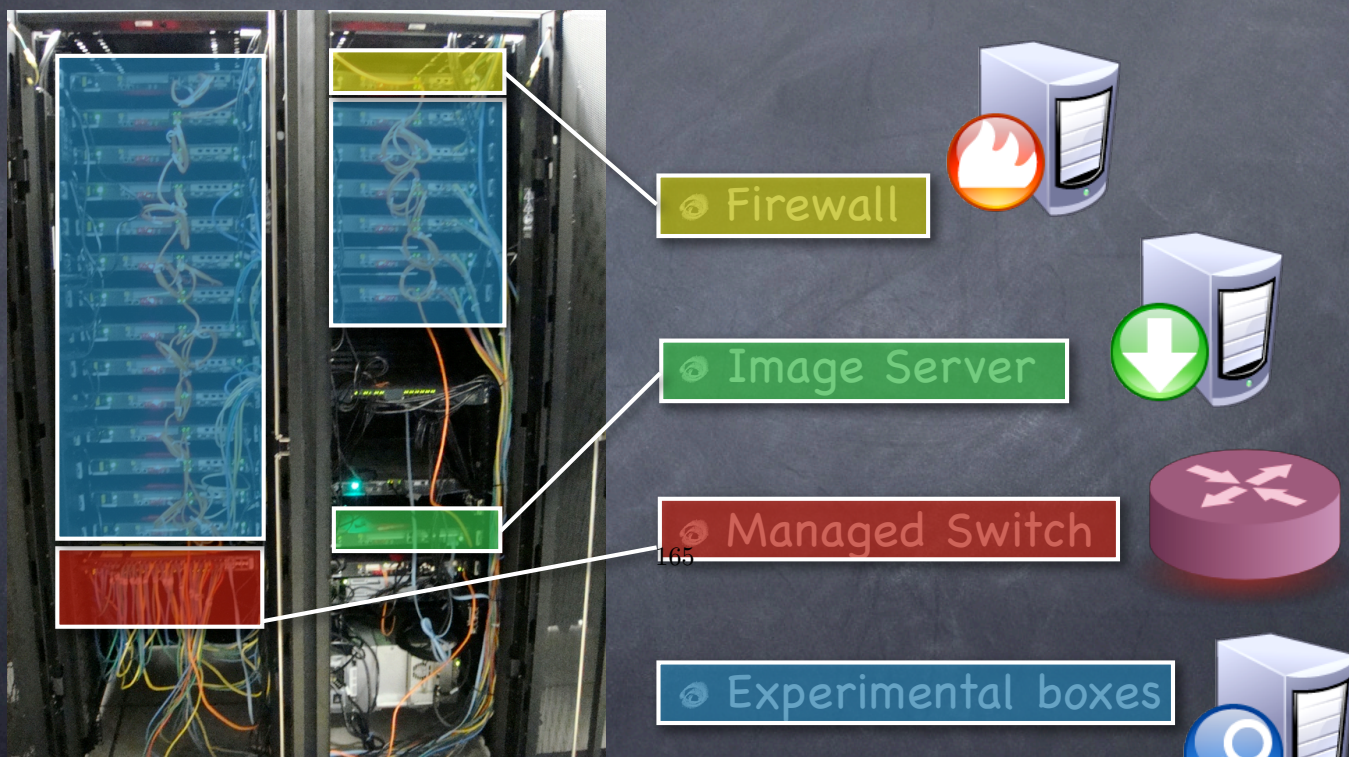


ReAssure & SELinux

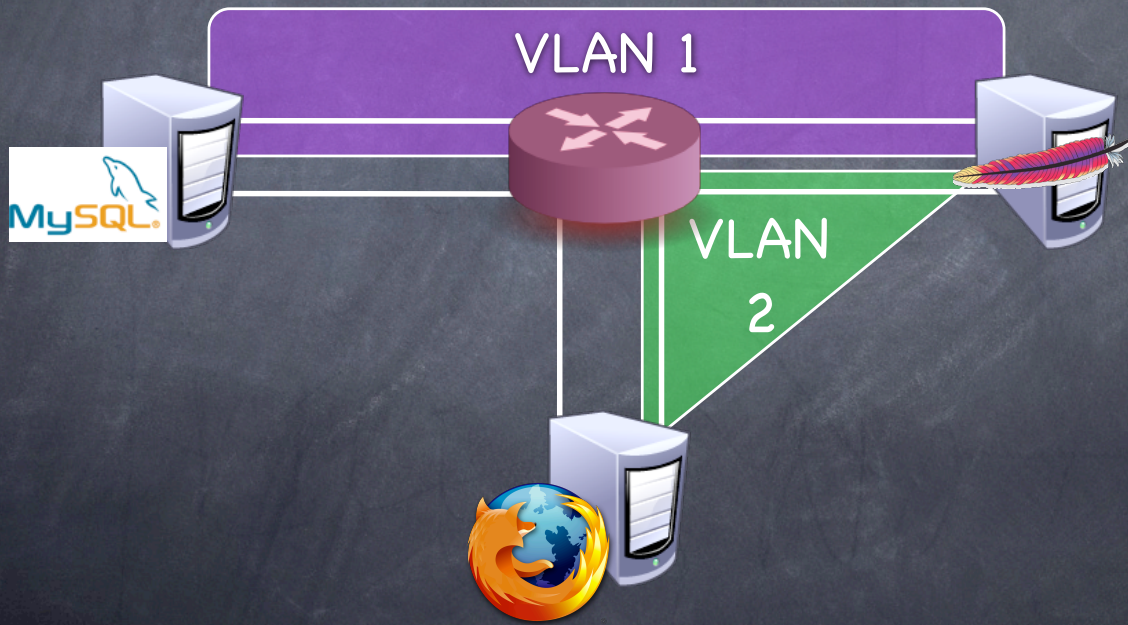
PURDUE
UNIVERSITY

Jacques Thomas, Pascal Meunier, Patrick Eugster,
and Jan Vitek

ReAssure



Deploying



Experimenting !



<https://reassure.cerias.purdue.edu>

The screenshot shows a web browser window with the URL <https://reassure.cerias.purdue.edu>. A blue speech bubble with the text "Request Account" is overlaid on the page. The website content includes a navigation menu with "System Help", "Logout", "Project Home Page", "Image Downloads", and "Contact us". A sidebar on the left lists "Your Account", "Experiments", "Files", "PCs", "Reservations", and "Testbed Status". The main content area features a section titled "Give Me A Computer!" with a form containing the following fields:

- Image file to transfer (optional): none
- How many hours? 24
- Create Experiment button

Below the form, there is a note: "If you only see 'none' as an image choice, you need to either copy public images to your directory or upload images." At the bottom of the page, contact information for CERIAS, Purdue University is provided: "CERIAS, Purdue University / Recitation Building / 656 Oval Drive / West Lafayette IN 47907-2086 phone (800)494-4419 / fax (765)496-3181 / Read our Privacy Policy".



Pascal Meunier

Virtual Machines,
Virtual Machines,
Virtual Security?



Accommodative Mandatory Access Control

- Capitalize on SELinux mechanisms
- Administrative model for SELinux
- Custom experiments without VM

Healthy Testbed



SELinux in. Malware out.

Questions

- When is a machine considered clean ?
- How does our SELinux admin model work ?
- Interesting for GENI ?
- Why SELinux ?

Thank you

jthomas@cs.purdue.edu

**3.25 Security for High-End CyberInfrastructure: Lessons Learned:
Randy Butler, Roy Campbell, Himanshu Khurana, Adam
Slagell, and Von Welch**

Security for High-end CyberInfrastructure: Lessons Learned

Randy Butler, Roy Campbell,
Himanshu Khurana, Adam Slagell, Von Welch
National Center for Supercomputing Applications
and
Information Trust Institute
University of Illinois



Lessons Learned from...

NSF TeraGrid
Extensible Terascale Facility

Los Angeles SDSC

UC/ANL Chicago

NCSA IU PU

ORNL

UT

PSC

NCSA

OCEAN OBSERVATORIES INITIATIVE
Science Plan
Revealing the Secrets of Our Ocean Planet

171
Open Science Grid

Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing

LCG



GENI and previous CI

- Some key differences.
 - Heavy use of VLANs and VMs.
 - Jobs are more "experimental" and "deeper" in nature.
 - e.g., the networking infrastructure itself is open to experimentation
- Many similar challenges and goals.
 - Multiple, distributed organizations.
 - Distributed user community.
 - Availability and Integrity of resources.
 - Keeping user "jobs" isolated.



Some Lessons GENI Can Build On

- Your biggest security problems are the ones you don't own.
- The hackers don't care about your software.
 - The hackers don't take the time to read the manual either.
 - It's all the usual stuff - Password theft, scans getting lucky, PHP, MySQL, kernel vulnerabilities, etc.
 - So far... the day may come, but it has been "coming" for a while.



Lessons

- Preparation and planning for incident response is critical.
 - Flowcharts.
 - Dry-runs and exercises.
 - Make sure you are doing the right logging and auditing.
- Plan for collaboration during an incident.
 - How will responders communicate with each other?
 - Who communicates with media? NSF? Users?
 - How do responders securely share data, correlate events, etc.

Lessons

- Getting agreement on security issues is hard
 - Need to include all the stakeholders.
 - Inevitably someone will have a problem with everything.
- Other Issues:
 - Handling software vulnerabilities is a constant distraction.
 - Don't underestimate value of training.
 - Of users, administrators and management.
 - Centralization versus decentralization of control.
 - Often move to the former as trust grows.

Opportunities with Virtualization

- VMs:
 - Better job isolation and lower level monitoring.
 - Can suspend and capture suspicious jobs.
- VLANs:
 - Better isolation of job traffic from “Internet background noise” allowing for better IDS through reduced false positives.
- All require tighter integration of security tools with VM/VLAN technologies than is typical today.



GENI Security Workshop (Jan 2009)

Von Welch <vwelch@ncsa.uiuc.edu>

Thank You.



3.26 Supporting Study of High-Confidence Criticality-Aware Distributed CPHS in GENI: Sandeep K. S. Gupta

Supporting Study of High-Confidence Criticality-Aware Distributed CPHS in GENI

Sandeep K. S. Gupta

Impact Lab (<http://impact.asu.edu>)
 Computer Science and Engineering
 Affiliated with EE, BMI, BME
 Arizona State University
sandeep.gupta@asu.edu



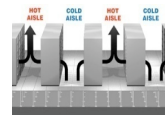
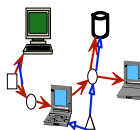
Workshop on GENI and Security – Jan 22-23, 2009



Sandeep K. S. Gupta, IEEE Senior Member

• Heads  @ School of Computing & Informatics 

Use-inspired, Human-centric research in distributed cyber-physical systems




ID Assurance	Mobile Ad-hoc Networks	Pervasive Health Monitoring	Criticality Aware-Systems	Thermal Management for Data Centers	Intelligent Container



BEST PAPER AWARD: Security Solutions for Pervasive HealthCare – ICISIP 2006.

BOOK: Fundamentals of Mobile and Pervasive Computing, Publisher: McGraw-Hill Dec. 2004



176

• TCP Chair



<http://www.bodynets.org>

•TCP Co-Chair:

GreenCom'07

<http://impact.asu.edu/greencom>

• Area Editor



Motivation

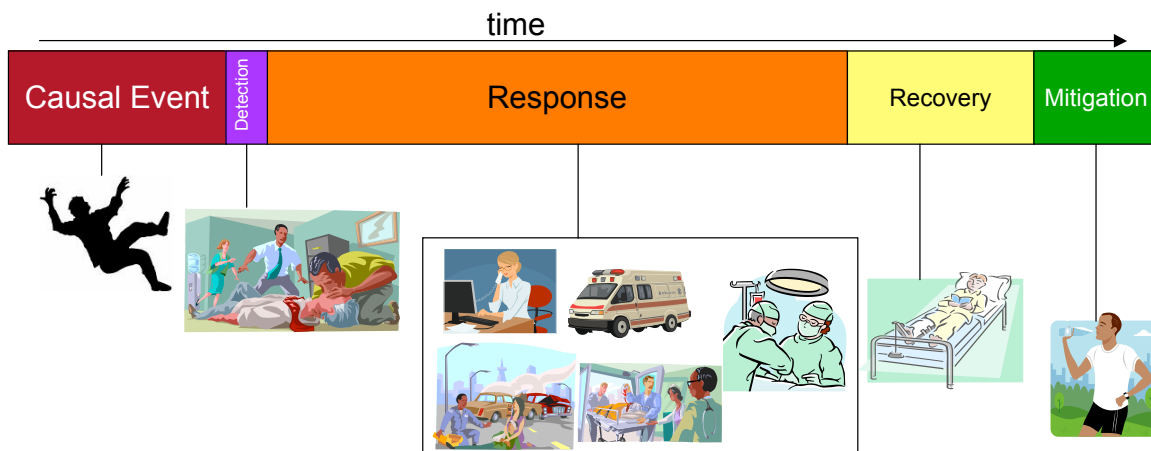
- Challenges – Traffic congestion, Energy Scarcity, Climate Change, Medical Cost ...
- Need Smart Infrastructure – distributed CPHS (Cyber-Physical-Human System (of systems))
- **Criticality-awareness**: the ability of the system to respond to unusual situations, which may lead to disaster (with associated loss of life and/or property)
 - How to design, develop, and test criticality-aware software for CPHS systems?
- **Unifying Framework** for Safe (Energy-Efficient) Spatio-Temporal Resource Management for CPHS
 - Thermal-Aware Scheduling for Data Centers and Bio Sensor Network (within Human Body)



Workshop on GENI and Security – Jan 22-23, 2009



Example Scenario



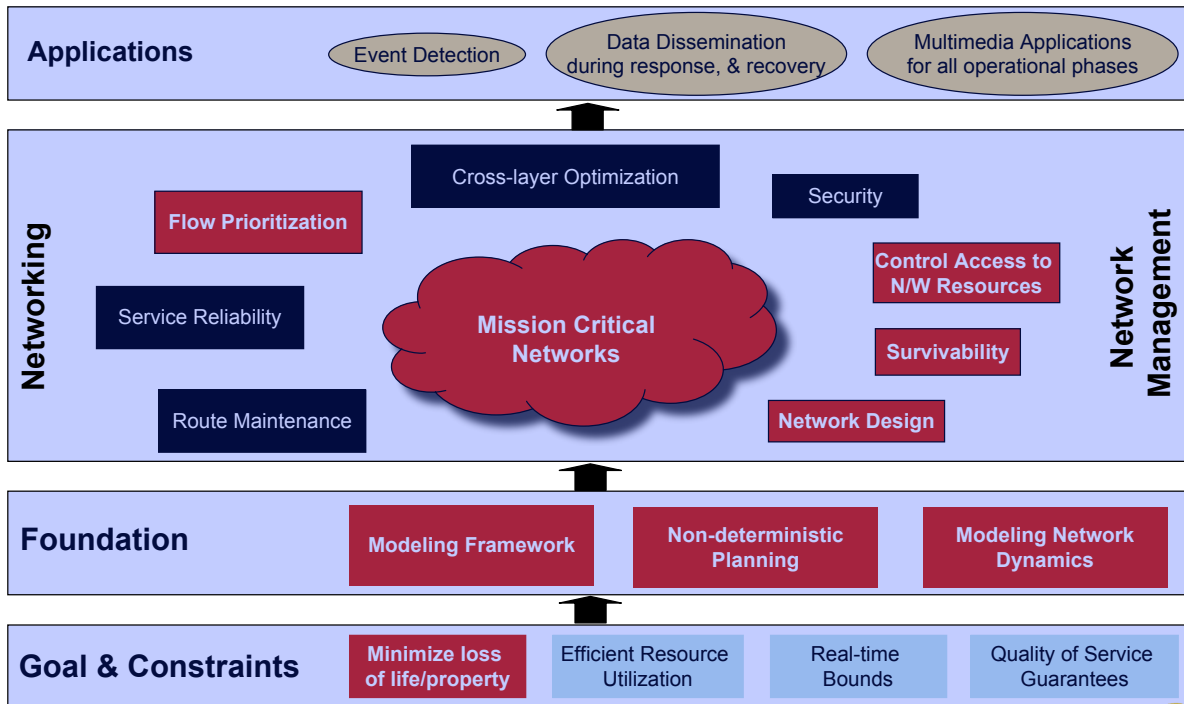
- Critical Event detected using **BSN** on the person - heart attack
- **BSN** provides patient's current health data to first responders
- Patient taken to hospital, **BSN** providing up-to-date information throughout the way.
- Information from **BSN** used by clinicians for diagnosis and treatment
- **BSN** helps in keeping track of patient recovery status
- Reduce hospital stay time.
- Control medicine dosage
- **BSN** tracks subject's health during normal times



Workshop on GENI and Security – Jan 22-23, 2009



Grand challenges for Distributed CPS



Workshop on GENI and Security – Jan 22-23, 2009



Recommendations from Real-time Embedded Systems GENI Workshop, Sep. 2006

- ▶ Recommendations for real-time and embedded networking infrastructure atop the GENI substrate
 - ▶ Uniform representation of **time and physical location information**,
 - ▶ **End to end timing predictability** across wired and wireless mobile networks,
 - ▶ **Co-existence of guaranteed, managed and best-effort QoS services**,
 - ▶ **Quantified** safety, reliability, availability, security and privacy,
 - ▶ **Scalability** across small deployments to national and world-wide deployments, and
 - ▶ Compatibility with regulatory organizations' requirements.

Properties - Cyber Physical Human Systems

- ▶ Tight coupling between physical and cyber-world
- ▶ Human-in-the-loop
- ▶ Heterogeneous entities with order of magnitude difference in capabilities, e.g. sensors, medical devices, servers, handheld computing devices, and Humans.



“HOT” Mission Critical Applications – Example of Environmental Effects on Networks

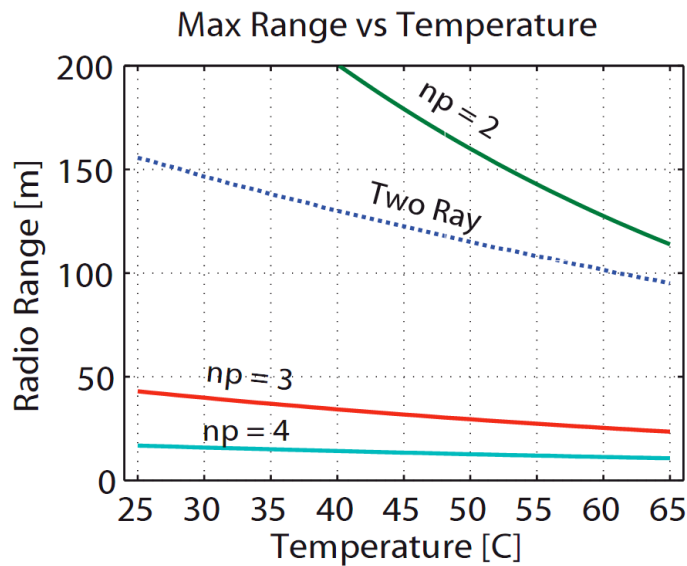


- Nodes exposed to the sun might easily reach 65C and above
- Temperature at nodes in a wildfire monitoring application have reported to reach 95C.

How to compensate for temperature effects at design/runtime?

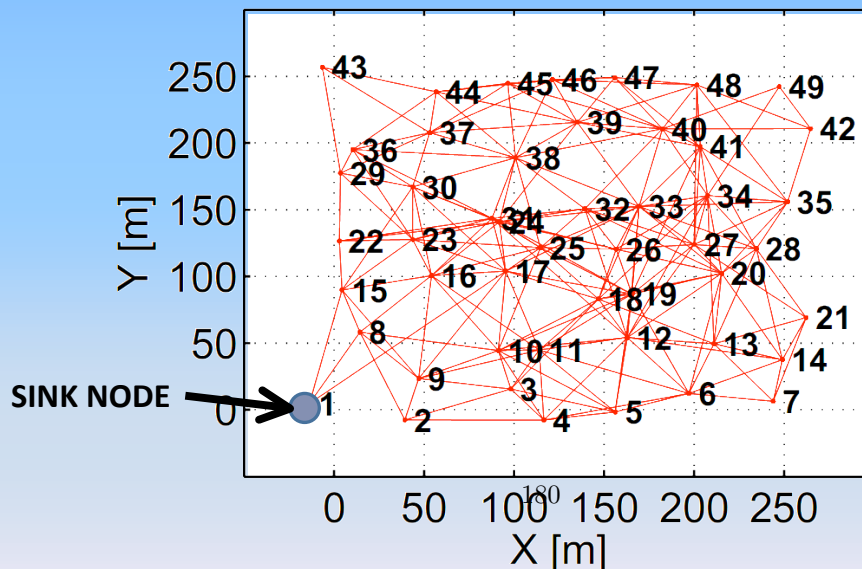


Communication Range



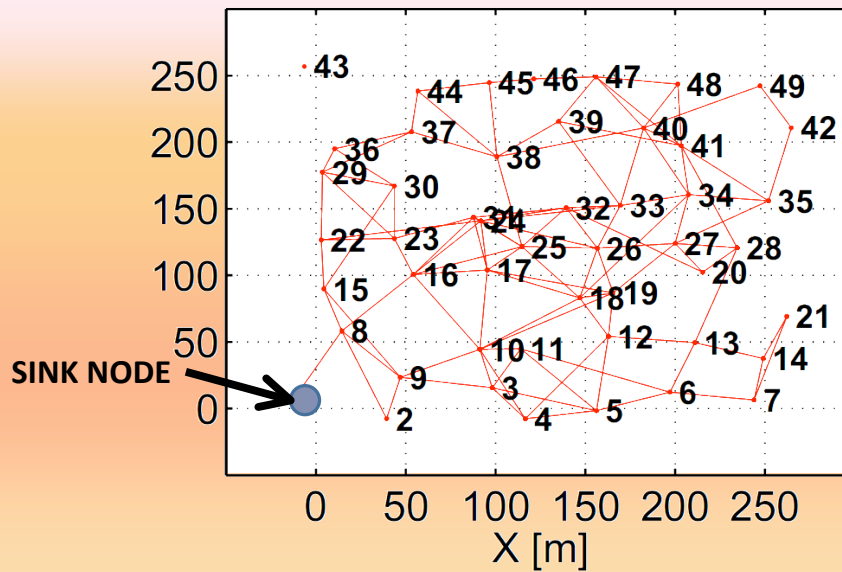
Depending on the path loss model, losses due temperature cause reduction in range comprised between **40%** and **60%** the max. value

Network Connectivity @ 25°C



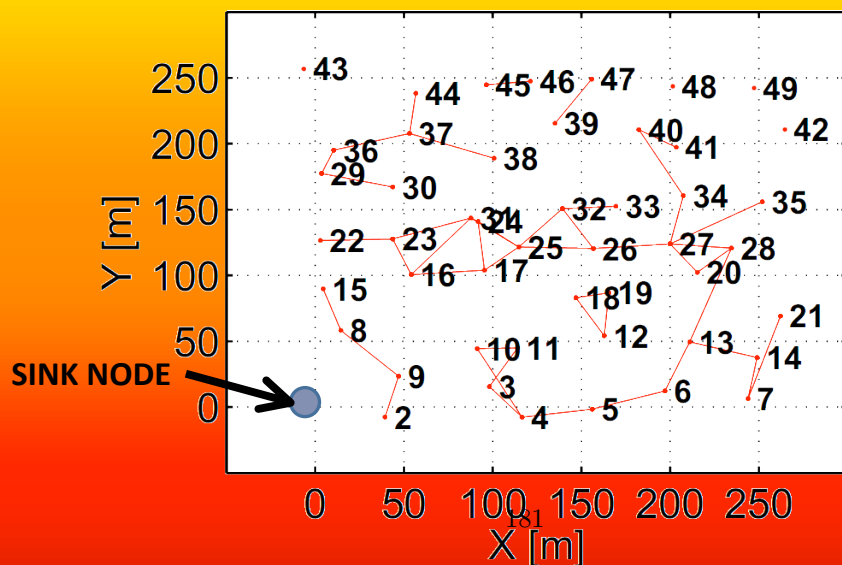
Average Connectivity = 8.94. Connected nodes = 100%.
Avg. Path Length = 2.95. **Network seems reliable.**

Network Connectivity @ 45°C



Average Connectivity = 4.57. Connected nodes = 98%.
Avg. Path Length = 4.93. **Few nodes are disconnected.**

Network Connectivity @ 65°C



Average Connectivity = 4.57. Connected nodes = 0%.
The sink is completely disconnected from the rest of the network!

Physical Aspects of CPS Security

- ▶ Modifying physical environment around the CPS can cause security breach
- ▶ Example –
 - ▶ Smart-car's theft protection system fails completely if it is fooled into thinking the car is on fire by trigger specific sensors.
 - ▶ No amount of securing all the other components will help
- ▶ The problem is compounded if security solutions for CPS depend on environmental stimuli for efficiency purposes
- ▶ Example –
 - ▶ Physiological value based security (PVS) utilizes common physiological signals from the body for key agreement
 - ▶ If one of the sensors is fooled into measuring incorrect physiological signals (by breaking the sensor-body interface), the whole process breaks down



Fundamental differences with Cyber Security

- ▶ Threat Model is fundamentally different
- ▶ The point of entry for traditional (cyber-only) is essentially cyber
 - ▶ Example – Attacker hacking a computing system through a network
- ▶ CPHS – it can be **cyber, environmental (physical), and human**
- ▶ CPHS system has several aspects each of which need to be secured–
 - ▶ Environment
 - ▶ Sensing
 - ▶ Communication
 - ▶ Processing
 - ▶ Feedback
 - ▶ Humans

Securing the environment and its interaction with other following unique to CPHS

Securing these addressed in traditional cyber security



GENI and CPHS Security Solutions

- ▶ GENI therefore needs to provide the ability –
 - ▶ To simulate/emulate diverse situations in which CPHS are deployed in real situations
 - ▶ To program the CPHS components to behave maliciously based on both cyber and environmental attacks.
 - ▶ Ability to sand-box cyber and physical components of the CPHS for evaluation various aspects of the attacks and defense mechanisms.
 - ▶ Collect feedback on security solutions' performance.
-

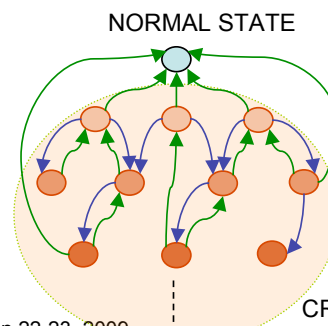
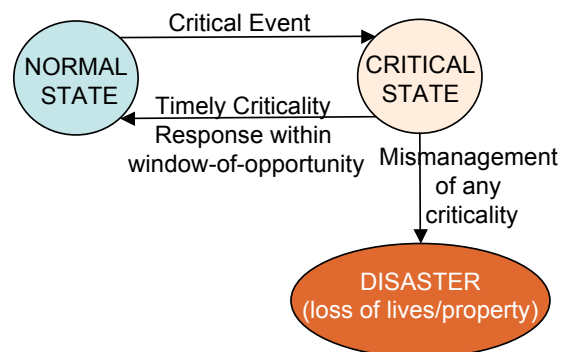
Some Results from IMPACT Lab

- ▶ [Analytical model to minimize energy overhead of pro-active protocols](#) for wireless networks
 - ▶ Classifies pro-active protocols based on periodic updates performed
 - ▶ Minimizes update overhead for all classes by finding [optimum update periods](#) based on link dynamics, network size, traffic intensity, and end-to-end reliability requirements
 - ▶ [Theory of criticality](#) capturing effects of critical events, which can lead to loss of lives/property.
 - ▶ [Probabilistic planning](#) of response actions for fire emergencies in off-shore oil & gas production platforms.
 - ▶ [Criticality-aware access control](#) policies for mission critical systems.
 - ▶ [Physiological Value](#) based security for Body Sensor Networks
 - ▶ Environment-aware [Communication Modeling & Network Design](#)
-

Our Approaches to Enable Criticality-Aware CPHS Study in GENI

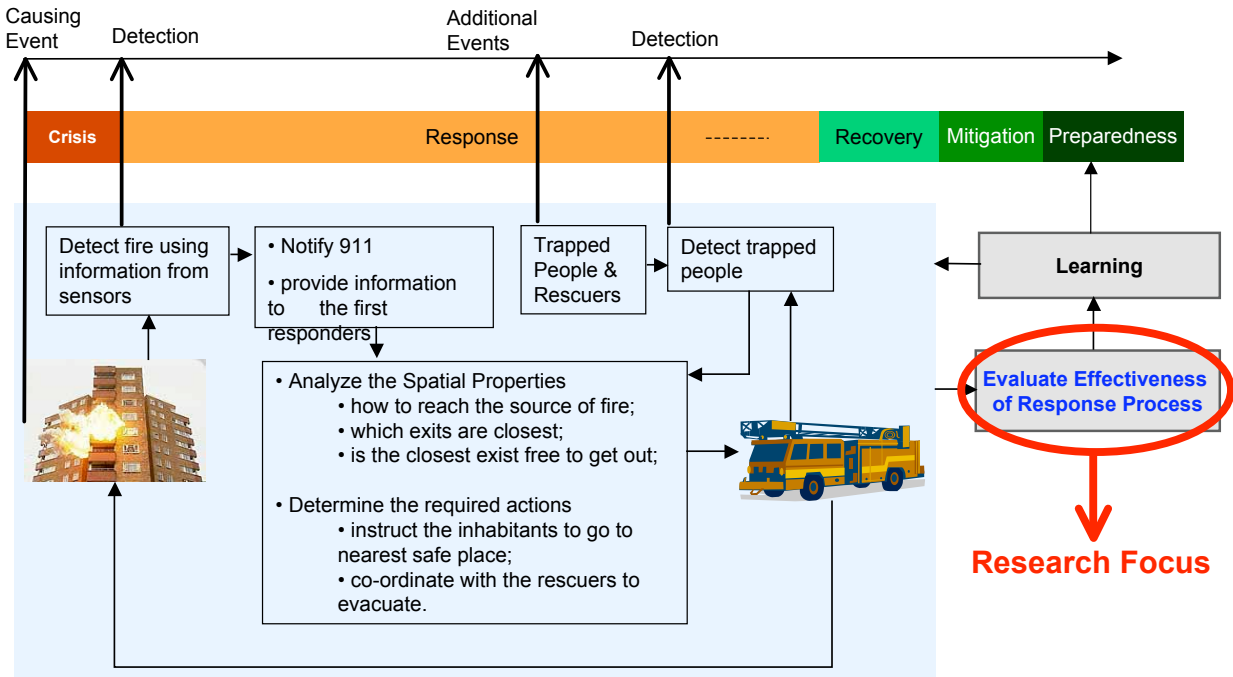
Theory of Criticality & Probabilistic Planning

- Critical events
 - Causes emergencies/crisis.
 - Leads to loss of lives/property
- Criticality
 - Effects of critical events on the smart-infrastructure.
 - Critical State – state of the system under criticality.
 - **Window-of-opportunity (W)** – temporal constraint for criticality.
- **Manageability** – effectiveness of the criticality response actions to minimize loss of lives/property.
- **State based stochastic model** capturing *qualifiedness* of the performed actions to improve manageability of critical events.
 - Probabilistic action planning to maximize manageability





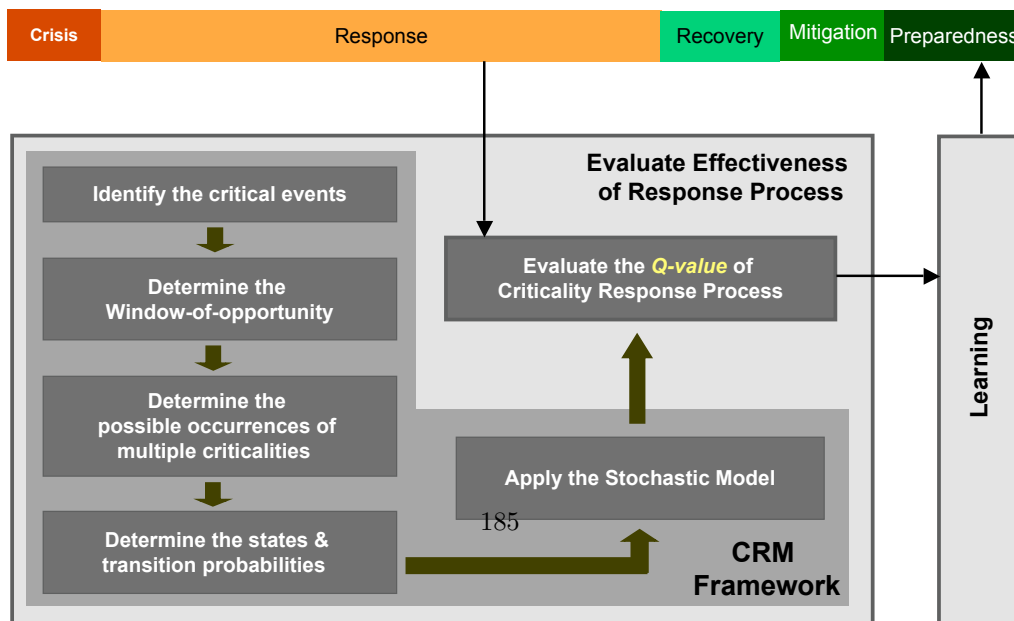
Crises Management – Fire in Smart-Building



Workshop on GENI and Security – Jan 22-23, 2009



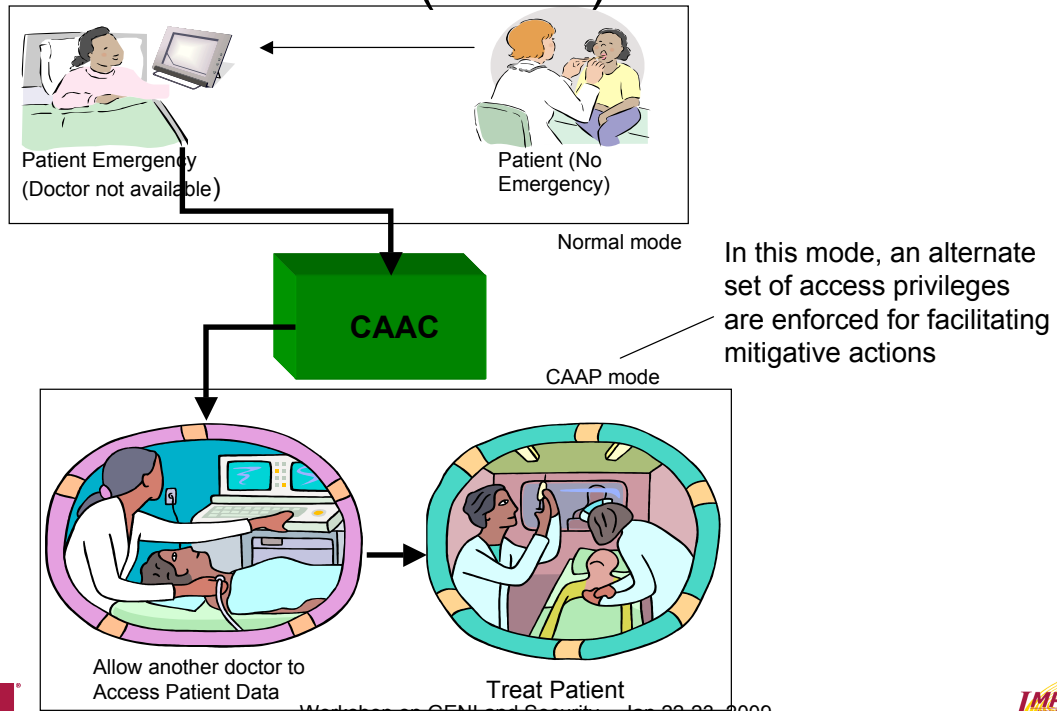
Criticality Response Modeling (CRM) Framework



Workshop on GENI and Security – Jan 22-23, 2009



Criticality Aware Access Control (CAAC)



Workshop on GENI and Security – Jan 22-23, 2009

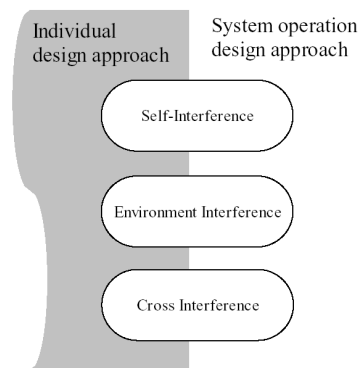
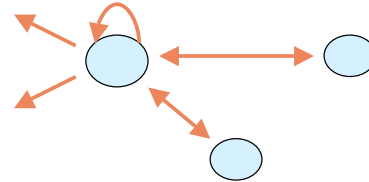


Unifying Framework for Modeling Spatio-Temporal Cyber-Physical Effects



Environmental Coupled Distributed CPS

- Terminologies
 - Self-interference
 - Environment –interference
 - Cross-interference
- Disturbance models
 - Quantitative model
 - Temporal model
 - Spatial model
 - Comprehensive model
- Individual design approach
- Network/system operation approach



23

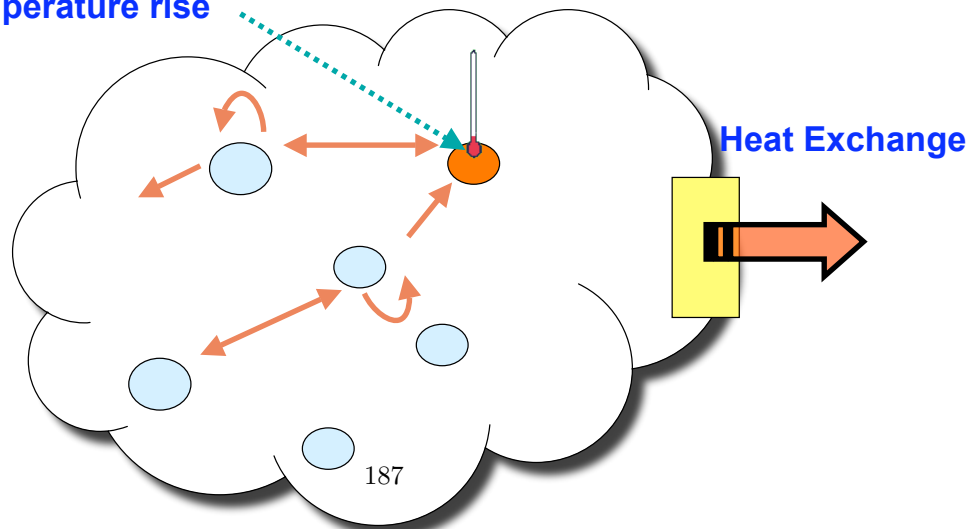


Workshop on GENI and Security – Jan 22-23, 2009



System Model

Interference cause undesired
Temperature rise



System performance depends on the thermal distribution

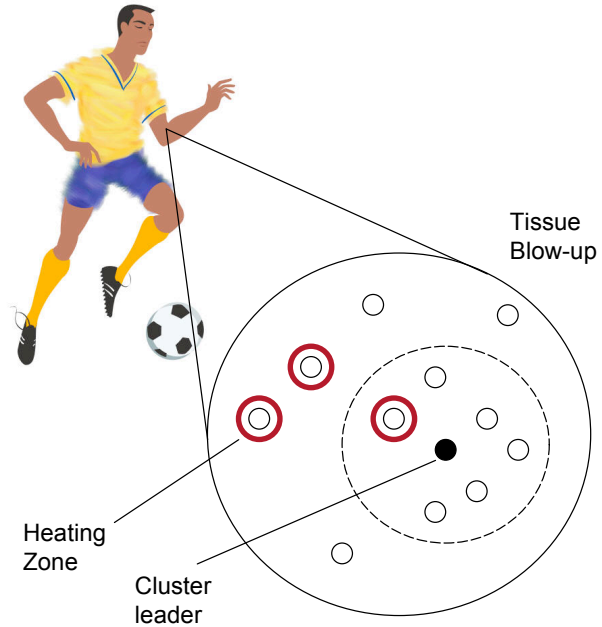


Workshop on GENI and Security – Jan 22-23, 2009



Tissue Heating

- Medical sensors implanted/worn by human need to be safe.
- Sensor activity causes heating in the tissue.
 - Heating caused by RF inductive powering
 - Radiation from wireless communication
 - Power dissipation of circuitry
- Goal: *minimize tissue heating*.
- Two solutions:
 - *Communication scheduling* for minimizing thermal effects:
 - Rotate cluster leader – balance energy usage + distribute heat dissipation
 - *Thermal aware routing*: route around thermal hotspots



Disturbance Minimization

Workshop on QoS and Security, Jan 22-26, 2009



BSN Scheduling

- Requirement**
- FCC Regulation

$$SAR = \sigma E^2 / \rho \text{ (W/kg)}$$

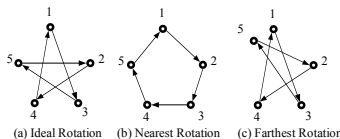
E = induced Electric Field
 P = tissue density
 σ = electric conductivity of tissue

IEEE Requirement (1g Tissue)

	Whole Body Average	SAR =	Peak Local	SAR =
CE		0.4W/Kg		8W/Kg
UCE		.08W/Kg		1.6W/Kg

Solution

- Random selection may lead to higher temperature rise
- Similar to Traveling salesman problem but with dynamic metric
- Heuristic: Leader selection based on sensor location, rotation history

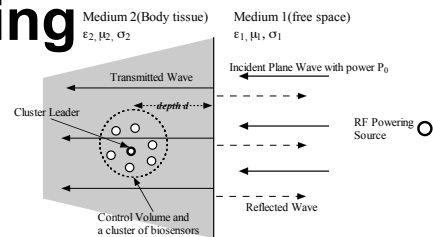


Four Approaches

- FDTD + enumeration
- FDTD + Genetic Algorithm
- TSP + enumeration
- TSP + Genetic Algorithm

System Model

- Consider only one cluster
- 2D Model
- Rotate cluster head - dist. energy consump. reduce heating



Temperature Rise: Pennes Bio-heat Equation

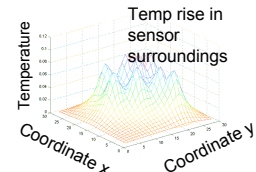
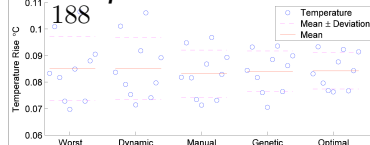
$$\rho C_p \frac{dT}{dt} = K \nabla^2 T + \rho SAR - b(T - T_b) + P_{circuitry} + Q_m$$

Heat accumulated by conduction Heat by radiation Heat by convection Heat by power dissipation Heat by metabolism

Results

Method	Result	Time
FDTD + enumeration	Optimal	720960 hrs (est.)
FDTD + Genetic Algorithm	Near Optimal	100 hrs (est.)
TSP + enumeration	Optimal	7.6 hrs
TSP + Genetic Algorithm	Near Optimal	5 min

Comparative Result



Data center Energy Consumption

What are datacenters

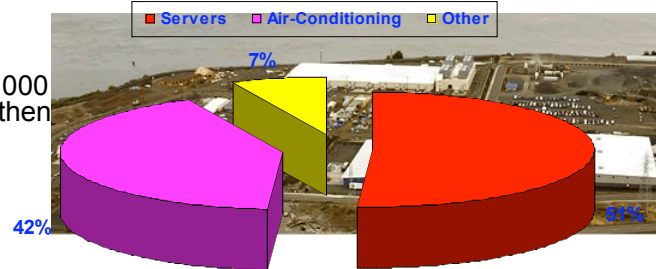
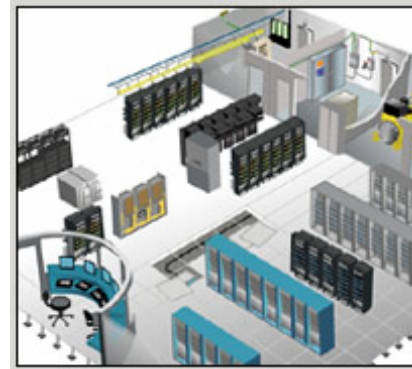
- Server farms, IT centers, computer rooms

Why they are important

- Centralized management, powerful computation capabilities
- Backbones of Internet Infrastructure

Why thermal management is important

- Improve reliability
- Reduce system down time
- Save energy cost !!
 - \$400,000 annually to power a 1,000 volume server-unit data center, then how much for this
- More than 40% is cooling cost



27

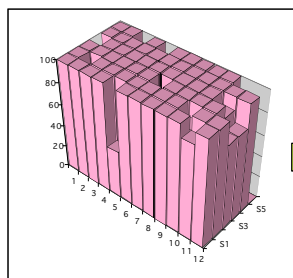


Workshop on GENI and Security – Jan 22-23, 2009

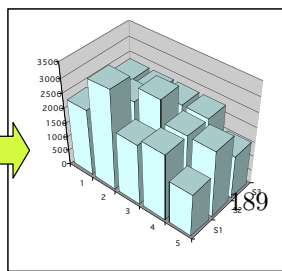


Ecosystem of Datacenters

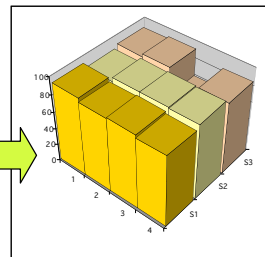
- ✓ Different task assignments lead to different power consumption distributions
- ✓ Different power consumption distributions lead to different temperature distributions
- ✓ Different temperature distributions lead to different total energy costs



Server load distribution



Power consumption distribution



Temperature distribution



Energy cost

28



Workshop on GENI and Security – Jan 22-23, 2009



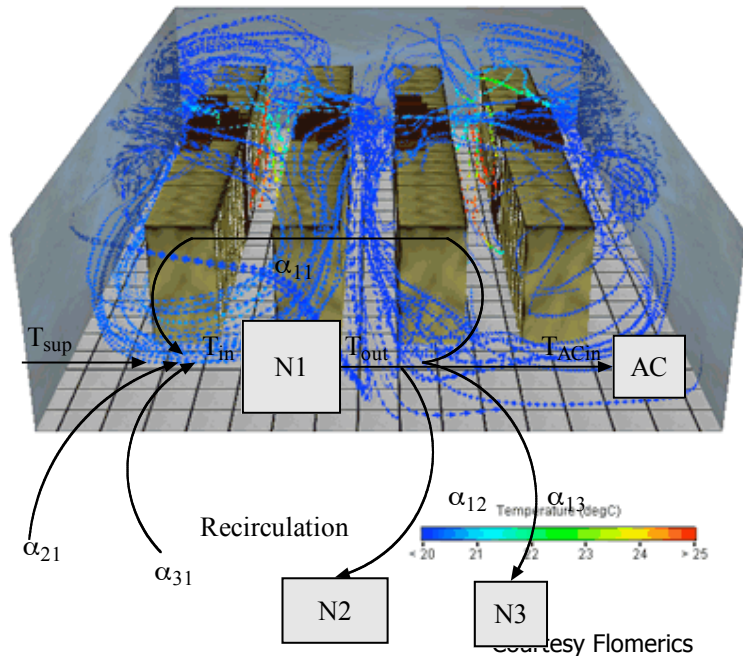
Interference in Datacenters

•Observation

- Airflow patterns are stable (confirmed through CFD simulations)

•Hypothesis

- The amount of recirculated heat is stable, can be quantified as recirculation coefficients
- Define α_{ij} as the percentage of recirculated heat from node i to node j



Workshop on GENI and Security – Jan 22-23, 2009



Two Studied CyberPhysical Applications

<i>Application scenario</i>	Implanted biomedical sensor networks	Computing nodes of data center clusters
<i>Objective</i>	find the best leadership sequence to minimize the temperature rise	find the best task assignment to minimize the energy cost
<i>Heat transfer mechanism</i>	Convection, conduction and radiation.	Convection
<i>Original numerical simulation</i>	Finite Difference Time Domain	Computational Fluid Dynamics
<i>Abstract Model: the function $F(\cdot)$</i>	Time-space function	Cross interference coefficients
<i>Placement or scheduling: the function $H(\cdot)$</i>	Temporal domain	Spatial domain

Conclusions

- Supporting interaction of Cyber and Physical Environment in GENI – essential to study important applications such as pervasive health monitoring, remote surgery etc.
- Makes GENI itself a CPHS system
- Would enable study of important issues such as subtle (or event emergent) interactions between Security and Safety

Questions ??



Impact Lab (<http://impact.asu.edu>)

*Creating Humane Technologies
for Ever-Changing World*

3.27 Privacy in the GENI Project: Robin Wilton

Privacy in the GENI project

Robin Wilton
Director

Future Identity Ltd

futureidentity@fastmail.fm
+44 (0)705 005 2931
<http://futureidentity.blogspot.com>

GENI and Privacy

GENI has visionary, constructive and beneficial aims, but some of the introduce privacy-related tensions:

- .Data-sharing, versus legitimate privacy expectations
- .Observability, versus sensitivity of real data and real traffic
- .Monitoring (and correlation) of usage patterns, versus user consent and control over collection and disclosure
- .Catering for cross-border differences in legislation and governance

Resolving these tensions can potentially

- directly benefit GENI participants*
- remove obstacles to end-user trust and adoption*
- show how to design privacy in from the outset, in the networks and distributed systems of the future*

1 – Why is privacy an issue?

There are both 'hard' and 'soft' factors -

- Identity data as a thing of value (and therefore a target)
- Cost of compliance, balance of technology and governance
- Respect for privacy as a cultural/trust/adoption factor

*When times are tough, governance ('people') measures get squeezed sooner – and harder – than technical measures
But technology isn't self-managing...*

2 – What tends to de-rail privacy discussions?

- Diversity of stakeholders and their views/goals (and who are GENI's stakeholders, exactly?)
- Lack of shared concepts and language
- “Techno-blinkers”: technology input is totally valid, but can narrow the discussion and distract attention from implementation and adoption factors
- Trying to apply a linear process instead of a cyclical and iterative (“bee-hive”) one

There are simple and well-tested models for building a common conceptual framework and creating a productive stakeholder dialogue

3 – Practicalities in the GENI context

Assume that:

- Privacy is really about disclosure... but with user control and consent;
- Privacy is about 'contextual integrity';
- Privacy is relationship-based, and can be multi-party, transitive and mediated...

(not that all those assumptions are easy to fulfil...)

- *How could privacy work if it were designed in?*
- *What outcomes could we expect?*
- *How would we get there from here?*

4 – Suggested approach

- Identify and engage with GENI stakeholders
- Establish conceptual framework for productive discussions (different perspectives, different levels)
- Contribute to a 'GENI position' on identity, privacy and governance, including cross-border use-cases
- Set goals for privacy-related GENI outcomes

How might we do GENI if we set the following goals?

- *Enhance the privacy of GENI participants*
- *Improve privacy for Internet users in general*
- *Minimize privacy-related risk to GENI*

Workshop on GENI and Security
January 2009
UC Davis

F

Thank you

Robin Wilton
Director

Future Identity Ltd

futureidentity@fastmail.fm
+44 (0)705 005 2931
<http://futureidentity.blogspot.com>

3.28 Secure Multi-Party Computation: Manoj Prabhakaran

Secure Multi-party Computation

What it is, and why you'd care

Manoj Prabhakaran
University of Illinois, Urbana-Champaign

SMC

SMC

- SMC conceived more than 30 years back

SMC

- SMC conceived more than 30 years back
- A very general concept that subsumes the bulk of theoretical cryptography

SMC

- SMC conceived more than 30 years back
- A very general concept that subsumes the bulk of theoretical cryptography
- Largely a well-kept secret

SMC: the question

SMC: the question

- Collaboration without trust?

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties
 - e.g. query with me, database with you

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties
 - e.g. query with me, database with you
 - e.g. query with me, encrypted database with you, key with someone else

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties
 - e.g. query with me, database with you
 - e.g. query with me, encrypted database with you, key with someone else
 - Goal: Nothing should be revealed “beyond the result”

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties
 - e.g. query with me, database with you
 - e.g. query with me, encrypted database with you, key with someone else
 - Goal: Nothing should be revealed “beyond the result”
 - “Ideally”: Use a trusted third party

SMC: the question

- Collaboration without trust?
 - Collaboration: compute on collective data belonging to different parties
 - e.g. query with me, database with you
 - e.g. query with me, encrypted database with you, key with someone else
 - Goal: Nothing should be revealed “beyond the result”
 - “Ideally”: Use a trusted third party
 - “Really”: Can’t agree on a trusted party. So...

SMC: the answer

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings
 - Tools: Verifiable secret-sharing, homomorphic encryptions, commitments, ZK proofs, oblivious transfer, ...

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings
 - Tools: Verifiable secret-sharing, homomorphic encryptions, commitments, ZK proofs, oblivious transfer, ...
- Simpler protocols if some trust already present

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings
 - Tools: Verifiable secret-sharing, homomorphic encryptions, commitments, ZK proofs, oblivious transfer, ...
- Simpler protocols if some trust already present
 - “Honest-but-curious”

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings
 - Tools: Verifiable secret-sharing, homomorphic encryptions, commitments, ZK proofs, oblivious transfer, ...
- Simpler protocols if some trust already present
 - “Honest-but-curious”
 - “Honest-majority”

SMC: the answer

- SMC protocol: among mutually distrusting parties, to emulate the presence of a globally trusted party
- Numerous protocols in literature for various functionalities, in various settings
 - Tools: Verifiable secret-sharing, homomorphic encryptions, commitments, ZK proofs, oblivious transfer, ...
- Simpler protocols if some trust already present
 - “Honest-but-curious”
 - “Honest-majority”
 - Simple (offline) trusted sources

SMC in GENI?

SMC in GENI?

- Where privacy is needed

SMC in GENI?

- Where privacy is needed
 - e.g. Measurement archives held by a *virtual* trusted party

SMC in GENI?

- Where privacy is needed
 - e.g. Measurement archives held by a *virtual* trusted party
 - Secure distributed storage and computation (secure unless all servers corrupt)

SMC in GENI?

- Where privacy is needed
 - e.g. Measurement archives held by a *virtual* trusted party
 - Secure distributed storage and computation (secure unless all servers corrupt)
 - May use “honest majority” in a federation

SMC in GENI?

- Where privacy is needed
 - e.g. Measurement archives held by a *virtual* trusted party
 - Secure distributed storage and computation (secure unless all servers corrupt)
 - May use “honest majority” in a federation
- Provide SMC as an “experiment support service”?

SMC in GENI?

- Where privacy is needed
 - e.g. Measurement archives held by a *virtual* trusted party
 - Secure distributed storage and computation (secure unless all servers corrupt)
 - May use “honest majority” in a federation
- Provide SMC as an “experiment support service”?
 - SMC offers a whole range of novel applications