

Dynamic Layer Instantiation as a Service

Flavio Esposito[§] Yuefeng Wang[§] Ibrahim Matta[§] John Day[‡]
flavio@bu.edu wyf@bu.edu matta@bu.edu day@bu.edu

[§]College of Arts & Sciences, Boston University, MA [‡] Metropolitan College, Boston University, MA

Motivation: One of many fundamental problems with the current Internet architecture is the lack of *scoping* of control and management functions, which makes it challenging to deliver a communication service with required characteristics when the range of operation (*e.g.*, capacity, delay, loss) is so wide. This deficiency, together with the desire to offer virtualized network services, have compounded existing network service management challenges. For example, in *Virtual Private Clouds* [1], some of these challenges include network isolation, custom addressing, and dynamic acquisition and release of virtual resources, with different Service Level Objectives (SLOs) and policies.

Background: Our Recursive InterNetwork Architecture (RINA) [2] is based on the fundamental principle that *networking is Inter-Process Communication (IPC)*. RINA recurses the IPC service over different scopes, *i.e.*, ranges of operation. Specifically, a collection of distributed IPC processes with shared states is called a *Distributed Inter-process communication Facility (DIF)*. A *DIF layer* provides data transport services over a certain scope. By *DIF layer instantiation*, we mean that a layer is built by customizing management and data transfer control policies, to deliver predictable services to applications. Stacking DIF layers on top of each other allows networks to be built from smaller and more manageable layers of limited scope. The DIF layer instantiation is *dynamic* as layers can be discovered and created on the fly, and their mechanisms are able to respond and support policy adaptation as network states (*e.g.*, quality of underlying services) change.

Our Contributions: RINA enables private (*e.g.*, virtual or content delivery) networks to be dynamically instantiated, by customizing network management policies, *e.g.*, authentication, routing, addressing, and resource allocation, into a single layer, without the shortcomings of the TCP/IP architecture. For example, since our DIF layer identifies applications (services) by location-independent names, they can move, migrate, or be multi-homed; since we support explicit (authenticated) enrollment into the DIF layer, communication through the layer is secure; and since DIF services are recursively built, they can be better coordinated and managed.

We demonstrate with our prototype how such dynamic instantiation of a DIF layer over lower-level service layers is a natural service that RINA can provide.

RINA subsumes the mechanisms and policies of all other architectures. For example, by layer isolation and explicit authentication of processes to join a layer, RINA supports secure address spaces as in XIA [3]; each DIF layer is a secure container. By naming application processes and not interfaces, RINA supports mobility and multihoming as in Serval [4] or




Fig. 1. Dynamic layer (DIF) formation: a higher level DIF (Virtual Private Cloud) is dynamically formed.

Mobility-First [5]. Thanks to its directory mechanism, RINA supports content discovery applications as in NDN [6]. By instantiating DIF layers with different scopes and policies, a DIF layer can be thought of as a virtual “slice” of a physical network, as envisioned by the Nebula Data Plane [7]. A DIF layer could also be a slice/layer over lower level slices. RINA references on mobility [8], security [9], and layer discovery [10], together with other detailed information can be found at [11].

Demonstration: We demonstrate the dynamic layer instantiation feature of RINA by creating on the fly a new, higher level *Virtual Private Cloud DIF*. The demonstration includes two IPC processes, VPC1 and VPC2, that initially use two separate private DIFs — an *Enterprise DIF* and a *Cloud Provider DIF* — to communicate with their respective local processes. Later on, an enterprise application process App1 asks for a flow service, so as to communicate with App2, a remote application process on the Cloud Provider DIF. Such request, handled by the underlying communication process VPC1, cannot occur unless there is a common underlying *Virtual Private Cloud DIF* to which both VPC1 and VPC2 subscribe. We demonstrate this dynamic instantiation of the DIF layer over the GENI testbed [12].

REFERENCES

- [1] Wood, T. *et al.*, “The Case for Enterprise-Ready Virtual Private Clouds,” in *Proceedings of HotCloud '09*.
- [2] J. Day, I. Matta, and K. Mattar, “Networking is IPC: a Guiding Principle to a Better Internet,” in *ReArch '08 (CoNEXT workshop)*, 2008.
- [3] D. Han *et al.*, “XIA: Efficient Support for Evolvable Internetworking,” in *NSDI*, 2012.
- [4] E. Nordstrom *et al.*, “Serval: An End-Host Stack for Service-Centric Networking,” in *NSDI*, April 2012.
- [5] I. Seskar *et al.*, “MobilityFirst Future Internet Architecture Project,” in *ACM AINTEC*, 2011.
- [6] V. Jacobson *et al.*, “Networking Named Content,” in *CoNEXT*, 2009.
- [7] Nebula project website: <http://nebula.cis.upenn.edu/>.
- [8] V. Ishakian *et al.*, “On Supporting Mobility and Multihoming in Recursive Internet Architectures,” *Comput. Commun.*, vol. 35, no. 13, 2012.
- [9] G. Boddapati *et al.*, “Assessing the Security of a Clean-Slate Internet Architecture,” in *Proc. of the NPSec*, 2012.
- [10] E. Trouva, E. Grasa, J. Day, and S. Bunch, “Layer Discovery in RINA Networks,” in *IEEE CAMAD*, 2012.
- [11] Recursive InterNetworking Architecture (RINA): <http://csr.bu.edu/rina/>.
- [12] Global Environment for Network Innovations <http://www.geni.net>.