



SmartAuth: Dynamic Context Fingerprinting for Continuous User Authentication

Davy Preuveneers
iMinds-DistriNet

Department of Computer Science, KU Leuven
Leuven, Belgium
davy.preuveneers@cs.kuleuven.be

Wouter Joosen
iMinds-DistriNet

Department of Computer Science, KU Leuven
Leuven, Belgium
wouter.joosen@cs.kuleuven.be

ABSTRACT

As recent incidents have shown, weak passwords are a severe security risk for authenticating users and granting access to protected resources. Additionally, strong passwords score low on usability, especially on mobile devices. In this work, we present SmartAuth, a scalable context-aware authentication framework built on top of OpenAM, a state-of-practice identity and access management suite. It uses adaptive and dynamic context fingerprinting based on Hoeffding trees to continuously ascertain whether a user's identity is authentic or not, and it respects privacy preferences by adopting consent-driven use of context information. We assess our approach from both an offensive and defensive security perspective. Our results show that dynamic context fingerprinting has good potential for a zero-interaction authentication scheme, with a minimal performance overhead compared to traditional authentication schemes.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection—*Information flow controls*; K.6.5 [Management of Computing and Information Systems]: Security and Protection—*Unauthorized access*; D.2.8 [Software Engineering]: Metrics—*complexity measures, performance measures*

Keywords

Authentication, fingerprinting, context, security, performance

1. INTRODUCTION

The main purpose of identity and access management (IAM) platforms is to address authentication, authorization, access and auditing as a common concern for online service providers. The added value that many state-of-practice identity management solutions have to offer is their capability of federated single sign on (SSO), simplifying access to partnering services with a single login. However, given

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SAC'15, April 13-17, 2015, Salamanca, Spain

Copyright 2015 ACM 978-1-4503-3196-8/15/04...\$15.00

<http://dx.doi.org/10.1145/2695664.2695908>

today's evolving threat landscape and recent attacks on on-line services, poor passwords have been shown over and over again to be a severe security threat [6]. Verizon's 2013 Data Breach Investigations Report [25] confirms that weak or default passwords, and stolen or reused credentials are still the main source of successful data breaches. Hence, traditional systems for identity and access management technologies that heavily rely on passwords no longer suffice. Indeed, the mainstream pin or username and password-based authentication combined with the long-lived user authenticated sessions do not offer the security guarantees for risk-sensitive services that require stronger continuous identity assurance. Furthermore, usernames and passwords are deemed inconvenient for mobile customers, and smartcard based authentication solutions – as in online banking applications – cause a skewed balance between security and usability, especially for online applications where a frictionless experience is paramount.

To address these concerns and trade-offs, we investigate reliable and less intrusive authentication techniques that can operate silently in the background based on additional context and behavioural information [24]. Continuous passive assessment of the login context through frictionless zero-interaction authentication [7] enables service providers to streamline access for trusted combinations of user accounts and consumer contexts. By leveraging information about device characteristics, time and location, etc. during authentication, we aim to leverage context to quantify the risk in zero-interaction user authentication, and reduce the risk of fraudulent activities through analysis of significant contextual deviations over time.

From a practical point of view, collecting context information for enhanced risk-based authentication raises several privacy concerns. To address this challenge, we specifically investigate context fingerprint-based authentication with minimal information disclosure, and techniques that require explicit consent on an individual basis about which context properties can be used for enhanced authentication.

SmartAuth, our approach for dynamic and adaptive context fingerprinting for continuous user authentication, uses Hoeffding trees [9] – a learning and mining technique for high-speed data streams – to continuously classify user identities and implement change detection. We use this technique to continuously assess the risk of fraudulent activities during long-lived user authenticated sessions. To address privacy concerns while collecting user context, our solution dynamically adapts to user consents. We share our expe-

rience with enhancing OpenAM¹ – a state-of-the-art open source federated identity and access management solution – for the implementation and validation of our work. The main contribution of this work is twofold:

- A user consent-driven probabilistic and adaptive authentication method to identify suspicious actions
- A scalable implementation on top of a state-of-practice identity and access management platform

After reviewing related work in section 2, we present our dynamic context fingerprinting method in section 3. The actual implementation on top of OpenAM is discussed in section 4. In section 5 we evaluate its performance, as well as the strengths and weaknesses of our implementation. We conclude in section 6 summarizing the main insights and identifying possible topics for future work.

2. RELATED WORK

Weak passwords [15] are a major cause of data and security breaches. SplashData reveals each year its annual *25 Worst Passwords of the Year* list². The top 3 most common passwords of 2013 were *123456*, *password*, and *12345678*, and these were also in the top 3 of the list of 2012 and 2011. SplashData compiles their top 25 based on files with millions of stolen passwords posted online by hackers. With dictionary attacks and optimized password cracking tools, users with simple or short (i.e. less than 8 characters) passwords are easy prey, especially if they use the same password for various services.

Efforts are ongoing to replace password-based authentication with better alternatives [1, 3, 11, 12]. With multi-factor authentication, users authenticate with a combination of authentication factors, i.e. *knowledge*, *biometrics*, and *possession*. A common example of two-factor authentication is accessing an ATM machine, where the credit card is something that you have, and the PIN code is something that you know. Knowledge-based factors include passwords, PIN codes and security questions. The advantage is that they do not need any equipment, but they are easily guessed, the user can be tricked into handing them to an attacker (i.e. phishing or social engineering), and complex passwords can be forgotten. Biometric factors like voice recognition, fingerprints or retina scans cannot be forgotten, but may require expensive equipment to implement. Possession factors like ownership of digital certificates, smart cards, USB tokens, or One Time Password generators are more resistant to guessing, but the hardware can be costly (e.g., in the case of tokens and smart cards), as can the process to implement it (e.g., in the case of certificate management. Additionally, possession factors can also be lost or stolen. Cost and ease-of-use, especially on mobiles, are two main reasons why such alternatives are still fighting an uphill battle to elevate authentication to the next level of protection.

Contemporary multifactor authentication schemes do augment security, but are often considered as cumbersome with a high impact on user experience. We need a passive assessment of user context through frictionless multi-factor authentication to reduce this inconvenience. Early works

¹<http://forgerock.com/products/open-identity-stack/openam/>

²<http://splashdata.com/press/worstpasswords2013.htm>

leveraged location- and proximity-based authentication [8, 16] to simplify and improve security. Nowadays, any context of the user [5, 13] is being considered for new types of online authentication and stronger continuous identity assurance. Bruce Schneier, an expert authority on security, commented on risk-based authentication as recently as November 2013 [23]:

I like this idea of giving each individual login attempt a risk score, based on the characteristics of the attempt: The risk score estimates the risk associated with a log-in attempt based on a user's typical log-in and usage profile, taking into account their device and geographic location, the system they're trying to access, the time of day they typically log in, their device's IP address, and even their typing speed ...

Browser fingerprinting schemes [10, 19] have been widely used to uniquely identify individuals for tracking purposes, but the same techniques can be applied for strengthening user authentication as well. In a Gartner 2013 report [14], Henry states that endpoint trustworthiness is especially an issue for bring-your-own-device (BYOD) scenarios. However, it is important to decide how much to trust the user and the contextual information that they are presenting. The quality of the used context information is key to objectively quantify the risk in zero-interaction authentication [7].

Manzoor et al. [18] argued on the importance of Quality of Context (QoC) for real-life applications, and presented QoC models to make effective use of context. However, current approaches often rely on ad hoc weighing functions to aggregate different types of context. There is no systematic approach to derive trust levels from the context and the quality of the context, nor to assess the required accuracy, precision and recency to reliably ascertain the risk in zero-interaction [7] or context-based user authentication [13].

Privacy and context are two intimately related [20] and often conflicting concerns. Chabridon et al. [4] survey existing works on the notions of privacy and QoC, and confirm that current solutions usually consider only one notion, and very few of them started to bridge privacy and QoC. Indeed, contemporary risk-based authentication systems have no support for end-user capabilities to provide explicit consent about which personal or transactional context is collected and used, nor are they able to quantify the risk in a personalized manner.

From a state-of-practice point of view, contemporary risk-based authentication systems rely on blacklists of IP addresses of malicious hosts (e.g. botnets), or leverage limited forms of geolocation (timezones or IPv4 databases to link addresses with locations) to restrict access. For mobile devices in a local region such databases offer subpar location quality and accuracy. Also, this technique will become far more complex when IPv6 is widely adopted [17] to get an accuracy comparable to these IPv4 databases.

3. CONTINUOUS IDENTITY ASSURANCE WITH CONTEXT FINGERPRINTS

The authentication mechanisms discussed so far are static, i.e. a user can always log in with the same credentials, independent of context. Adaptive authentication, on the other hand, incorporates a risk model per type of interaction with the system or service. It acknowledges that some situations are inherently more risky than others, and therefore require stronger guarantees. As an example, a user checking the

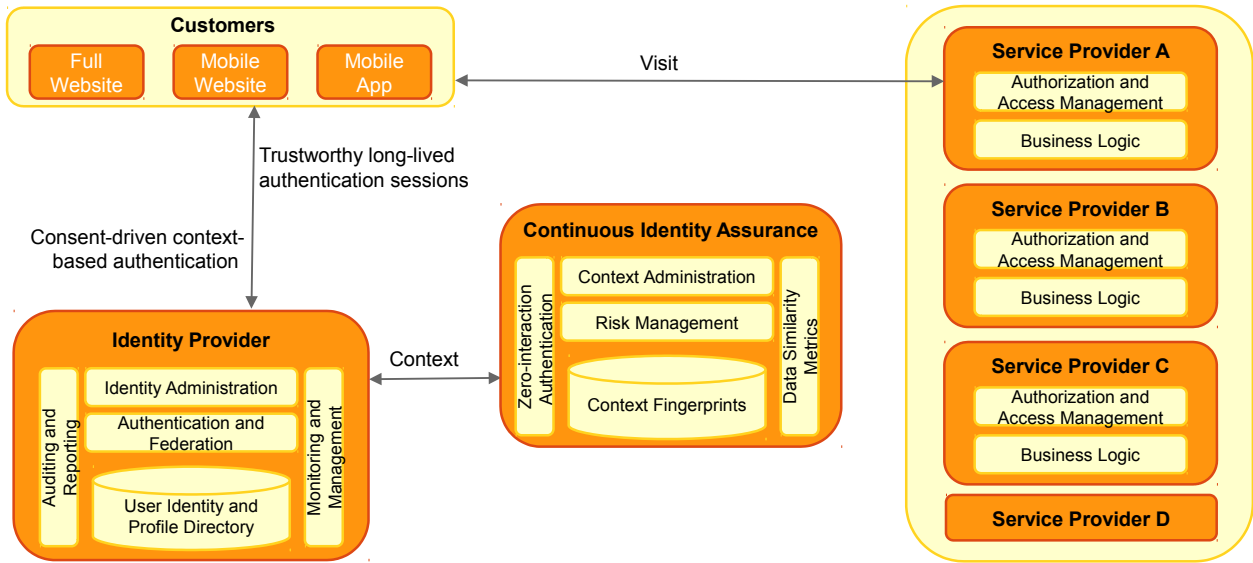


Figure 1: Context-based identity assurance in a single sign on ecosystem of identity and service providers

news from his trusted mobile device every morning in the vicinity of Brussels can be considered normal, while clearly something is odd when the same user makes a purchase from Russia, ten minutes later. For mobile applications, this risk model can be based on capabilities like location awareness (through WiFi, cell-tower triangulation, GPS) and mobile device identification (device model, language, and screen size). Anomalies such as locations or devices which are new to the user are deemed high risk. For payment transactions, the risk model can analyze time and day of access, typical transaction amounts, and frequency of payment transactions.

By continuously scrutinizing for deviations from normal interaction patterns, as depicted in Figure 1, the system can avoid interrupting the user with more complex authentication schemes up to a point where the system begins to doubt the person’s identity. In the following subsections, we will discuss both server-side and client-side fingerprints, and how we tackle their dynamic behavior at runtime.

3.1 Security requirements for server- and client-side context fingerprints

The context fingerprint based solution is subject to several security requirements, which we will list below:

1. Ensure that context fingerprints cannot be compromised
2. Prevent replay attacks of context fingerprints
3. Support revocability of context fingerprints
4. Fingerprints should have strong similarity checks

Our identity and access management solution offers adaptive consent-driven intelligence to protect against risk-based threats. It assesses the risk during authentication based on a configurable score system, and determines whether to require the user to complete other authentication steps (such as HOTP, requiring the user to authenticate with a one-time

password delivered by email or by SMS) when the sum of the scores exceeds a certain threshold. The scores are based on both server- and client-side context fingerprints, including historic IP addresses and ranges, geolocation, time since last login, device, software and transaction fingerprints.

Contrary to existing systems, our solution offers 2 enhancements to preserve privacy on sharing sensitive context information: (1) for each context-attribute used for risk-assessment, the user has to provide *consent*, and (2) client-side components apply an *efficient similarity preserving hash function* on sensitive context-information rather than sending and storing the context information in the clear.

3.2 Dynamic context fingerprints

As no particular fingerprinting algorithm is universally better than all the others, combining multiple models generally offers some kind of performance improvement. We therefore use a dynamic fingerprinting algorithm that combines many fingerprints into one that is more accurate than the individual ones.

Furthermore, we also take user consent into consideration by dynamically adapting and optimizing the weight of multiple context fingerprints. The proposed approach can be formulated as follows:

$$d = \sum_{i=1}^C w_i(x) \times h(x, D) \quad (1)$$

C is the number of context types that can be used for fingerprinting. x is the context value, and h computes a hash value of this context value and compares it with entries in its database D . The hash value has two purposes, i.e. (1) it reduces the amount of data to be stored and communicated, and (2) it obfuscates sensitive information when sent across the network. The weights are first initialized according to user consents, i.e. $w_i(x)$ is 0 if the user has not provided consent to collect and use context parameter i .

We use Hoeffding trees [9], a mining technique for high-speed data streams, to continuously test and classify adap-

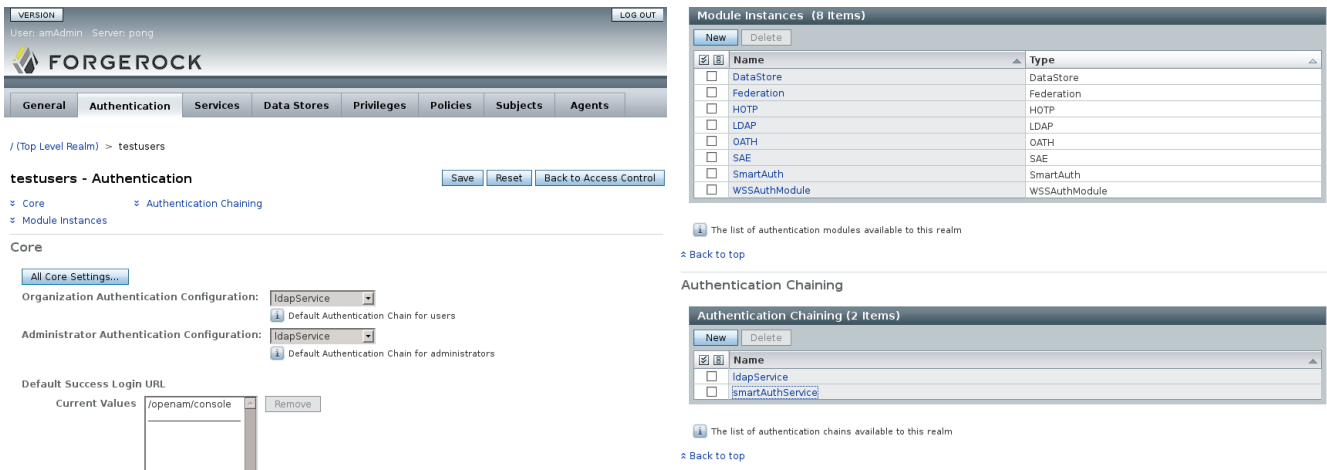


Figure 2: SmartAuth: Dynamic context fingerprinting as an authentication plugin in ForgeRock’s OpenAM

tation rules and actions. The added value of using Hoeffding trees is manifold:

- They operate in a limited amount of memory and time.
- They are ready to predict and classify at any time.
- They support an interleaved test-then-train setting.

During the training phase, our algorithm collects all fingerprints, and classifies the fingerprint combinations as a positive or negative training instance depending on whether they match with the actual authentication of the individual.

During the testing phase, the Hoeffding tree is used to ascertain whether a given fingerprint will be effective or not given the possible fingerprints the user granted to be collected. Our solution builds upon Hoeffding trees to make sure that fingerprints with a negative utility are never used, and to be able to adapt to concept drift to handle dynamically evolving fingerprints (e.g. current location or browser version). Hoeffding trees are driven by the Hoeffding bound ϵ that decides how many instances are needed to achieve a certain level of confidence on the best attribute to split the tree.

$$\epsilon = \sqrt{\frac{R^2 \ln(1/\delta)}{2n}} \quad (2)$$

The Hoeffding bound ϵ defines that with probability $1-\delta$, the true mean of the real-valued variable r with range R is at least $\bar{r} - \epsilon$, with n being the number of independent observations of r . For example, assuming $R = 1.0$, then for $\delta = 0.95$ and $n = 10$ samples, ϵ evaluates to 0.05.

For our work, this means that one context fingerprint is superior compared to other fingerprints when the difference of information gain is greater than ϵ . The real-valued variable r in our solution is based on the distance function of the similarity preserving hash function.

4. IMPLEMENTATION

This section covers the implementation of our approach towards dynamic context fingerprinting for continuous user authentication.

4.1 Dynamic hash-based context fingerprints

Our current solution combines a variety of fingerprints, collected from both the user device as well as parameters collected at server side:

- **Client side:** language, color depth, screen resolution, timezone, platform, plugins, etc.
- **Server side:** IP address range, time of access, geolocation, request headers, etc.

Fingerprints that are collected at client side are hashed so that these are not sent through the network in the clear. We apply similarity preserving hash functions like `sdfhash` [22] and `t1sh` [21]. The key benefit is that such hash functions ensure that similar inputs yield similar hash values. The Hoeffding decision tree implementation is provided by the Massive Online Analysis (MOA) framework [2].

4.2 Extending OpenAM with a new authentication module

We implemented the above solution on top of ForgeRock’s OpenAM version 11.0 by integrating our solution as a new authentication plugin (see Figure 2), called *SmartAuth*. OpenAM offers the added value that authentication plugins can be chained, so that a risk-based assessment of the context fingerprints can fall back on a stronger authentication method if need be. This is the case when the collected fingerprints are not recognized or offer insufficient distinguishing features. Depending upon which fingerprints are observed frequently, the Hoeffding classification tree adapts by updating its optimal splitting attributes.

The context fingerprint also embeds additional fields to prevent phishing or replay attacks:

- **Token:** This identifier represents the authenticated session after initial login. This token can be revoked by logging out.
- **Counter:** With each submission of the fingerprint within a session, this counter is incremented with 1. This value is compared with the value of the previous fingerprint stored at server side.

- **Timestamp:** This field is initialized with a random value at initial login, and incremented with 1 every second during the lifetime of the authenticated session.
- **Random number:** A random number is added to increase the entropy of the context fingerprint.
- **Checksum number:** This number is a CRC-32 checksum of all the fingerprints and the previous fields.

The fingerprint is sent encrypted from the client to the SmartAuth framework. The receiving party decrypts the message, verifies the checksum, the validity of the authentication token, whether the counter and timecounter is greater than the last stored value. The timestamp is additionally used to check whether the timestamp and the time elapsed since the previously submitted fingerprint (of the same authenticated session) corresponds with the timestamp of the current fingerprint. If any of these checks fails, the fingerprint is rejected and a fallback to a stronger authentication method is initiated.

Revisiting the security requirements of section 3.1, we aim to prevent compromising the fingerprints (1) using public key encryption, and prevent replay attacks (2) by adding counters and timestamps to the fingerprints. The SmartAuth authentication framework stores the last received and valid context fingerprint. These fingerprints are revoked (3) with every new submission of the fingerprint, and by logging out hereby invalidating the authentication token.

For each individual, our framework keeps track of the 1024 fingerprints. When the user logs in again, the authentication token will be different, the counter will be reset, and the timestamp will be reinitialized. However, the actual fingerprints are compared with those previously stored. To compare the similarity of the fingerprints (4), the corresponding hash function should produce exactly the same value or be within certain limits when using similarity preserving hash functions.

5. EVALUATION

Contrary to browser fingerprinting techniques like Panoptick³ that can uniquely identify a browser in more than 1 million entries, our approach allows individuals to selectively enable and disable context fingerprints. As a result, our approach will not achieve the same uniqueness rates if fewer fingerprints are used. However, the Panoptick method relies on static attributes, whereas our solution supports dynamic fingerprinting and can recognize situations where an individual appears to be at 2 different locations at the same time by leveraging the similarity preserving hash functions.

To evaluate our methodology, we collected 2000 different system and connectivity configurations based on 10 different types of context fingerprints. These are based on known and valid combinations of browser user agents, software versions and device models using the WURFL mobile device description database⁴.

In a second stage, we replayed these fingerprints but simulated time and location variations, and applied similarity based hash functions on the latter two attributes. The results are shown in Table 1.

³<https://panopticklick.eff.org/>

⁴<http://wurfl.sourceforge.net/>

# Fingerprints	Correct	Incorrect
2	63%	37%
4	79%	21%
6	85%	15%
8	91%	9%
10	97%	3%

Table 1: Correct and incorrect classification of individuals based on their context fingerprints.

In a second experiment, we made the context fingerprinting dynamic. Both the sending and receiving side know which fingerprints to exchange in a particular context. The fact whether or not context fingerprints are being communicated, not only depends on whether the user has given consent to compute them, but also on other constraints and preferences the user may wish to impose. The simple examples below illustrate the adaptive fingerprint is made time and/or location dependent.

```

1 if (location ≈ 'work') {
2   sendFingerPrint(hash(location))
3   sendFingerPrint(hash(screensize))
4 }
5
6 if (time ≈ 'morning') {
7   sendFingerPrint(hash(time))
8   sendFingerPrint(hash(plugins))
9 }
10
11 ...

```

Note that the actual values used in the conditions have to be sent as part of the fingerprints too. When sending the similarity preserving hash of the location, the receiving party can recognize *work* as the current location, and can infer that the screensize must be submitted too. If the last one is missing, or other hashes have been sent, we can identify a mismatch for this particular individual. In our current implementation, we manually configured a fixed scheme of which fingerprints to send under which conditions, but we hope in the future to revise this approach with a negotiation protocol between the client and server to improve the usability of this dynamic fingerprinting.

By making the exchange of context fingerprints dynamic, we add additional differentiating features. We generated for 10000 users a collection of random context fingerprint exchange scenarios. The classification accuracy is shown below in Table 2:

# Fingerprints	Correct	Incorrect
2	87%	13%
4	91%	9%
6	95%	5%
8	98%	2%
10	99%	1%

Table 2: Correct and incorrect classification of individuals based on dynamic context fingerprints.

In this adaptive context fingerprinting scenario, we see a significant improvement in classification accuracy. However,

as the complexity increases, we also measured the performance of dynamic context fingerprint on OpenAM in comparison with a simple login/password based authentication scheme.

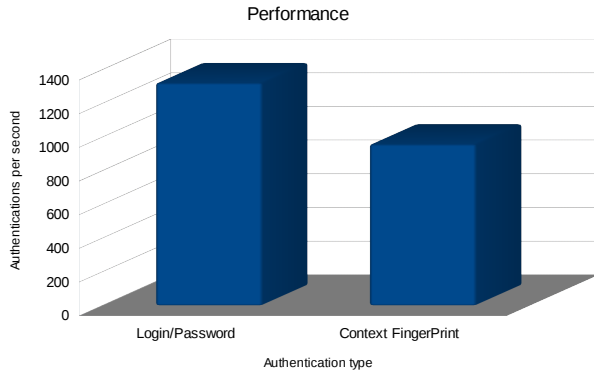


Figure 3: Performance of fingerprinting-based authentication

The results in Figure 3 show that there is a performance penalty of using dynamic context fingerprints for continuous user authentication. Most of the hashes of the context fingerprints are computed at client side, and as such do not cause any performance impact on the server. However, our solution implements continuous authentication where these fingerprints are repeatedly compared. With a simple login/password authentication scheme, only the authentication token received after a first login needs to be verified. This validation is a very lightweight operation. Additionally, building upon capabilities of OpenAM and OpenDJ (i.e. the LDAP backend of OpenAM), we can deploy multiple instances of our solution to achieve horizontal scalability, with replication between the LDAP instances to guarantee failover in case an OpenAM instance goes down, as depicted in Figure 4 below.

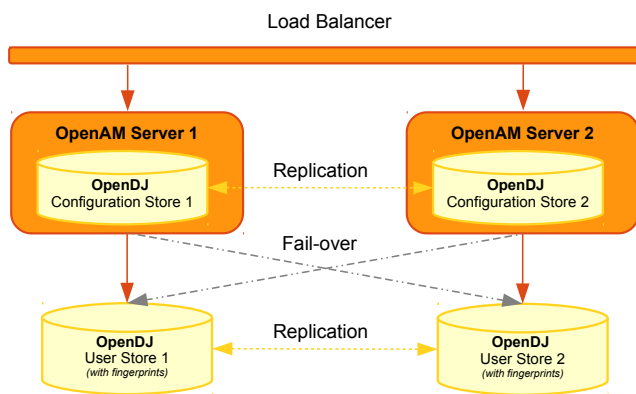


Figure 4: OpenAM scalability with fail-over support

We also carried out a user study with 6 individuals and 9 concrete devices (smartphones and tablets), mainly to test the SmartAuth framework from a usability point of view and to compare it with OpenAM’s built-in Device-Print Adaptive Authentication module. Similar to our approach, this

module detects the type of device requesting the authentication. However, it is up to the administrator (and not the user) to enable which features should be part of the device fingerprint. If there is a mismatch with one of the attributes (user agent, installed fonts, browser plugins, screen resolution or depth, timezone and geolocation, etc.), a pre-configured number of penalty points (per attribute) is added to a global score, and if the total penalty points exceeds a configured threshold, the user may be asked to verify his identity with stronger authentication schemes.

When the users were confronted with the facts that (1) their fingerprints were sent in full (and not hashed), that (2) users were not able to opt-in or opt-out of which fingerprints were collected, and that (3) the weighted scoring function is identical for each individual rather than personalized, and (4) does not evolve with their login behavior, the test subjects understood at least at a conceptual level the added value of the SmartAuth method. However, the test subjects in this study had mostly a technical background and the number of participants in this preliminary usability assessment was too small to produce results that are statistically meaningful. Nonetheless, context fingerprints offer a better user experience compared to traditional logins and passwords or even multi-factor authentication schemes. However, similar to biometrics, context fingerprints are not secrets, and as such they may be easily accessible to hackers. Additionally, some fingerprints are more tightly linked to the identity than others. Therefore, we recommend to use context-based authentication as a first level of authentication in a multi-layered risk-based approach, where high risks require a greater degree of certainty that the identity is indeed the one user’s claim to be.

6. CONCLUSION

The key contribution of this work is SmartAuth, a flexible and non-intrusive authentication scheme that it offers increased and adaptive security with support for transparent authentication, with the ability to authenticate the user periodically throughout the day in order to maintain confidence in the identity of the user. This authentication scheme leverages context fingerprints as a key enabler for long-lived authenticated users and risk-sensitive services.

From a user perspective, SmartAuth delivers ease-of-use through context-based zero-interaction authentication with support for end user consent about which contextual information attributes are gathered and processed. The main advantage of contextualizing user authentication and the ability to detect deviations from normal login behaviour is that it can be almost impregnable by malicious users who attempt to get personal details through various intrusion techniques or by plain theft of a mobile device.

In this work, we demonstrated how such an adaptive authentication system can be inception on top of OpenAM, a state-of-practice federated single sign-on solution. After extending this identity and access management system with context-based enabling services and concepts, a performance assessment was carried out to gain insights on the impact of these incremental security enhancements onto the identity management system. Our results show that dynamic context fingerprinting has good potential for a zero-interaction authentication scheme, with a minimal performance overhead compared to traditional authentication schemes.

7. REFERENCES

- [1] A. Bhargav-Spantzel, A. Squicciarini, and E. Bertino. Privacy preserving multi-factor authentication with biometrics. In *Proceedings of the Second ACM Workshop on Digital Identity Management, DIM '06*, pages 63–72, New York, NY, USA, 2006. ACM.
- [2] A. Bifet, J. Read, B. Pfahringer, G. Holmes, and I. Zliobaite. Cd-moa: Change detection framework for massive online analysis. In A. Tucker, F. Höppner, A. Siebes, and S. Swift, editors, *IDA*, volume 8207 of *Lecture Notes in Computer Science*, pages 92–103. Springer, 2013.
- [3] J. Bonneau, C. Herley, P. C. v. Oorschot, and F. Stajano. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy, SP '12*, pages 553–567, Washington, DC, USA, 2012. IEEE Computer Society.
- [4] S. Chabridon, R. Laborde, T. Desprats, A. Oglaza, P. Marie, and S. M. Marquez. A survey on addressing privacy together with quality of context for context management in the internet of things. *Annales des Télécommunications*, 69(1-2):47–62, 2014.
- [5] B. Chen, L. H. Nguyen, and A. W. Roscoe. When context is better than identity: Authentication by context using empirical channels. In B. Christianson, B. Crispo, J. A. Malcolm, and F. Stajano, editors, *Security Protocols XIX - 19th International Workshop, Cambridge, UK, March 28-30, 2011, Revised Selected Papers*, volume 7114 of *Lecture Notes in Computer Science*, pages 115–125. Springer, 2011.
- [6] W. Cheswick. Rethinking passwords. *Commun. ACM*, 56(2):40–44, Feb. 2013.
- [7] M. D. Corner and B. D. Noble. Zero-interaction Authentication. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking, MobiCom '02*, pages 1–11, New York, NY, USA, 2002. ACM.
- [8] D. E. Denning and P. F. MacDoran. Internet besieged. chapter Location-based Authentication: Grounding Cyberspace for Better Security, pages 167–174. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1998.
- [9] P. Domingos and G. Hulten. Mining high-speed data streams. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '00*, pages 71–80, New York, NY, USA, 2000. ACM.
- [10] P. Eckersley. How unique is your web browser? In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies, PETS'10*, pages 1–18, Berlin, Heidelberg, 2010. Springer-Verlag.
- [11] E. Grosse and M. Upadhyay. Authentication at scale. *IEEE Security and Privacy*, 11:15–22, 2013.
- [12] R. P. Guidorizzi. Security: Active authentication. *IT Professional*, 15(4):4–7, 2013.
- [13] E. Hayashi, S. Das, S. Amini, J. Hong, and I. Oakley. CASA: Context-aware Scalable Authentication. In *Proceedings of the Ninth Symposium on Usable Privacy and Security, SOUPS '13*, pages 3:1–3:10, New York, NY, USA, 2013. ACM.
- [14] T. Henry. Adaptive Access Control Brings Together Identity, Risk and Context. Gartner Report. <https://www.gartner.com/doc/2578515/adaptive-access-control-brings-identity>, Aug 2013.
- [15] M. Jakobsson and M. Dhiman. The benefits of understanding passwords. In *Proceedings of the 7th USENIX conference on Hot Topics in Security, HotSec'12*, page 10, Berkeley, CA, USA, 2012. USENIX Association.
- [16] W. Jensen, S. Gavrila, and V. Korolev. Proximity-Based Authentication for Mobile Devices. In *Proceedings of The 2005 International Conference on Security and Management*, pages 398–404, June 2005.
- [17] R. Koch, M. Golling, and G. D. Rodosek. Geolocation and Verification of IP-Addresses with Specific Focus on IPv6. In G. Wang, I. Ray, D. Feng, and M. Rajarajan, editors, *CSS*, volume 8300 of *Lecture Notes in Computer Science*, pages 151–170. Springer, 2013.
- [18] A. Manzoor, H. L. Truong, and S. Dustdar. Quality of context: models and applications for context-aware systems in pervasive environments. *Knowledge Eng. Review*, 29(2):154–170, 2014.
- [19] N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy, SP '13*, pages 541–555, Washington, DC, USA, 2013. IEEE Computer Society.
- [20] H. Nissenbaum. A Contextual Approach to Privacy Online. *Daedalus*, 140(4):32–48, 2011.
- [21] J. Oliver, C. Cheng, and Y. Chen. Tlsh – a locality sensitive hash. In *Proceedings of the 2013 Fourth Cybercrime and Trustworthy Computing Workshop, CTC '13*, pages 7–13, Washington, DC, USA, 2013. IEEE Computer Society.
- [22] V. Roussev. Data fingerprinting with similarity digests. In *Advances in Digital Forensics VI - Sixth IFIP WG 11.9 International Conference on Digital Forensics, Hong Kong, China, January 4-6, 2010, Revised Selected Papers*, pages 207–226, 2010.
- [23] B. Schneier. Schneier on Security. Risk-Based Authentication. https://www.schneier.com/blog/archives/2013/11/risk-based_auth.html, Nov. 2013.
- [24] E. Shi, Y. Niu, M. Jakobsson, and R. Chow. Implicit authentication through learning user behavior. *Information Security*, pages 99–113, 2011.
- [25] Verizon. 2013 Data Breach Investigations Report. http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2013_en_xg.pdf, 2013.