Remote access, built to be locked down

How MantisRMM connects to your endpoints, who can use that connection, what limits the damage if an account is ever abused, and how you can see every action afterwards.

1 THE CONNECTION

Outbound only. Each agent opens its own encrypted connection to the platform. The agent does not run a listening service for remote access, so there is no inbound port to scan or attack.

Relay by default. Remote screen and terminal traffic runs through the platform, scoped to your tenant. Direct peer connections are an option, not the default.

Your tenant, isolated. Each customer has its own database, and each device connection is bound to its tenant. An ID from another tenant returns not found.

2 THE DEVICE IS VERIFIED, NOT JUST THE USER

Per-device secret. Every enrolled device holds its own secret. The platform stores only a hash of it, and a wrong secret is refused before it can replace the real connection.

Clone detection. The agent reports a hardware fingerprint. A copy of a device that presents the right secret from different hardware is held: it gets no jobs, no remote desktop and no terminal, and a high-priority security event is raised. Reset key is the approval path.

Require a key. Admins can switch a tenant to refuse devices that have no key.

Temporary enrollment. Installer codes always expire (you choose 1 hour to 30 days) and can be single-use. Each one is exchanged for a one-time key valid for one device and one hour, after which the device holds its own secret; the key is not kept on the device.

3 WHO CAN USE IT

MFA for every password sign-in. No user can opt out. Authenticator app, emailed code or a passkey; admins can require passkeys. Single sign-on counts as the second factor only when your identity provider asserts it.

Step-up on servers. Opening a remote desktop or terminal on a Windows Server, Linux or macOS machine asks for a fresh authenticator code, even inside a signed-in session.

Granular permissions. Terminal, remote desktop, scripts, files and software are separate permissions, enforced at the API rather than hidden in the interface, and can be limited to specific companies, sites or device groups.

IP allowlists. Admins can restrict console sign-in and sessions to approved networks, and can apply the same restriction to agent and probe connections. Saving a list that would lock you out is refused, so is switching on agent or probe restriction while it would cut off agents that are online, and blocked connections are logged.

4 LIMITS ON THE BLAST RADIUS

Bulk actions escalate. Acting on more than 10 devices in a short window requires a fresh authenticator code. Acting on more than 50, or across companies, requires a second administrator to approve. Approvals expire after an hour, and thresholds are tenant settings that need step-up to loosen.

Destructive actions always escalate. Agent removal, removing security software and scripts matching known destructive patterns (shadow-copy deletion, disk wipe, encoded payloads) always need step-up.

Sensitive secrets are gated twice. Revealing a managed local-admin password needs its own permission plus a fresh authenticator code. BitLocker keys need step-up and are redacted from logs.

API keys stay narrow. Keys are shown once and stored hashed, can expire, and can be limited by IP and company. They cannot run scripts, open terminals or start remote sessions.

5 KILL SWITCHES

Pause all dispatch. Any admin can stop all new work for the tenant in one click: queued jobs are cancelled, running ones are told to stop and open terminal and remote desktop sessions are ended. Resuming needs the admin's password and an authenticator code.

Freeze a user. Takes effect immediately, ends that user's live sessions, and unfreezing needs another administrator.

Revoke access. A password or MFA reset, role change or removal signs the user out of every MantisOps product.

6 VISIBLE TO THE PERSON AT THE MACHINE

Always notified. Every remote desktop session shows a notice with the operator's name and organization when it starts and ends. It cannot be turned off; you can brand and reword it.

Windows and Linux also show a persistent session indicator. Windows adds a privacy screen and local-input block, and each use is logged.

7 SEEING WHAT HAPPENED

Every action is logged with who, from where, which devices and a hash of what was sent; remote desktop sessions also record when they ended and for how long. The log is chained so edits are detectable, and it is kept outside your tenant database.

Stream it to your own tools. Webhook, Splunk HEC or generic HTTP, each line signed, so your SIEM holds the retention you choose.

8 SOFTWARE YOU CAN TRUST

Signed updates. The agent installs an update only if its manifest is signed by a release key pinned in the agent and every file matches its size and SHA-256. Windows builds are Authenticode-signed and macOS builds are Apple-notarized.

QUESTIONS TO ASK ANY REMOTE ACCESS VENDOR

- ✓ Does the agent listen for inbound connections, or only connect out?
- ✓ Is each device verified with its own secret, and what happens to a cloned device?
- ✓ Is MFA mandatory for every user, and can you require passkeys for admins?
- ✓ Does opening a terminal on a server ask for a fresh code, even in a signed-in session?
- ✓ Are terminal, remote desktop, scripts and files separate permissions, enforced at the API?
- ✓ What stops one compromised account from running a script on every device at once?
- ✓ Can you stop all remote work for the whole tenant in one step, and who can resume it?
- ✓ Is the person at the machine told when someone connects, and can the vendor switch that off?
- ✓ Can you restrict console and agent connections to your own networks?
- ✓ Can you stream the audit log to your own SIEM, and are the lines signed?
- ✓ How long do install codes and enrollment keys live?
- ✓ How does the agent verify that an update really came from the vendor?

MantisRMM answers each of these on pages 1 and 2.

Ask us to walk your team through it.

https://mantisops.net/security · security@mantisops.net

This brief describes how MantisRMM behaves at the time of writing. Defaults can be changed by tenant administrators, and some controls depend on agent version and platform. It is not a certification, an audit report or a guarantee against misuse. Security questions: security@mantisops.net.