Keynote

An End to the End-to-End Arguments

Hisashi Kobayashi
Professor Emeritus, Princeton University,
Princeton, NJ 08544, USA and
Executive Advisor, NICT, Tokyo, Japan
hisashi@princeton.edu

EXTENDED ABSTRACT

The so-called "end-to-end (E2E) argument" advocated by Saltzer, Reed and Clark [1] seems to have served as a major design guideline throughout the evolution of the Internet. Their argument goes, in the context of a network design, that such communication functions as error control and security should be implemented not within the network, but at the end nodes, since these functions can be completely specified only at the end nodes that run applications, and any partially implemented functions within the network will be redundant, waste network resources and degrade the system performance.

Proponents of the E2E principle advocate protection of innovation and "network neutrality" [2] as its rationale. I believe that some of the advantages associated with E2E design can be satisfied by ingenious use of network virtualization techniques.

I revisit the E2E arguments here, because I disagree with some technical arguments made in the original paper, and the end-to-end design should merely be one of many design options and should not be labeled as a "principle." It is unfortunate that some Internet experts hold a dogmatic view of this design option.

In the problem of file transfer the authors of [1] wrongly argue that a checksum should be done only at the end-points. It is important to understand that E2E retransmission makes sense only if the channels are clean and the file size is not so large. There are many situations where reliability and delay can be improved by applying localized retransmission.

In view of shrinking costs of storage, processing power and bandwidth, and because of the recent trend where the end systems are small mobile devices, I argue that error control functions should be placed within the network in order to attain the level of reliability required by a majority of applications. A network should provide various options depending on the need of end users and applications.

Network-based design, as opposed to E2E design, considers sharing of network resources as a primary design goal. Thus, a richer set of functions to monitor and control these resources is inherent to the basic design of the network. E2E design, on the other hand, is not in a position to effectively take advantage of functions and services provided by resource sharing. As such, there are severe limitations to functions and services that E2E based applications can receive from multiplexers, switches and routers.

The notion of sharing should apply not only to physical resources but also to functions and services, such as *security functions*. Applications that require the same security function should share such functions and avoid duplicating the function end-to-end. Authentication, for instance, need not be the sole responsibility of end users. End users may delegate much of the typical authentication procedures to an "agent" residing in some network nodes that they can trust. Such an arrangement can be made more secure than E2E authentication, since an end node can depend on an agent that can exert strong security against the other end or its agent.

The issue of trust is increasingly important. Security measure should not be retrofit into an existing architecture as an afterthought. Adding some cryptographic protection to protocols after they are designed is a flawed approach. Instead we must define trust relationships between the network elements and include them as part of the network architecture. The trust boundaries should be defined, and when communication occurs across a trust boundary, cryptographic or other security protection may be necessary.

Effective resource sharing leads to *economies of scale*. The E2E design proponents often argue that a simple or dumb network is more scalable because there is little processing going on within the network. A counter argument would be that a significant portion of the network resources is wasted because of unnecessary retransmissions and congestion created by the dumb network. It is true that the recent success of peer-to-peer networking and its applications (e.g., Napster) can be ascribed to the simplicity of the core network of the Internet. However, such techniques as NAT and caching require more processing within the network than the end-to-end design approach can allow.

I argue that we should return to first principles and re-examine what the end-to-end design approach can and cannot accomplish, and understand the limitations and constraints imposed by this design approach.

[1] J. H. Saltzer, D. P. Reed and D. D. Clark, "End-to-End Arguments in System Design," *ACM Trans. Comp. Sys.*, 2 (4), pp. 277-288, Nov. 1984

[2] http://en.wikipedia.org/wiki/Network_neutrality

BIOGRAPHY

Hisashi Kobayashi is the Sherman Fairchild University Professor Emeritus of Princeton University and an Executive Adviser of the National Institute of Information and Communications Technology (NICT), Japan. He earned his BE (1961) and ME (1963) degrees from the University of Tokyo and his Ph.D. from Princeton University (1967). He worked for Toshiba Corporation as a radar designer (1963-65) and for the IBM US Research Division for 15 years (1967-82), and then as the founding director of IBM Tokyo Research Laboratory (1982-86). He then joined the Princeton faculty as Dean of the School of Engineering and Applied Science (1986-91) and taught in the electrical engineering department until June 2008.

In 1970 he invented a high-density digital recording scheme, widely known as "partial-response coding and maximum-likelihood decoding" (PRML) (see the "History of Communications" column, IEEE Communication Magazine, March 2009, pp. 14-17), and was awarded the 2005 Eduard Rhein Technology Award. He also received Humboldt Prize of Germany (1979) and IFIP's Silver Core Award (1980). He was elected a Fellow of IEEE (1977), a Fellow of IEICE, Japan (2004). He is a member of the Engineering Academy of Japan — Japan's National Academy of Engineering — since 1992.

He authored "Modeling and Analysis"(Addison Wesley 1978) and coauthored with Brian Mark "System Modeling and Analysis" (Pearson/Prentice Hall, 2008). He is also working on textbooks on "probability and random processes" (to be published from Cambridge University Press) and "network protocols, performance and security."

During his 19 year association with the IBM Research (1967-86) Kobayashi was granted several sabbatical leaves to teach at major universities as a visiting professor: UCLA (1969-70), the University of Hawaii (1975), Stanford University (1976), Technische Hochschule Darmstadt (1979-80), and at Universite Libre de Bruxelles in Belgium (1980). While he was with Princeton University, he took a

sabbatical leave to the RCAST (Research Center for Advanced Science and Technology) of the University of Tokyo as an NEC C&C Visiting Professor (1991-92) and to the University of Victoria, B.C., Canada as an ASI (Advanced System Institute) Fellow (1998-99).

He served as a scientific advisor for numerous organizations in the United States, Japan Canada, Singapore and Hong Kong. Currently he serves on the Board of Directors of the "Friends of Todai (The University of Tokyo) Foundation-USA, Inc.," and the "Edwin Howard Armstrong Foundation" of Columbia University. He is an executive advisor to the NICT Japan on their future Internet project AKARI. Most recently, he has been appointed as a reviewer of the "G-Lab" of Germany for their Future Internet design, and as an advisor for President of Toyota Technological Institute, Japan.

For additional information, please see http://www.princeton.edu/kobayashi