



Addressing GENI's Risks


Chip Elliott
GENI Project Director

www.geni.net

Clearing house for all GENI news and documents



“Hey, how difficult could it be?”






What's the Goal of this Meeting?

- Agree on some of GENI's significant risks (there are plenty)
- Begin to agree on how they can be addressed
- Discuss how GENI's prototyping efforts will drive down risks in building & operating GENI



Brief Introductions 1



Key Roles and Responsibilities

GSC

- Definitive source of “what we need in GENI”
- Authors of GENI Research & Education Plan
- Technical advisory & oversight to GPO

GPO

- Project management and execution
- GENI architecture and system engineering
- Cost & schedule estimates for construction
- Authors of GENI facility construction plan
- Home for Working Groups



Brief Introductions 2



Chip Elliott
Project Director



Aaron Falk
(Community Nominee)
Engineering Architect



Kristin Rauschenbach
Substrate Architect



Henry Yeh
Project Manager




Heidi Picher Dempsey
Operations &
Integration Director



Craig Partridge
Outreach Director



Current Timeline for GENI Planning and Construction





Current Funding Status

- GSC – ✓ up and running
- GPO – ✓ up and running
- Working Groups – ✓ up and running
- System engineers – expect funding Q1 2008
 - 1 currently at work (John Jacob, risk funding)
 - 3 more will be hired
 - Also interns, outreach efforts, etc.
- Community D&P – expect funding Q1 2008
 - About \$7.5 million / year
 - Proposals due ~ mid Feb, money flows ~ mid May
 - Will try to raise more, from other agencies



Major Categories of Risk

- **Technical Risks**
 - Getting our arms around this beast
 - Keeping up with advances in technology and research interests
- **Funding Risks**
 - Is GENI eating into research funds?
 - Will Congress fund the project?
- **Getting the Scope Right**
 - Narrow vs. Broad Participation
 - Exactly what Communities are we talking about?



How We'll Build GENI

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

Please read the GENI Project Development Plan (PDP) and Project Execution Plan (PEP) for detailed planning information.



An ambitious goal

The GENI facility will allow experiments to incorporate **all the key technologies for global networks and distributed services within a 10-20 year time frame** – specifically CPU & disk farms, programmable ‘routers’, optical networks, and wireless access.

That's way too ambitious!

Exactly what wireless? or optics?

Technology becomes obsolete fast!

Overlays are all you'll ever need!

Nobody will use it – it's a white elephant!





Managing real risks

You are identifying important **risks**.

A typical “blueprint then execute” process suitable for building many kinds of predictable engineering projects (such as chemical plants) will lead to extremely high levels of risk if used for planning and building GENI.

Our plan for building GENI successfully relies on two main risk-management techniques:

- **Spiral development**
- **Federation**





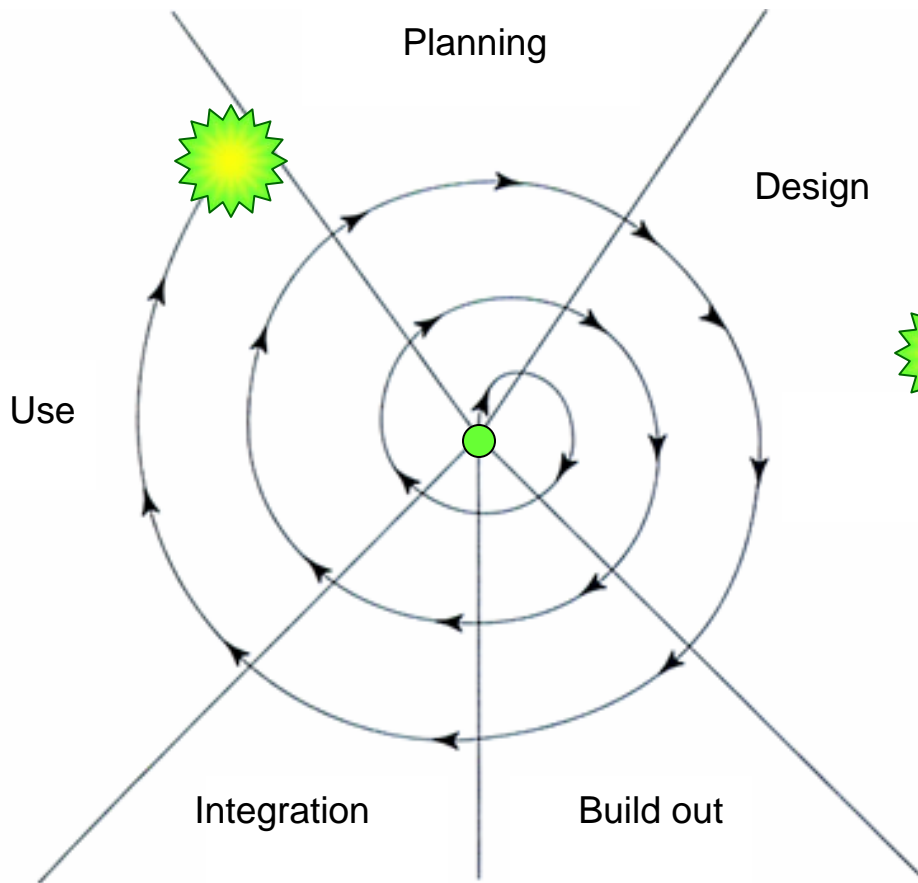
Our plan for building GENI

- Start with a clear, achievable starting point and an envisioned “ultimate goal”
- Begin prototyping and trials immediately
 - Gain **practical experience** with prototypes, and adjust “wishlists” and requirements as we go
 - Make **realistic estimates** of cost and operational complexity based on early experience with prototype systems, rather than guess-work
 - Add features, complexity, and new technologies incrementally, based on experience to date
- Repeatedly assess GENI’s current risk and usefulness as planning and construction unfold, and adjust plans accordingly



Spiral Development

GENI grows through a well-structured, adaptive process



Strawman GENI Construction Plan

- An achievable **starting point**
Example: Rev 1 “narrow waist”, federation of multiple substrates (clusters, wireless, regional / national optical net with early GENI ‘routers’, perhaps some existing testbeds), Rev 1 user interface and instrumentation.



Envisioned **ultimate goal**

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

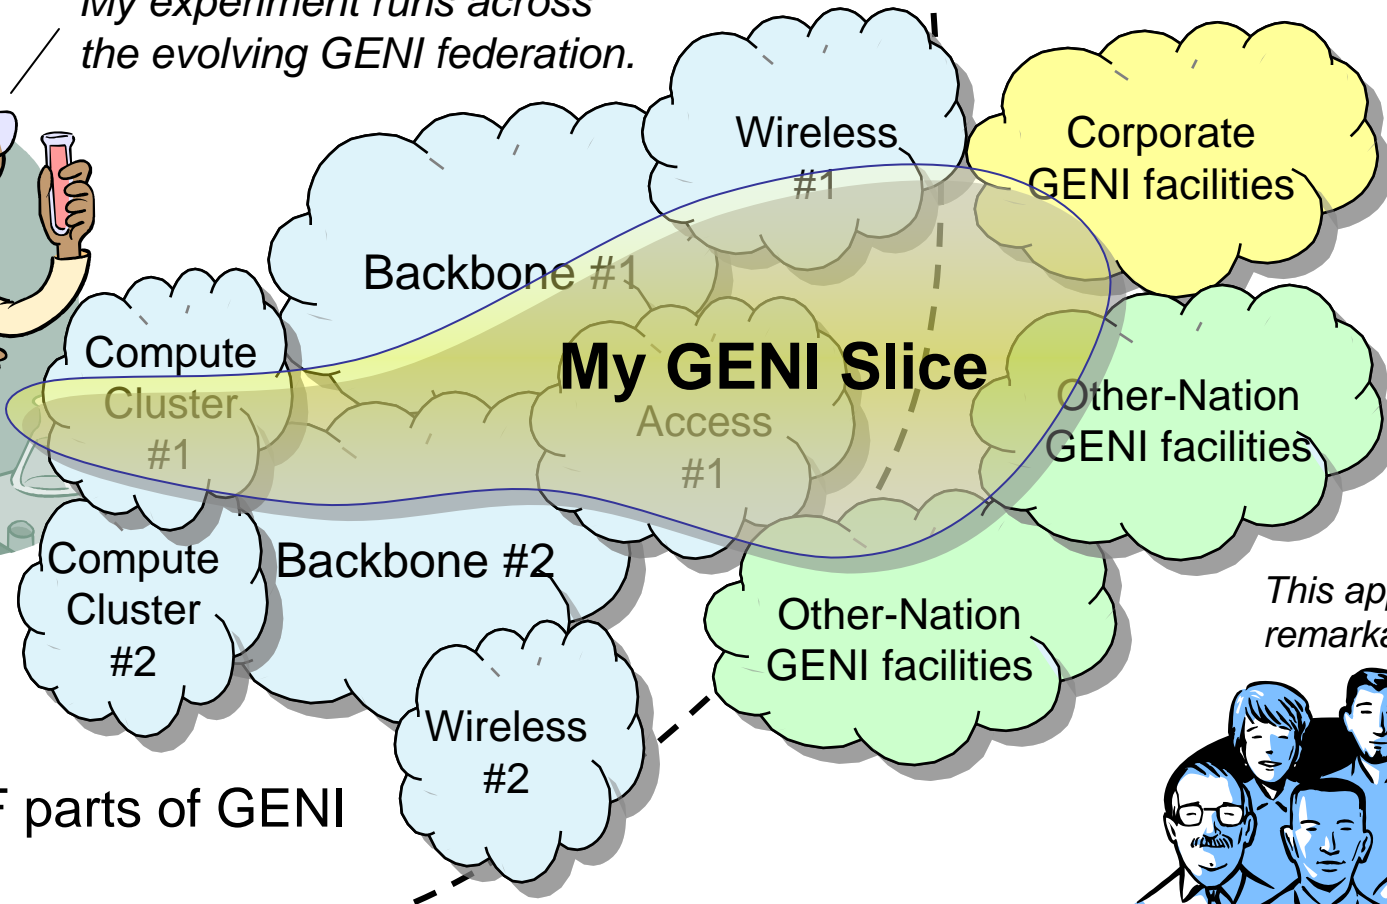
- **Spiral Development Process**
Re-evaluate goals and technologies yearly by a systematic process, decide what to prototype and build next.



Federation

GENI grows by “gluing together” heterogeneous facilities over time

My experiment runs across the evolving GENI federation.



This approach looks remarkably familiar . . .



NSF parts of GENI

Goals: avoid technology “lock in,” add new technologies as they mature, and potentially grow quickly by incorporating existing facilities into the overall “GENI ecosystem”



It's all about managing risks

The Central Goal of GENI Planning and Construction



We'll take it little by little. Those parts of GENI that are widely used will grow; those that aren't, won't get more funding. But it won't be impromptu or ad hoc – we will follow a well-defined, formal process throughout: spiral development.

I see. We are avoiding an “all or nothing” gamble – we don't try to specify all of GENI right now, then live with it for the next 20 years. Thank heavens!





Moral of this story

- GENI has real risks . . .
 - Many have been identified already
 - Others will emerge as prototyping / construction get underway
 - The “white elephant” risk is certainly real, as are many technological risks including rapid obsolescence
 - Accurate understanding of operating expenses will be critical
- Risk management is central to GENI planning
 - “Winging it” would almost surely lead to disaster
 - Systematic, formal processes must be used identify and drive down risks throughout planning and construction
 - Spiral development and federation greatly reduce risk
- **Rapid prototyping should begin immediately**, as a key technique for risk reduction in GENI’s planning phase



Major Categories of Risk

- Technical Risks
 - Getting our arms around this beast
 - Keeping up with advances in technology and research interests
- Funding Risks
 - Is GENI eating into research funds?
 - Will Congress fund the project?
- Getting the Scope Right
 - Narrow vs. Broad Participation
 - Exactly what Communities are we talking about?



Is GENI eating into network research funds?

- **NO** – these are “infrastructure” funds. GENI is not shrinking the available funds for network research.
- **JUST THE OPPOSITE** – the GPO is working hard to get additional money from outside the NSF to use for GENI prototyping, thereby **growing** the amount of money for the CS community.

CISE research \$\$\$ —————> research

CISE infrastructure \$\$\$ ———> GENI planning & prototyping (now)
GENI operations (later)

MREFC \$\$\$ —————> GENI construction (later)

Note that GENI operational costs must fit into CISE’s infrastructure budget, so they will be a key design driver for GENI.



Will Congress Fund GENI ?

- You are smarter than I am if you can predict what Congress will do in the future.
- NSF has already set aside funds for the GENI Planning Phase; that's why you're here today.
- When there is a coherent, realistic GENI construction plan, NSF and Congress will decide if and when it should be funded; that point is at least 3 years in the future. The \$350 million estimate was just that – an estimate.



Major Categories of Risk

- **Technical Risks**
 - Getting our arms around this beast
 - Keeping up with advances in technology and research interests
- **Funding Risks**
 - Is GENI eating into research funds?
 - Will Congress fund the project?
- **Getting the Scope Right**
 - Narrow vs. Broad Participation
 - Exactly what Communities are we talking about?



Getting the Scope Right

- Narrow vs. Broad Participation
 - The GPO strongly believes in a “big tent” GENI, but understands the potential for loss of decisiveness.
 - We think our plans for community involvement offer many levels of participation, and provide good tools for bringing new people into GENI while maintaining brisk forward progress.
 - How can we improve our plans?
- Exactly what Communities are we talking about?
 - At the very least, the “hard core techie” communities of distributed systems, networking research, security and privacy.
 - The GPO strongly believes that the following kinds of communities will also be extremely important: theory, economics, ethics, law and society
 - Others? What do you think?
 - How can we best structure GENI planning to usefully engage these communities?



We need good, creative ideas . . .

- We have structured GENI's planning to identify and drive down technical risk through analysis, prototyping, and trial integration.
 - How can these plans be improved?
- We plan for early operational experience with GENI so its construction and operational costs don't break the bank.
 - Are there additional ways to quickly gain solid operational experience?
- How best to usefully engage additional research communities that can really add value to GENI?
- Where / how to find near-term, additional funding from outside the NSF for more community prototypes?

The key is to identify risks and then ***do something*** about them!

- Bringing in other communities
 - GENI may not deliver on security needs, risk of too high expectations; need more engagement from security & privacy communities
- Real user traffic
 - Attractive apps: GENI (spam-free) email?
- Too successful
 - Burden of supporting a production network (rigidity)
- Insufficient momentum
- Operational risks
 - Challenges in keeping network operational as it grows; expectations for robustness will be low during prototyping
 - Need clarity in interface between experimental services & operations
 - Need buy-in from regional nets, CIOs for deploying GENI (e.g., non-IP) last-miles
- Funding
 - Need economic & resource mgmt models that are self-sustaining
- Engagement
 - Risk of lack of industry involvement or opposition