Target Research Area: Generic Virtualization Architecture (GVA) for Cyber-infrastructure
Reporter: Jerry Sobieski, Chief Research Officer, NORDUnet
Collaborators: Dr. Bijan Jabbari (George Mason Univ.), Ciprian Popoviciu (East Carolina Univ.), Dr. Susanne Naegele-Jackson (Frederick Alexander University (Erlangen DE)), others.

Importance: An "architecture" specifies how things work, and how they work together. While there are many technologies and services that claim to be "virtual", most such products are simply a software version of a product or technology. There is no ***architecture***. A *virtualization* architecture defines how virtual objects behave *independently* of their underlying technology or implementation. Virtualization will allow users to design applications in the "concept plane" on a "glass canvas" and then reliably infer and predict how those distributed applications will perform when deployed into infrastructure. A formal Generic Virtualization Architecture (GVA) offers a common generic object model for virtual resources, and a common set of life cycle primitives to manipulate those virtual resources. These aspects enable intelligent automated software agents to map the user's concepts to appropriate, optimized, and efficient infrastructure components. Groups of virtual objects can be composed to build [virtualized] services and applications. And such virtualized CI slicing can re-invigorate production networks by providing a means of deploying new services within fully functional yet insulated virtual environments that effectively manage risk across a wide variety of services and maturity levels. But a thorough treatment of virtualization introduces complex issues: efficient conforming placement of resources in the infrastructure; isolation of virtual objects; deterministic performance; multi-domain construction of composite virtual environments; performance monitoring & verification; scaling and security of virtual objects and protocols; virtualization of sensors and instruments; virtualization of mobile services and rf spectrum; etc.

Existing Facilities: Research infrastructures such as Chameleon/CloudLab, FABRIC, and PAWR offer facilities to develop and test specific software functionality. These will be part of the virtualization strategy. But they are research oriented and it is unclear whether they can provide multi-domain scale "pilot" services with real users, networks and applications.

The Case for a MERI. We need a Virtualization MERI strategy that addresses two key requirements: First, it provides low level access to hardware to develop the underlying virtualization software - infrastructures like FABRIC and cloud labs to fill this need. Second, it must provide the long-term experimental service environment to deploy, refine, and demonstrate the higher level "rubber meets the road" benefits to applications – and an [international] pilot capability is needed. This latter long-term multi-domain pilot component is currently missing.

Key Components: First, the default action will be to utilize existing NSF investments to do software development and testing. Second, we need to engineer/deploy the pilot infrastructure. This should span the US, and should reach down to incorporate regional and campus facilities, and reach internationally to incorporate virtualization programs in other regions of the world. The ***US*** MERI component will deploy four to six "pods," connected at 100 Gbps to provide a mesh infrastructure. This US domain can use FABRIC if the long-term allocation of transport, hardware, or colo are available. The ***International*** MERI will be a trans-Atlantic "bridge" infrastructure with 100Gbps waves: Washington to Paris, and NYC to AMS. We would place pods at the four terminals, and would close the ring on each side to create a resilient and flexible cross ocean network research facility. This ocean "bridge" could offer low level access to FABRIC researchers in return for FABRIC enabling the longer-term higher level pilot infrastructure. This MERI project would promote collaborative US/EU virtualization research.

We estimate the total to be ~$3M for equipment. Annual recurring to be ~$500K/yr. and~$1M/yr for personnel. This for a 10 pod infrastructure spanning the US to EU.