MCP SERVER | NO CODE | CLOUD HOSTED

Care Communication & Consent Map MCP

Enforce strict healthcare privacy and consent rules automatically.

Care Communication & Consent Map MCP acts as a decision engine for managing sensitive healthcare data disclosures. It uses least-disclosure logic to determine exactly what information can be shared with specific recipients, ensuring your communications stay within legal and ethical boundaries. You can connect it to your AI client to handle complex permission mapping and compliance monitoring.

A+ Quality Score 100/100

consent | data-privacy | healthcare-tech | compliance-engine | permission-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026

vinkius.com/en/ai-agent-connect/care-communication-consent-map

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

> *The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.*
>
> — Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT

```
https://edge.vinkius.com/{token}/mcp
```

Claude · Cursor · VS Code · Windsurf · Grok · Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1) Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

> Get started at **cloud.vinkius.com** — connect your AI agent in under 60 seconds.

Care Communication & Consent Map MCP

4 tools available
Cloud-hosted on Vinkius

Navigating the overlap between patient privacy and the need for relational transparency is difficult. This MCP handles that complexity by applying strict consent precedence and least-disclosure rules. Instead of guessing who can see what, you let your agent check the specific permissions tied to a patient or subject.

You can use it to map out exactly which communication methods are allowed for specific people or to pull approved messaging templates that follow role-based guidelines. It also keeps you on track by identifying privacy barriers and calculating when you need to perform your next compliance review. It essentially acts as a guardrail for your AI client, making sure that every piece of information sent or received respects the established consent framework.

Core Capabilities

01 — Permission Mapping

Your agent uses this to find out who can receive specific data and how.

02 — Restriction Checking

The AI checks for privacy barriers before suggesting any communication.

03 — Template Generation

Your agent pulls pre-approved text to ensure messaging stays within role-based guidelines.

04 — Compliance Scheduling

The tool calculates when the next review is due to maintain regulatory standards.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/care-communication-consent-map** — connect your AI agent in three steps.

01 Connect the MCP to your preferred client like Claude or Cursor via Vinkius.

02 Ask your agent to check permissions for a specific patient or recipient.

03 The agent queries the communication matrix to find allowed data and methods.

04 The agent uses restriction flags to confirm no privacy rules are being broken.

05 The agent generates or sends the approved communication using the provided templates.

Connecting this MCP to your AI client gives your agent immediate access to your consent and communication rules.

Built For

This MCP is built for professionals managing sensitive healthcare data who need to automate consent verification.

Healthcare Administrators

They hand off the task of verifying patient consent and communication permissions to the agent.

Compliance Officers

They use the tool to monitor review schedules and ensure data disclosure stays within legal limits.

Care Coordinators

They use the approved templates to communicate updates to caregivers and family members.

What Changes When You Connect

01 Enforces least-disclosure rules to prevent unauthorized data sharing.

02	Uses role-based templates to keep messaging consistent and compliant.
03	Automates the calculation of required compliance review dates.
04	Identifies privacy restrictions before any communication is sent.

Real-World Applications

Verifying Recipient Permissions

Check if a specific family member or contact has the legal right to see a patient's medical updates.

Standardizing Care Updates

Use approved templates to send consistent, role-appropriate messages to primary caregivers.

Compliance Auditing

Monitor the review schedule to ensure all consent permissions are up to date.

Preventing Data Leaks

Identify restriction flags to stop unauthorized information from being shared with non-authorized parties.

Patterns to Avoid

Manual permission guessing

AVOID

A coordinator assumes a family member has access to medical updates because they are listed in the contact directory.

✓ INSTEAD

Always use `query_communication_matrix` to confirm exactly what information a specific person is allowed to receive.

Unstructured messaging

AVOID

An agent drafts a custom update for a caregiver that includes sensitive details not permitted for that specific role.

✓ INSTEAD

Use `get_update_templates` to pull pre-approved text that fits the recipient's role and permission level.

Ignoring privacy blocks

AVOID

Sending a patient update without checking if there are active restrictions on the subject.

✓ INSTEAD

Run `check_restriction_flags` first to see if the communication is blocked before you attempt to send it.

The Right Fit

Connect this MCP if your team manages sensitive healthcare data and needs to automate consent verification. Do not connect it if your data environment is public or lacks strict role-based privacy requirements. The decision rests on whether you need to enforce least-disclosure rules during patient communication.

Care Communication & Consent Map Toolset

These tools manage the intersection of patient privacy and active communication. They handle everything from permission lookups to compliance scheduling.

#	TOOL	DESCRIPTION
01	`calculate_review_schedule`	This tool determines the date for the next required compliance review.
02	`check_restriction_flags`	This tool checks if a specific communication is blocked based on the subject and the recipient.
03	`get_update_templates`	This tool provides approved text templates for sending specific information to different roles.
04	`query_communication_matrix`	This tool identifies exactly what information can be shared with a person and which methods are permitted.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U Who is allowed to receive medical information for patient ID 123?

Based on current consents, the following individuals are allowed to receive medical information: Legal Guardian (ID: 456) and Primary Caregiver (ID: 789).

U Generate a message for the Primary Caregiver about the patient's recent checkup.

Dear Primary Caregiver, this is an update regarding the patient's recent checkup. The results indicate that all vitals are within normal ranges.

U Is the neighbor allowed to receive emergency contact details?

No, the neighbor (ID: 001) does not have permission to receive emergency contact details under the current consent configuration.

Frequently Asked Questions

01 **What does this MCP do?**

It acts as a decision engine for managing healthcare privacy and communication permissions. It helps determine what information can be shared with specific people based on consent rules.

02 **How does it handle privacy?**

The MCP uses least-disclosure rules and explicit consent precedence. It checks restriction flags to ensure no unauthorized data is shared.

03 **Which AI clients can use this?**

You can use this with any MCP-compatible client, including Claude, Cursor, Windsurf, and VS Code.

04 **Can it help with messaging?**

Yes, it provides approved text templates for communicating specific information to different roles.

05 **Does it manage compliance?**

It helps with compliance by calculating required review schedules to ensure permissions remain valid.

06 **How does the system handle conflicting permissions?**

The system follows the Explicit Consent Precedence rule: if a person has explicitly denied permission for a specific category, that denial overrides any general permission.

07 **What is the 'Least-Disclosure Rule'?**

In the absence of explicit permission, the system defaults to the most restrictive state. Information is only shared if there is a clear, valid permission associated with the recipient.

08 **Can I check if a specific person is blocked from receiving data?**

Yes, you can use the `check_restriction_flags` tool to identify privacy barriers or 'No-Go' zones for a specific subject and recipient.

Go Live in 60 Seconds

Get your connection token from **cloud.vinkius.com**, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

```
https://edge.vinkius.com/[TOKEN]/mcp
```

CLIENT	WHERE TO CONFIGURE
Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add `"care-communication-consent-map": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → Add endpoint URL
ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

- Ask ChatGPT
- Ask Claude
- Ask Perplexity
- Ask Gemini
- Ask Grok

READY TO CONNECT

Care Communication & Consent Map is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Care Communication & Consent Map. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	September 2026
MCP Server	Care Communication & Consent Map MCP
Server ID	`01a0d525-21a0-7048-81af-889d3785d34d`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/en/ai-agent-connect/care-communication-consent-map.