



SWITCH

The Swiss Education & Research Network

GENI Network Virtualization Concepts

Alexander Gall <gall@switch.ch>

22.2.2007

EFNI Workshop, Amsterdam

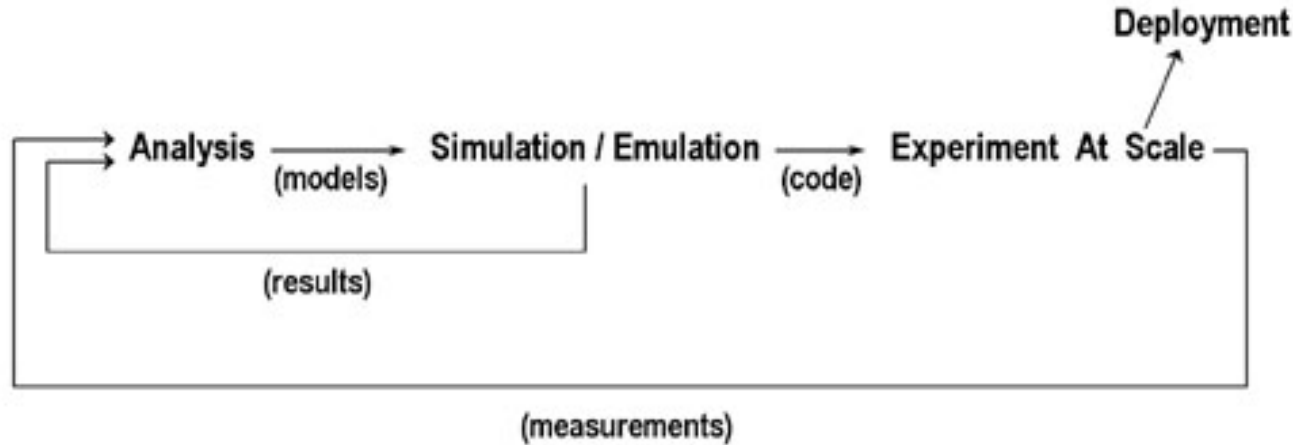
GENI: Global Environment for Network Innovations.

- Initiative planned by the NSF CISE (Directorate for Computer and Information Science and Engineering). First workshops in early 2005, first “Project Execution Plan” (PEP) from January 10 2006.
- New PEP planned for early 2007, draft not yet available
- Large-scale experimental facility to help developing the “future Internet”. Includes
 - Research program (strong coupling to FIND “Future INternet Design” NSF project)
 - Global experimental facility
- Funded by NSF through grants for US research community
- Long-term project (15-20 years, construction phase 5-7 years), budget 370M USD (Jan 2006)
- <http://www.geni.net/>, <http://www.nsf.gov/cise/cns/geni/>

- Security and robustness
 - Eroding trust, may end up with “gated communities”
 - No “security architecture”
 - Must work in times of crisis, “at least as reliable as telephone system”
- Wireless communication, mobility of hosts and networks
 - No true mobility (addressing/routing)
 - Little location awareness (resource location, peer discovery)
 - Security difficult in changing hostile environments
- PAN, sensors
 - Not reachable from the Internet, hard to extend across the Internet
- Network management
 - Very heterogeneous
 - Easy to make mistakes, hard to debug

Incremental changes are not good enough, need “clean slate” design

Research cycle

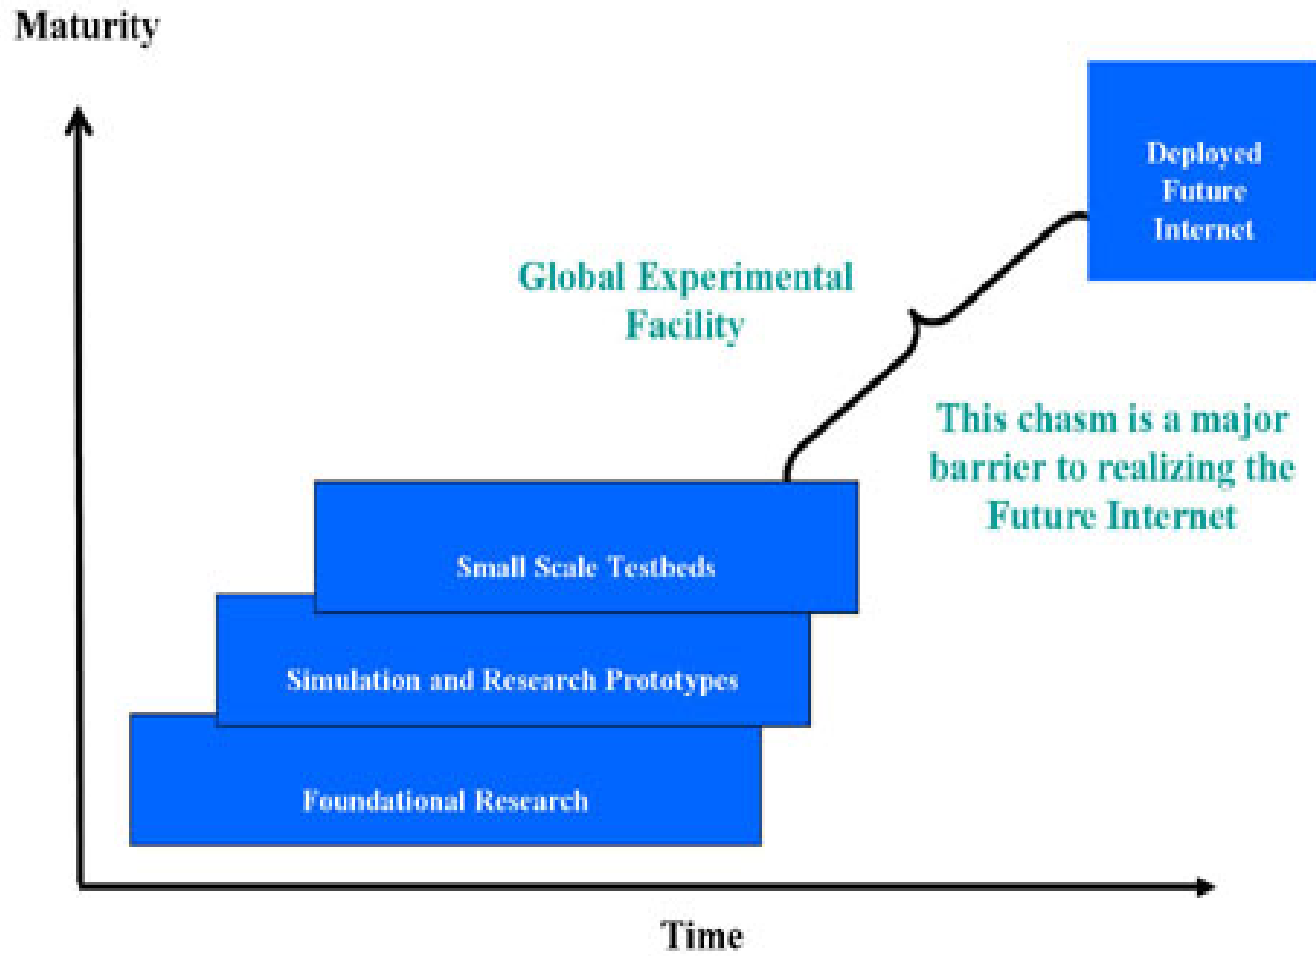


- Simulation/Emulation
 - Backed by conjecture rather than empirical data
 - Overly simplistic modelling (topology, workload, device failure etc.)
- Theoretical Analysis
 - Mostly even more detached from reality

Extensive experimentation is required to understand complex new protocols. **Need real users, real traffic, real implementations, real measurements.** However, ...

Two types of testbeds

- Production-oriented
 - Only incremental changes to existing technologies
 - Stable implementations
 - Real traffic from real users
 - Conservative
- Research-oriented
 - Radical new designs (“clean-slate”)
 - Bleeding edge implementations
 - No real users/small group of dedicated testers
 - Synthetic traffic



- Collection of globally distributed hosts (700+ hosts, 300+ sites, 30+ countries)
- Run common software for management
- Key concepts:
 - Distributed virtualization
 - Each physical node hosts any number of virtual nodes
 - A collection of virtual nodes forms a *slice*
 - Nodes within a slice form an overlay network on top of the Internet
 - Programmability: virtual nodes are freely programmable
 - Opt-in per-user/per-application (through proxies, gateways etc.)
 - Attracts real users (3-4 TB/day)
 - Demand drives deployment
- New paradigm
 - Simultaneous support for real users and clean-slate design
 - Provide plausible deployment path

Limitations

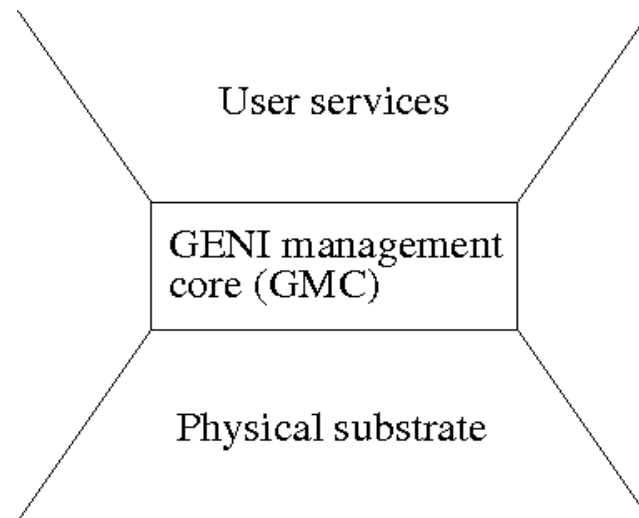
- Commodity PC hardware, no way to incorporate other node or link technologies
- Artefacts from underlying network may affect experiments
- Notoriously underprovisioned (node hardware, network capacity at hosting organizations)
- Steep learning curve
- Largely designed and supported by graduate students

GENI extends PlanetLab idea to

- Support richer set of node technologies
- Expose low-level link behaviour by slicing through to layers 1/2
- Instrumentation of all components for measurement
- Supports federation

- Multiple simultaneous experiments, isolated or interconnected
 - virtualization
- Generality: must be able to change all aspects of networking
 - programmability
- Diversity in technology
 - extensibility
- Support real users
 - must reach to the edge of the network, support long-term experiments
- Fidelity: expose functionality at the right level of abstraction
 - real hardware and topology
- Must run for a long time
 - federation, sustainability

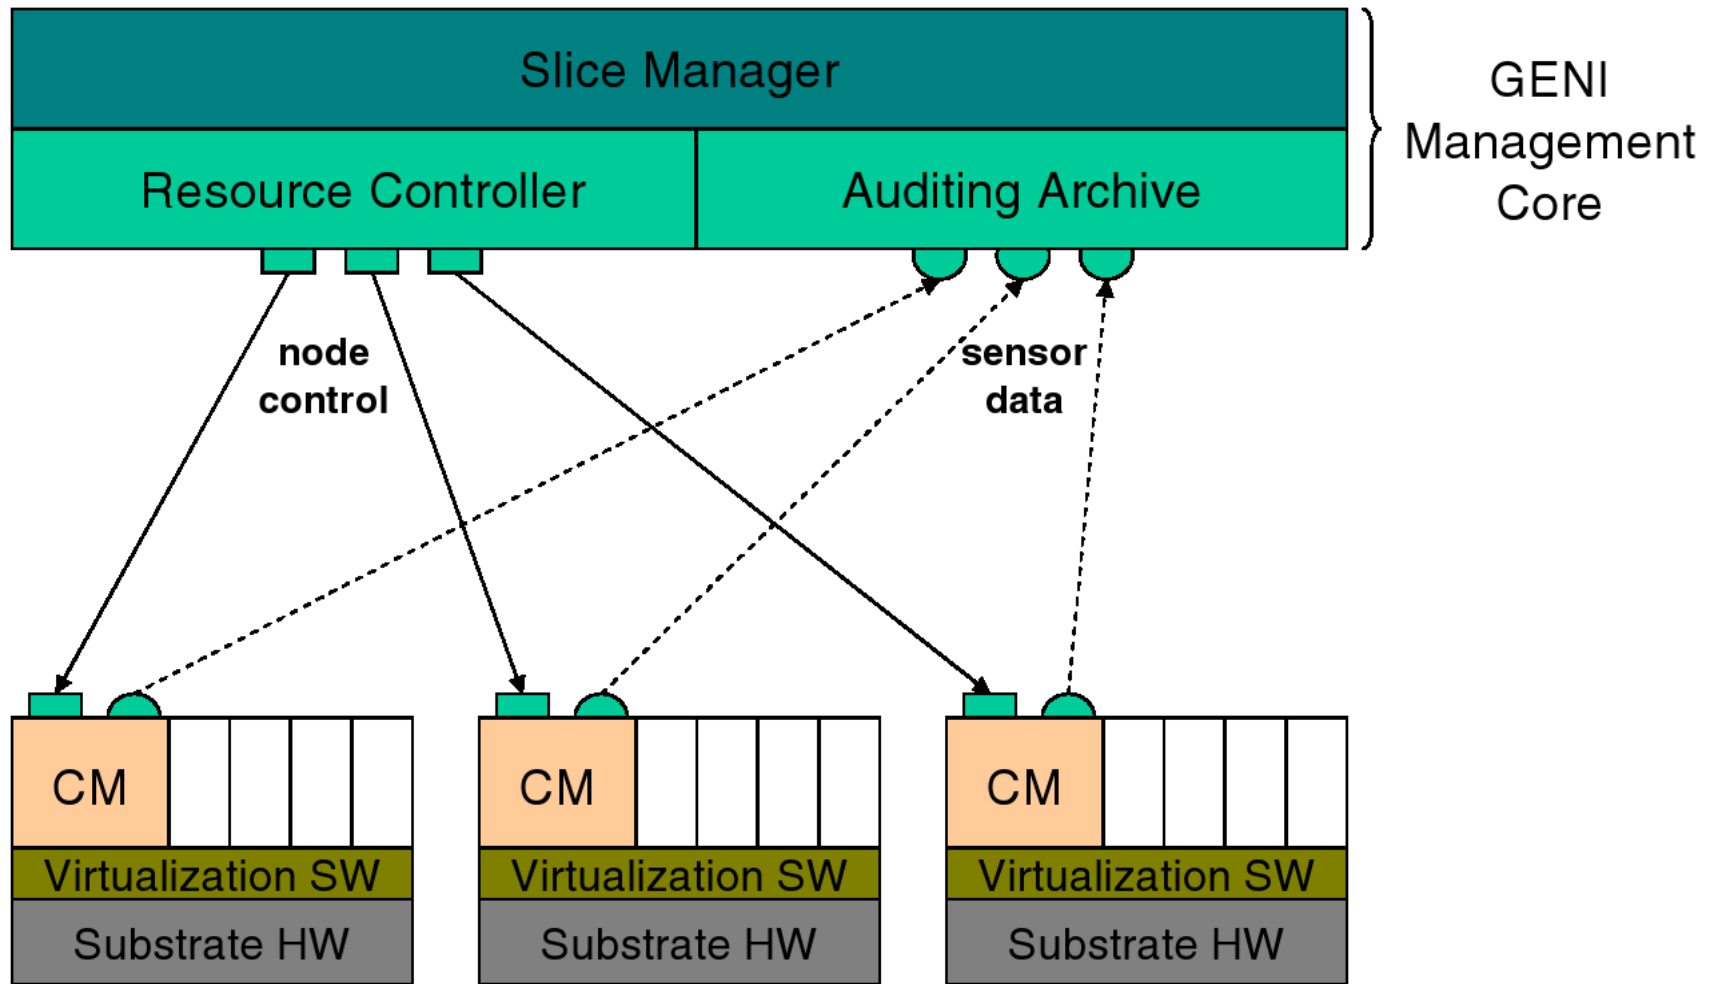
Layered design



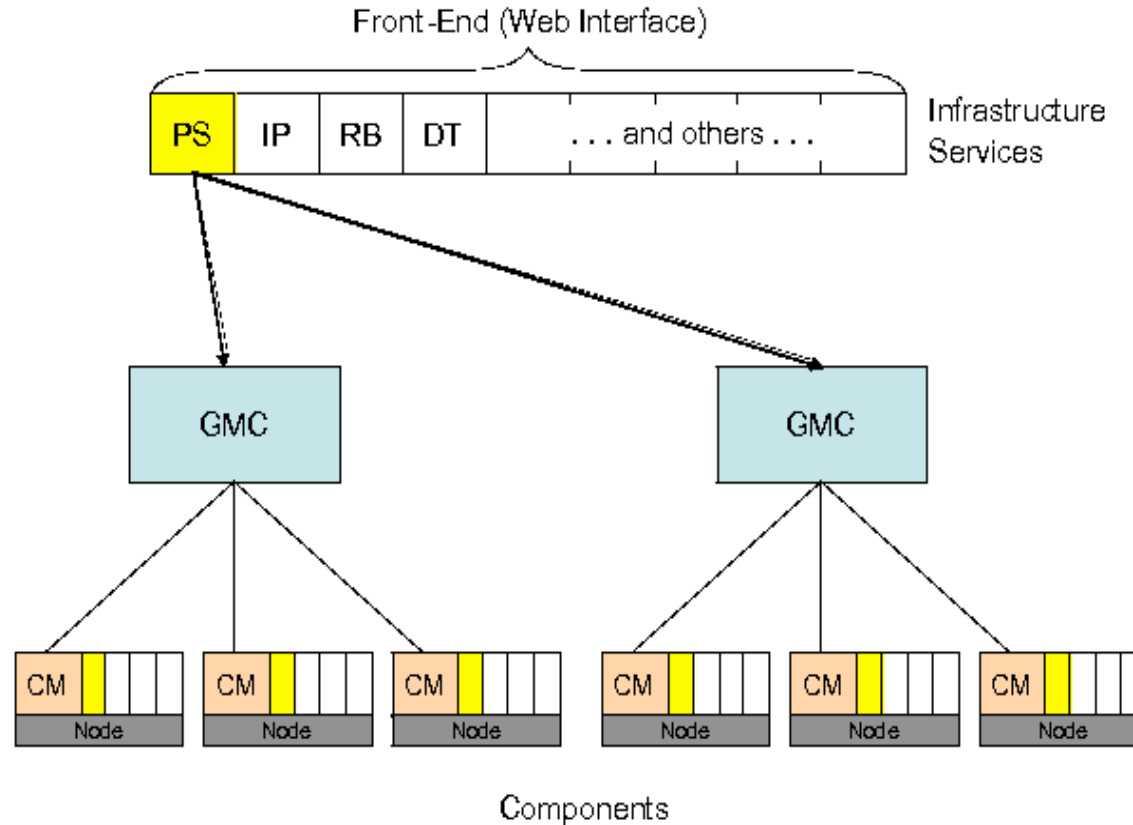
- Physical substrate: collection of network and computing devices (links, routers, processors, storage, ...)
- User services provide access to the GENI facility to researchers
 - Manage slices
 - Collect and retrieve measurement data

Defines stable, long-lived framework while user services and physical substrate can change and evolve (“waist of the hourglass”)

- Namespace
 - Unique naming of objects
 - Basis of authorization
- Abstractions
 - Component: virtualizable collection of resources that form an entity (e.g. a computer hosting virtual machines)
 - Component Manager: controls a component, exports well-defined interface, provides access to instrumentation
 - Aggregate: set of components
 - Sliver: virtual partition of a component
 - Slice: set of slivers that contains an experiment
- Operations on components, aggregates, slices (create, destroy, ...)
- Data objects




GENI supports federation: independent infrastructures can be plugged together through the GMC.

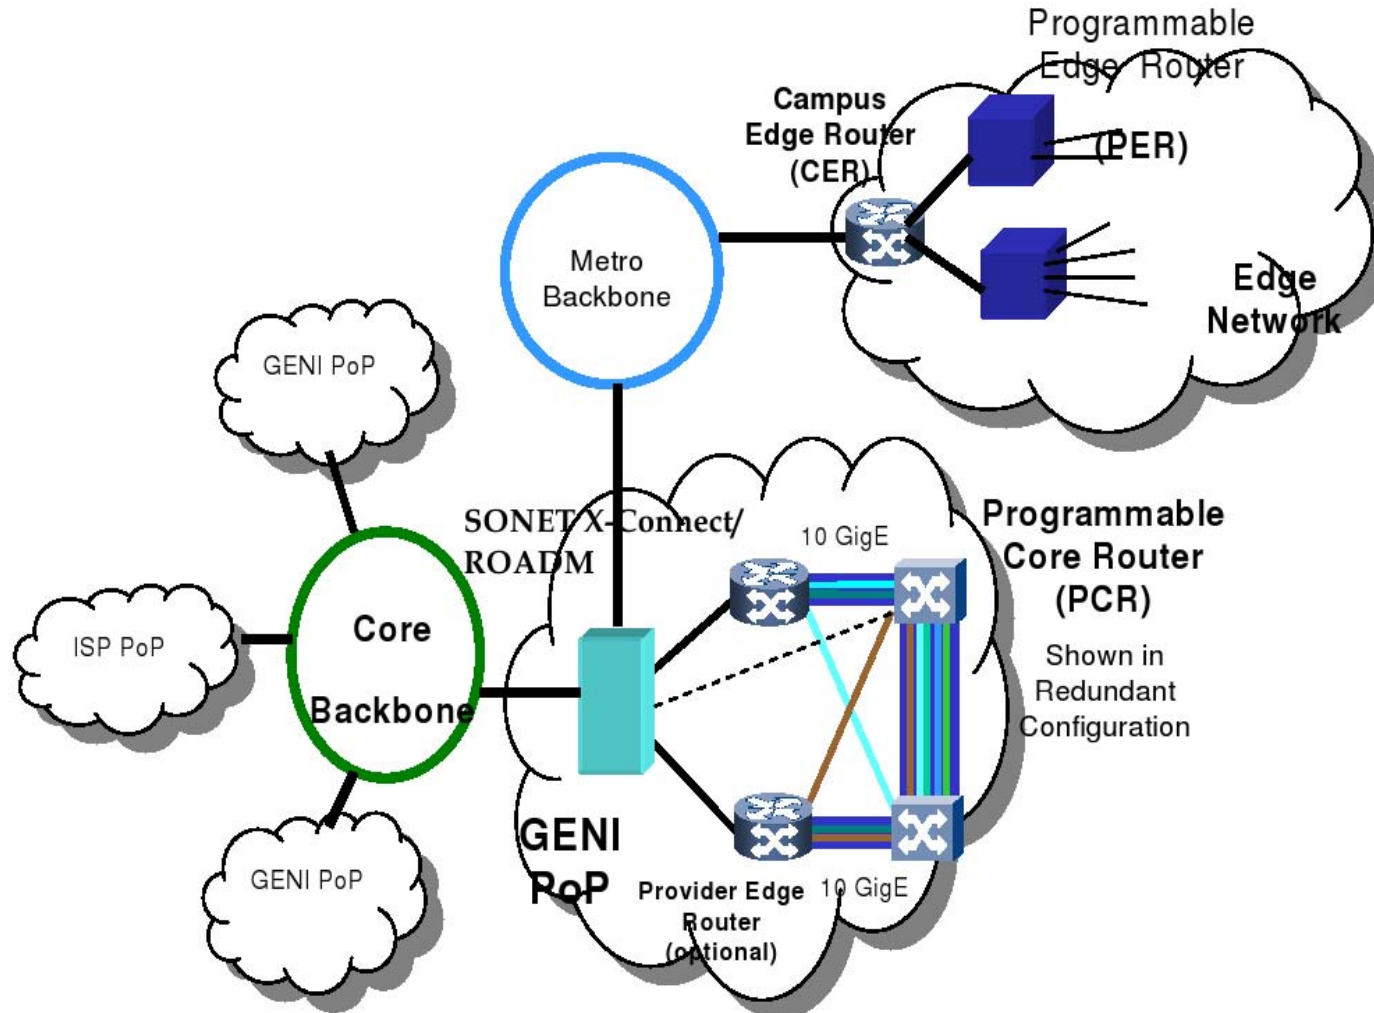



PS: Provisioning Service, IP: Information Plane, RB: Resource Broker, DR: Development Tools

Instantiation of GENI within the NSF project is called “GENI proper” or “NSF GENI”, consists of

- Links
 - Fiber backbone with ~25 PoPs (National Lambda Rail/Abilene?)
 - Tail circuits (lambdas, MPLS, tunnels, ...) to connect edge sites and interface with the legacy Internet
- Nodes
 - Programmable core routers (developed as part of the project)
 - Optical switches (e.g. SONET cross-connect, Reconfigurable Add-Drop Multiplexer)
 - Programmable edge devices (routers, hosts, access points)
- Wireless subnets
 - 802.11, WiMax (802.16), Cognitive Radio Subnets, Sensor Subnets, programmable devices







- New network technology cannot be adequately tested in traditional test facilities
- PlanetLab introduces new paradigm that help bridging the “chasm”
 - open (not restricted to project participants)
 - realistic (real users, real traffic)
 - based on virtualization and programmability, allows parallel experimentation with vastly different technologies on all layers
 - opens deployment path (via long-lived experiments)
- GENI extends this idea
 - more general
 - extensible
 - support of federations
 - includes lower layers



SWITCH

The Swiss Education & Research Network