ATTESTATION LETTER

Third-Party Offensive Security & AI-Agent Boundary Attestation

CLIENT ORGANIZATION	STAGING ENVIRONMENT	ASSESSMENT WINDOW	ISSUED
Meridian Financial Systems, Inc.	`staging.meridian.example`	2 to 13 November 2026	17 November 2026

01 Executive summary

August Cyber carried out an independent offensive security assessment of the Meridian Financial Systems staging environment under a signed authorization. **Four findings were reproduced, fixed by Meridian engineering, and retested against the patched build.** No reproduced finding remains open.

VERDICT
ALL CRITICAL & HIGH FINDINGS REMEDIATED AND RETESTED · 0 OPEN EXPOSURES

02 Scope & methodology

SURFACES TESTED
- S1 Multi-tenant REST & GraphQL APIs
- S2 AI Support Agent & MCP Tool Server
- S3 Session, OAuth & Role Boundaries
- S4 Cloud Storage & Secrets Egress

METHODOLOGY
- M1 Non-refusal agentic boundary probing inside kernel scope-locked August Sandboxes
- M2 Manual Senior Security Engineer verification of every finding
- M3 Byte-exact post-patch retest of every recorded request

03 Findings by severity

SEVERITY	FOUND	REMEDIATED	RETESTED	OPEN	STATUS
Critical	1	1	1	0	✓ Remediated & Retested
High	2	2	2	0	✓ Remediated & Retested
Medium	1	1	1	0	✓ Remediated & Retested
Total	**4**	**4**	**4**	**0**	0 open exposures

04 Attestation

August Cyber attests that this assessment was performed independently and within the authorized scope, that every finding counted above was reproduced by a Senior Security Engineer before it was reported, and that each fix was confirmed by replaying the recorded request against the patched build. This letter describes the in-scope systems during the assessment window only.

ATTESTED BY
Lead Reviewer
August Cyber Offensive Engineering

CONTACT
security@august.so
august.so

VERIFICATION HASH
`sha256:9f4c82b1...e81a`
Evidence record #03, page 2

F-01 CRITICAL REPRODUCED FINDING

VERIFIED · 403 BLOCKED
RETESTED 12 NOV 2026

Cross-Workspace Secret & Payout Read (BOLA)

A signed-in member of one workspace could read another workspace's stored secrets, including its payout credentials, by changing the workspace ID in the URL. The route checked for a session but never checked membership.

`GET /v1/workspaces/{id}/secrets` / CWE-639 / OWASP API1:2023 Broken Object Level Authorization / Patched build **4e1f9a2**

05 HTTP wire reproduction

Actor A belongs to ws_3b91d7. Actor B owns ws_8f2c41.

RECORDED REQUEST · ACTOR A · NON-OWNER — 04 Nov 2026 · 14:12:47 UTC

```
GET /v1/workspaces/ws_8f2c41/secrets HTTP/1.1
Host: staging.meridian.example
Authorization: Bearer eyJhbGciOiJFUzI1NiJ9.[actor-a.redacted]
```

BEFORE FIX · ACTOR A · NON-OWNER

```
HTTP/1.1 200 OK
Content-Type: application/json

{
  "workspace_id": "ws_8f2c41",
  "secrets": {
    "PAYOUT_API_KEY": "[redacted]",
    "LEDGER_DB_URL": "[redacted]"
  }
}
```

AFTER FIX · ACTOR A · NON-OWNER

```
HTTP/1.1 403 Forbidden
Content-Type: application/json

{
  "error": "forbidden"
}
```

Same request bytes, replayed against patched build 4e1f9a2.

OWNER CONTROL · ACTOR B · OWNER

```
HTTP/1.1 200 OK
Content-Type: application/json

{
  "workspace_id": "ws_8f2c41",
  "secrets": {
    "PAYOUT_API_KEY": "[redacted]",
    "LEDGER_DB_URL": "[redacted]"
  }
}
```

06 Minimal source patch

src/routes/workspaces/secrets.ts:44 · @@ -42,3 +42,7 @@

```
42  42  router.get("/v1/workspaces/:workspaceId/secrets", requireSession, async (req, res) => {
43  43    const { workspaceId } = req.params;
44      -   const secrets = await db.secret.findMany({ where: { workspaceId } });
        44 +  const member = await db.membership.findFirst({
        45 +    where: { workspaceId, userId: req.session.userId },
        46 +  });
        47 +  if (!member) return res.status(403).json({ error: "forbidden" });
        48 +  const secrets = await db.secret.findMany({ where: { workspaceId: member.workspaceId } });
```

07 SHA-256 chained evidence log

#01	Scope Verified 2026-11-02T08:30:11Z	prev hash	sha256:00000000...0000 **sha256:41d7a90e...6c2b**
#02	Exposure Reproduced 2026-11-04T14:12:47Z	prev hash	sha256:41d7a90e...6c2b **sha256:b86e03f5...12d9**
#03	Patch Verified 403 2026-11-12T10:05:33Z	prev hash	sha256:b86e03f5...12d9 **sha256:9f4c82b1...e81a**

08 Related finding

F-02 HIGH — VERIFIED · BLOCKED

MCP Tool Scope Escalation

The AI support agent could be steered into calling an MCP tool with another workspace's ID. The tool server now binds every call to the workspace in the session. On retest the same prompt was refused.