

SDI Research Challenges and Opportunities

Ilya Baldin <ibaldin@gmail.com>, Jeff Chase <chase@cs.duke.edu>

Challenges

We believe one of the central challenges and opportunities for future SDI research comes from investigating scaling. A large number of problems left untouched by GENI relates to the various scaling challenges in building federated, deeply-networked, software defined infrastructures. While many SDI/N/X solutions are being developed, none of them have a proven scale. Those scaling challenges have multiple dimensions, which include:

- The scaling properties of the federation mechanisms (trust, resource management etc.) as the number of participants and the respective amounts of contributed resources grow.
- The scaling properties of resource orchestration mechanisms operating in an environment with large number of federated providers that support varying degrees of visibility of their resources to each other and their customers.
- The scaling properties of the resource provisioning mechanisms (e.g. developing better cloud solutions with tighter network integration) as the size of the federated infrastructure increases.
- The scaling properties of the underlying low-level virtualization mechanisms (e.g. in supporting the fluid migration of computational activities across this federated infrastructure).
- The scaling properties of resource monitoring and measurement systems, necessary for day-to-day operations, experiments, computational steering and other tasks.

Another set of challenges comes in the area of resource orchestration, especially as it relates to the intersection of topology embedding and NFV. The deep integration of computational, storage and networking resources brings new challenges for orchestrating resources from different providers and “stitching” these resources together to achieve usable (and useful!) arrangements. While NFV has been identified as a solution for the problems facing telecom operators, there are opportunities to expand this concept to support a broader set of applications, e.g. in supporting computational domain sciences.

A related set of challenges lies in designing and implementing generic abstraction, algorithms and data representation and handling mechanisms that support a broad range of operations, like search, subgraph embedding, view-based models, merging of distributed representations of resources multiple providers into a consistent whole that can be used for embedding, provisioning and monitoring in an SDI system. This problem has a particular bearing on SDXs, which must be able to inter-connect multiple providers in flexible ways, yet without compromising their privacy.

Similarly, SDXs bring a host of authorization issues for the operations they must support in order to interconnect providers in a dynamic fashion, allowing for interposition of peer-specific policies to control traffic crossing their boundaries.

Testbed Requirements

- Wide distribution of sites
- A mixture of scales of deployment
- Heterogeneity of hardware
- Connectivity to non-federated resources – instruments, supercomputers, data repositories

Prior Background

Chase and Baldin have served as PIs in a number of GENI projects, which included the development of the ORCA control framework for GENI and the deployment of ExoGENI testbed as part of GENI federation. Chase is currently a Co-PI in the SILVER NSF Frontiers award focusing on novel cloud security models. Baldin is currently a PI in an NSF NeTS project titled SERPENT designing novel representation and query techniques for semantically rich representations of networked infrastructure.