

M C P S E R V E R

N O C O D E

C L O U D H O S T E D

Local Storage Upgrade Planner MCP

Calculate exact capacity needs and hardware costs in seconds.

Local Storage Upgrade Planner MCP handles the math for your next infrastructure expansion. It calculates required capacity based on your current usage, growth rates, and backup needs. You can then test different hardware configurations to find the best balance between cost and headroom without manual spreadsheets.

A+ Quality Score 100/100

storage

capacity-planning

hardware

backup

cost-analysis



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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

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

Firewall & DLP

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

Agent Activity

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

Tool Health

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

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.

Local Storage Upgrade Planner MCP

4 tools available

Cloud-hosted on Vinkius

Planning a storage upgrade usually involves messy spreadsheets and constant guesswork regarding growth rates and backup overhead. This MCP changes that by giving your AI client the ability to run precise capacity models. You can input your current data footprint, projected annual growth, and required safety buffers to get a clear picture of your future needs.

Once you know your target capacity, you can use the tool to vet specific hardware options. Instead of wondering if a certain drive array is enough, you can have your agent check it against your requirements. It helps you avoid the two biggest mistakes in infrastructure planning: buying too little and running out of space too soon, or overspending on massive arrays you won't fully utilize for years. You get clear headroom analysis and cost-effectiveness rankings to justify your hardware decisions.

Core Capabilities

01 — Capacity Modeling

Your agent calculates total raw capacity by accounting for growth and safety buffers.

02 — Hardware Matching

The AI compares your required capacity against available hardware drives.

03 — Cost Optimization

Your agent ranks different hardware setups by cost-effectiveness.

04 — Headroom Calculation

The tool identifies the exact amount of surplus capacity in a specific setup.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/local-storage-upgrade-planner — connect your AI agent in three steps.

- 01 Connect the MCP to your client like Claude or Cursor via Vinkius.
- 02 Provide your current usage and growth parameters to your agent.
- 03 Input available hardware options for comparison.
- 04 Review the ranked configurations and headroom analysis provided by the agent.

Get your storage math done by connecting this MCP to your preferred AI client.

Built For

This MCP is built for technical professionals managing data infrastructure and hardware procurement.

Systems Administrator

Hands off growth projections and current usage to the agent for capacity planning.

IT Manager

Uses the tool to compare hardware costs and justify budget requests.

Data Center Engineer

Runs headroom analysis to ensure long-term system stability.

What Changes When You Connect

- 01 Eliminates manual capacity math using growth and buffer variables.
- 02 Identifies the most cost-effective hardware configurations automatically.

- 03

Provides clear headroom metrics to prevent premature storage exhaustion.
- 04

Reduces procurement errors by validating hardware against actual needs.

Real-World Applications

Annual Capacity Planning

Project future storage needs based on current usage and annual growth rates.

Hardware Procurement

Compare different drive options to find the best price-to-capacity ratio.

Backup Strategy Modeling

Account for backup overhead when calculating total required raw capacity.

Infrastructure Auditing

Check existing hardware setups to see how much headroom remains before an upgrade is required.

Patterns to Avoid

Guessing growth rates

❌ AVOID

You tell your agent you need 10TB of storage because it feels right for next year.

✓ INSTEAD

Input your actual current footprint and projected annual growth into `calculate_storage_needs` to get a math-based requirement.

Ignoring backup overhead

❌ AVOID

You plan a new array based only on your primary data size and forget about the extra space needed for backups.

✓ INSTEAD

Use `calculate_storage_needs` to factor in backup requirements so you don't run out of space immediately after deployment.

Buying based on lowest price alone

❌ AVOID

You pick the cheapest drive array available without checking if it actually meets your long-term capacity needs.

✓ INSTEAD

Run `rank_configurations` to see which hardware options provide the best balance of utility and cost for your specific needs.

The Right Fit

Systems administrators and IT managers should connect this MCP if they manage physical or virtual storage infrastructure. If you are just a casual user with a single external hard drive, you don't need this. The decision depends on whether you need to justify hardware spend with hard data rather than gut feeling.

4 Tools for Precision Storage Planning

These tools handle the heavy lifting of capacity modeling, hardware vetting, and cost analysis.

#	TOOL	DESCRIPTION
01	calculate_storage_needs	This tool determines the total raw capacity you need by factoring in growth, safety buffers, and backup requirements.
02	evaluate_configurations	Use this to match your required capacity against specific hardware drive options to see which setups actually work.
03	get_headroom_analysis	This tool calculates exactly how much surplus space remains in a chosen hardware configuration.
04	rank_configurations	This tool sorts your valid hardware options by how much utility and cost-effectiveness they provide.

See It in Action

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

U I have 500GB used, a 10% annual growth rate, want 20% free space, and a 1.5x backup overhead. How much total capacity do I need?



Based on your parameters, the total required capacity is 750GB.

U I need 1000GB of storage. I have two options: a 1TB drive for \$100 or a 2TB drive for \$180. Which is better?



The 1TB drive is the most cost-effective option at \$100, as it meets your minimum requirement of 1000GB.

U If I buy a 4TB drive but only need 3TB, how much extra space will I have?



You will have 1TB of headroom, which represents a 33.3% excess capacity.

Frequently Asked Questions

01 How does this MCP calculate storage needs?

It uses your current usage, projected growth rates, safety buffers, and backup overhead to determine the total raw capacity required.

02 Can I use this to compare different hard drives?

Yes, you can provide a list of available hardware drives and the agent will match them against your requirements and rank them by cost-effectiveness.

03 What is headroom analysis?

Headroom analysis calculates the specific surplus capacity you will have left over after choosing a specific hardware setup.

04 Which AI clients can use this MCP?

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

05 Does this help with budget planning?

Yes, by ranking hardware configurations based on cost-effectiveness, it helps you identify the most economical way to meet your capacity needs.

06 How do I calculate my total storage needs?

Use the ``calculate_storage_needs`` tool. You will need to provide your current used capacity, expected growth rate, desired free-space ratio, and the backup overhead multiplier.

07 Can I compare different hardware drives?

Yes. After calculating your needs, use ``evaluate_configurations`` with a list of available drives to find valid setups, then use ``rank_configurations`` to find the most cost-effective option.

08 What is headroom in the context of storage planning?

Headroom is the surplus capacity available in a hardware configuration after meeting all requirements. You can analyze this specifically using the ``get_headroom_analysis`` tool.

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 <code>"local-storage-upgrade-planner": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

Local Storage Upgrade Planner 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 Local Storage Upgrade Planner. 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	October 2026
MCP Server	Local Storage Upgrade Planner MCP
Server ID	01a0edde-2c36-7221-a738-a003674f27db
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/local-storage-upgrade-planner.