

Risk-Based Authentication for OpenStack: A Fully Functional Implementation and Guiding Example

Data/Toolset paper

Vincent Unsel
H-BRS University of Applied Sciences
Sankt Augustin, Germany
vincent.unsel@h-brs.de

Nils Gruschka
University of Oslo
Oslo, Norway
nilsgrus@ifi.uio.no

Stephan Wiefling
Ruhr University Bochum
Bochum, Germany
stephan.wiefling@rub.de

Luigi Lo Iacono
H-BRS University of Applied Sciences
Sankt Augustin, Germany
luigi.lo_iacono@h-brs.de

ABSTRACT

Online services have difficulties to replace passwords with more secure user authentication mechanisms, such as Two-Factor Authentication (2FA). This is partly due to the fact that users tend to reject such mechanisms in use cases outside of online banking. Relying on password authentication alone, however, is not an option in light of recent attack patterns such as credential stuffing. Risk-Based Authentication (RBA) can serve as an interim solution to increase password-based account security until better methods are in place. Unfortunately, RBA is currently used by only a few major online services, even though it is recommended by various standards and has been shown to be effective in scientific studies. This paper contributes to the hypothesis that the low adoption of RBA in practice can be due to the complexity of implementing it. We provide an RBA implementation for the open source cloud management software OpenStack, which is the first fully functional open source RBA implementation based on the Freeman et al. algorithm, along with initial reference tests that can serve as a guiding example and blueprint for developers.

CCS CONCEPTS

• **Security and privacy** → **Authentication; Access control; Web application security; Usability in security and privacy; Intrusion detection systems**; • **Computer systems organization** → *Cloud computing*; • **Software and its engineering** → *Designing software*.

KEYWORDS

Risk-Based Authentication, Implementation Challenges, OpenStack

ACM Reference Format:

Vincent Unsel, Stephan Wiefling, Nils Gruschka, and Luigi Lo Iacono. 2023. Risk-Based Authentication for OpenStack: A Fully Functional Implementation and Guiding Example: Data/Toolset paper. In *Proceedings of the Thirteenth ACM Conference on Data and Application Security and Privacy (CODASPY '23)*, April 24–26, 2023, Charlotte, NC, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3577923.3583634>

1 INTRODUCTION

Passwords are still the predominant authentication method for most online services [27], despite their long-known weaknesses [21] and the continuous emergence of new attacks such as credential stuffing and password spraying [1]. To effectively protect their users, online services must use alternative or additional measures to passwords. Other authentication factors using special user-owned devices or physical biometrics, are generally impractical for online services, as they require additional hardware and active user enrollment. This is why they are generally not used in practice [10]. The composition of two different user authentication factors, usually a password combined with something the user possesses or is, suffers from the same acceptance problems. Apart from few applications such as online banking [8, 28, 33], Two-Factor Authentication (2FA) [25] is not yet widely accepted by users in practice either [19, 24, 31].

Risk-Based Authentication (RBA) [36] is an online-service-side complement to authentication systems such as password-based authentication that does not require direct user interaction. The user performs the log-in process simply by entering their login credentials (i.e., username and password). In many cases this is sufficient to authenticate to the service. Only when the online service's RBA component detects a deviation from the usual log-in behaviour (e.g., different user location), a further authentication factor is requested [33]. This is also true when a correct username-password combination is provided, i.e., an attacker used leaked credentials. Therefore, to keep it with the security principle of “good security now” [11], RBA can be used as an immediate additional security measure for password-protected user accounts in online services. It enhances the security of such online accounts right away until alternative and more secure authentication methods become established in the mainstream.

CODASPY '23, April 24–26, 2023, Charlotte, NC, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the Thirteenth ACM Conference on Data and Application Security and Privacy (CODASPY '23)*, April 24–26, 2023, Charlotte, NC, USA, <https://doi.org/10.1145/3577923.3583634>.

Research Hypothesis. Although the use of RBA is recommended in the literature [33–35] and in national policies [4, 5, 13, 23], the adoption of RBA in practice is still very limited [12, 33]. Recent research found that 78% of 235 popular online services inside the Tranco 5K [17] still do not use any form of RBA [12]. As Internet users have an average of 92 to 130 online accounts [16], services of the major Internet companies protect only a fraction of these accounts with RBA. One reason for this low RBA adoption rate could be a lack of guidance on the implementation of RBA. Insights from the literature show that effective RBA implementations can be very complex [6, 9, 35]. Therefore, we expect that many developers cannot estimate the challenges they face when implementing RBA in their online services. Furthermore, there are almost no measures available that could help developers with the implementation or testing of a custom RBA implementation.

Contributions. To close this gap, we provide the following: (i) We introduce a conceptual model of RBA that provides a generic view on how to implement and integrate RBA in online services. (ii) We then instantiate the model as a fully functional open source RBA plug-in for the cloud computing software OpenStack. To the best of our knowledge, this is the first open RBA implementation based on the algorithm of Freeman et al. [9]. (iii) Finally, we provide a reference test and reference values based on real world login data that can be used to test RBA implementations based on the Freeman et al. algorithm.

Overall, our work aims to help developers, administrators, and service owners to strengthen password-based authentication by providing guidance on how to implement, integrate, and test RBA for their online services. This should make it easier to put RBA into practice. Our code repository should also serve as a research environment to study how such example implementations help developers to bring security measures into software products.

2 RISK-BASED AUTHENTICATION (RBA)

Figure 1 shows the common architecture of an RBA system [9, 34, 36], typically used in addition to password-based authentication. The system logs contextual features (e.g., network and browser information) for each successful login attempt and compares them to the previously observed feature values. In the case of RBA-enhanced password authentication, this means that access to the service is only granted if the login context is not too different from the previously observed ones. This comparison of the current login context with a recorded history of login contexts is done by calculating a so-called risk score. The score is a number indicating the deviation from the expected values, i.e., the higher the number, the higher the risk of account takeover. Only if this risk score is below a certain threshold, the whole login process behaves like regular password-based authentication from the user’s perspective. If the threshold is exceeded, the RBA-based authentication system requires a re-authentication step, i.e., requesting a second authentication factor. Typically, this factor is a verification code sent to a registered email address or phone number [12, 36]. Only in case this additional factor is successfully verified, the user is authenticated to the service. In case of a very high risk score, e.g., when we are very sure that it is an automated attack, the online service can block access. Although

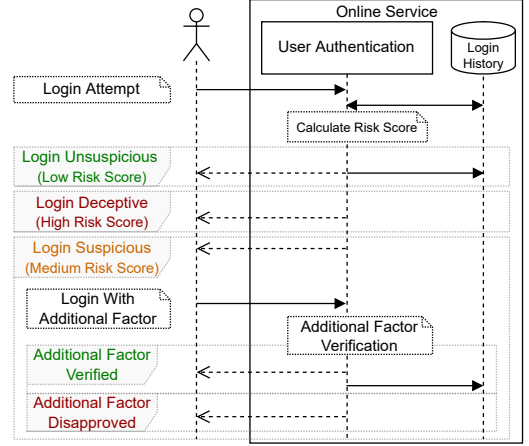


Figure 1: Overview of an RBA system, showing the communication between user and online service in different scenarios (low, medium, and high risk)

this use case is frequently mentioned in literature, it seems to be uncommon in practice [12, 36].

The contextual features used in the risk score calculation can be defined by the person configuring the RBA system, and can vary from network (e.g., IP address or RTT) and device (e.g., user agent string), to behavioral biometric information (e.g., login time). Related research suggests that network-based (IP address) and device information (user agent string) are sensible features providing high security with high usability in practice [9, 34, 36]. When used in combination, these features protected against sophisticated attackers that know the users’ passwords, locations, and devices, without legitimate users noticing any differences in the online service’s behavior [34]. Therefore, we focused on these features in our OpenStack implementation.

To the best of our knowledge, Freeman et al. [9] published the first RBA algorithm for computing a risk score for a login attempt. The algorithm was also suspected to be used in some form at the popular online services Google, Amazon, and LinkedIn [34]. As it also showed good performance in a practical evaluation and outperformed an algorithm used in the open source single sign-on solution OpenAM [34], we selected this algorithm for our implementation. The algorithm calculates the risk score S for a user u and a given feature set $FV = (FV^1, \dots, FV^d)$ with d features as:

$$S_u(FV) = \left(\prod_{k=1}^d p(attack|FV^k) \frac{p(FV^k)}{p(FV^k|u, legit)} \right) \frac{p(u|attack)}{p(u|legit)} \quad (1)$$

In the formula, we use the probabilities $p(FV^k)$ for the feature value appearing in the global login history of all users, and $p(FV^k|u, legit)$ for the feature value being used by the legitimate user trying to sign in. For some features, comparison of occurrence rarely result in an exact match. For example, nowadays, IP addresses are typically assigned dynamically and devices (especially mobile devices) are roaming in different networks. In these cases, when using a unseen feature, $p(FV^k|u, legit)$ would be zero, resulting to an undefined value. To compensate for such frequently

changing feature values, sub-feature derivation techniques can be applied to some features as so-called smoothing [9]. Therefore, the IP address feature can be split at three granular levels. Including the Autonomous System Number (ASN), to which the address belongs, smooths the unseen feature value to the granularity level of the Internet Service Provider (ISP)'s ASN range [14]. The third and least granular level is the corresponding country to even compensate frequent rotations of different Access Points (APs) [15]. Likewise for smoothing the User-Agent (UA) string, sub-features are extracted to include browser, operating system and device information into the risk score. These derived sub-features are aggregated in the probability calculation using a linear interpolation [9], in which the history is divided into all entries containing each sub-feature value to estimate the originated feature value. Furthermore, $p(u|attack)$ estimates how likely a user is being attacked and $p(u|legit)$ how likely the user is signing in on this online service, i.e., $p(u|legit) = \frac{\text{Number of user logins}}{\text{Number of all logins}}$. Some online service might also have additional attack data, e.g., a list of IP addresses that were previously used in attacks. In this case, $p(attack|FV^k)$ estimates how likely the feature value has been previously seen in attacks. Related work indicates, however, that attack data should be used with caution or not at all, as it also can negatively influence the RBA system's security and usability properties [35]. When not using attack data, this term can be neglected, i.e., $p(attack|FV^k) = 1$ [34, 35, 38].

3 RELATED WORK

The Freeman et al. paper [9] can be considered the academic birth of RBA. Follow-up academic research provides scientific evidence on useful features [2, 3, 34, 36], user perspective [7, 18, 33, 37], suitable RBA algorithms [34, 38], and a large-scale evaluation of effective RBA protection based on login data of a real online service [35]. To the best of our knowledge, there is no research addressing the reasons for the low adoption of RBA in practice or supporting measures for software developers. When analyzing the repositories of relevant open source online software, and identity and access management (IAM) solutions, we found that there are virtually no projects among them that contain RBA capabilities. The IAM software OpenAM¹ is one notable exception. Still, its RBA implementation is very limited in terms of monitored features and risk score calculation. The feature set only includes the IP address, and the current IP address is only compared with those already included in the login history. This approach does not meet the security or usability goals that RBA can otherwise provide [34–36].

Achieving a functional RBA implementation is a very complex task, as it requires a variety of skills and knowledge. To even complicate this, the available scientific literature offers little support for implementation. For instance, the paper by Freeman et al. presents the theory and the algorithm, but many implementation-relevant aspects are left unconsidered and the authors themselves do not provide a reference implementation. The lack of a reference implementation of the Freeman et al. algorithm has since been resolved [32], but this is still insufficient as a developer support. The availability of test cases with corresponding test data is often another necessary prerequisite. Therefore, to facilitate the adoption of RBA, we provide a fully functional implementation of RBA in a

¹<https://www.openidentityplatform.org/>

relevant open source software project. The implementation puts all the necessary pieces together to provide a complete guiding example, including testing capabilities.

4 OPENSTACK

OpenStack² is an open source software suite for building cloud computing platforms. It offers different Infrastructure-as-a-Service (IaaS) cloud services, primarily virtual machines (VM) and storage services. Commercial cloud service providers can use OpenStack to build public cloud platforms and organisations can build private clouds with it.

The OpenStack software is designed as a modular framework composed of different application services. As a central component for IaaS systems, the *Nova* compute service manages and hosts virtual instances. Users can administrate *Nova* through the *Horizon* web dashboard. In addition to *Nova*, a minimal configuration of OpenStack contains the following modules: the image service *Glance* to discover, register, and retrieve VM images; the *Placement* service to enable other services track their own resources; the Software Defined Network (SDN) component *Neutron* to create and attach virtual network infrastructure devices to instances from other modules; and the *Keystone* identity service to enable user authentication and access management.

Relevant service components for integrating RBA functionality are *Horizon* and *Keystone*. The *Horizon* web dashboard allows administrators and users to access and manage cloud computing resources. Authentication requests are delegated to the corresponding identity management service. *Horizon* offers two possibilities to render the received authentication result in the frontend: When the login credentials were incorrect or the user has to provide an additional additional factor, it can display this information as an error, i.e., a red box in the login form. When access was granted, *Horizon* redirects to the user's dashboard.

Keystone is the access control component for the framework's services. It can identify users and verify their authorization on managed assets. On successful user authentication, *Keystone* issues a unique login session identifier (session token) to the user. *Keystone* then constantly verifies whether this token is valid and whether user actions are within their permitted scope.

We chose OpenStack as the example software for providing an RBA implementation. We made this decision because many clouds are built on OpenStack [20]. They can therefore immediately benefit from increased account security by installing and using RBA. In addition, OpenStack has a modern microservices-based architecture and built-in extensibility. For this reason, RBA-related source codes are highly decoupled from the other components and can be read and understood with limited knowledge of the overall system. With Python as the underlying programming language, which is also the most popular programming language [29], the RBA source code should be accessible and comprehensible to most software developers.

5 OPENSTACK RBA EXTENSION

To enable RBA in OpenStack, we must extend the *Horizon* frontend component with RBA-specific login user interfaces and implement

²<https://www.openstack.org/>

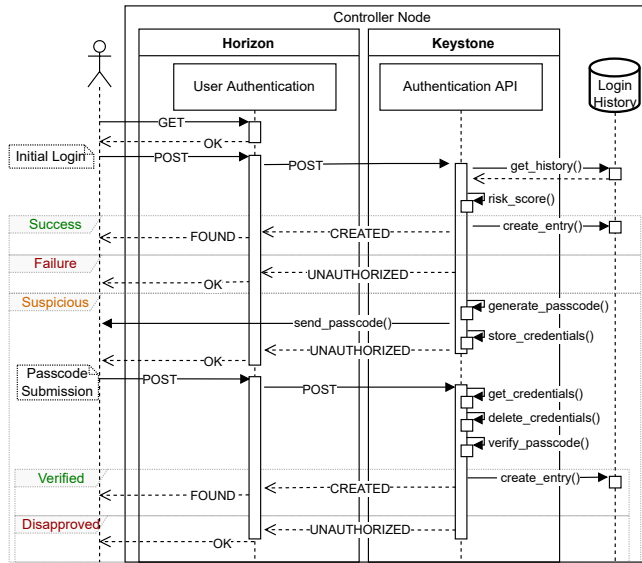


Figure 2: Authentication Flow in OpenStack using the RBA Extension

the risk score calculation in the Keystone backend component. We leveraged OpenStack’s modular microservice architecture and extension interfaces and developed two extension plug-ins for these two components as follows. They work at least with the Wallaby release up to the actual released stable series Zed.

5.1 Horizon RBA Extension (Frontend)

The Horizon RBA plug-in³ extends the default password-based authentication by unobtrusively collecting the client’s RBA features at the login screen. In the current version, the IP address, user agent string, and the Round-Trip-Time (RTT) features are supported by the extension. We chose these features as previous research identified them as the most effective ones to identify users [34, 35]. The first two features are obtained directly from the client Hypertext Transfer Protocol (HTTP) request. The last one required an additional Horizon extension that uses the asynchronous WebSockets protocol for measuring the transmission time of messages from the Horizon service to the client and back. The RTT is a strong and privacy-preserving indicator to distinguish legitimate login attempts from those spoofing the “correct” location using a Virtual Private Network (VPN) with an egress point in the same region as the legitimate user [34, 38]. To compensate the network delay jitter, these round-trip measurements are performed five times, and the shortest time measurement is selected as the RTT feature in RBA authentication [34].

Figure 2 shows the RBA-enabled authentication flow in OpenStack. When the user performs a login attempt, Horizon forwards the RBA features together with the password to the authentication API of Keystone. The RBA evaluation (see Figure 1) can have three different results where the Horizon frontend initiates different actions. These are (i, success) the user is authenticated, (ii, failure)

³Available online at <https://github.com/das-group/password-rba-horizon>




Figure 3: Implemented re-authentication prompt in the Horizon frontend for a suspicious login attempt

the authentication has failed, and (iii, suspicious) the RBA feature verification calculated a suspicious risk score that requires a re-authentication factor to be requested. In our implementation, we use a verification code (passcode) generated by a HMAC-Based One-Time Password (HOTP) [22] generator, and sent over a separate channel to the user as a re-authentication factor. We chose email based verification with a six digit code in subject line and body, as this variant performed best in a user study [33]. Also, verification codes are common practice in real-world RBA systems [12, 36].

In the *suspicious* case, the login dialog requests the verification code that was sent to a previously deposited user contact address, e.g., email address (see Figure 3). We based the dialog and verification email on common patterns found in RBA deployments of popular online services, and recommendations of an email verification study [33, 37]. We did this to create an RBA user interface that was tested successfully in multiple usability studies. To be compatible with different contact points in the future (e.g., phone number), we changed the term “email address” to “contact address”. For technical reasons, the user’s contact address was not shown in the dialog (see the discussion in Section 7). The subsequent login attempt using the RBA method now contains the verification code instead of features and it needs to match the verification code issued by Keystone to succeed.

5.2 Keystone RBA Extension (Backend)

The Keystone RBA plug-in⁴ implements the backend RBA evaluation based on the features delivered by the Horizon RBA frontend, the stored login history, and the Freeman et al. algorithm (see Figure 2).

Data Collection. The RBA features received at the beginning of a login attempt are validated and normalized to ensure comparability during further risk score processing. This includes sub-feature extraction to extend the acceptance range for unseen feature values.

⁴Available online at <https://github.com/das-group/keystone-rba-plugin>

Finally, the RTT feature is normalized by rounding to the nearest ten milliseconds, as suggested by related work [34].

Risk Score Calculation. Our implemented risk score algorithm uses the derived feature values with a linear interpolation probability estimation, as proposed by Freeman et al. [9], for the calculation of appearance in the user’s login history (see Section 2). Although smoothing including sub-features is possible for the probability estimation on the global history, the decision to compare only the originated features was made to limit the needed resources. Otherwise, it would be necessary to either keep a consistent copy of the whole history in-memory or query all entries from the login history database for each login attempt. Both variants seemed not acceptable with a growing user base. For the same reason, a size limit is configurable to even cap the amount of entries for each user. When the limit is exceeded, the oldest entry will be replaced in favor of a new one. Instead of narrowing down entries of the whole login history, only the currently attempting user’s entries are used and requested from the database. The other estimations can operate on lookup dictionaries or hash tables containing the actual value occurrences in the login history, as suggested by related work [35].

Furthermore, the risk score can optionally include the attack probability for a given IP address. We realized this as a third party reputation system, as proposed by Freeman et al. [9], by using a list containing recently reported malicious network addresses [30]. Such lists are common for firewall services to automatically block access to protected systems. In our case, we used the IP reputation database from the FireHOL project [30], that provides a daily updated collection from several sources.

Risk Classification. Three different results can be derived from the calculated risk score (see Section 5.1). Therefore, two threshold values can be configured, allowing to adjust the occurrence of these cases. Risk scores below the lower threshold are considered successful authentication and a new entry with the attempt’s features is stored in the database.

By exceeding the lower threshold, but still below the upper one, the response indicates Horizon to request an additional verification code. The RBA extension sends the verification code to the registered user via email by default, but this behavior can be changed by setting a different messenger in the configuration (e.g., to send text messages to mobile phone numbers).

In case even the upper threshold was exceeded, the authentication attempt is rejected. Nevertheless, it is quite common in practice to disable the rejection case [12, 36] and let RBA request the re-authentication factor only. Our implementation supports this by setting the rejection threshold to an unreachable high value.

Re-Authentication Request. In the re-authentication case, Horizon will include the verification code entered by the user in the following authentication request to Keystone. If the transmitted verification code could be verified, then a new record will be added to the login history. These previously labeled suspicious feature values will now be taken into account as *“already seen”* on further login attempts.

5.3 Extending the Feature Set

We integrated the IP address, user agent string, and RTT into the RBA extension, as they proved to be effective to identify users [34, 35]. Nevertheless, developers can change the feature set collected by the RBA extension in the code.

To achieve this, they first need to change the list of collected features and their values by the Horizon RBA extension. As Horizon sends these values to Keystone, they also need to adjust this feature set in the Keystone RBA extension. Furthermore, in case of new features beside the three implemented ones, they need to write and connect new validation functions to evaluate the new features.

After that, they need to adjust the two risk score thresholds to values reflecting the risk score values using the new feature set. Adding and removing new features will change the potential range of risk scores, as the amount of multiplications, mostly consisting of probabilities p with $\{p \in \mathbb{R} \mid 0 <= p < 1\}$, will change (see Equation 1). Therefore, we can assume that more features will likely lower the risk scores values. To get an idea of potential values, a reference test can be helpful.

6 REFERENCE TEST

In order to test a self-developed RBA implementation, some kind of reference test is required. Such a test could provide risk score values obtained by an openly available RBA reference implementation and an openly available login data set. Developers could calculate the risk scores based on the login data set using their own RBA implementation and compare the calculated risk scores to those of the reference implementation. Fortunately, both a reference implementation of the Freeman et al. algorithm [32] and a RBA login data set [35] got recently publicly available.

To obtain a reference test from the available resources, one must first determine the reference risk scores. The RBA reference implementation can be used for this purpose. It makes use of the Python pandas [26] library that is often used in big data and data science use cases. The reference implementation operates directly on a pandas DataFrame object that is initialized with the successful login attempts from the login data set. It provides a test function that calculates risk scores based on all preceding entries at a specified starting point. The result of the test function call contains all risk scores in the interval between the starting entry and the amount of entries to be considered. Note that the reference implementation allows to calculate risk scores just for slices of the data set. This is a useful feature in early development and testing stages, as it allows to focus the assessments on parts of the data set and reduce high computational costs. The feature also enables scaling the risk score computation across multiple servers and processor cores, e.g., for high performance computing clusters.

The results do not contain risk scores of login attempts by so far unseen users, as these would be zero anyway. Hence, as soon as a user has logged in more than once, the corresponding risk score is calculated and becomes part of the output. Table 1 shows the first calculated risk scores when going through the data set from the beginning. The 64th entry in the data set represents the first recurring login attempt of the same user. Thus, it is the first output of the test function of the reference implementation.

Table 1: Sample of selected risk scores calculated sequentially for successful login attempts of the RBA data set and compared to the risk scores calculated using the Keystone plugin.

Login Attempt		Reference Risk Score	Plug-in Risk Score
Global	User		
64	2	0.0105382376	0.0105382376
67	2	0.0005499333	0.0005499333
89	2	0.0024253951	0.0024253951
...			
10000	2	0.0184015408	0.0184015408
10004	2	0.0000824648	0.0000824648
10008	4	0.0063803620	0.0063803620
10012	3	0.0002190521	0.0002190521
10013	9	0.0000297278	0.0000297278
10014	2	0.0035949540	0.0035949540
...			
29328	4	0.0009868327	0.0009868327
29331	3	0.0000687314	0.0000687314
29333	3	0.0000862387	0.0000862387

As shown in the table, the implementation of the risk score algorithm in the OpenStack RBA plug-in calculates the same values as the reference implementation. It should be noted that the reference implementation uses the IP address and user agent string and their sub-features ASN, country code, browser name, browser version, operating system name, operating system version and device type from the data set to calculate the risk values. As the validation, normalization, and sub-feature derivation process could result in deviations of feature values from the data set entries used by the implementation to be tested, it is important to circumvent these processes and calculate the risk scores directly on the same feature values contained in the data set.

7 DISCUSSION

Our guiding example shows that integrating RBA into software projects can be rather complex. We outline some of the issues we faced during implementation in the following.

Showing the Contact Address in Horizon. RBA dialogs typically show the user’s email address in censored form [33, 37]. This could help users to determine which email address received the notification while reducing attack surface for attackers, in case they do not know the full email address of the target. Unfortunately, the current code base of Horizon makes it difficult to forward the user’s contact address to the dialog. The only way to forward information from the plugin logic to the user interface seems to be via Python exceptions, where the contact address has to be attached to the error message. However, this might entail security and compatibility issues, as we have to parse the contact address correctly while keeping the parsing mechanisms compatible to other variants in the future (e.g., phone numbers). Therefore, to reduce complexity for future implementation versions, we had to decide against showing the contact address in the dialog. Future work should enhance OpenStack’s code base to allow forwarding variables from the plugin logic to the user interface to solve this problem.

Complexity in Implementing RBA. To integrate RBA into their own software projects, developers need to implement the RBA algorithm in the backend and test the algorithm with a data set and compare the output with the reference implementation. When the comparison was successful, they can integrate the data collection in the frontend, and implement the communication between frontend and backend. After that, they can implement the login flow including the re-authentication prompt. Finally, the developers have to set the different thresholds to classify the different risk categories. This can be complex in practice, however.

As derived by the surrounding risk scores in Table 1, a score above 0.003 could be medium risk and 0.018 high risk, as these risk scores are not too common in the sample. We assume that developers will not find this intuitive, as they rather expect numbers like 0.5 for medium risk and 1.0 for high risk. Ways on how to calibrate these scores to understandable values were not described in the Freeman et al. paper. To solve this problem, Wiefeling et al. [35] suggested a machine learning based algorithm, which returns a risk score threshold for medium risk. This threshold could be used as a baseline to calibrate all risk scores to more readable values.

Another source of complexity could be the integration of privacy into RBA. To protect collected features from attacks, several privacy-enhancing methods for the Freeman et al. algorithm were introduced [35, 38]. While most of the methods will not change the risk score, few of them will. Therefore, developers need to check with reference tests whether the privacy enhancements will keep RBA’s usability and security properties.

8 CONCLUSION AND OUTLOOK

RBA offers good security that service owners should deploy now to increase account protection for their users, until secure and usable alternatives to passwords for online services become a reality. Technical standards, political instruments, and scientific literature support this. What is still lacking, however, is widespread use beyond the few major online services that have been early adopters. Similarly, there is a lack of open implementations that software developers can use as a source of information to tackle the complex task. To address this gap and facilitate the use of RBA in practice, we provided a first fully functional open source RBA implementation based on the Freeman et al. algorithm for the OpenStack cloud management software. On the one hand, the OpenStack RBA plugin can immediately secure many cloud computing platforms and their resources, as well as users. On the other hand, it can serve as a guiding example for developers. To further assist developers in the complex task of implementing RBA for their online services, we provided a way to test the algorithm implementations of Freeman et al. Finally, we want to engage with software developers through our repository who made use of our open RBA implementation in some way. Our goal is to collect qualitative and quantitative empirical data on how useful complete examples are perceived by developers and what other factors play a role in the adoption of security technologies.

Following our OpenStack implementation, we plan to integrate RBA into more open source software projects. This should help to impact a widespread use of RBA in practice to protect more users from attacks like credential stuffing and password spraying.

REFERENCES

- [1] Akamai. 2019. Credential Stuffing: Attacks and Economies. *[state of the internet] / security* 5, Special Media Edition (April 2019). <https://web.archive.org/web/20210824114851/https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/soti-security-credential-stuffing-attacks-and-economies-report-2019.pdf>
- [2] Furkan Alaca and P. C. van Oorschot. 2016. Device Fingerprinting for Augmenting Web Authentication: Classification and Analysis of Methods. In *32nd Annual Computer Security Applications Conference (ACSAC '16)*. ACM, 289–301. <https://doi.org/10.1145/2991079.2991091>
- [3] Nampoina Andriamilanto, Tristan Allard, and Gaëtan Le Guelvouit. 2021. “Guess Who?” Large-Scale Data-Centric Study of the Adequacy of Browser Fingerprints for Web Authentication. In *Innovative Mobile and Internet Services in Ubiquitous Computing*. Springer, Cham, 161–172. https://doi.org/10.1007/978-3-030-50399-4_16
- [4] Australian Cyber Security Centre. 2021. *Australian Government Information Security Manual*. Technical Report. [https://web.archive.org/web/20210830131917/https://www.cyber.gov.au/sites/default/files/2021-06/01.%20ISM%20-%20Using%20the%20Australian%20Government%20Information%20Security%20Manual%20\(June%202021\).pdf](https://web.archive.org/web/20210830131917/https://www.cyber.gov.au/sites/default/files/2021-06/01.%20ISM%20-%20Using%20the%20Australian%20Government%20Information%20Security%20Manual%20(June%202021).pdf)
- [5] Joseph R. Biden Jr. 2021. Executive Order on Improving the Nation’s Cybersecurity. *The White House* (May 2021). <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>
- [6] Anne Bumiller, Olivier Barais, Nicolas Aillery, and Gael Le Lan. 2022. Towards a Better Understanding of Impersonation Risks. In *15th International Conference on Security of Information and Networks (SIN '22)*. IEEE, Sousse, Tunisia. <https://doi.org/10.1109/SIN56466.2022.9970540>
- [7] Periwinkle Doerfler, Kurt Thomas, Maija Marincenko, Juri Ranieri, Yu Jiang, Angelika Moscicki, and Damon McCoy. 2019. Evaluating Login Challenges as a Defense Against Account Takeover. In *The World Wide Web Conference (WWW '19)*. ACM, 372–382. <https://doi.org/10.1145/3308558.3313481>
- [8] Jonathan Dutton, Danny Allen, Dennis Eggett, and Kent Seamons. 2019. “Don’t punish all of us”: Measuring User Attitudes about Two-Factor Authentication. In *4th European Workshop on Usable Security (EuroUSEC '19)*. <https://doi.org/10.1109/EuroSPW.2019.00020>
- [9] David Freeman, Sakshi Jain, Markus Dürmuth, Battista Biggio, and Giorgio Giacinto. 2016. Who Are You? A Statistical Approach to Measuring User Authenticity. In *NDSS*, Vol. 16. 21–24.
- [10] Ajit Gaddam. 2019. Usage of Behavioral Biometric Technologies to Defend Against Bots. In *Enigma 2019*. USENIX Association.
- [11] Simson L. Garfinkel. 2005. Design principles and patterns for computer systems that are simultaneously secure and usable.
- [12] Anthony Gavazzi, Ryan Williams, Engin Kirda, Long Lu, Andre King, Andy Davis, and Tim Leek. 2023. A Study of Multi-Factor and Risk-Based Authentication Availability. In *32nd USENIX Security Symposium (USENIX Security '23)*. USENIX Association, Anaheim, CA, USA.
- [13] P. A. Grassi, M. E. Garcia, and J. L. Fenton. 2017. *Digital Identity Guidelines*. NIST Special Publication 800-63-3. National Institute of Standards and Technology, Gaithersburg, MD 20899-2000. <https://doi.org/10.6028/NIST.SP.800-63-3>
- [14] J. Hawkinson. 1996. Guidelines for creation, selection, and registration of an Autonomous System (AS). RFC 1930. <https://www.rfc-editor.org/rfc/rfc1930.html>
- [15] ISO 3166 Maintenance Agency. 2020. ISO 3166-1:2020(en) Codes for the representation of names of countries and their subdivisions – Part 1: Country code. ISO 3166-1. <https://www.iso.org/obp/ui/#iso:std:iso:3166:-1:ed-4:v1:en>
- [16] Tom Le Bras. 2015. Online Overload – It’s Worse Than You Thought. <https://web.archive.org/web/20150919202348/https://blog.dashlane.com/infographic-online-overload-its-worse-than-you-thought/>
- [17] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *2019 Network and Distributed System Security Symposium (NDSS '19)*. Internet Society. <https://doi.org/10.14722/ndss.2019.23386>
- [18] Philipp Markert, Theodor Schnitzler, Maximilian Golla, and Markus Dürmuth. 2022. “As soon as it’s a risk, I want to require MFA”: How Administrators Configure Risk-based Authentication. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. USENIX Association, Boston, MA, 483–501. <https://www.usenix.org/conference/soups2022/presentation/markert>
- [19] Grzegorz Milka. 2018. Anatomy of Account Takeover. In *Enigma 2018*. USENIX Association. <https://www.usenix.org/node/208154>
- [20] Paul Miller and Lauren E Nelson. 2015. Brief: OpenStack Is Now Ready For Business. *Forrester Report Brief* (Sept. 2015).
- [21] Robert Morris and Ken Thompson. 1979. Password security: A case history. *Commun. ACM* 22, 11 (Nov. 1979), 594–597. <https://doi.org/10.1145/359168.359172>
- [22] D. M’Raihi, M. Bellare, F. Hoornaert, D. Naccache, and O. Ranen. 2005. HOTP: An HMAC-Based One-Time Password Algorithm. RFC 4226. <https://doi.org/10.17487/RFC4226>
- [23] National Cyber Security Centre. 2018. *Cloud security guidance: 10, Identity and authentication*. Technical Report. <https://www.ncsc.gov.uk/collection/cloud-security/implementing-the-cloud-security-principles/identity-and-authentication>
- [24] Lily Hay Newman. 2021. Facebook Will Force More At-Risk Accounts to Use Two-Factor. <https://web.archive.org/web/20211212185008/https://www.wired.com/story/facebook-protect-two-factor-authentication-requirement/>
- [25] Thanasis Petsas, Giorgos Tsirantonakis, Elias Athanasopoulos, and Sotiris Ioannidis. 2015. Two-factor Authentication: Is the World Ready?: Quantifying 2FA Adoption. In *Eighth European Workshop on System Security (EuroSec '15)*. ACM. <https://doi.org/10.1145/2751323.2751327>
- [26] PyData Development Team. 2020. pandas documentation. <https://pandas.pydata.org/pandas-docs/version/1.1.5/>
- [27] Nils Quermann, Marian Harbach, and Markus Dürmuth. 2018. The State of User Authentication in the Wild. In *Who are you?! Adventures in Authentication Workshop 2018 (WAY '18)*. <https://wayworkshop.org/2018/papers/way2018-quermann.pdf>
- [28] Ken Reese, Trevor Smith, Jonathan Dutton, Jonathan Armknecht, Jacob Cameron, and Kent Seamons. 2019. A Usability Study of Five Two-Factor Authentication Methods. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS '19)*. USENIX Association, 357–370. <https://www.usenix.org/conference/soups2019/presentation/reese>
- [29] IEEE Spectrum. 2022. Top Programming Languages 2022. <https://spectrum.ieee.org/top-programming-languages-2022>
- [30] Costas Tsaoasis. 2022. All Cybercrime IP Feeds. <https://iplists.firehol.org/>
- [31] Twitter. 2022. Account Security - Twitter Transparency Center. <https://web.archive.org/web/20220211182429/https://transparency.twitter.com/en/reports/account-security.html#2021-jan-jun>
- [32] Stephan Wiefeling. 2022. Basic Algorithm for Risk-Based Authentication. <https://github.com/das-group/rba-algorithm>
- [33] Stephan Wiefeling, Markus Dürmuth, and Luigi Lo Iacono. 2020. More Than Just Good Passwords? A Study on Usability and Security Perceptions of Risk-based Authentication. In *36th Annual Computer Security Applications Conference (ACSAC '20)*. ACM, 203–218. <https://doi.org/10.1145/3427228.3427243>
- [34] Stephan Wiefeling, Markus Dürmuth, and Luigi Lo Iacono. 2021. What’s in Score for Website Users: A Data-Driven Long-Term Study on Risk-Based Authentication Characteristics. In *25th International Conference on Financial Cryptography and Data Security (FC '21)*. Springer, 361–381. https://doi.org/10.1007/978-3-662-64331-0_19
- [35] Stephan Wiefeling, Paul René Jørgensen, Sigurd Thunem, and Luigi Lo Iacono. 2023. Pump Up Password Security! Evaluating and Enhancing Risk-Based Authentication on a Real-World Large-Scale Online Service. *ACM Transactions on Privacy and Security* 26, 1, Article 6 (2023). <https://doi.org/10.1145/3546069>
- [36] Stephan Wiefeling, Luigi Lo Iacono, and Markus Dürmuth. 2019. Is This Really You? An Empirical Study on Risk-Based Authentication Applied in the Wild. In *34th IFIP International Conference on ICT Systems Security and Privacy Protection (IFIP SEC '19)*. Springer, 134–148. https://doi.org/10.1007/978-3-030-22312-0_10
- [37] Stephan Wiefeling, Tanvi Patil, Markus Dürmuth, and Luigi Lo Iacono. 2020. Evaluation of Risk-based Re-Authentication Methods. In *35th IFIP International Conference on ICT Systems Security and Privacy Protection*. Springer, 280–294. https://doi.org/10.1007/978-3-030-58201-2_19
- [38] Stephan Wiefeling, Jan Tolsdorf, and Luigi Lo Iacono. 2021. Privacy Considerations for Risk-Based Authentication Systems. In *2021 International Workshop on Privacy Engineering (IWPE '21)*. IEEE, 320–327. <https://doi.org/10.1109/EuroSPW54576.2021.00040>