To Session or not to Session — Design Issues for the Future Internet

C. M. Gauger, B. Gloss, C. Hauser, J. Jähnert, S. Kiesel, M. Köhn,
P. J. Kühn, M. C. Necker, A. Reifert, D. Sass, M. Scharf
Institute of Communication Networks and Computer Engineering,
University of Stuttgart, Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Corresponding author: marc.necker@ikr.uni-stuttgart.de

Today's communication networks are based on one of two fundamentally different paradigms. On the one side are the traditional circuit-switched telecommunication networks controlled by telecom operators. On the other side is the packet-switched Internet, composed of a large number of independent domains. The paradigm in telecom area comprises closed networks under tight control, which achieves a high level of security and privacy. Additionally, its service-specific infrastructure provides means for accounting, user management and Quality of Service (QoS). In contrast, the Internet employs a simple and service-unaware infrastructure for data transport. Services need to be realized in the end-systems, making it relatively easy to introduce new services. Compared to the telecommunication paradigm, the network is open and in a constant change. The origins are in the military and academic area, originally driven without commercial interests.

In the past decade, the Internet has undergone a rapid change towards a huge commercial platform. Along with this came the introduction of new as well as the migration of traditional services, such as voice or television. As this process continues, the Internet is taking over more and more functionality from traditional circuit-switched networks with an increasing user penetration. This raises questions on both the technical and the business side. First, the Internet lacks mechanisms that support certain services, such as the already mentioned QoS and accounting support. Second, it is questionable whether the existing business models can directly be taken over from the telecom world. Therefore, it can be argued that the current Internet is overwhelmed by the present developments.

As a consequence, telecom operators have started to adapt IP-based infrastructure for their networks while maintaining basic concepts of their traditional circuit-switched networks. This development can best be observed in the ongoing activities in the area of 3GPP cellular networks. The IP Multimedia Subsystem (IMS) is a first approach to facilitate new services within an IP-platform. In fact, it is an attempt to establish a service delivery platform using IP technology, which implements the above identified missing functionality. However, while building on Internet-technology, operators are trying to keep these domains closed and under tight control, which is in contrast to the prevailing Internet philosophy.

Telecommunication networks base their functionality on keeping states within the network and establishing a network-based session for each communication activity. In contrast, the Internet holds state information only in the end systems, limiting the session concept to relations between these end systems. Maintaining per-session state makes networks aware of the characteristics and functional and non-functional requirements of sessions. This additional knowledge has substantial consequences on the type of services offered by a network: First, session information allows guaranteed QoS by means of resource reservation and traffic engineering. While there is only limited support of QoS mechanisms in current IP networks, the migration of telephony to IP technology and the introduction of new real-time services will benefit from QoS support.

Second, and even more important, security mechanisms in the network such as access control require state information, e.g., for authentication and authorization purposes. The major security problems of today's Internet, such as denial-of-service attacks and spam, originate from its state-less end-to-end design. Furthermore, many privacy issues arise from end-to-end sessions, which can be addressed by breaking up sessions within the network. Third and finally, session information is required for service differentiated accounting and pricing, as they are predominant in current fixed and mobile telecommunication networks.

However, maintaining state within networks comes at the cost of complexity and limited robustness. Even though processing capabilities have significantly increased during the last decades, stateful network components still have a limited scalability. They are also more vulnerable to failures and inconsistencies, which is particularly critical in case of dynamic reconfigurations. Furthermore, the required control plane signalling introduces overhead and requires some kind of cooperation even over competing administrative domains.

The question of stateful vs. stateless network design is not only driven by technical problems. There are also economic and legal aspects to be considered. In traditional telecommunication networks, the operator is involved in service delivery. In contrast, Internet network operators only provide connectivity. Thus, the convergence of networks may require that operators redefine their role and may even give raise to new business models. For example, current discussions in the US on net neutrality reveal a certain economic interest in per-session QoS mechanisms. In addition to this, regulation requirements such as legal intercept may have a strong impact on the design of communication networks.

With respect to the above discussion, there is a clear tradeoff regarding the introduction of session state in a future Internet. An alternative to a stateful network is to keep the core Internet as simple as it is today and provide all necessary intelligence in the end-systems or access networks. In order to realize real-time services with such an approach, the core network would need to be overprovisioned.

It is likely that neither of these two concepts will dominate a future Internet. For instance, access networks may handle sessions in a different way than core networks. Furthermore, cross-layer solutions may emerge, combining different concepts that may be required on different layers. In this talk, we will analyse the tradeoff between a stateful and a stateless future Internet on different protocol layers and for different functional aspects. We will contrast these two concepts for access and core networks taking into account both wireless and optical transport technologies. We will show how functionalities like QoS, privacy, security or network management can benefit from either concept. Finally, we will converge on some basic conclusions for future Internet design.