ASSURED

SECURITY CONSULTANTS

Report

Guardian app web, API and cloud security assessment

Thomas Stacey, Patrik Aldenvik, Wictor Olsson

Project	Version	Date
DNS006	1	2025-09-04

ASSURED
SECURITY CONSULTANTS

REPORT

Project	Version	Date
DNS006	1	2025-09-04

Executive summary

Between 2025-05-26 and 2025-06-12 Assured Security Consultants performed a security assessment on behalf of Guardian Firewall. Further, between 2025-08-13 and 2025-08-29 a verification test was performed, where the findings from the original report has been reassessed and given a status of: FIXED, PARTIALLY FIXED, ACCEPTED or REMAINING.

In scope was the Guardian Firewall web frontend and API's as well as its Google Cloud deployment. The test aimed to assess the application's general security posture, new functionality and cloud infrastructure.

The assessment covered many aspects of the Guardian Firewall web and API services, including authentication, authorization and data handling. Additionally, a thorough review of the Google Cloud deployment was conducted to identify potential security vulnerabilities in the underlying infrastructure.

A combination of manual testing, automated tools, and vulnerability scanning techniques was employed during the assessment.

This report lists the security issues found, along with recommendations for remediation or mitigation. In our conclusions we discuss the issues and address apparent patterns in areas where security is lacking.

Observations were made with the following risk severity assessments (number of issues):

Critical 0 High 2 Medium 6 Low 7 Note 0

A CSRF vulnerability was identified in the "update email" functionality of the Guardian API which enables a unauthenticated attacker to take control of other users accounts.

Insufficient error handling in the API backend leads to crashes. These flaws can be exploited by an unauthenticated attacker to disrupt the Guardian backend API.

Inadequate cloud network segmentation and default VM service accounts create opportunities for lateral movement within the Google Cloud environment, potentially allowing an attacker to access sensitive data or disrupt services.

In summary a few common high risk issues were identified but nothing critical, some areas could be improved to increase robustness and defence in depth.

The verification test shows that Guardian has prioritized mitigating issues both in their application implementation as well as infrastructure. A few lower risk issues has the status of ACCEPTED for now, they are planned to be implemented in the future. With this the Connect API, web frontend and cloud setup has been thoroughly tested and reviewed. No issues related to customer traffic integrity were identified and issues with potential privacy impact was dealt with swiftly.

Assured would like to thank the Guardian team for their support during this penetration test. We are happy to answer any questions and provide further advice.

Contents

1 Observations

1.1 Web and APIs

This section contains observations related to the guardian web and its related APIs.

1.1.1 HIGH FIXED Cross-Site Request Forgery in update email

Likelihood: MEDIUM (5), Impact: HIGH (7)

Verification note: CSRF token has now been implemented on the `https://guardianapp.com/user/update-email` endpoint, adding unpredictable data to the request, and therefore preventing an attacker from performing a successful CSRF attack. Furthermore, the `grd-auth-cookie` session cookie's `SameSite` attribute has been set to `Strict` which adds an extra layer of protection.

1.1.2 MED FIXED Inadequate error handling in REST API, leading to unhandled crashes

Likelihood: MEDIUM (3), Impact: MEDIUM (3)

Verification note: The vulnerabilities leading to crashes/service disruption have been corrected, Guardian has also made improvements in its development process to catch such issues earlier

1.1.3 MED FIXED User Enumeration via Password reset

Likelihood: MEDIUM (3), Impact: MEDIUM (3)

Verification note: The application no longer responds with any information that would allow an attacker to determine whether or not the user exists. Furthermore the response time for this endpoint is consistent, regardless of whether or not the user exists, additionally preventing any form of time-based user enumeration.

1.1.4 MED FIXED Time-based user enumeration during login

Likelihood: MEDIUM (3), Impact: MEDIUM (3)

Verification note: Response time for the login endpoint is not deterministic, mitigating time-based user enumeration.

1.1.5 MED ACCEPTED No bruteforce protection or account lockout

Likelihood: MEDIUM (3), Impact: MEDIUM (3)

There were no protections or limitations identified when attempting to bruteforce an account password, and the account was not locked during testing.

This can be abused by an attacker to guess common credentials of users in the system to gain access. Credential based attacks are one of the most common.

We recommend implementing a limit on how many login attempts that can be performed within a given timeframe from a single source. Implement an account lockout policy for failed authentication attempts.

1.1.6 LOW FIXED Input validation of strings missing

Likelihood: MEDIUM (4), Impact: LOW (2)

Verification note: The offending API endpoint/functionality has been removed.

1.1.7 LOW ACCEPTED Insufficient anti-automation for email sending

Likelihood: LOW (2), Impact: MEDIUM (3)

The application's password reset functionality implements no rate-limiting of CAPTCHA of any kind, allowing an attacker to repeatedly trigger a password reset request for arbitrary users. Each reset will trigger an email towards the victim user, spamming their inbox and increasing the likelihood of the application's legitimate emails being marked as spam. Furthermore, any mail servers are likely to eventually become strained possibly denying other users' reset requests from being sent or received.

We recommend implementing a rate-limiting or CAPTCHA on the password-reset endpoint as it is very unlikely that any legitimate usage of the password reset functionality requires many requests within a short time frame.

1.1.8 LOW FIXED Authenticated session not invalidated on logout

Likelihood: LOW (2), Impact: MEDIUM (3)

Verification note: The authentication cookie is now properly invalidated upon sign out.

1.1.9 LOW ACCEPTED Missing Multi-Factor Authentication

Likelihood: LOW (2), Impact: MEDIUM (3)

Multi-factor authentication (MFA) enhances security by requiring users to provide at least two separate forms of authentication, typically something they know (e.g., a password), something they have (e.g., a hardware token or a smartphone), and/or something they are (e.g., a biometric identifier like a fingerprint). By not implementing MFA, the application becomes more vulnerable to unauthorized access, as attackers only need to compromise a single authentication factor to gain entry.

Missing or lacking MFA increases the likelihood of unauthorized access to user accounts and sensitive data. If an attacker manages to obtain a user's login credentials, either through phishing, brute force attacks, or credential stuffing, they can easily access the user's account without the need to bypass any additional authentication layers. This can lead to data breaches, identity theft, unauthorized modifications, and financial losses, depending on the privileges associated with the compromised account.

We recommend implementing MFA as part of user authentication. This can be achieved by using various authentication methods, such as one-time passwords (OTP) generated by a dedicated app, hardware tokens, or biometric authentication. MFA should be mandatory at least for administrators and forced on logins from new devices and for certain sensitive actions.

1.2 Google Cloud Platform

This section contains observations related to the hosting environment in Google Cloud Platform.

1.2.1 HIGH PARTIALLY FIXED Network segmentation

Likelihood: MEDIUM (4), Impact: HIGH (6)

Verification note: Coarse network segmentation has been implemented at the GCP level. The more granular firewall rules are applied at the host level, using UFW.

1.2.2 MED FIXED Default VM service account

Likelihood: MEDIUM (4), Impact: MEDIUM (3)

Verification note: There is no active default service account used by any of the VMs within the environment.

1.2.3 MED PARTIALLY FIXED Default network resources in use

Likelihood: MEDIUM (3), Impact: MEDIUM (3)

Verification note: The default resources for the Guardian firewall production environment have been removed. There are still default resources in other non-production environments, which are planned to be removed.

1.2.4 LOW ACCEPTED DNSSEC

Likelihood: LOW (2), Impact: MEDIUM (3)

Verification note: DNSSEC adds complexity and Guardian has decided to not implement it to reduce possible errors.

1.2.5 LOW FIXED Eligible project owner

Likelihood: LOW (2), Impact: MEDIUM (3)

Verification note: The user has been removed.

1.2.6 LOW PARTIALLY FIXED Overpriviliged service account

Likelihood: LOW (2), Impact: MEDIUM (3)

Verification note: GCP does not provide an out-of-the box solution for this service account. A working solution is in development.