

The Need for a "Communicating System"

Joe Touch, USC/ISI

August 15, 2005

touch@isi.edu

**NSF Workshop on Grand Challenges in Distributed Systems
Cambridge, MA – Sept. 28-29, 2005**

Computer systems today benefit from the shared resource management, coordinated control, abstraction, and protection enforcement of operating systems (OS). A corresponding system for communications, a "communicating system" (CS), is needed to allow distributed systems to take advantage of the emerging capabilities of peer and virtual networks without reimplementing them inside each application. Like an OS, a CS would consist of an architecture, an API, and an implementation of the underlying services that support them.

OSs apply abstractions like virtual memory (VM), threads, and processes to provide an environment where programmers are freed from the nuisance of managing their own resources and mapping them to a specific hardware platform. A similar CS would free distributed systems programmers from their current focus on application-layer forwarding (in peer nets), or tunnel deployment and management (in VPNs).

Many of the preliminary pieces already exist to enable us to conceive of a CS concept. Transport connections and messages correspond to processes and threads, and virtual networks extend the Internet and correspond to virtual memory [1]. Support for multinetting - akin to multiprocessing - is being developed [2]. Although there are APIs for deploying virtual nets, they lack the uniformity of the sockets API for conventional communication.

The virtual communications model and its supporting management systems are lacking the CS equivalent of the BSD reference implementation and APIs. Like BSD, a CS reference system design would require a multifaceted approach, developing the variety of network resource management and coordinated control components which need to be integrated these with existing capabilities for virtual networking, virtual hosts, and virtual routers [3].

Specific challenges are based on the networking equivalent of OS capabilities. OSs provide shared resource coordination, which in networking would include QoS, resource discovery and reservation, and coordinated configuration. OSs provide a virtual machine abstraction, similar to what virtual networks provide, especially when they support configurable forwarding engines [4]. OSs also provide protection, which requires more pervasive support throughout the OS and network API and in forwarding inside the network [1].

Although some of these issues have been defined or explored as instances in existing work, a general system for providing this capability, a "CS", has yet to be developed. Just as in OSs, a base CS would provide a further platform for exploring rich variations, e.g., optimization, "do what I mean" models, and pluggable frameworks. As VM is a simpler, linear address space that enables much more complex application manipulation of data structures, so too a CS would enable distributed applications to take full advantage of peer networking and virtualization to decouple their capabilities and semantics from the underlying communications infrastructure.

It is as challenging to describe the potential impact of such a CS, as it would have been to describe that of OSs 40 years ago. The primary goal is obvious - of emancipating application designers and system architects from low-level, lab-coat incantations of network primitives; of reducing the recapitulation of protocol design that encumbers peer-to-peer nets; and of enabling users to think, design, and operate on distributed objects and groups without assembler-level skills.

- [1] Touch, J., Wang, Y., Eggert, L., Finn, G., "Virtual Internet Architecture," Future Developments of Network Architectures (FDNA) at Sigcomm, August 2003. Available as ISI-TR-2003-570.
- [2] Marko Zec, "Implementing Clonable Network Stacks in the FreeBSD Kernel," USENIX Annual Technical Conference (FreeNIX), San Antonio, TX, June 2003.
- [3] Kamp, P., Watson, R., "Jails: Confining the omnipotent root," Proc. 2nd International System Administration and Networking Conf. (SANE), May 2000.
- [4] Tennenhouse, D., et al., "A Survey of Active Network Research," IEEE Comm. Mag., Jan. 1997, pp. 80-86.