Implications and Control of Middleboxes in the Internet

Xiaoming Fu Martin Stiemerling Henning Schulzrinne

Middleboxes in the Internet have been explored, sometimes quite controversially, in operations, standardization, and the research community for more than 10 years. The main concern in the past has been their contradicting nature to the Internet's end-to-end principle. In the past, many have expressed concerns that middleboxes contradict the Internet's end-to-end principle that is often understood to posit that "intelligence" is placed in end system and network elements just forward packets. Middleboxes introduce functions beyond forwarding in the data path between a source and destination, as described, for example, in RFC 3234. RFC 3234 describes a wide range of middle boxes, from TCP performance enhancing proxies to transcoders.

On the other hand, middleboxes were introduced in the Internet for various reasons: NATs intend to decouple the internal IP addressing from the public address space while allowing multiple hosts to share a single public IP address, for the purpose of preserving the IP address space; firewalls are used for administrators to enforce policies on the data traffic at administrative borders with the intention of preventing their networks from being attacked or monitored; application level gateways (ALGs) are typically used to assist applications in their operations.

The implications of the emergence and popularity of middleboxes are complicated. With middleboxes it is difficult to even provide basic end-to-end connectivity for many applications. For example, Internet hosts behind NATs can only initiate a TCP connection with another host, but cannot accept a connection request. Unlike in the past, when the vast majority of applications followed the client-server design pattern, and most hosts behind NATs were clients anyway (e.g., your browser accessing a Web server), a variety of new applications today, such as voice-over-IP, gaming, and peer-to-peer file sharing cause an enormous list of issues. Hosts behind NATs are not reachable from any other host anymore, which become particularly troublesome for VoIP and other peer-to-peer applications. Likewise, firewalls are usually statically configured to block certain TCP ports or do not understand non-TCP protocols, making it difficult to deploy new applications and protocols. This results in a number of issues to be considered in the design and development of new protocols and applications.

To mitigate the negative impacts of these issues, quite a number of techniques have been developed, which can be categorized as explicit control and implicit control of firewalls and NATs. For explicit control, an entity, either the end host or a proxy in the network, has a relationship with the middlebox and controls its behavior (e.g., the set of policies or filter rules loaded). Examples of explicit control are universal plug and play (UPnP), Internet Engineering Task Force (IETF) Middlebox Communications (MIDCOM), and IETF Next Steps in Signaling (NSIS). On the other hand, implicit control is the traditional way of traversing middleboxes. Implicit control does not have any control relationship with the middlebox, because end hosts, probably with the support of other end hosts, are using hole punching techniques to get a working middlebox traversal. Examples of implicit control are the IETF's Session Traversal Utilities for NAT (STUN), Traversal Using Relays around NAT (TURN), and Interactive Connectivity Establishment (ICE). In addition, there have been some recent attempts to design or use certain types of middleboxes, such as various application proxies.

In this special issue we are pleased to introduce a series of state-of-the-art articles on this specific area. These articles cover the subject from a variety of perspectives, offering the readers an understanding of the issues and implications of various middleboxes in the Internet, including their control mechanisms. A total of eigh articles, selected from 26 submissions based on a strict peer review process, cover a broad range in the field of implications and control of middleboxes in the Internet. While some articles present more general issues with middleboxes, understanding their behaviors and implications, others focus on new approaches to controlling and usiing middleboxes.

NATs, an unplanned reality, have posed complications to the Internet architecture and applications. The first article, “A Retrospective View of NAT” by Lixia Zhang, takes readers back to the early days of middleboxes. It gives a historic review of NATs and the lessons learned, including how they impeded standardization and deployment of IPv6, and an expected solution for addressing the Internet address depletion problem. Without a timely standardization of NAT, today there have been a number of different NAT implementations, and it is vital to understand their behaviors due to their nearly ubiquitous presence.

The second article, “Behavior and Classification of NAT Devices and Implications for NAT Traversal” by Andreas Müller, Andreas Klenk, and Georg Carle, provides a comprehensive overview of NAT behaviors and currently available

NAT traversal techniques. The article presents a new categorization approach based on an analytical abstraction of NAT traversal, which classifies NAT traversal services into four distinct types and deduces the corresponding NAT behaviors. This may help developers of new protocols and applications to determine applicable techniques for NAT traversal.

While the first two articles describe the history, behavior, and classification of NAT, the next article by Dilip Joseph and Ion Stoica, "Modeling Middleboxes," proposes a formal and generic model for deducing middlebox functionalities and behaviors. Using this model, the article illustrates how different middleboxes process packets, and how four common middleboxes — firewall, NAT, layer 4, and layer 7 load balancers — may be depicted. As such, the article provides an initial step for relevant designers, users, and researchers to understand and refine the behaviors and implications of various middleboxes.

Existing middleboxes mostly consider TCP and UDP in their implementations, and typically do not support other protocols, such as the Stream Control Transmission Protocol (SCTP). In the fourth article, Michael Tüxen *et al.* describe the extensions required to support NAT for SCTP. The analysis presented in this article may be useful as a general lesson in the near future, as several other protocols after SCTP, including DCCP, XCP, and HIP, use similar techniques such as multihoming, rehoming, and handshake cookies.

Applications using the Session Initialization Protocol (SIP) or peer-to-peer way of operation (P2PSIP or just normal P2P applications) are among those that suffer most from the middlebox traversal issue. The fifth article, "Distributed Connectivity Service for a SIP Infrastructure" by Luigi Ciminiera *et al.*, examines this issue and presents an alternative approach to the current STUN/TURN/ICE approach to middlebox traversal. The approach distributes the rendezvous and relay functions among SIP user agents, which discover their peers autonomously and maintain a P2P overlay to ensure connectivity across NATs and firewalls in a SIP infrastructure without relying on a centralized server.

The remaining three articles address new applications of middleboxes. The sixth article, "Dial M for Middlebox Managed Mobility" by Stephen Herborn and Aruna Seneviratne, describes a new usage type of middleboxes for mobility support via the concept of virtual private "personal networks." Such a network is created and maintained by way of HIP combined with IPsec and supported by middlebox state drop "(at least to some extent)" plus middlebox state, which may be interesting (at least to some extent) for the recent research efforts on network virtualization, as they use today's technologies directly.

An increasing number of home users today are using NATs to connect their home IP devices with the Internet. Choongul Park *et al.* discuss this issue in their article "Issues in the Remote Management of Home Network Devices." By extending SNMP and using additional management objects (MOs) to gather NAT binding information, the authors attempt to address the NAT traversal problem under a symmetric NAT, based on their observations in Korea. While the success rate of NAT traversal could be a potential issue outside Korea, the article provides an insight of what home networking standards may have to deal with.

Yet another type of middlebox function, intelligent route control (IRC) for multihomed sites and subscribers, has been recently identified as a key issue in efficient network operations. The final article, "Improving the Performance of Route Control Middleboxes in a Competitive Environment" by Marcelo Yannuzzi *et al.*, addresses this issue and introduces an IRC approach for competitive environments, by blending randomization with adaptive filtering techniques.

We hope that these articles will help to clarify and explain the state-of-the-art advances on middlebox issues in the Internet, providing current visions of how the behaviors, implications, and control of middlboxes may be analyzed, encompassed, and utilized. In preparing this special issue, we wish to thank all the peer reviewers for their efforts in carefully reviewing the manuscripts to meet the tight deadlines. We are grateful to our liaison editor Jon Crowcroft for his constructive feedbacks, and Editor-in-Chief Ioanis Nikolaidis for his timely and critical suggestions.

Biographies

XIAOMING FU [M'02] (fu@cs.uni-goettingen.de) received his Ph.D. degree in computer science from Tsinghua University, Beijing, China, in 2000. After almost two years of postdoctoral work at Technical University Berlin, he joined the University of Göttingen as an assistant professor, leading a team working on networking research. Since April 2007 he has been a professor and head of the Computer Networks Group at the University of Göttingen. During 2003–2005 he also served as an expert on the ETSI Specialist Task Forces on Internet Protocol Testing; he was also a visiting scientist at the University of Cambridge and Columbia University. In the research fields of architectures, protocols, and applications for QoS, firewalls, p2p overlay, and mobile networking as well as related security issues, he (co-)authored more than 50 referred papers as well as several RFCs/I-Ds. He has served as TPC member and session chair for several conferences, including IEEE INFOCOM, ICNP, ICDCS, GLOBECOM, and ICC. He was also founding chair of the ACM Workshop on Mobility in the Evolving Internet Architecture (MobiArch) and is TPC Co-Chair of IEEE GLOBECOM 2009 Next Generation Networking and Internet Symposium. He is currently a member of the editorial board of *Computer Communications Journal* (Elsevier).

MARTIN STIEMERLING [M'00] (stiemerling@cs.uni-goettingen.de) received his M.Sc. degree (Diploma) in electrical eengineering with a focus on IP networking technologies from the Polytechnic University of Applied Sciences in Cologne in 2000. After that he joined the NEC Laboratories Europe, Heidelberg, Germany, where he is currently a senior researcher. His areas of research interest are Internet architecture, Internet signaling protocols, network management, and overlay/peer-to-peer systems. He has published several papers in these areas, and served as a TPC member of IEEE IPOM 2007. In the IETF he is active as working document editor in the MIDCOM, MMUSIC, and NSIS working groups, as well as in other IETF working groups and IRTF research groups. He is co-chair of the IETF Next Steps in Signaling (NSIS) working group, and secretary of the IP over DVB (IPDVB) working group, and a co-author of RFC 3816, RFC 3989, and RFC 4540, as well as RTSPng.

HENNING SCHULZRINNE [F'06] (hgs@cs.columbia.edu) received his Ph.D. from the University of Massachusetts in Amherst, Massachusetts. He was a member of technical staff at AT&T Bell Laboratories, Murray Hill, New Jersey, and an associate department head at GMD-Fokus (Berlin) before joining the Computer Science and Electrical Engineering Departments at Columbia University, New York. He is currently a professor and chair of the Department of Computer Science. He has been a member of the Board of Governors of the IEEE Communications Society and is vice chair of ACM SIGCOMM, former chair of the IEEE Communications Society Technical Committees on Computer Communications and the Internet, has been technical program chair of Global Internet, INFOCOM, NOSSDAV, and IPTCOMM, and was General Chair of ACM Multimedia 2004. He has also been a member of the Internet Architecture Board. Protocols co-developed by him, such as RTP, RTSP, and SIP, are now Internet standards, used by almost all Internet telephony and multimedia applications. His research interests include Internet multimedia systems, ubiquitous computing, mobile systems, quality of service, and performance evaluation.