

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

Backup Storage Capacity Forecast MCP

Stop guessing when your backup storage will run out.

Backup Storage Capacity Forecast MCP gives your AI agent the math needed to manage data growth. It calculates monthly storage footprints, predicts exact exhaustion dates, and determines the specific hardware capacity you need to add to stay ahead of your retention policies.

A+ Quality Score 100/100

backup

capacity-planning

storage-management

data-growth

retention



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.

Backup Storage Capacity Forecast MCP

4 tools available

Cloud-hosted on Vinkius

Managing backup infrastructure shouldn't feel like a guessing game. This MCP gives your AI client the ability to run precise capacity planning models based on your current growth rates and retention schedules. Instead of manually calculating spreadsheets, you can ask your agent to project your storage footprint months or years into the future. It handles the heavy lifting of checking if your retention policies actually align with your version counts, ensuring your math holds up before you commit to new hardware. You can use it to pinpoint exactly when a storage array will hit its limit or to figure out how much extra capacity is required to hit a specific target date. It turns raw growth data into actionable hardware procurement plans.

Core Capabilities

01 — Growth Projections

Your agent calculates monthly storage footprints over a defined horizon.

02 — Exhaustion Prediction

The AI identifies the exact date your storage capacity hits zero.

03 — Hardware Planning

Your agent determines the specific capacity needed to meet future targets.

04 — Policy Auditing

The tool verifies that retention schedules and version counts are mathematically sound.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/backup-storage-capacity-forecast — connect your AI agent in three steps.

- 01 Connect your preferred client like Claude or Cursor to Vinkius.
- 02 Provide your current storage stats and growth rates to your agent.
- 03 Ask the agent to run specific tools like exhaustion or expansion calculations.
- 04 Receive precise capacity forecasts and hardware requirements in your chat.

Connect your AI client to Vinkius and start running capacity models immediately.

Built For

This MCP is built for infrastructure teams managing large scale data backups and storage arrays.

Storage Administrator

Hands off capacity planning and growth modeling to the agent.

Backup Engineer

Uses the agent to validate retention policies and predict exhaustion.

IT Infrastructure Manager

Asks the agent for hardware expansion requirements to build budget requests.

What Changes When You Connect

- 01 Eliminates manual spreadsheet calculations for storage growth.
- 02 Provides exact dates for storage exhaustion to prevent emergency upgrades.
- 03 Matches hardware procurement to actual mathematical needs.

- 04 Validates that retention policies won't cause unexpected storage spikes.
-

Real-World Applications

Budget Forecasting

Ask your agent how much storage you'll need in two years to justify a hardware purchase.

Capacity Auditing

Check if your current retention settings are actually compatible with your storage limits.

Growth Monitoring

Track how monthly growth rates impact your long term storage availability.

Risk Mitigation

Identify the exact month you will run out of space to avoid backup failures.

Patterns to Avoid

Manual spreadsheet guessing

❌ AVOID

You try to manually calculate next year's storage needs by looking at last month's growth in a spreadsheet. This ignores how retention policies change your footprint over time.

✓ INSTEAD

Let your agent handle the math. Use ``get_monthly_usage_forecast`` to see how your footprint evolves month by month based on real data.

Reactive hardware buying

❌ AVOID

You wait until a storage array hits 95% capacity before you start looking for new hardware. This leads to emergency procurement and potential backup failures.

✓ INSTEAD

Stay ahead of the curve. Use ``calculate_exhaustion_date`` to find the exact day you'll run out of space so you can plan upgrades months in advance.

Ignoring policy math

❌ AVOID

You assume your retention settings are fine because they worked last year, even though your version counts are climbing. This causes unexpected storage spikes.

✓ INSTEAD

Verify your logic first. Use ``validate_retention_policy`` to ensure your retention schedules and version counts are actually mathematically consistent.

The Right Fit

Infrastructure teams managing large scale data backups should connect this MCP. If you are just managing a single laptop or a small local drive, you don't need this. The deciding factor is whether your storage growth is predictable enough to require mathematical forecasting rather than just watching a dashboard.



Backup Storage Capacity Forecast 4 Tools

These tools handle the math for storage growth, exhaustion dates, and policy validation.

#	TOOL	DESCRIPTION
01	estimate_expansion_requirement	This tool calculates the exact amount of additional storage you need to reach a specific capacity target by a certain date.
02	get_monthly_usage_forecast	Use this to project your storage footprint month by month over any timeframe you specify.
03	validate_retention_policy	This tool checks that your retention schedules and version counts are mathematically consistent.
04	calculate_exhaustion_date	This tool identifies the specific date when your current backup storage will run out of space.

See It in Action

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

U When will my 100TB backup storage run out if I grow at 5% monthly with 10 versions and 12 months retention?



Based on a 5% monthly growth rate, your storage capacity will be exhausted in 24 months.

U How much extra storage do I need for month 36?



To satisfy requirements for month 36, you will need an additional 450TB of capacity.

U Show me the monthly storage usage for the next 12 months.



Month 1: 105TB, Month 2: 110.25TB, Month 3: 115.76TB, Month 4: 121.55TB, Month 5: 127.63TB, Month 6: 134.01TB, Month 7: 140.71TB, Month 8: 147.75TB, Month 9: 155.13TS, Month 10: 162.89TB, Month 11: 171.03TB, Month 12: 179.58TB.

Frequently Asked Questions

01 What can this MCP do for my backup strategy?

It predicts when your storage will run out, calculates monthly growth, and determines how much new hardware you need to buy.

02 How does it handle retention policies?

It uses a validation tool to ensure your retention schedules and version counts are mathematically consistent.

03 Can I use this with Claude or Cursor?

Yes, this MCP works with any MCP-compatible client including Claude, Cursor, and Windsurf.

04 Do I need to host this myself?

No, Vinkius hosts and manages the MCP so it is ready to use as soon as you connect.

05 How accurate are the expansion estimates?

The estimates are based on the specific growth rates and target dates you provide to the tools.

06 How can I predict when my storage will be full?

You can use the `calculate\_exhaustion\_date` tool to determine the specific month when your projected data volume will exceed your total capacity.

07 Can I plan for future hardware purchases?

Yes, the `estimate\_expansion\_requirement` tool calculates exactly how much additional capacity you need to meet a specific target date.

08 How does compression affect my forecast?

The compression ratio is a key input for all forecasting tools. It reduces the physical footprint of the raw data, allowing for more accurate capacity planning.

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>"backup-storage-capacity-forecast": { "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

Backup Storage Capacity Forecast 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 Backup Storage Capacity Forecast. 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	Backup Storage Capacity Forecast MCP
Server ID	01a0ed98-0764-7326-8d96-666ad8710ba9
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/backup-storage-capacity-forecast.