

Suppressing Bot Traffic with Accurate Human Attestation

Muhammad Jamshed
Computer Science
Department
University of Pittsburgh
210 S. Bouquet St
Pittsburgh, PA 15260
USA
ajamshed@cs.pitt.edu

Wonho Kim
Computer Science
Department
Princeton University
35 Olden Street
Princeton, NJ 08544
USA
wonhokim@cs.princeton.edu

KyoungSoo Park
Electrical Engineering
Department
KAIST
335 Gwahangno, Yuseong-gu
Daejeon, 305-701
Republic of Korea
kyoungsoo@ee.kaist.ac.kr

ABSTRACT

Human attestation is a promising technique to suppress unwanted bot traffic in the Internet. With a proof of human existence attached to the message, the receiving end can verify whether the content is actually drafted by humans. This technique can significantly reduce bot-generated abuse such as spamming, password cracking or even distributed denial-of-service (DDoS) attacks. Unfortunately, existing methods rely on the probabilistic characteristics of attestations and can be exploited by smart attackers.

In this paper, we propose deterministic human attestation based on trustworthy input devices. By placing the root of trust on the input device, we tightly bind the input events to the content for network delivery. Each input event is generated with a cryptographic hash that attests to human activity and the message consisting of such events gets a third-party verifiable digital signature that is carried to the remote application. For this, we augment the input device with a trusted platform module (TPM) chip and a small attester running inside the device. We focus on trustworthy keyboards here but we plan to extend the framework to other input devices.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and protection; D.4.6 [Operating Systems]: Security and Protection

General Terms

Security, Reliability

Keywords

Human Attestation, Bot Traffic Suppression, Trusted Computing, Networked System Security

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

APSys 2010, August 30, 2010, New Delhi, India.

Copyright 2010 ACM 978-1-4503-0195-4/10/08 ...\$10.00.

1. INTRODUCTION

Unsolicited messages have widely permeated into our daily Internet life. Despite decade-long efforts, the volumes of email, blog, and instant messaging (IM) spam are continually on the rise, and millions of infected botnet machines aggravate the situation. 87.7% of the total emails, that is, 107 billion messages being sent globally on a daily basis, are estimated to be spam in 2009 [7]. University of Pittsburgh received 51 million spam emails in November 2009 alone while only 17 million emails were delivered as legitimate to the users [10]. While best practices such as content-based filtering [2, 19, 21], DNS blacklisting [12] and network-based fingerprinting [4, 11] greatly reduce the spam delivery, they often create false positives - legitimate emails classified as spam, making the Internet message delivery less reliable.

Human attestation is a promising technique that can potentially exterminate unwanted bot traffic. By carrying a non-forgable proof of human existence with the message, the receiving end can reliably determine the identity of the traffic source and adjust her filtering policy to better accommodate human traffic. Existing methods typically infer the human activity from key clicks or mouse events and use TPM-generated signatures as human attestation. This approach is shown to be effective in reducing spam, DDoS attacks, click frauds, and so on [3, 8]. However, one serious problem is that they depend on the probabilistic characteristic of the attestation. For example, Not-a-Bot [3] allows any bot to get a human attestation for its own content within an allowed timing window after key or mouse clicks are generated. Though the rate of illegitimate attestations is bounded by that of human activity, smart attackers can deterministically bypass any spam filter based on human attestation. That is, bots could generate sophisticated spam messages for targeted users or act like humans when they access Internet banking or E-commerce sites if they adopt existing methods as human identification.

The key question this paper poses is how strongly one can bind human activity to the content that is sent over the network. We argue that the content should carry a signature that can be independently verified that its content was actually created by a human. We build a framework in which each element in the content is confirmed to be human-generated. In our framework, only the part of the content that is made of keystrokes gets a human attestation and the rest of it is tagged as unattested before delivering to the remote application. Based on the accurate human content attestation, the remote application can make an informed

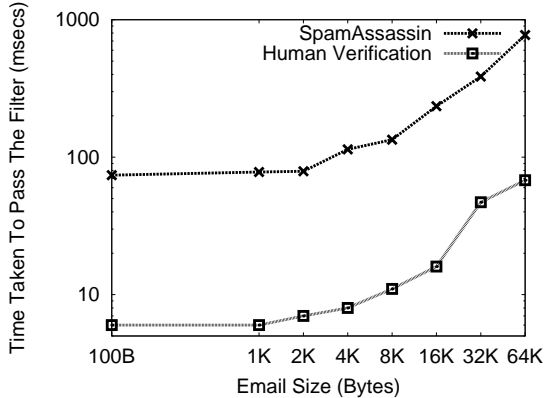


Figure 2: Time to verify human-attested emails. On 2.50 GHz Intel Pentium single-core machine with 6 GB RAM.

terface. We use Linux 2.6.21 and an on-board TPM from STMicroelectronics that ships with Dell Optiplex 780, an 2.70 GHz Intel dual-core machine with 3 GB memory.

We first measure the TPM function latency and show the results in Table 2. TPM_Extend loads the data to be signed to Platform Configuration Register (PCR) and TPM_PCR_Reset flushes the data. TPM_Quote signs the PCR value using the private AIK loaded by TPM_LoadAIK_Key. TPM_Quote (2048-bit RSA) takes 746 ms, which is the slowest operation in getting the attestation. A typical 2048-bit TPM_Quote size is 580 bytes long consisting of an AIK (256B), a TPM-constructed data blob containing essential PCR attributes (48B), the RSA signature (256B) of that data blob and the SHA1 hash (20B) of the actual data. The current TPM specification do not allow AIK certification for key sizes smaller than 2048 bits. This is due to the concern that smaller RSA keys can be broken in the future.

4.1 Human-typed Email Verification

We have implemented a Mozilla Thunderbird 2.0.0.23 extension that attaches Base64-encoded human attestation signature to every human-typed email. We also implement a verification filter that works with Postfix 2.5.5-1 [9]. The filter scans the human attestation signature attached to emails and flags them as human-generated or unattested.

Figure 3 shows our Postfix filter performance. All emails in the graph take less than 70 ms for human verification. We compare it with the filtering latency of SpamAssassin 3.2.5, a popular content-based spam filter, with the default configuration. Our verification performance is comparable to that of SpamAssassin for small content, but as the size of the email increases, the performance gap widens. Our filter is 12 times faster than SpamAssassin at 64 KB email.

4.2 SSH Command Attestation

We have added human attestation to the Dropbear 0.52 [1] SSH client/server suite so that each ssh command carries an attestation signature and the server verifies the signature before launching the command. That is, for every enter key, we send the signature to the server for human verification.

Figure 3 shows the attestation signature generation time as the command size increases. We downgraded the CPU clock frequency of Dell Optiplex 780 to 1.20 GHz to simulate a low-power processor that the attester is likely to be using.

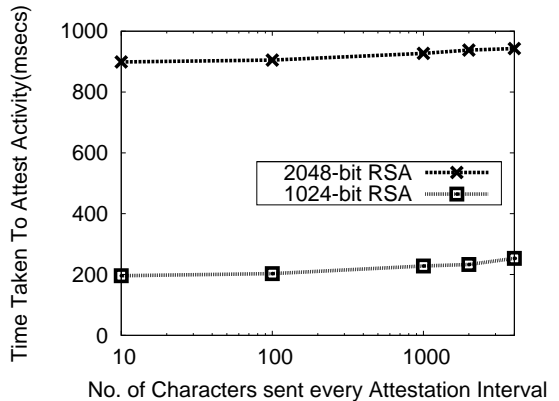


Figure 3: Time taken to attest a group of characters on a Dropbear ssh session. On a 1.20 GHz Intel Pentium single core machine with 3GB RAM.

We see relatively high latency for 2048-bit RSA signature generation, reflecting the slow operation of TPM_Quote. In contrast, 1024-bit RSA signatures show much smaller latency, which makes it practical for interactive environments. Given that 1024-bit RSA is widely used in the SSL protocol, we believe 1024-bit RSA can be a reasonable choice for the current day use if the the TPM specification allows it. One nice trend is that the response time stays more or less the same regardless of the command size. We hope that the performance of TPM_Quote improves as the demand for faster TPM grows.

5. RELATED WORK

Not-a-Bot (NAB) [3] motivated our work. NAB proposes a simple human attestation framework that can be used to reduce spam, click frauds, DDoS attacks, etc. They infer the human existence from input device events, and allow human content attestation with a TPM within an acceptable time window (they use one second) from the last key click. However, their scheme can be abused by smart attackers, which is noted in their work. We share the same goal of suppressing the bot traffic, but provide a stronger attestation framework that eliminates the bot abuse by tightly binding the key events to the content. Our earlier work explored bot detection in the Web environment by having obfuscated Javascript code catch the input events of the users [8]. It has shown to be effective to block most Web abuse on CoDeeN content distribution networks [18], though it does not prevent sophisticated bots that generate software interrupt-based input events.

Sailer *et al.* developed the first framework that employs a TPM for static software stack-based remote attestation [13]. McCune *et al.* built a *trusted path* between sandboxed software and the remote application by using dynamic root of trust [6] and reduced the trusted computing base to a single secure hypervisor having a small code footprint. Our framework adds human attestation and can further improve the level of trust in human-interactive applications as it eliminates the dependence of a trusted software stack from the scheme.

Modern spam filters [2, 15] typically rely on content-based filtering coupled with IP-based [14] and URL-based [16] black-

listings. There has been significant amount of research on devising accurate detection techniques to defend against spammers primarily originating from botnets [4, 5, 12, 17, 20]. They usually add a set of features to spam-detecting classifiers that capture the behavior of botnets. While the spam detection accuracy has improved greatly over time, the inherent probabilistic nature of learning-based filtering often creates false positives, which makes the legitimate email delivery unreliable. Human attestation can potentially bring the false positive rate to zero.

6. CONCLUSION

We have developed an accurate human attestation framework that draws the trust from the input device. Without any dependency on the software stack integrity beyond trustworthy input devices, remote applications can verify if the content is actually created by humans. We have also shown that it is straightforward to employ our scheme into existing applications that deal with human-typed content.

The immediate next step would require adding other input devices such as mouse. We are working on capturing the sequence of mouse events and connecting them to the content creation context. Remote verifier policy is another interesting area to explore. With the timestamp of each character, one can calculate the fraction of the words created from left-to-right and use the information to block any content that is mixed and matched by sophisticated bots. User privacy can be a concern, and minimizing the exposure of the timing information while maintaining a good filtering quality should be further studied.

7. ACKNOWLEDGEMENT

We thank Daniel Mosse and Sunghwan Ihm for valuable discussion. We also thank anonymous APSys reviewers for their insightful comments. This research was funded by KAIST award G04100004.

8. REFERENCES

- [1] Dropbear SSH server and client. <http://matt.ucc.asn.au/dropbear/dropbear.html>.
- [2] Google Postini Services. <http://www.google.com/postini>.
- [3] R. Gummadi, H. Balakrishnan, P. Maniatis, and S. Ratnasamy. Not-a-Bot (NAB): Improving service availability in the face of botnet attacks. In *Proceedings of USENIX NSDI*, 2009.
- [4] S. Hao, N. A. Syed, N. Feamster, A. G. Gray, and S. Krasser. Detecting spammers with SNARE: Spatio-temporal network-level automatic reputation engine. In *Proceedings of the USENIX Security Symposium*, 2009.
- [5] J. P. John, A. Moshchuk, S. D. Gribble, and A. Krishnamurthy. Studying spamming botnets using botlab. In *Proceedings of USENIX NSDI*, 2009.
- [6] J. McCune, B. Parno, A. Perrig, M. Reiter, and H. Isozaki. Flicker: An execution infrastructure for TCB minimization. In *Proceedings of EuroSys*, 2008.
- [7] MessageLabs Intelligence: 2009 Annual Security Report. http://www.messagelabs.com/mlireport/2009MLIAnnualReport_Final_PrintResolution.pdf.
- [8] K. Park, V. S. Pai, K.-W. Lee, and S. Calo. Securing web service by automatic robot detection. In *Proceedings of the USENIX Annual Technical Conference*, 2007.
- [9] Postfix Mail Transfer Agent. <http://www.postfix.org/>.
- [10] Private conversation with Ben Carter, IT staff member for the University of Pittsburgh, 2009.
- [11] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *Proceedings of ACM SIGCOMM*, 2006.
- [12] A. Ramachandran, N. Feamster, and S. Vempala. Filtering spam with behavioral blacklisting. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2007.
- [13] R. Sailer, X. Zhang, T. Jaeger, and L. v. Doom. Design and implementation of a TCG-based integrity measurement architecture. In *Proceedings of the USENIX Security Symposium*, 2004.
- [14] Spam and Open Relay Blocking System (SORBS). <http://www.au.sorbs.net/>.
- [15] The Apache SpamAssassin Project. <http://spamassassin.apache.org/>.
- [16] URL Blacklist. <http://urlblacklist.com/>.
- [17] S. Venkataraman, S. Sen, O. Spatscheck, P. Haffner, and D. Song. Exploiting network structure for proactive spam mitigation. In *Proceedings of the USENIX Security Symposium*, 2007.
- [18] L. Wang, K. Park, R. Pang, V. Pai, and L. Peterson. Reliability and security in the CoDeeN content distribution network. In *Proceedings of the USENIX Annual Technical Conference*, 2004.
- [19] Y. Xie, F. Yu, K. Achan, R. Panigrahy, G. Hulten, and I. Osipkov. Spamming botnets: Signatures and characteristics. In *Proceedings of ACM SIGCOMM*, 2008.
- [20] Y. Zhao, Y. Xie, F. Yu, Q. Ke, Y. Yu, Y. Chen, and E. Gillum. BotGraph: Large scale spamming botnet detection. In *Proceedings of USENIX NSDI*, 2009.
- [21] L. Zhuang, J. Dunagan, D. R. Simon, H. J. Wang, and J. D. Tygar. Characterizing botnets from email spam records. In *Proceedings of the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET) Botnets, Spyware, Worms, and More*, 2008.