ULTRARED

ULTRA RED MCP: AI-POWERED ATTACK SURFACE OPERATIONS

Closed Beta | Model Context Protocol Integration

What is ULTRA RED MCP?

ULTRA RED MCP is a Model Context Protocol integration that connects AI assistants—Claude, GitHub Copilot, custom agents—directly to ULTRA RED's live attack surface intelligence.

It's an operational bridge that lets you embed ULTRA RED's full intelligence engine into your AI environment. AI assistants retrieve data, correlate findings, execute remediation workflows, and surface critical decisions. Your team makes the calls. The AI does the heavy lifting.

Why MCP Matters Now

The security industry is in the middle of a fundamental shift. Every major vendor—Tenable, Qualys, Rapid7, CrowdStrike—is racing to ship MCP integrations. This is the new standard for AI-powered security operations.

ULTRA RED's MCP is already deployed.

The new model: AI agents that work on your attack surface proactively. They triage findings, surface shadow IT, track remediation progress, and flag what matters. All without you logging in.

The Result

Every security analyst becomes a **10x operator.**

Instead of clicking through dashboards, you're asking natural language questions to an agent that understands your entire attack surface and executes at scale.

Strategic Multiplier

100x more intelligence

at target scope. ULTRA RED's capability—already the industry leader in speed and depth—becomes an embedded intelligence layer in your operations.

Available Tools

ULTRA RED MCP provides 15 tools across the full attack surface management lifecycle. Two categories: **intelligence** and **response.**

Read Operations — Intelligence & Monitoring

TOOL	WHAT IT DOES	OPERATIONAL VALUE
get_targets	List all monitored organizations/ scopes with asset and vector counts	Dashboard-free visibility into your monitored scope
list_assets	List all assets for a target with technologies, tags, sensitivity ratings, vector counts	Fast asset inventory, prioritized by risk
get_asset	Deep asset details: technologies, enrichments, remediation steps, IP/port data	Full context for remediation decisions
get_asset_timeline	Chronological event history for any asset	Track changes, understand discovery/remediation velocity
list_vectors	List security findings (proven vulnerabilities with exploit demonstrations) with status/priority filtering	Triage by severity, see what's verified vs archived
get_vector	Full vector details including exploit HTTP request/response, CVSS, EPSS, tags	Exploit proof for validation and reporting
get_discovery_data	Attack surface discovery: related domains, discovered vs assigned assets	Shadow IT detection—surface unknown infrastructure
get_scanner_proxy_ips	Current scanning infrastructure IPs for firewall allowlisting	Operational—keep your WAF/ firewall aligned

Write Operations — Response & Management

TOOL	WHAT IT DOES	OPERATIONAL VALUE
add_asset	Add new IPs, domains, FQDNs to monitoring	Onboard discovered assets without leaving your AI environment
add_asset_enrichment	Add contextual metadata observations to assets	Tag assets with business context, owner, risk notes
add_asset_tags / remove_asset_tags	Tag-based asset grouping and filtering	Organize assets by team, environment, risk tier
set_asset_monitoring	Enable/disable continuous scanning per asset	Control scanning cadence based on business need
update_vector_status	Move vectors through remediation lifecycle *(verified → archived → accepted_risk)*	Track remediation progress, record risk decisions
create_vector	Inject findings from external tools into ULTRA RED	Unify security data—integrate Burp, Nessus, custom tools into one platform

Use Cases

Automated Triage

AI agent monitors new vectors, prioritizes by Priority/Sensitivity/EPSS and cross-correlates with your context, flags critical findings to Slack. Your team wakes up to a highly contextualized, prioritized summary.

Attack Surface Audit

"Show me all assets with critical vectors that haven't been scanned in 30 days." Natural language query. Instant answer. No manual spreadsheet work.

Shadow IT Detection

Compare discovered vs assigned assets. Surface unknown infrastructure before it becomes a breach vector. ULTRA RED finds it; MCP flags it for you.

Remediation Tracking

AI agent tracks vector status changes, reports on MTTR, identifies stale findings. Know exactly where remediation stands, who owns what, and what's overdue.

Custom Security Workflows

Build org-specific agents that combine ULTRA RED data with Jira, ServiceNow, Slack, email, or internal tools. Your security process, automated.

Executive Reporting

Natural language queries to your attack surface data. Instant posture summaries, trend analysis, risk narratives—no manual report generation.

Pricing

ULTRA RED MCP is available as an add-on to existing ULTRA RED subscriptions.

Closed Beta	Production Pricing
Complimentary access for select customers. Includes full MCP tool access, dedicated onboarding, and direct feedback channel to the product team.	TBD

Closed Beta Invitation

We're opening limited seats in the ULTRA RED MCP closed beta.

What You Get

- Full access to all 15 MCP tools
- Dedicated onboarding and configuration support
- Direct feedback channel to the ULTRA RED product team
- First-mover advantage as MCP becomes the standard

Requirements

- Active ULTRA RED subscription
- AI assistant environment: Claude Desktop, GitHub, Copilot, or custom agent compatible with MCP

How to Join

Contact your ULTRA RED account manager or reply to this document with your interest. Limited seats available.