# GONE: an Infrastructure Overlay for Resilient, DoS-Limiting Networking

Xiaoming Fu
University of Göttingen, Germany
fu@cs.uni-goettingen.de

Jon Crowcroft
University of Cambridge, UK
jon.crowcroft@cl.cam.ac.uk

## ABSTRACT

With today's penetration in volume and variety of information flowing across the Internet, data and services are experiencing various issues with the TCP/IP infrastructure, most notably availability, reliability and mobility. Therefore, a critical infrastructure is highly desireable, in particular for multimedia streaming applications. So far the proposed approaches have focused on applying application-layer routing and path monitoring for reliability and on enforcing stateful packet filters in hosts or network to protect against Denial of Service (DoS) attacks. Each of them solves its own aspect of the problem, trading scalability for availability and reliability among a relatively small set of nodes, yet there is no single overall solution available which addresses these issues in a large scale.

We propose an alternative overlay network architecture by introducing a set of generic functions in network edges and end hosts. We conjecture that the network edge constitutes a major source of DoS, resilience and mobility issues to the network, and propose a new solution to this problem, namely the General Internet Signaling Transport (GIST) Overlay Networking Extension, or GONE. The basic idea of GONE is to create a half-permanent overlay mesh consisting of GONE-enabled edge routers, which employs capability-based DoS prevention and forwards end-to-end user traffic using the GIST messaging associations. GONE's use of GIST on top of SCTP allows multi-homing, multi-streaming and partial reliability, while only a limited overhead for maintaining the messaging association is introduced. In addition, upon the services provided by GONE overlays, hosts are identified by their unique host identities independent of their topologies location, and simply require (de-)multiplexing instead of the traditional connection management and other complex functionality in the transport layer. As a result, this approach offers a number of advantages for upper layer end-to-end applications, including intrinsic provisioning of resilience and DoS prevention in a dynamic and nomadic environment.


*NOSSDAV* '06 Newport, Rhode Island USA


## Keywords

Overlay Networking, Denial-of-Service, Resilience, General Internet Signaling Transport, Host Identity Protocol

## 1. INTRODUCTION

The original TCP/IP architecture did not deliberately consider path instability, middleboxes, security, and device mobility. To dynamically adapt the variance of Internet topology and path characteristics, end systems simply utilize TCP to react to network congestions and routers implement routing protocols to disseminate and construct new path information over time, in addition to best effort IP forwarding which has been recently enhanced with e.g., differentiated services [1] to satisfy the needs of increasing amount of real-time multimedia applications. With today's popularity of information flowing across the Internet, these issues have become essential that impair the availability of the Internet services worldwide. Service providers are suffering failures in providing effective measures to resolve certain routing pathologies in their infrastructure, such as link or node failures or temporary routing loops. For example, some measurements performed in 2000 have shown that the chance of encountering end-to-end path failures in Internet communications were around 3.3% [2], far higher than the level of the PSTN network (typically within $10^{-5}$). This can result in packet losses and connection failures for end-to-end applications. Such deteriorated quality could be intolerable for most video or audio streams. In addition, a denial of service (DoS) attacker can compromise a victim's network service availability, typically by flooding the victim with a huge number of useless requests thus exhausting its bandwidth or computational resources. A quantitative estimation of worldwide DoS attack frequency found 12,000 attacks against more than 5,000 distinct targets over a 3-week period in 2001 [3]. The issue about DoS limiting is even more crucial for multimedia than for classic Internet adaptive or asynchronous applications like web and emails. As a result, firewall middleboxes have been emerged rapidly to reduce the volume of malicious connections. In addition to network address translators (or NATs, another common type of middleboxes), these middleboxes have largely changed the Internet end-to-end principle and become a challenging issue for services and applications [4, 5]. Moreover, the proliferation of wireless devices and need of mobility has posed a critical challenge for the conventional Internet to support seamless mobility for user applications [5].

Studies over the last decade have attempted to address these issues by a variety of means, including content repli-

cation (e.g. [6]), host-, site- or ISP-level multi-homing [7], resilient overlay routing [8], mobility using tunneling and redirection techniques [9, 10] or an identifier/locator split [11, 12, 13], DoS prevention by installing filters either in the network alone or also in receivers to filter out unwanted traffic (e.g. [14, 15, 16, 17]). While these approaches deal with their respective functional aspects, typically one approach simply addresses a certain specific problem space and may not best serve or even work for other scenarios. For example, solution for resilient overlay routing does not consider mobility, whereas multihoming approaches usually do not address DoS issues. Furthermore, most approaches are suitable in scenarios with a relatively small number of nodes, and do not consider large networks and many end nodes.

The above approaches fall into either an application-layer solution [8, 15, 6], which operates only in end hosts, or an infrastructure-based solution [14, 16, 17], which involves intermediate nodes in addition to end hosts. This paper examines the potential of GONE (the GIST [18] Overlay Networking Extension), a generic infrastructure-based overlay architecture for improving availability, reliability and supporting mobility. By using existing well-specified standards (SCTP, HIP [19] and GIST), this approach provides a fairly easy means to specify and implement a network edge with desired functions and software components. On one hand, as a DoS-limiting infrastructure is supplied with GONE, the strength of building such systems based on IETF standards will be more outstanding. On the other hand, the reuse of the common and fundamental component in the next generation signaling framework [20] enables us to build an IETF standard-based platform for realizing ideas like Plutarch [21].

After a short discussion of related work in Section 2, we elaborate the GONE design overview in Section 3, followed by more detailed discussions in Section 4. We briefly review our ongoing research status in Section 5 before concluding in Section 6.

## 2. RELATED WORK

A recent advance is the introduction of the *Host Identity Protocol (HIP)* [19, 12]. HIP attempts to resolve the issue of separating host locator from identifier, allowing end host authentication, device mobility and multihoming [22], as well as reducing DoS attacks. However, HIP does not completely address the issues of resilience and path availability at the ISP/AS level. As resource exhausting DoS attacks usually take advantage of the cost of setting up a state for a protocol on the responder compared to the 'cheapness' on the initiator, HIP intentionally let a responder impose an increased cost for the start of state on the initiator, thus reducing the cost for the responder. This is done by having the responder start the authenticated Diffie-Hellman exchange instead of the initiator, which includes a puzzle (a cryptographic challenge that the Initiator must solve before continuing the exchange) based DoS reduction scheme. HIP mobility support [23] allows a host to dynamically change the primary locator that it uses to receive packets. For resolving the mapping between identifiers and locators, some centralized entity, the so-called *rendezvous server (RVS)*, may be used. The *Host Identity Indirection Infrastructure (Hi3)* [24] extends HIP based on $i3$ [14], which presents a distributed scheme for routing HIP handshake messages based on host identities.

The concept of capability-based DoS prevention proposed by Yang *et al.* [17] refrains the DoS attacks by limiting the sender to send only traffic permitted by the receiver. Different from HIP, which relies on some mechanisms in end hosts, this method introduces some kind of capability filters in certain routers in the data path. Another similar work is the *Secure Overlay Service (SOS)* [16] architecture, which constructs an overlay using a combination of secure overlay tunneling, routing via consistent hashing, and filtering in the network edge. However, SOS does not consider the path resilience issue; the pre-established SOS edge nodes are assumed to be known to the public without taking account of nomadic or mobile users.

Feamster *et al.* [25] analyzed the path failure issue and suggested that it can be improved by using reactive routing such as RON [8], especially when hosts have multiple connections to the Internet. Guo *et al.* [26] shows that performance gains can be achieved if an access router is connected to several neighboring ISP networks (i.e., multi-homed).

The *General Internet Signaling Transport (GIST)* [18] is a general purpose signaling transport protocol currently developed by the IETF NSIS working group. GIST provides a soft state mechanism and richer security than RSVP [27] for delivering any kind of path/flow-coupled state in IP-based networks. GIST can use reliable stream- or message-oriented protocols such as TCP or SCTP, or unreliable transport protocols such as UDP to deliver the required signaling message. Readers not familiar with GIST are suggested to take a look at [18, 20]. Our GONE approach is built on GIST over SCTP [28]. Different from other usages of GIST, GIST here is not only for GONE control message signaling, but also for end-to-end data traffic forwarding.

## 3. GONE OVERVIEW

In this section we present an overview of GONE. We start with the description of a general communication scenario, then outline the GONE design.

### 3.1 A General Communication Scenario

As shown in Fig. 1, generally a network communication between two end hosts (here, $H1$-$H3$ or $H2$-$H4$) encompasses two access networks (here, ISPs 1/3, 5/7) and backbone networks (here, ISPs 2/4, 4/6). In this example, both ISP1 and ISP3 have two links connecting to its two adjacent backbone networks (ISP2 and ISP4) via dual connectivity between their edge routers (ERs), respectively. Therefore, there is an ISP-level multi-homing from the $H1$-$H3$ communication's point of view. For the $H2$-$H4$ communication, additionally there is a host-level multi-homing, where $H2$ and $H4$ are connected to two access routers (ARs), respectively. Therefore, when one path encounters failures (such as a link failure in either $a$ or $b$, or a node failure in some routers in ISP2 along the path indicated by the direct line between $H1$ and $H3$), by applying routing algorithm the network shall resume after some time the ongoing connectivity between $H1$ and $H3$ and an alternative path (the dash line). Furthermore, a host can move from one point of attachment to another, e.g., $H1$ can move to the area where $H2$ is shown, thus mobility shall be supported.

For the convenience of our discussion, the following assumptions are made:

- Multi-homing is common for ISP networks;
- Strong DoS protection is needed in access networks;

- Hosts support HIP locator-independent identifier;
- Hosts and ARs support GONE. To make the best benefits, (especially multi-homed) ERs may support GONE, but this is not mandatory for possible incremental deployment.

Next we will describe the overall design of GONE.

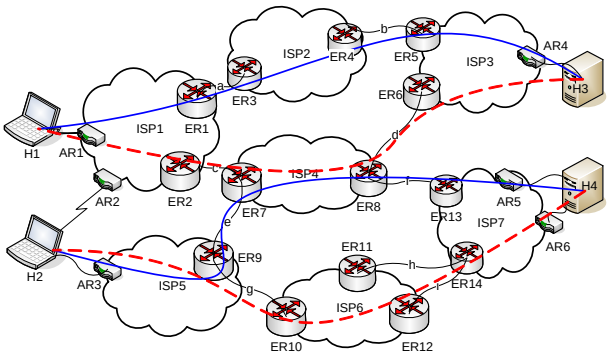


**Figure 1: Network scenario**

## 3.2 GONE Strawman Design

We believe there is a need for an effective edge-based overlay routing and DoS-prevention system – one that enhances service availability, improves end-to-end path stability and is resistant to DoS attacks. In particular, for multimedia systems, instances, as they are often coupled with resource reservation, call control and other signaling functions, all being critical parts that converge in network edges, it is reasonable to assume such generalized edge to avoid a multitude of complexity. Such a system is thus designed to provide a new view of the network architecture, as shown in Fig. 2. GONE consists of a GONE base protocol layer, which works upon GIST over SCTP for overlay routing and forwarding, and consists of a GONE control protocol and a GONE data protocol. GONE control protocol maintains messaging associations and GONE overlay routing state, and exchanges capability information between GONE aware ARs/ERs. GONE data protocol delivers end user's traffic over the messaging associations to the next GONE node, known from the mapping of the host identity and next GONE node's IP address in the GONE overlay routing state. End hosts simply need to support the multiplexing and de-multiplexing function in addition to the GONE stack.

In addition, for easy, ubiquitous Internet access and supporting mobility for end devices, HIP host identity is used to identify and authenticate a host without the need to change its identify for ongoing communications. The mapping of a host's identity and IP address is maintained in GONE nodes with a Distributed Hash Table (DHT), such as Chord [29] or OpenDHT [30]. DHT is chosen here mainly due to its known better searching performance than unstructured peer-to-peer techniques. Maintaining the DHT in edge devices is particularly useful in nomadic and mobile environments, since a host identity remains unchanged while a host moves and is still quickly retrievable for its IP address.

Here, ISP edge-level multi-homing together with benefits of HIP and capability-based DoS prevention are considered our motivating design background. However, different from previous approaches, which utilize different mechanisms to maintain overlay routing state, we construct and maintain an SCTP overlay mesh between network edges for delivering end users' traffic based on permitted capability, minimizing the complexity of maintaining various functional boxes while providing desired features. The idea is to use nearly "always-on" SCTP associations between edge routers for fast path failure recovery and load balancing (due to its multi-homing and multi-streaming support), and to apply capability based DoS prevention directly in edge nodes. SCTP associates are created and dynamically maintained by the IETF GIST protocol [18], which has been initially designed for control plane signaling, but here we extend as GIST Overlay Networking Extension (GONE). GONE intends to provide better DoS protection than HIP and better path failure recovery than traditional overlay networks (e.g., RON [8], *i*3 [14]).

Let us explain this approach in an example. In a simple $H1$-$H3$ communication scenario shown in Fig. 1, assume only $H1$, $H3$, $AR1$ and $AR4$ support GONE, and there is an existing GIST messaging association between $AR1$ and $AR4$ (which runs over SCTP). When the host H1 expects to deliver upper-layer application data traffic to its communicating peer $H3$, it initializes a GIST message routing state establishment process while discovering GONE intermediaries, namely $AR1$ and $AR4$ in this example. This is done by applying a GONE capability negotiation procedure. The result are established messaging associations between $H1$ and $AR1$, and between $AR4$ and $H3$, each using an SCTP association respectively, as well as established messaging routing states and remembered capability in $H1$, $AR1$, $AR4$ and $H3$. The SCTP associations (and GIST messaging associations in turn) are then maintained using soft state provided by GIST protocol to minimize GONE setup overhead. The multi-homed SCTP association between $AR1$ and $AR4$ allows path failure recovery. GONE overlay nodes keep track of the negotiated capability and filter out any traffic that does not conform to the capability.

GONE extends the IETF GIST protocol and defines a new NSLP application. It consists of two protocols:

- GONE control protocol: capability negotiation and setting up overlay routing state.
- GONE data protocol: transmission of GONE data between adjacent GONE nodes and capability-based DoS traffic filtering and rate limiting.

Any GONE message has the following format:

GONE message := [GONE header] [GONE payload]
GONE header := [Type] [Length] [NSLPID="GONE"]
GONE payload := [Ctrl_Req] | [Ctrl_Resp] | [Data]

### 3.2.1 GONE Control Protocol

The purpose of GONE control protocol is to establish and maintain necessary states in GONE edge routers, access routers as well as end hosts for construct overlay routing and legal traffic pattern for the end-to-end communication. A key concept used in GONE control protocol is "capability". Following the work in *Traffic Validation Architecture (TVA)* [17], a capability comprises a router timestamp and a keyed hash. A sender can request via a GONE *Ctrl_Req* message towards the receiver, expecting the receiver to grant this sender with a capability associated with rate limit to

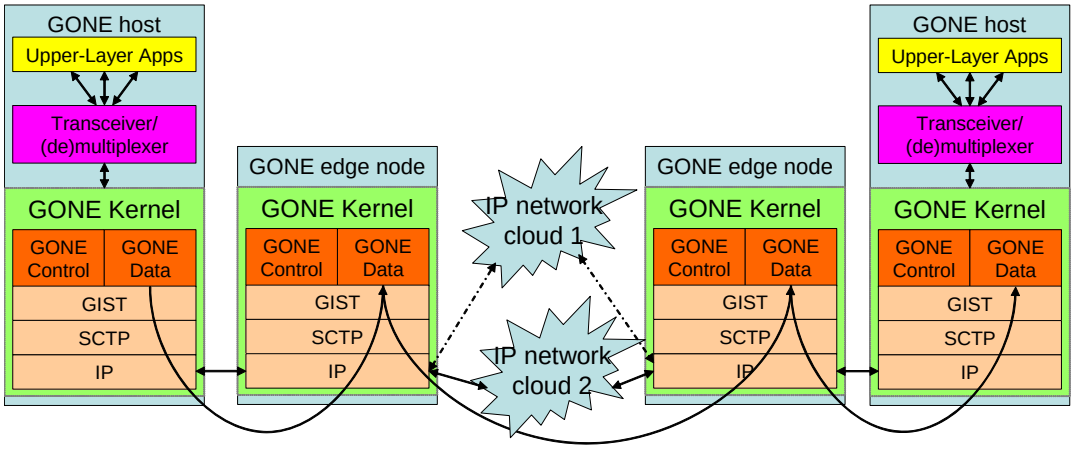


**Figure 2: GONE protocol stack**

send traffic. A Ctrl_Req message is delivered by GIST (which in turn creates or updates GIST messaging association in the background), directing to all the GONE nodes in the path. The traversed GONE nodes will in turn generate a *pre-capability* and add it into the Ctrl_Req message, before forwarding it to the next GONE node. When the receiver sees the Ctrl_Req message, it checks whether the sender is allowed to send the corresponding traffic. If yes, it grants this traffic with a capability with certain rate limit. While sending backwards along the GONE node chain a *Ctrl_Resp* message, a capability token is installed in the GONE nodes, allowing the GONE data protocol to check the validity of the traffic and whether a given traffic exceeds its rate limit.

Different from TVA, where the sender and receiver's IP addresses (in addition to timestamps and private secret) are used as inputs for keyed hash computation, GONE uses the hosts' location-independent identifier (known as HIT or Host Identity Tag in HIP [12]) instead, to allow the seamless mobility of end hosts. The capability is installed in the GONE intermediaries (ARs/ERs) and maintained by GONE soft state mechanism, automatically removing inactive GONE overlay routing state (e.g., due to host mobility or abnormal session termination). A host can also explicitly remove the GONE overlay routing state and associated capability from GONE intermediaries.

GIST has a three-way handshake/discovery component and a data transport component. Thus, there are two possible ways to build the control protocol: either to extend the discovery component to allow collecting the pre-capabilities and installing capabilities in the GONE nodes, or to implement it as part of GONE as an "NSIS Signaling Layer Protocol" (NSLP) [31, 20]. The simplicity of the NSLP-based approach seems to outweigh the relative performance benefit in the discovery extension approach, since the GONE data protocol can be easily implemented by the same NSLP, which results in a unified design for both GONE protocol components.

To summarize, GONE control protocol extends the GIST data transport to enable end-to-end capability negotiation. It consists of two messages:

Ctrl_Req := [SenderHIT] [ReceiverHIT] [SessionID]
[Flag="Install/Remove"] [Pre-Capability list]
Ctrl_Resp := [Capability]

Where, a probabilistically unique SessionID is generated per 5-tuple <upper-layer protocol, sender and receiver's port numbers and host identities>. This is used in the GIST layer for overlay routing information (i.e., GIST's *message routing state*) and also used as part of identification metric of end-to-end traffic.

### 3.2.2 GONE Data Protocol

There is only one message type for the GONE data protocol:

Data := [UserData] [SessionID]

When a Data message is received, a GONE node checks the validity of the capability corresponding to SessionID.

GONE's use of GIST over SCTP for transport brings additional benefit: the header-of-line blocking avoidance, potential of customized reliability level for data traffic, and failover handling due to multi-homing. Moreover, the reuse of messaging associations allows frequently used networks to establish GIST overlay states for new arriving traffic. After the GIST messaging associations are established during the GONE setup phase, they should be set with a longer lifetime (e.g., several hours), whereas the messaging associations between end hosts and GONE access routers should be set with a shorter lifetime (e.g., several minutes) to allow efficient resource usage in the end hosts and minimal messaging association setup overhead in GONE routers.

## 4. FURTHER DISCUSSIONS

### 4.1 Host Identity and Capability-based DoS Prevention

The host identity can be the public key of a host or HIP tag (e.g., can be IPsec SPI if ESP is used), which allows uniqueness of a host at the given time in a node and easy

authentication of the sender from the receiver (e.g., if the HI namespace is based on public-key cryptography).

Another approach of capability generation and validation is based on a pre-capability offered by each GONE routers (e.g., a hash function of the incoming 48-bit probabilistically unique interface ID and a timestamp).

## 4.2 Host Addressing and Application Interface

In GONE, hosts' IP addresses are only meaningful for the last/first-hop communication (between GONE hosts and GONE access routers). Due to the separation of locator and identifier, a host can choose any available means to obtain IP addresses, including but is not limited to manual configuration, stateless autoconfiguration and DHCP.

From the viewpoint of the high layer application interface with GONE, only (de)multiplexing is needed, since the GONE infrastructure overlay also provides the other desired transport layer functionality, such as fragmentation, connection management, congestion control and flow control. Multiplexing above GONE in host level ensures multiple upper-layer applications can use the same GONE overlay infrastructure service.

## 4.3 Mobility Considerations

When a GONE host roams from one network to another, it needs to update the DHT for the entry containing its host identity with its new locator (IP address). If this host is the data sender, it then initializes a new capability negotiation and GONE intermediary discovery, by sending a GONE_Req message towards the receiver and the receiver in turn generates a new capability. If the host is the data receiver, it notifies the sender to initialize capability negotiation.

## 4.4 Security Considerations

The security properties of GONE inherit GIST security and are extended with the additional capability-based denial of service prevention mechanism. Similar to [17, 18], the strength of the pre-capabilities and GIST discovery phase cookie determines the security level which GONE can achieve. It is conceivable that with the introduction of other discovery mechanisms in GIST, stronger or weaker messaging association security will be inferred.

Another issue is associated with GONE message security. It is arguable that there can be some scenarios requiring hop-by-hop security over end-to-end IPsec ESP-encrypted data. If such cases are necessary, IPsec and TLS may be used in securing hop-by-hop GONE messages for achieving both end-to-end confidentiality and hop-to-hop secure transport. GIST using TLS over SCTP is discussed in [28].

## 4.5 Performance and Deployment Considerations

One important aspect with any overlay solution is its performance. In the GONE design, the use of GIST over SCTP, locator-independent host identifier and capability-based DoS prevention allow flexible and generic resilient overlay with high availability and support for mobility. In a more realistic system implementation, developers need to carefully consider the way to move the GONE kernel stack to the OS kernel instead of user space, and avoid unnecessary data replication in a single GONE processing of data traffic.

Above all, GONE represents a way of building customized network edge using existing (well-specified) standards, for achieving various numerous fancy features such as a higher availability, reliable and DoS resilient network infrastructure as presented above. The market takeup of such a solution may involve several key steps: GIST availability, host level protocol stack update in hosts, edge's GONE support, and also importantly, considerations of its integration with the charging and AAA infrastructure.

# 5. IMPLEMENTATION STATUS AND FUTURE WORK

We have implemented a prototype of the GONE system in Linux [32], which supports any number of GONE intermediaries to provide soft state overlay routing and data delivery while conserving resilience. We are performing performance and scalability studies in an experimental testbed.

We have also used GONE to support several applications. One application for GONE is similar to RTP, which is designed to extend the GONE application interface to provide end-to-end transport functions suitable for applications transmitting stored multimedia data. Data can be accessed by any end system and made available at any time. GONE provides the resilient routing, DoS-preventing forwarding, and mobility functionality that a user desires. In particular, GONE efficiently and robustly routes messages across the wide-area by routing aross less loaded or secondary paths. Finally, the the simple (de-)multiplexing layer in GONE hosts allows easy distribution and collection of end-to-end user traffic.

Initial measurements show that GONE provides certain scalability in terms of the number of intermediaries, while leveraging GONE to provide fault-tolerant on-time packet delivery and minimal duplication of packets. Our next priority is on further performance analysis under a variety of conditions and parameters and evaluation on PlanetLab. This would help us better understand GONE's position in the overlay research space, as well as how it compares to other approaches such as RON, $i3$/H$i3$ and SOS, and possibly allow us to define a taxonomy of the research space. On the application side, we are developing intelligent network applications that exploit network-level statistics and utilize GONE routing to minimize data loss and improve latency and throughput.

Transport mechanisms in GONE play an essential role for overall network performance. One interesting topic here is fairness. We are currently studying the feasibility of applying concepts like multiTCP [33] in achieving fairness for data belonging to different sessions in a same messaging association.

# 6. CONCLUSION

In this paper, we presented GONE, an overlay architecture intended to be self-organized, scalable, DoS-limiting and robust wide-area infrastructure that efficiently routes traffic in the presence of path faults and node mobility. We showed how a GONE overlay network can be efficiently constructed and employ capability-based DoS prevention to enhance resilience and availability in dynamic and mobile environments. While GONE shows some similarities to RON [8], SOS [16], $i3$ [14] and HIP approaches [12, 24], we have embedded mechanisms that leverage soft state information and

provide self-management, robustness, dynamic routing detection and recovery in the presence of failures and high load by lower layer functions – GIST and SCTP, while eliminating the shortcoming of a lack of detailed protocol specification in some overlay systems and providing reusable software components for various services.

Moreover, GONE provides a plausible solution for customizing the network edge, where most fancy functions such as peer-to-peer, VoIP or NAT traversal are located. This paper presents such a use for dynamic overlay routing that need to deliver messages across ISP networks in a location independent manner, using usually pre-established messaging associations and without centralized services. GONE does this, in part, by using HIP host identifiers, capability concepts, as well as soft state and reuse of standard common signaling component in the network edge to achieve both mobility and enhanced service availability and network resilience.

## Acknowledgment

The authors would like to thank Jan Demter, Christian Dickmann and Henning Peters for their insightful comments and implementation efforts which helped the GONE design. In addition, Andreas Pashalidis and Hannes Tschofenig provided helpful feedbacks to the initial version of this paper.

## 7. REFERENCES

[1] S. Blake, D. L. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated service," RFC 2475, Dec. 1998.

[2] Y. Zhang, V. Paxon, and S. Shenker, "The Stationarity of Internet Path Properties: Routing, Loss and Throughput," ACIRI, Tech. Rep., May 2000.

[3] D. Moore, G. Voelker, and S. Sava, "Inferring Internet Denial-of-Service Activity," in *Proc. Usenix Security Symposium*, 2001.

[4] J. Kempf and R. Austein, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture," RFC 3724, Mar. 2004.

[5] D. Clark, C. Partridge, R. Braden, B. Davie, S. Floyd, V. Jacobson, D. Katabi, G. Minshall, K. Ramakrishnan, T. Roscoe, I. Stoica, J. Wroclawski, and L. Zhang, "Making the World (of Communications) a Different Place," *Computer Communication Review*, vol. 35, no. 2, pp. 91–96, July 2005.

[6] L. Wang, K. Park, R. Pang, V. Pai, and L. Peterson, "Reliability and Security in the CoDeeN Content Distribution Network," in *Proc. USENIX Annual Technical Conference*, Boston, MA, June 2004.

[7] M. Bagnulo, A. Garcia-Martinez, A. Azcorra, and D. Larrabeiti, "Survey on proposed IPv6 multi-homing network level mechanisms," Internet draft (draft-bagnulo-multi6-survey6), work in progress, July 2001.

[8] D. Andersesn, H. Balakrishnan, M. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Proc. SOSP*, 2001.

[9] C. Perkins, "IP Mobility Support for IPv4," Internet Engineering Task Force, RFC 3344, Aug. 2002.

[10] D. B. Johnson, C. E. Perkins, and J. Arkko, "Mobility support in IPv6," Internet Engineering Task Force, RFC 3775, June 2004.

[11] B. Aboba, "IAB Considerations for the Split of Identifiers and Locators," draft-iab-id-locsplit-00.txt, work in progress, Mar. 2004.

[12] R. Moskowitz and P. Nikander, "Host Identify Protocol Architecture," Internet draft (draft-ietf-hip-arch-03), work in progress, June 2005.

[13] E. Nordmark and M. Bagnulo, "Level 3 multihoming shim protocol," Internet draft (draft-ietf-shim6-proto-03), work in progress, Sept. 2005.

[14] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet Indirection Infrastructure," in *Proc. SIGCOMM*, 2002.

[15] R. Mahajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling high bandwidth aggregates in the network," *Computer Communication Review*, vol. 32, no. 3, pp. 62–73, 2002.

[16] A. Keromytis, V. Misra, and D. Rubenstein, "SOS: Secure Overlay Services," in *Proc. SIGCOMM*, 2002.

[17] X. Yang and D. Wetherall and T. Anderson, "A DoS-limiting Network Architecture," in *Proc. SIGCOMM*, 2005.

[18] H. Schulzrinne and R. Hancock, "GIST – General Internet Signaling Transport," Internet draft (draft-ietf-nsis-ntlp-09), work in progress, Feb. 2006.

[19] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, "Host Identify Protocol," Internet draft (draft-ietf-hip-base-04), work in progress, Oct. 2005.

[20] X. Fu, H. Schulzrinne, A. Bader, D. Hogrefe, C. Kappler, G. Karagiannis, H. Tschofenig, and S. Van den Bosch, "NSIS: A New Extensible IP Signaling Protocol Suite," *IEEE Communications Magazine*, vol. 43, no. 10, pp. 133–141, Oct. 2005.

[21] J. Crowcroft, S. Hand, R. Mortier, T. Roscoe, and A. Warfield, "Plutarch: An Argument for Network Pluralism," in *SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, Aug. 2003.

[22] P. Nikander, J. Ylitalo, and J. Wall, "Integrating Security, Mobility, and Multi-homing in a HIP Way," in *Proc. NDSS*, 2003.

[23] P. Nikander, J. Arkko, and T. Henderson, "End-Host Mobility and Multi-Homing with the Host Identity Protocol," Internet draft (draft-ietf-hip-mm-01), work in progress, Feb. 2005.

[24] P. Nikander, J. Arkko, and B. Ohlman, "Host Identity Indirection Infrastructure (Hi3)," in *Proc. 2nd Swedish National Computer Networking Workshop*, Karlstad, Sweden, Nov. 2004.

[25] N. Feamster, D. Andersen, H. Balakrishnan, and M. Kaashoek, "Measuring the Effects of Internet Path Faults on Reactive Routing," in *Proc. SIGMETRICS*, 2003.

[26] F. Guo, J. Chen, W. Li, and T. Chiueh, "Experiences in Building a Multihoming Load Balancing System," in *Proc. INFOCOM*, 2004.

[27] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," RFC 2205, Sept. 1997.

[28] X. Fu, C. Dickmann, and J. Crowcroft, "General Internet Signaling Transport (GIST) Over SCTP," Internet draft, work in progress, Feb. 2006.

[29] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-Peer Lookup Service For Internet Applications," MIT, Tech. Rep. TR-819, Jan. 2002.

[30] S. Rhea, B. Godfrey, B. Karp, J. Kubiatowicz, S. Ratnasamy, S. Shenker, I. Stoica, and H. Yu, "OpenDHT: A Public DHT Service and Its Users," in *Proc SIGCOMM*, 2005.

[31] R. Hancock, G. Karagiannis, J. Loughney, and S. V. den Bosch, "Next Steps in Signaling (NSIS): Framework," RFC 4080, June 2005.

[32] "GONE Implementation." [Online]. Available: http://user.informatik.uni-goettingen.de/~fu/gone

[33] J. Crowcroft and P. Oechsli, "Differentiated End-to-End Internet Services using a Weighted Proportional Fair Sharing TCP," *Computer Communication Review*, vol. 28, no. 3, pp. 53–69, 1998.