ASSURED

SECURITY CONSULTANTS

Report

Mullvad VPN - Email server audit

Wictor Olsson, Benjamin Svensson

Project	Version	Date
MUL017	1	2024-02-12

ASSURED
SECURITY CONSULTANTS

REPORT

Project	Version	Date
MUL017	1	2024-02-12

Executive summary

Between 2024-01-08 and 2024-01-12 Assured Security Consultants performed a security assessment on behalf of Mullvad VPN.

The mail servers `mail.mullvadvpn.net` and `mail2.mullvadvpn.net` were in scope.

The overall security level of the in-scope servers is considered to be good, only a few issues were identified in this assessment. No identified issue poses a major risk to Mullvad VPN. Some listed observations are rated as notes and are suggestions to improve the security configuration of the services or systems further.

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

Issues were found with the following risk severity assessments (number of issues):

Critical 0 High 0 Medium 2 Low 3 Note 3

Our recommendations are to patch and reconfigure the services and system according to the suggestions in this report.

Assured would like to thank Richard and Victor for their support during this assessment. We are happy to answer any questions and provide further advice.

Contents

1 Introduction

1.1 Background

Assured AB (Assured) was contracted to perform an audit on Mullvad VPN:s support mail server. Assured was given access to the code repository and SSH access to the servers in scope. A communications channel was setup to allow for immediate support and reporting.

1.2 Constraints and disclaimer

This report contains a summary of the findings found during the project period. This report should not be considered as a complete list of all vulnerabilities, security flaws and/or misconfigurations.

1.3 Project period and staffing

Assured started the project on 2024-01-08 and finished on 2024-01-12.

This report was last reviewed on 2024-02-12.

Involved in the penetration testing was Assured consultant Wictor Olsson and Benjamin Svensson.

1.4 Risk rating

1.4.1 OWASP Risk Rating Methodology

In this report we have assessed the severity of issues and identified vulnerabilities according to the OWASP Risk Rating Methodology [1].

Table 1: OWASP Risk Rating overall severity model

Overall risk severity				
Impact	HIGH	Medium	High	Critical
	MEDIUM	Low	Medium	High
	LOW	Note	Low	Medium
		LOW	MEDIUM	HIGH
	Likelihood			

As Table 1 visualizes, the overall risk assessment is determined from a combined likelihood and impact of an identified vulnerability or security issue. A value from 0 to 9 is assessed for each variable, where 0-2 is determined LOW, 3-5 is MEDIUM and 6-9 is HIGH.

Likelihood is dependent on attributes related to threat actors and the identified vulnerability, with factors such as: the skill level and motivations of the threat agents; how easily the vulnerability can be found and exploited, and; how likely an exploit may be detected.

Impact depends on technical and business factors, such as: level of loss of confidentiality, integrity, availability and accountability; potential financial damage; potential brand damage, and; potential violations of privacy.

Please note that the severity assessment is made by Assured consultants and ratings may differ from the resource owners' ratings.

2 Scope and methodology

2.1 Scope

The scope of the assessment included two servers, `mail.mullvadvpn.net` and `mail2.mullvadvpn.net`. The main testing effort was focused on `mail.mullvadvpn.net` because it was installed with all the services for email (Postfix, Dovecot, Postgresql, rspamd) where `mail2.mullvadvpn.net` was only configured as an MX (secondary SMTP server). Ansible source code and configuration files were included in the scope for the test.

2.2 Methodology

The assessment started with review of the Ansible code used to setup and install the servers. All configuration and documentation were included in the ansible repository. Testers were provided with SSH access to the servers to login and review the running configuration and any operating system issues. Wiregurad keys were also provided to access the network where IMAP client access is allowed. Focus was primarily on the email services and their configuration.

2.2.1 Tools used

- lynis
- nmap
- ansible-lint
- testssl.sh
- swaks
- ismtp
- thunderbird
- claws

3 Observations

3.1 FIXED MEDIUM Postfix SMTP smuggling CVE-2023-51764

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

The Postfix installation was found to be vulnerable to a recently disclosed attack vector known as SMTP Smuggling. The attack involves a COMPOSITION of two email services with specific differences in the way they handle line endings other than <CR><LF>. This can be abused by causing misinterpretations of the standard SMTP end-of-message marker <CR><LF>.<CR><LF>.

As an example, consider two mail services A and B. Service A does not recognize malformed line endings in SMTP commands, such as in <LF>.<CR><LF>. If such a sequence is included in an email message from an authenticated attacker to a recipient at email service B, the malformed sequence is forwarded from service A to service B.

If service B does support malformed line endings in SMTP such as in <LF>.<CR><LF>, the service is vulnerable. An attacker crafts a payload where the malformed ending is followed by ”smuggled” SMTP MAIL/RCPT/DATA commands and message header plus body text, email service B is tricked into receiving two email messages.

- one message with the content before the <LF>.<CR><LF>
- one message with the ”smuggled” header plus body text after the ”smuggled” SMTP commands.

Service A believes only one message was forwarded, while service B believes it has received two messages. If Mullvad’s email server is service B in this example, Mullvad may receive and process spoofed email.

The attacker can use the ”smuggled” SMTP MAIL/RCPT/DATA commands and header plus body text, to spoof an email message from any MAIL FROM address whose domain is also hosted at email service A, to any RCPT TO address whose domain is also hosted at email service B.

The spoofed email message will pass SPF-based DMARC checks at email service B, because the spoofed message has a MAIL FROM address whose domain is hosted at email service A, and because the message was received from an IP address for email service A.

We recommend applying the mitigating configuration below and/or upgrade to Postfix 3.8.4, 3.7.9, 3.6.13 or 3.5.23.

```
smtpd_data_restrictions = reject_unauth_pipelining
smtpd_discard_ehlo_keywords = chunking, silent-discard
```

For further reference see Postfix advisory [2] and NIST CVE [3].

3.2 FIXED MEDIUM Postfix SMTP denial of service

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

The SMTP configuration is missing options which can increase the service resistance against denial of service attacks.

An attacker could send a large number of messages in a short time, causing the SMTP server to refuse new incoming messages, and effectively block other legitimate servers from sending messages to the service.

We recommend adding configuration options to limit the size of message headers and in the rate messages can be sent from the same source.

```
smtpd_client_connection_rate_limit =
smtpd_client_event_limit_exceptions =
header_size_limit =
```

You can find more information on how to tune performance and harden the configuration for denial of service attacks at `https://www.postfix.org/TUNING_README.html` and `https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/security_guide/sect-security_guide-server_security-securing_postfix`.

3.3 FIXED LOW Postfix TCP 587 submission exposed

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

The TCP port 587 is usually used for SMTP submission, which is a mail submission agent port. It is used by mail clients to connect and authenticate when sending emails.

Mullvad's configuration forces the clients to access the mail server from a Wireguard tunnel, and does not require the submission service to be listening to an external interface. The snippet below shows a `nmap` scan identifying TCP port 587 as open.

```
# Nmap 7.80 scan initiated Mon Jan 8 08:56:59 2024 as: nmap -v -sV -sC -p- -oN nmap_scan_service.nmap -Pn mail.mullvadvpn.net
Nmap scan report for mail.mullvadvpn.net (185.213.154.124)
Host is up (0.0091s latency).
Not shown: 65532 filtered ports
PORT   STATE SERVICE VERSION
25/tcp open smtp Postfix smtpd
|_smtp-commands: mail.mullvadvpn.net, PIPELINING, SIZE 52428800, VRFY, ETRN, STARTTLS, ENHANCEDSTATUSCODES, 8BITMIME, DSN, SMTPUTF8, CHUNKING,
| ssl-cert: Subject: commonName=mail.mullvadvpn.net
| Subject Alternative Name: DNS:autoconfig.mullvadvpn.net, DNS:imap.mullvadvpn.net, DNS:mail.mullvadvpn.net, DNS:mta-sts.mullvadvpn.net, DNS:rspamd.mullvadvpn.net, DNS:smtp.mullvadvpn.net
| Issuer: commonName=R3/organizationName=Let's Encrypt/countryName=US
| Public Key type: rsa
| Public Key bits: 4096
| Signature Algorithm: sha256WithRSAEncryption
| Not valid before: 2023-12-21T09:53:50
| Not valid after: 2024-03-20T09:53:49
| MD5: cd8f 229f d571 707b 13e5 0954 5212 42b0
|_SHA-1: 90df 897e b25e 3b43 127b 40b6 22e4 9c6f 80fe def7
|_ssl-date: TLS randomness does not represent time
587/tcp open smtp Postfix smtpd
|_smtp-commands: mail.mullvadvpn.net, PIPELINING, SIZE 52428800, VRFY, ETRN, STARTTLS, ENHANCEDSTATUSCODES, 8BITMIME, DSN, SMTPUTF8, CHUNKING,
| ssl-cert: Subject: commonName=mail.mullvadvpn.net
| Subject Alternative Name: DNS:autoconfig.mullvadvpn.net, DNS:imap.mullvadvpn.net, DNS:mail.mullvadvpn.net, DNS:mta-sts.mullvadvpn.net, DNS:rspamd.mullvadvpn.net, DNS:smtp.mullvadvpn.net
| Issuer: commonName=R3/organizationName=Let's Encrypt/countryName=US
| Public Key type: rsa
| Public Key bits: 4096
| Signature Algorithm: sha256WithRSAEncryption
| Not valid before: 2023-12-21T09:53:50
| Not valid after: 2024-03-20T09:53:49
| MD5: cd8f 229f d571 707b 13e5 0954 5212 42b0
|_SHA-1: 90df 897e b25e 3b43 127b 40b6 22e4 9c6f 80fe def7
|_ssl-date: TLS randomness does not represent time
1022/tcp open ssh OpenSSH 8.9p1 Ubuntu 3ubuntu0.5 (Ubuntu Linux; protocol 2.0)
Service Info: OS: Linux; CPE: cpe:/o:linux:linux_kernel
Read data files from: /usr/bin/../share/nmap
Service detection performed. Please report any incorrect results at https://nmap.org/submit/ .
# Nmap done at Mon Jan 8 08:59:10 2024 -- 1 IP address (1 host up) scanned in 131.40 seconds
```

This exposed service has authentication enabled which an attacker could abuse and perform for example credential based attacks against to gain access to the Mullvad email service.

We recommend configuring the firewall to block incoming traffic to TCP port 587 and configuring Postfix to bind the port on the Wireguard interface.

3.4 FIXED LOW Postfix SMTP user enumeration (VRFY)

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

The VRFY SMTP command is enabled which makes it possible for an unauthenticated external attacker to enumerate existing email addresses configured on the system.

The following systems are affected `mail.mullvadvpn.net`, `mail2.mullvadvpn.net`.

An attacker can verify if a email exists on the system, this could then later be abused through credential based attacks on other services.

We recommend to disable the command, example below.

```
disable_vrfy_command=yes
```

3.5 FIXED LOW Ansible secrets in clear-text

Likelihood: LOW (1), Impact: MEDIUM (3)

The application source code contains hardcoded credentials or other secrets.

The following files are configured with hard-coded secrets in ansible.

```
/inventory/group_vars/all/global.yml - __global_vmail_db_password: "password goes here"
ansible/inventory/host_vars/mail.mullvadvpn.net/main.yml - Wireguard private key
ansible/inventory/host_vars/mail2.mullvadvpn.net/main.yml - Wireguard private key
mail/ansible/inventory/group_vars/all/secrets.yml -password: "{SHA512-CRYPT}
    XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX"
```

Compromise of the source code can occur in multiple ways, such as malware on a developer workstation, compromise of an external consultant, or leak of repository credentials.

An attacker who gains access to the source learns hardcoded secrets such as credentials and may use them to attack the infrastructure. If the credentials are for the production environment, the impact is the compromise of the service associated with the leaked credentials. Even if the credentials are for a development environment, the attacker may be able to spot credential reuse for other services.

We recommend removing the affected credentials from the source code, implementing a secure secret storage for runtime access to secrets such as `ansible-vault`. The specific credentials in place should be revoked and replaced. If the credentials are for a production environment, take steps to verify that they have not already been compromised.

3.6 FIXED NOTE IPv6 enabled but not used

IPv6 protocol is enabled on the server but is not expected to be used. As seen in figure 1, several services are listening on all interfaces, including IPv6 loopback.

from 2024-01-11 10-56-40.png

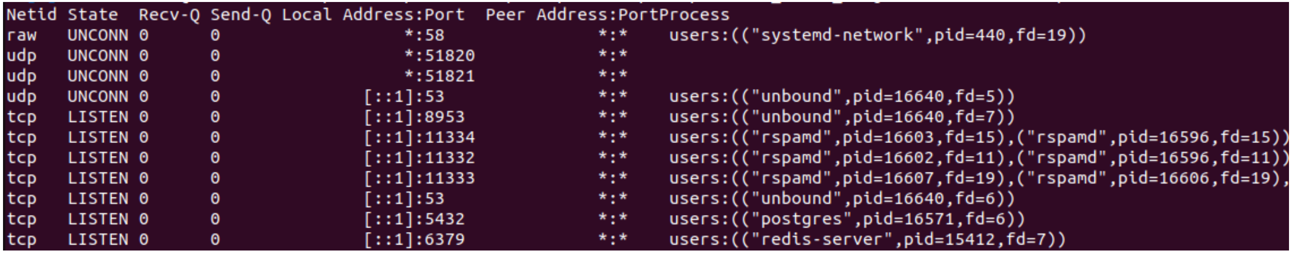

```
Netid State  Recv-Q Send-Q Local Address:Port  Peer Address:PortProcess
raw   UNCONN 0      0                  *:58               *:*    users:(("systemd-network",pid=440,fd=19))
udp   UNCONN 0      0                  *:51820            *:*
udp   UNCONN 0      0                  *:51821            *:*
udp   UNCONN 0      0              [::1]:53               *:*    users:(("unbound",pid=16640,fd=5))
tcp   LISTEN 0      0              [::1]:8953             *:*    users:(("unbound",pid=16640,fd=7))
tcp   LISTEN 0      0              [::1]:11334            *:*    users:(("rspamd",pid=16603,fd=15),("rspamd",pid=16596,fd=15))
tcp   LISTEN 0      0              [::1]:11332            *:*    users:(("rspamd",pid=16602,fd=11),("rspamd",pid=16596,fd=11))
tcp   LISTEN 0      0              [::1]:11333            *:*    users:(("rspamd",pid=16607,fd=19),("rspamd",pid=16606,fd=19),
tcp   LISTEN 0      0              [::1]:53               *:*    users:(("unbound",pid=16640,fd=6))
tcp   LISTEN 0      0              [::1]:5432             *:*    users:(("postgres",pid=16571,fd=6))
tcp   LISTEN 0      0              [::1]:6379             *:*    users:(("redis-server",pid=15412,fd=7))
```

Figure 1: Services listening on IPv6

We recommend disabling IPv6 on the server using *sysctl*. Add the following lines to */etc/sysctl.conf*. And to reconfigure the services to only listen on the expected interfaces.

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
```

3.7 FIXED NOTE Dovecot config cleanup

Dovecot is configured to listen on two ports: 143 and 993. Both ports enforce TLS. Typically 143 is used for unencrypted IMAP but in this case TLS is used for both 143 and 993. The services are only used by Mullvad employees where Mullvad themselves choose their mail clients, and thus two ports are not needed.

We recommend removing the configuration for the TCP port 143 and only listening on TCP port 993.

3.8 FIXED NOTE Dovecot server ciphers not preferred

The option `ssl_prefer_server_ciphers` is set to `no` in the dovecot configuration. This means that the client will be choosing the ciphers used for the TLS connection. Even though Mullvad sets a good suite of supported ciphers, it is recommended to have the server choose the order in which ciphers are supported.

We recommend setting the option `ssl_prefer_server_ciphers` to `yes`. See `https://doc.dovecot.org/configuration_manual/dovecot_ssl_configuration/#ssl-security-settings` for more information.

4 Conclusions and recommendations

Assured's assessment of Mullvad VPN's email servers is that their security posture is good. We have reviewed configuration and performed dynamic testing against the Dovecot and Postfix services. Most issues that have been identified are recommendations and suggestions for improvements rather than actual vulnerabilities. A few are considered to be risk, albeit low severity; primarily unneccesary exposure of Postfix, a known vulnerability and rate-limit hardening.

Access to the setup is very limited both in terms of network access as well as administrative functions. The services configuration follows security best practices in regards to hardening for the most parts.

The most severe issues identified are CVE-2023-51764 and a lack of rate limiting which could make it easier to run denial of service attacks agains the service. We identified secrets in clear text in the source code, which is not public and an attacker would therefore need to steal the source code from a Mullvad employee to access the secrets.

The SMTP submission port is publicly exposed which opens up for credentials based attacks, the observed credential quality in general is deemed very strong making this unlikely but not impossible. The SMTP command VRFY is enabled, which enabled an attacker to enumerate email users, but since this type of deployment only has one user this is not considered to pose any direct risk to Mullvad currently.

We recommend Mullvad VPN to:

- Patch for CVE-2023-51764 and/or apply mitigating configuration
- Address SMTP options for resistance against denial of service attacks
- Block TCP 587 from public exposure
- Remove support of the VRFY command
- Use ansible-vault (or equivalent) for secrets
- Disable IPv6
- Configure Dovecot to prefer server ciphers
- Configure Dovecot to only listen to TCP 993 with TLS
- Use a linter such as ansible-lint to check for risky behaviours or other errors

References

[1] OWASP, “OWASP Risk Rating Methodology.” `https://owasp.org/www-community/OWASP_Risk_Rating_Methodology`, 2023.

[2] Postfix, “SMTP Smuggling.” `https://www.postfix.org/smtp-smuggling.html`, 2024.

[3] NIST, “CVE-2023-51764.” `https://nvd.nist.gov/vuln/detail/CVE-2023-51764`, 2023.