

Extended Abstract: Authentication on the Edge - Distributed Authentication for a Global Open Wi-Fi Network

Nathanael A Thompson, Zuoning Yin,
Haiyun Luo
Dept. of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL, USA
{nathomps,zyin2,haiyun}@cs.uiuc.edu

Petros Zerfos, Jatinder Pal Singh
Deutsche Telekom AG Laboratories
Ernst-Reuter-Platz 7 Berlin, Germany
{petros.zerfos,jatinder.singh}@telekom.de

ABSTRACT

A global-scale low cost outdoor Internet access infrastructure is finally attainable. Emerging projects are leveraging the proliferation of private Wi-Fi networks to build a global-scale ubiquitous access infrastructure from autonomous, independently owned Internet connections at homes and other private properties. To ensure the traceability and accountability required by the broadband ISPs and private owners of these Wi-Fi networks, reliable authentication and authorization are needed. This paper describes *authentication on the edge*, a localized and distributed authentication method for open Wi-Fi networks. Three main ideas are used to adapt to the variability and unreliability of these networks: the use of certificate-based authentication, the distribution of a segmented certificate revocation list to all entities and the self organization of access points into a social lookup network. These methods achieve the scalability needed for the overwhelming size and volume of a global network and increase resiliency against temporary failures in the infrastructure.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.2.3 [Network Operations]: Network management

General Terms

Management, Performance, Reliability

Keywords

Authentication, EAP-AGE, Social Network, Wi-Fi Networks

1. INTRODUCTION

High speed low cost ubiquitous Internet access has been a long-standing vision for years, attracting major effort from both academia and industry. As shown in recent studies [6, 8], indoor Internet connectivity is becoming pervasive in the U.S. as a result of the steady growth of broadband penetration to the home and the proliferation of 802.11-based wireless local area networks (WLANs) deployed in households and other popular indoor locations. On the other

hand, outdoor coverage for broadband Internet access is seriously lagging behind despite the explosion of wireless network technologies and deployments over the last decade.

We see a fundamentally new approach emerging to bridge this gap which has recently been conceived in both academia (Extended HotSpots [4], OBAN [12], P2PWNC [13], PERM [16]) and industry (ABitCool [2], Fon [5]). Following the peer-to-peer spirit, individual users share their subscribed broadband access over their private wireless routers under the supervision of an authoritative party. The approach is born with unprecedented *scalability* and *cost-effectiveness* when compared to other alternatives because each new user adds the necessary infrastructure to support herself. We call this global-scale Internet access infrastructure based on *privately* owned wireless Internet access points (either residential or small business) “GIANT”. The method of deployment and maintenance of a GIANT network is simple: a user opens up her own Internet connection through her wireless router to other participating users who in turn grant her access when she is away from her own fixed Internet connection. To satisfy legal and regulatory requirements of the ISPs over which the GIANT network will run, some *operator-assistance* is required. This trusted third-party operator is responsible for managing user credentials, handling billing and also enforcing fair and controllable sharing.

In this paper we present our design and implementation of a scalable and resilient service for *authentication on the edge* (AGE) of GIANT networks that is used for access control of the nomadic users. A set of *semi-distributed* algorithms and protocols operate under light-weight centralized coordination to authenticate users for controlled access. Our design of AGE seeks to strike a balance between fully centralized approaches and fully distributed approaches in order to capitalize on the simplicity and ease of management of the former and scalability and ideal robustness of the latter. AGE’s mechanisms allow the bottleneck and single point of failure found in existing authentication schemes to be avoided.

AGE supports a single authentication authority allowing clients to access the service anywhere in the world with the same user id and authentication credentials with as little user interaction as possible. The key components of AGE are the use of certificates and public/private keys, segmentation of certificate revocation lists (CRLs) based on the subscriber’s primary physical location and a social overlay formed among the access points of friends in the GIANT network, which is used to support trusted verification of foreign CRL segments. We implement AGE and verify the

feasibility of our approach using a large social graph of SMS users extracted from traces obtained from a national cellular provider. We then evaluate AGE in a GIANT setting and demonstrate its performance enhancements over EAP-TLS by showing that it is able to achieve an average 47.9% reduction in authentication delay.

2. MOTIVATION AND CHALLENGES

Recent collaborative Internet access efforts such as PERM [16], MoB [11], and P2PWNC [13], as well as metropolitan community groups [9] are creating open Wi-Fi networks. They are mostly focusing on the individual sharing relationships among peers and do not address the fundamental issue of controlled access to those networks. In order for GIANT networks to see wide adoption, access control across all wireless access points is needed to ensure user traceability and accountability. The latter is also required in order to satisfy legal and regulatory requirements posed to ISPs over which GIANT will operate. Leaving authentication and access control to individual end users is not a viable solution as most of them are not technically adept or willing to handle such task.

On the other hand, centralized authentication approaches such as those based on the EAP/802.1X family of standards [1, 15] are ill-suited for the GIANT network which changes many assumptions about their operating environment. Firstly, many of the existing methods are not designed to scale to the size of GIANT. As the authentication server involved in EAP/802.1X is usually intended for deployment in enterprise-scale LANs with hundreds or maybe thousands of clients, scaling it to handle potentially millions of active clients would become prohibitively expensive, challenging GIANT's cost-effective method of deployment.

Secondly, because the existing frameworks are designed for LAN authentication, it is assumed that all entities are deployed in a well-managed local enterprise network. The GIANT system though acts as an overlay operator (or an overlay wireless ISP) on an access infrastructure that is collectively owned by its actual users, without centralized management. This leaves the authentication function vulnerable to general Internet outages and distributed DoS attacks against the authentication server. The access points should be able to authenticate clients reliably and autonomously, even when access to functions provided by a centralized authentication server is temporarily unavailable.

Thirdly, network anomalies, latency and intermittent loss of connectivity are frequently encountered in wide area networks at the Internet scale. The entire authentication process using EAP/802.1X methods involves several rounds of communication between the client and the authentication server. Traversing the Internet leads to high and variable delay in the completion time, and, in the case of packet loss, it becomes difficult to gauge an appropriate timeout before restarting the authentication process. However, low authentication delay is critical for providing micro-mobility management and handoffs between neighboring GIANT access points.

The current design tenet for existing community Wi-Fi sharing networks that operate under centralized authorities such as private hotspots [3, 7], or FON [5] and aBitCool [2] is to use captive portals to perform authentication. Users are required to type a password whenever associating to a new access point, which is then forwarded to a centralized au-

thentication server for processing. Typically these networks provide no wireless encryption because the authentication is performed at higher levels in the network stack. Due to lack of security, potentially long authentication delays and the centralized authentication server becoming a single point of failure, captive portals are also not appropriate for securing GIANT networks.

From the above discussion, it is clear that centralized approaches for authentication and access control would be inadequate for the GIANT environment. Deploying intermediate servers or replicating servers, like OBAN [12] and cellular networks, faces the challenge of optimally placing servers for effective coverage of a dynamically growing, opportunistically deployed GIANT network. Having additional servers also leads to increased maintenance and financial cost, failing to capitalize on the inherent scalability of GIANT networks. An ideal system should allow access points to operate independently regardless of the state of the backend management system, while staying within the control of a central authority.

3. AUTHENTICATION ON THE EDGE

To address the challenges of scale, outages and delay we propose *Authentication on the Edge* (AGE) for GIANT networks. AGE localizes the authentication process on the access points and avoids the unpredictability and potential unreliability of Internet-based authentication. AGE is comprised of three components: (1) mutual authentication between an access point and a client using certificates, (2) distribution of different CRL segments to participants at different locations and (3) fallback p2p lookup of the CRL freshness. AGE achieves scalability and resiliency by excluding the on-line involvement of the back-end authentication/directory server as illustrated in Figure 1.

The AGE network is composed of two types of devices, access points (APs) and wireless clients. Every device is owned by a single AGE participant. Typically each participant owns only a single AP but might own many clients. Certificates are issued by a trusted central authority (CA) and are verifiable given the CA's public key. By installing the CA public key on each GIANT device access to the CA is not necessary for most authentications. Certificate based authentication also has the advantage that it requires no user interaction during authentication as compared to password based schemes.

During a certificate's lifetime, which is usually on the order of months or even years, it may be necessary to revoke the certificate because of key compromise, cancellation of service or for punitive reasons. The AGE CA maintains a certificate revocation list (CRL) to track the validity of certificates. To prevent the need to download the CRL from the CA during authentication, each GIANT device carries the CRL along with its certificate and exchanges both as its credentials. Given GIANT's global nature the number of entries in the CRL could be quite large resulting in CRLs that exceed the storage capacity of the GIANT APs and clients. Therefore the CRL is partitioned and managed in segments. Every certificate falls in a subspace encoded by a CRL segment called the "containing CRL" for that certified name. A certificate's non-existence in its containing CRL proves the validity of the certificate up to the time when the CRL segment was generated.

In most cases devices can verify each other using the lo-

cally stored certificates and CRL segments. Occasionally, however, the presented CRL segment might be deemed too old for the AP’s freshness criteria. In a global heterogeneous system of autonomous entities like GIANT, it is difficult to reach consensus regarding a freshness cutoff. Rather than enforce a single threshold, AGE APs are organized into a peer-to-peer overlay to find the most recent version of a CRL segment. A p2p network is favored over centralized approaches like OCSP [14] because of its resiliency and zero infrastructure cost. To reduce network load the GIANT CA maintains and publishes CRL segment timestamps signed by the CA confirming the last update time for the segment. Rather than fetching the entire CRL segment only the timestamp need be transferred.

Assigning the containing CRL for a given participant is an important task to reduce the number of CRL segment lookups and the corresponding network overhead. If an authenticating client has the same containing CRL as the AP (a local authentication) then no lookup is required. Therefore, the assignment of participants to CRL segments should be done to maximize the number of local authentications. Users typically spend most of their time in a small number of physical locations (at work, at home, at a favorite cafe). Thus the majority of authentications can be expedited if devices in the same physical locations share the same containing CRL. To map locations to CRL segments we divide physical areas into non-overlapping zones which completely cover the area. Each zone is represented by one CRL segment which is assigned a unique id. All participants located in the same zone have their certificates assigned to the same CRL segment.

In order to verify the freshness of a given CRL segment the AGE p2p network must provide a trustworthy lookup function to map a CRL id to a CRL segment. Although distributed hash tables solve this problem the lookup can be manipulated by any node in the routing path making the results untrustworthy. To avoid untrusted links the access points in AGE instead form a social network overlay based on the trusted friendship of the owners. Using an interface on the AP each AGE participant maintains a friends list as she would for an instant messaging client. Trusted overlay links are built based on the friends list. To convert between a friend’s name and the current IP address of the friend’s AP, each AP publishes its current IP address in a separate overlay. Updates are verified using the publishing AP’s certificate so that an AP can only update entries for its own certified name. After finding an IP address for a friend, TLS authentication is performed between the APs to ensure the remote AP’s identity.

The CRL segment query is unique in that the query needs to be forwarded only to an area, not a specific AP. Any AP with a containing CRL that matches the request can reply. Routing through the social network then is a matter of finding a friend or a friend of a friend, etc. who is in the target zone. Using the intuition that a friend closer to the target zone is more likely to know someone in the target zone, greedily forwarding the request to the friend closest to the target should usually minimize the routing hops.

The AGE social network uses the location of each zone to forward messages between friends. Each AP keeps a listing signed by the AGE CA of all known zones and their coordinates. During routing an AP’s location is the same as its containing zone. This location based approach to routing

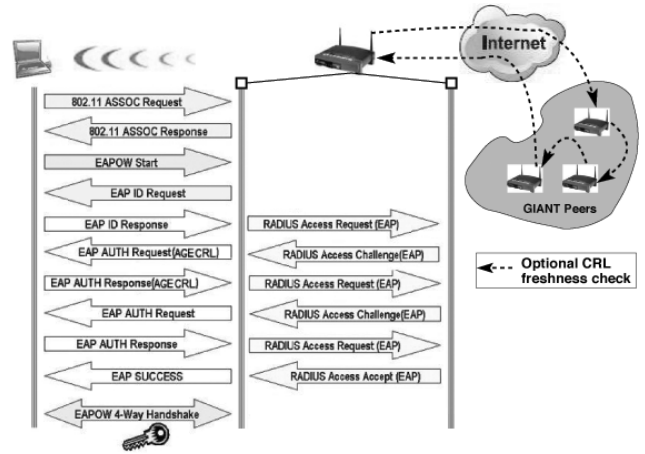


Figure 1: AGE authentication is fully localized between the access point and the client. *On-line* CRL freshness lookups are optional and rarely necessary.

is the same as greedy-face forwarding in ad hoc and sensor wireless networks, with the addition of long distance links. In these routing protocols a node greedily forwards a message to the neighbor closest to the destination. If the current node is closest amongst its neighbors to the destination the algorithm switches to some variant of face routing until greedy forwarding can be resumed. We use the GFG algorithm [10] to perform routing in the AGE social network.

The basic GFG routing will fail in AGE in the situation that an AP only has links to other APs in the same zone. To resolve the issue, AGE uses *intra-zone* connectivity of its friends to select the best next hop. Selecting the friend with the most intra-zone connections ensures routing can resume and increases the likelihood of finding a long distance link toward the destination. In the extreme case that an AP has no intra-zone friends and no friends with intra-zone friends the message is forwarded to a randomly selected friend.

When a new participant joins the GIANT network she first registers with the GIANT provider either through a dedicated registration server or using off-line methods. A unique username and temporary password are granted. She then installs the AGE software on her AP and client devices. A special registration application on the AP generates a new private/public key pair. Using the temporary password the public key is submitted to the AGE CA and a new certificate binding the public key to the new username is generated and transferred to the AP. The participant’s client devices synchronize with the AP to get new certificates and CRL segments. Any future updates to the CRL segment or user certificate are pushed from the CA server to individual APs.

4. SOCIAL NETWORK ANALYSIS

AGE’s ability to correctly authenticate users depends on the connectivity of the social network. If an AP cannot forward requests to a zone then it may reject valid clients from that zone. To verify the feasibility of AGE’s social network we extract the social graph from the SMS traces of a large nationwide cellular provider. A cellular network is a good approximation for a GIANT network because of its large user base and geographic coverage. Also each cellular user is assigned to a home mobile switching center (MSC)

	EAP-TLS/LAN	EAP-TLS/GIANT	EAP-AGE/GIANT
Client Processing Time	0.021556 [1.00x]	0.021171 [0.98x]	0.028369 [1.32x]
Server Processing Time	0.033561 [1.00x]	0.013828 [0.41x]	0.485165 [14.46x]
Network Delay	0.065948 [1.00x]	0.964469 [14.62x]	0.007132 [0.11x]
Total Auth. Delay	0.121065 [1.00x]	0.999468 [8.26x]	0.520666 [4.30x]

Table 1: Latency breakdown in seconds for EAP-TLS in a baseline LAN deployment, EAP-TLS in GIANT and AGE in GIANT. The change compared to EAP-TLS/LAN is shown in brackets.

identified in her telephone number, mirroring AGE’s CRL segment zones. Because SMS exchange is a social activity between friends the connections between SMS users form a social graph across zones. We generate the SMS social graph from our traces by forming a link between two users if at least one of the pair sent the other at least one SMS message. Our traces cover the individual SMS messages from over 6 million users and 22 different MSCs.

Because routing in AGE’s social network depends on intra-zone and long distance links we examined the SMS social graph for such relationships. One third of the users have only one friend, likely infrequent SMS users. In contrast, some users have tens of thousands of friends. These users likely are bots or services from the cellular provider. Over 99.99% of users have fewer than 200 friends. We filter out the users with fewer than 2 friends (infrequent users) and more than 200 friends (bots). In the filtered set there is an average of 4.2 friends per user. We define intra-zone friends as friends with a different MSC and long distance friends as intra-zone friends in non-adjacent MSCs as determined by the GPS coordinates of the MSCs. In the filtered user set, 94% of all users have at least one intra-zone friend and 33% of users have at least one long distance friend.

To decide if these connectivity properties could support the forwarding of CRL segment lookups to all zones we evaluated the reachability from every user to each of the 22 MSCs. An MSC is reachable if any member in the MSC can be found following only the friendship links in the SMS social graph. Our analysis shows that the distribution of reachable MSCs is bimodal - 70% of users can reach all 22 MSCs while the remaining can reach only 5 or less. Among the users with incomplete reachability the average number of friends per user was between 1.71 and 3.17. Among the well connected users the average number of friends was 5.02. This result suggests that if the average degree in the AGE social network is 5 or higher than the reachability to all AGE zones can be maintained. In the SMS traces 22% of users had 5 or more friends.

5. EVALUATION

We implemented AGE as a new EAP type and compared its authentication delay due to both network and processing latencies with EAP-TLS. We compared our AGE implementation to the existing IEEE 802.1x EAP-TLS because it is most similar in features to AGE. All EAP-TLS requests were handled by a workstation running **FreeRADIUS**. We compare EAP-AGE to two EAP-TLS deployments, one in a fast and reliable LAN and one in a GIANT network where the authentication server is only reachable over the Internet. In all experiments a Linux laptop in Berlin, Germany was authenticated using `wpa_supplicant` and internal 802.11b/g wireless card. The AP was the Linksys WRT54G running a modified version of the OpenWRT software. In the GIANT

EAP-TLS scenario a slightly more powerful authentication server was located in Urbana, IL, USA. Table 1 shows the breakdown in processing and network overhead for each scenario. The average AGE delays were about 4 times longer than LAN EAP-TLS but half as long as GIANT EAP-TLS. EAP-AGE reduces the network delay to only 7ms compared to 65ms for LAN EAP-TLS and 900ms for GIANT EAP-TLS. However, AGE suffers a 14x increase in processing delay, accounting for 94% of the authentication time, because of the limited power of the AP’s 200mhz processor.

6. CONCLUDING REMARKS

GIANT networks offer a promising new approach for large-scale, outdoor Internet access. Authentication on the Edge (AGE) running in GIANT networks will pave the way for wide scale adoption of GIANT, as users can trust the security and performance enabled by AGE while operators can be confident in having the traceability and accountability they require. AGE’s mechanisms of localizing and distributing authentication enable security and accountability to easily scale with the enormous growth potential inherent in GIANT networks.

7. REFERENCES

- [1] 802.1x. <http://standards.ieee.org/getieee802/download/802.1X-2004.pdf>.
- [2] Abitcool wi-fi community. <http://www.abitcool.com/>.
- [3] Boingo wireless service. <http://boingo.com/>.
- [4] Extended HotSpots. http://www.extended-hotspots.net/opencms/opencms/index.html?_locale=en.
- [5] Fon. <http://www.fon.com/>.
- [6] June 2006 bandwidth report. <http://www.websiteoptimization.com/bw/0606/>.
- [7] T-mobile hotspots. <http://hotspot.t-mobile.com/>.
- [8] Wi-Fi Surpasses Ethernet for Home Networking. http://www.parksassociates.com/press/press_releases/2005/gdl2.html.
- [9] M. Bina and G. Giaglis. Emerging issues in researching community-based w lans. *Journal of Computer Information Systems*, Fall 2005.
- [10] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [11] R. Chakravorty, S. Agarwal, S. Banerjee, and I. Pratt. Mob: a mobile bazaar for wide-area wireless services. In *Proceedings of ACM MobiCom*, 2005.
- [12] E. Edvardson. Fixed and mobile convergence. In *Proceedings of BroadBand Europe*, 2004.
- [13] E. C. Efstathiou, P. A. Frangoudis, and G. C. Polyzos. Stimulating participation in wireless community networks. In *Proceedings of IEEE INFOCOM*, 2006.
- [14] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. Online certificate status protocol - OCSP. IETF Request For Comments (RFC2560), 1999.
- [15] D. Stanley, J. Walker, and B. Aboba. RFC 4017: Extensible authentication protocol (EAP) method requirements for wireless LANs, Mar. 2005.
- [16] N. Thompson, G. He, and H. Luo. Flow scheduling for end-host multihoming. In *Proceedings of IEEE INFOCOM*, 2006.