Position Paper for the
Software Defined Infrastructure / Software Defined Exchange Workshop
Submitted by
Ibrahim Matta, Lou Chitkushev, John Day
http://www.cs.bu.edu/~matta
http://www.bu.edu/met/academic-community/faculty/full-time/lou-chitkushev/
https://en.wikipedia.org/wiki/John_Day_(computer_scientist))
Boston University

   Software-defined (programmable) networks have received much attention recently. The basic premise is to separate longer timescale control aspects from short timescale data forwarding aspects, (logically) centralize the control outside of network switches (elements) and run it on commodity hosts where it can be programmed in software, and create two interfaces: one between the controller and network elements, and another between the user (manager) and controller. Very quickly the community realized that centralized control does not scale and proposals for distributed controllers emerged. Furthermore, the underlying mechanisms remain tied to the TCP/IP architecture and hence inherit all its known inadequacies in dealing with security, multihoming mobility, QoS, *etc*. We posit that this view takes us down a path that stifles new thinking and innovation. Recently, we have been studying a foundational approach to networking called Recursive Internetwork Architecture (RINA) [1]. The key idea to programmability enabled by RINA revisits the Internet traditional view of a "layer" as simply modularity. Since the early ARPANET, it was understood that networking was Inter-Process Communication (IPC) and a "layer" essentially is a distributed application (resource allocator) that provides IPC; in RINA, it is called a Distributed IPC Facility (DIF). The IPC processes of a DIF are configured for a certain scope (range) of operation in terms of scale and performance characteristics. The DIF requires only two protocols, which are defined to be invariant with respect to syntax: one for data transfer among IPC processes, and an application protocol for autonomic management within the DIF. In search for commonalities, it turns out that a DIF is a special case of a DAF (Distributed Application Facility), where a DAF performs an application-specific function, *e.g.* weather forecast/analytics, or an operating system as opposed to IPC. A DIF can be recursively built over lower-level (smaller scope) DIFs. The functions of the DIF distinguish mechanisms and policies. This allows the policies of data transfer and management (including addressing, error and flow control, routing and resource allocation, access control) to be optimally configured for the range of operation of each DIF and across different scopes. This recursive construction of DIFs enables us to insert network policies at *any* point in the network, without the need for special network appliances like firewalls or load balancers. Thus, we believe that RINA provides a more promising solution to envisioned SDN/NFV scenarios.
   We imagine a future Internet that is dynamic in the face of application requirements and network management goals; a future network that can be programmed to build layers of communication to aggregate traffic (to reduce traffic burstiness and resource consumption), to create trusted communities (to provide security and contain malicious traffic and attacks), to enable cooperation as well as a competitive marketplace of providers at different levels (from infrastructure to service and brokers), to enable Internet-of-Things (IoT) scenarios with appropriate scoping for data aggregation and real-time communications without the need for a global address space (a la IPv6) or special "middleboxes", *etc*. This requires a disciplined approach to networking that decomposes the problem, not as a software engineering one but as multi-layered distributed computing. Toward this vision, a DIF is a secure container, by construction, as its IPC processes are authenticated when they join the DIF, and addresses local to the DIF must be unique but can be re-used across DIFs. DIFs can be dynamically instantiated and stacked on top of each other; this is akin to providers (e.g., brokers) implementing a higher-level service using the services of lower providers (e.g., infrastructure providers). Traffic from a higher-level DIF can be aggregated and sent to a lower-level DIF as a single aggregate transport flow (and not just a routed tunnel). Much research needs to be done that includes designing policies that are "compatible" across layers (scopes), operating systems that support multi-layered IPC, advanced optimization, control and game theoretic analysis, and wide-area experiments (*e.g.*, over GENI and FIRE) to validate the design and test interoperability across different implementations. We look forward to discussing these issues with other attendees.
   If accepted, we intend to apply for travel support.
[1] RINA: http://www.pouzinsociety.org, csr.bu.edu/rina/; http://www.irati.eu; http://www.ict-pristine.eu