murvord



CYBER AUDIT REPORT · PUBLIC SUMMARY

mssgs

mssgs, the chat service at mss.gs

REFERENCE	DATE OF THE REPORT	VALID UNTIL	VERSION
MV-BCW6E4R8	28 November 2026	28 November 2027	1.0

This report can be verified at https://murvord.com/verify/MV-BCW6E4R8, where Murvord confirms the reference, the scope and whether the badge is still valid, and lists the SHA-256 checksum of this file.

PART 1

Summary

Murvord audited the public side of mssgs and the parts of its source code that protect messages and accounts. The work ran from 5 to 28 November 2026; every external measurement in this report was repeated on 28 November 2026.

We found no critical and no high-severity issues. Two points are open, both of low severity. Eight issues were fixed during the audit: two of medium severity found by the external vulnerability scan on 5 November, and six from this report's own measurements, five fixed and measured again on 28 November and one on 29 November.

The basics are in good order. Connections are encrypted with current TLS and a post-quantum key exchange, only the web ports answer from the internet, mail sent in the name of mss.gs is protected against spoofing, card payments never touch mssgs, and one-to-one chats between the current apps are end-to-end encrypted.

0	0	0	2	0
CRITICAL	HIGH	MEDIUM	LOW	INFO

Open findings by severity on 28 November 2026. Resolved during the audit: 8.

Who wrote this, and what it is

Murvord works with RUAL, the company that makes mssgs, on the security of the platform on a continuing basis, and had access to the source code and the hosting for this audit. This is a review by a party that knows the system from the inside.

Scope

The public web front of mssgs (mss.gs, www.mss.gs, app.mss.gs): transport encryption, response headers, DNS and mail authentication, network exposure and the external vulnerability scan; and a source review of message encryption, account protection, dependencies and secrets.

The badge that refers to this report may be shown on: mss.gs, www.mss.gs, app.mss.gs.

PART 2

What we did

WHEN	WORK
5 and 6 November 2026	External vulnerability scan by a PCI Approved Scanning Vendor, two runs. We reviewed the results, the two fixes and the passing rescan.
18 and 19 November 2026	External configuration review of the web front: open ports, TLS, certificates and response headers.
22 November 2026	A filter at the network edge was put in front of the hosting. Checked from measurement points outside the network that only the intended services answer.
24 November 2026	Source review: how messages are encrypted, how passwords and two-factor secrets are stored.
28 November 2026	All external checks measured again; dependency and secret scans of the server code. Five findings fixed the same evening and measured again (R-03 to R-07); this report.
29 November 2026	Pages can no longer be framed by other sites: the fix measured in a browser (R-08). The script part of that finding stays open (F-02).

METHOD	HOW
External checks	Read-only measurements from the internet against mss.gs, www.mss.gs and app.mss.gs: protocol and cipher enumeration (OpenSSL, nmap ssl-enum-ciphers), certificate chain, HTTP response headers, redirects, DNS records read from a resolver outside the office network, and the service's security.txt.
Vulnerability scan	A full external scan of the three web addresses and the voice node by a PCI Approved Scanning Vendor (Qualys scanners). The scan was ordered by mssgs; Murvord reviewed the 157-page result, the fixes and the rescan.
Network exposure	A scan of all 65,535 TCP ports on the three public web addresses, and checks from independent measurement points in Europe and North America when the edge filter went live.
Source review	Reading the code of the chat backend, the apps and the platform underneath for the encryption of messages and the storage of passwords and two-factor secrets. Call-graph analysis of known advisories in dependencies (govulncheck) and a secret scan of the tracked files (gitleaks) on the chat backend, the website and the platform core.

PART 3

Results by area

AREA	RESULT	WHAT WE MEASURED
Transport encryption	IN ORDER	Only TLS 1.3 and TLS 1.2 are accepted; TLS 1.0 and 1.1 are refused. TLS 1.2 offers three cipher suites, all AEAD with forward secrecy, in the server's order (R-04); nmap grades every suite A. TLS 1.3 negotiates the hybrid post-quantum key exchange X25519MLKEM768. Plain HTTP is redirected to HTTPS with a 301, and HSTS is set for one year including subdomains, with preload.
Certificate	IN ORDER	An ECDSA certificate from Let's Encrypt for mss.gs and *.mss.gs, valid from 22 October 2026 to 20 January 2027. The chain verifies.
Response headers	ATTENTION	X-Content-Type-Options, Referrer-Policy and HSTS are set on every response, and since 29 November a Content-Security-Policy that keeps the pages out of frames on other sites (R-08). The policy does not yet restrict where scripts come from (F-02) and one header names the platform version (F-03).
DNS and mail authentication	IN ORDER	Five name servers under four top-level domains. SPF ends in -all and DMARC is p=reject with strict DKIM alignment, so mail that falsely claims to come from mss.gs is refused by receivers that check. A CAA record limits certificate issuance to one authority (R-03), and an MTA-STS policy with TLS reporting is published in testing mode (R-05). The zone is signed with DNSSEC and the registry publishes its key (R-07).
Network exposure	IN ORDER	On the three public web addresses only ports 80 and 443 and one webhook listener answer; every other TCP port is filtered. Since 22 November a filter at the network edge drops traffic to management and internal ports before it reaches a server. Confirmed from outside on the same day.
External vulnerability scan	IN ORDER	PCI ASV scan: compliant on 6 November 2026, with no finding at high or medium severity after two fixes (R-01, R-02). The scan is repeated every quarter.
Payments	IN ORDER	Card payments are a full redirect to the bank's payment page, so card data never passes through or is stored by mssgs (PCI DSS SAQ A). Purchases in the apps go through Apple and Google.
Message encryption	IN ORDER	Checked in the source code. A one-to-one chat between the current apps for iPhone, iPad, Android and desktop is end-to-end encrypted with Olm (Double Ratchet); the keys stay on the devices. A one-to-one message that cannot go end-to-end, because one side uses the web app or an old version, is encrypted on the server with AES-256-GCM before it is stored, under a key per chat that is kept in a different database. Messages in community channels are not end-to-end encrypted, by design.

Account protection	IN ORDER	Checked in the source code of the platform. A password is never stored, only a slow, salted hash of it. Two-factor secrets are stored encrypted.
Secrets in source code	IN ORDER	No production secret was found in the tracked files of the three code bases. What the scan flagged were test fixtures and two tokens that are public by design.
Software dependencies	IN ORDER	The chat backend, the platform core, the website, the voice service and the link-preview service were updated on 28 November and no longer call code with a published advisory. The running programs were read back as the updated builds (R-06).
Vulnerability disclosure	IN ORDER	mss.gs publishes a security.txt with a contact address and an expiry date, so a researcher knows where to report.

PART 4

Open findings

Severity is Murvord's judgement of the risk to mssgs and its users: critical, high, medium, low, or informational for points that are worth knowing but carry no direct risk.

F-02 The Content-Security-Policy does not restrict where scripts come from

LOW

OBSERVED Since 29 November pages on mss.gs and app.mss.gs carry a Content-Security-Policy that limits framing, the base address and plugin content (R-08). It has no script-src: the pages use inline scripts, which such a rule would block as they are written now.

RISK A script rule is a second line of defence: it limits what an injected script can do. We found no injection point; without the rule one would have more effect.

ADVICE Give the inline scripts a nonce, then add a script-src to the policy, starting in report-only mode.

F-03 The platform version is in a response header

LOW

OBSERVED Every response carries an X-Powered-By header with the platform's name and version number.

RISK It tells an attacker which version to look up. It is not a weakness by itself.

ADVICE Remove the header or drop the version from it.

PART 5

Resolved during the audit

R-01 Web addresses answered ICMP timestamp requests

MEDIUM FIXED

OBSERVED Found by the external scan on 5 November on all three web addresses (CVE-1999-0524): the servers returned their clock to anyone who asked.

FIX The firewall drops timestamp requests since 6 November. Re-tested the same day: no reply on any of the three, and normal reachability unchanged. The rescan passed.

R-02 Voice node closed TLS connections that sent no server name

MEDIUM FIXED

OBSERVED Reported by the external scan on 5 November as possible scan interference on port 443 of the voice node: a TLS client that named no host was handed to a service that does not speak TLS.

FIX Since 6 November every TLS client reaches the web service whatever name it sends. Re-tested the same day and passed in the rescan.

R-03 No CAA record

LOW FIXED

OBSERVED Measured on 28 November: the zone mss.gs had no CAA record, so any certificate authority could issue for the domain if it accepted a forged proof of control.

FIX A CAA record now names the one authority in use, on mss.gs and on the 13 alias domains of mssgs. Published and read back through two public resolvers on 28 November.

R-04 TLS 1.2 offered CBC cipher suites

INFO FIXED

OBSERVED Measured on 28 November: of the 13 suites offered on TLS 1.2, six used CBC mode, and the server followed the client's order of preference. No practical attack is known against them as they were configured.

FIX TLS 1.2 now offers three AEAD suites only and the server chooses. Changed on one web address first, measured, then on the other two; all three measured again on 28 November.

R–05 No MTA-STS policy for incoming mail

INFO FIXED

OBSERVED Measured on 28 November: the domain published no MTA-STS policy and no TLS reporting record, so a sending mail server could not be told to insist on encryption.

FIX A policy and a TLS reporting address are published since 28 November. The policy is in testing mode: senders report failures and still deliver. It moves to enforce once the reports show no failures.

R–06 Libraries with published security advisories

MEDIUM FIXED

OBSERVED Call-graph analysis on 28 November found advisories in code the services actually call: one in the chat backend, six in the platform core, and two libraries in the website, which also did not pin the Go version it was built with. A check of the running programs then showed two supporting services on older builds: the voice service on a build from 21 October, and the link-preview service, which had not been part of the first analysis.

FIX All five were updated the same evening and the build toolchains are pinned. The same analysis on the updated source of each finds no advisory in code the services call. The running programs were read back as the updated builds: the platform core, the chat backend, the voice service on both of its nodes and the link-preview service on both gateways. The chat backend now runs this check before every deploy build.

R–07 The zone was not signed with DNSSEC

INFO FIXED

OBSERVED Measured on 28 November: no DNSKEY at mss.gs and no DS at the parent, so resolvers could not validate answers for the domain.

FIX mss.gs is signed since 28 November and every answer validates against the zone's key on all five name servers. The .gs registry published the key the same evening; public resolvers that had picked it up returned validated answers, the others still held their earlier answer and none failed. Ten alias domains of mssgs were signed the same way, and their registries publish their keys.

R–08 **Pages could be shown in a frame on another site** LOW FIXED

OBSERVED Measured on 28 November: pages on mss.gs and app.mss.gs, the sign-in page among them, were served without a Content-Security-Policy and without X-Frame-Options, so another site could show them inside a frame under its own content.

FIX Since 29 November every page carries a Content-Security-Policy with frame-ancestors limited to mssgs's own hosts, base-uri 'self' and object-src 'none', and X-Frame-Options for older browsers. Measured in a browser on 29 November: another site that tried to frame the home page, the sign-in page and the web app was refused on all three, and mssgs's own pages loaded with nothing blocked. The clip player stays embeddable on other sites, on purpose.

PART 6

What this audit did not cover

- No penetration test of the application with user accounts: what one signed-in user can reach of another's data was not tested.
- The app packages for iPhone, iPad, Android and desktop were not analysed as shipped; the encryption was read in their source code.
- Servers and the internal network were assessed through their configuration and from outside, not by an attacker's foothold inside.
- No denial-of-service testing, no social engineering and no physical test of the data centre.
- The secret scan covered the files as they are now, not the full history of the repositories.
- Services of third parties (Apple, Google, the bank's payment page, outgoing mail delivery) are outside the scope.
- An audit describes the state on its dates. Changes after 29 November 2026 are not covered.

Verification

Reference **MV-BCW6E4R8**, version 1.0, 28 November 2026. Murvord confirms this report and the status of its badge at `https://murvord.com/verify/MV-BCW6E4R8`. The reference is derived from the content of the report, so it belongs to this text and to no other.