MCP SERVER | NO CODE | CLOUD HOSTED

Kaggle Market Intelligence MCP

Track pain points and trends across global data science forums.

Kaggle Market Intelligence connects your agent directly to Kaggle's entire ecosystem. It lets you scan for trending datasets, audit competition discussions, and find specific technical pain points across massive data science communities. Use it to track what developers are struggling with or what models are gaining traction in real-time.

A+ Quality Score 100/100

kaggle | growth-hacking | community-management | data-science | reconnaissance

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/kaggle-market-intelligence

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

Kaggle Market Intelligence MCP

10 tools available
Cloud-hosted on Vinkius

This MCP gives your AI agent the ability to manage all your interactions within the Kaggle platform. Instead of manually navigating dozens of forums, searching through datasets, and reading hundreds of comments, you can ask your agent to intercept key conversations. Your agent scans trending competitions for keywords like 'error' or 'missing,' flagging exactly where a data scientist is stuck. Need to see what people are building? You can pull code from notebooks or push your own analysis back out to the community. This capability turns complex reconnaissance into a simple conversation with your AI client, letting you act as an instant growth hacker. When you connect this toolset via Vinkius, you gain immediate access to deep market intelligence on ML models and data prep pain points across the entire catalog.

By using it, you can search for niche datasets or track specific model architectures being deployed, giving you a rapid understanding of where the community focus lies.

Core Capabilities

01 — Find Data Sources

Search and list available data sets, models, and entire competitions within Kaggle.

02 — Track Code and Results

Pull code from existing notebooks or get the current leaderboard status for a competition.

03 — Engage Community Discussions

Search discussion threads, read comments, and automatically post technical replies to guide users.

04 — Update Content

Create new datasets or push your own analyzed code back into the community for visibility.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/kaggle-market-intelligence** — connect your AI agent in three steps.

01 Subscribe to this Vinkius integration and enter your Kaggle API token (your username and key).

02 Connect your preferred AI client or compatible agent to access the full suite of tools.

03 Ask your agent specific questions, like 'Find all datasets discussing data prep errors,' and get actionable intelligence back.

The bottom line is you stop searching Kaggle manually; you just ask your agent what it finds.

Built For

This MCP is essential for anyone embedded in the data science or AI space. It helps founders and technical marketing leads track community pain points, allowing them to strategically introduce their product as the solution.

Developer Advocate

Using this MCP, you find discussions related to your specific tech stack so you can naturally engage with developers who are struggling and guide them toward your tools.

Growth Lead

You run rapid audits of trending datasets and competitions to spot community gaps, then deploy content or suggestions that increase platform usage.

Technical Sales Engineer

When faced with a new client problem, you use the agent to perform quick searches on Kaggle models and data sets to understand the current technical landscape of their industry.

What Changes When You Connect

01 Find immediate technical gaps: Use the agent to search datasets for keywords like 'error' or 'missing.' This lets you pinpoint exactly where users are running into problems, giving you a clear talking point for your product.

02 Understand competitive landscape: By checking the leaderboard for a competition or searching models, you see who is winning and what architectures are currently popular in the data science community.

03 Stay current on best practices: You can pull code from existing notebooks to quickly read developer strategies. This helps you understand how others solve problems before you build your own solution.

04 Directly influence conversations: If you find a discussion thread where users need help, your agent posts technical replies automatically, positioning your infrastructure as the expert solution right where they are looking for it.

05 Maintain visibility: Use the push_notebook tool to share your latest analysis or code directly back onto Kaggle. This keeps you visible in niche communities and establishes thought leadership.

Real-World Applications

Identifying a new enterprise pain point

A sales engineer needs proof that data teams struggle with cleaning messy inputs. They use the agent to search datasets for 'missing' or 'error.' The agent reports 3 active discussions mentioning these issues, allowing the engineer to schedule a demo focused on your pre-processing layer.

Launching a new feature set

A developer advocate wants to prove their platform solves model deployment complexity. They use the agent to search notebooks and pull code from winning examples, identifying common gaps in current ML workflows they can target with an update.

Monitoring a competitor's success

A product manager needs to know if a rival is gaining traction. They use the agent to search models and get the competition leaderboard, seeing that a specific architecture has recently been deployed by multiple users, signaling market interest.

Quickly validating data requirements

A founder needs to know if enough clean data exists for their next product iteration. They use search_datasets to verify the availability of niche datasets and list_dataset_files to check the schema, confirming viability before spending engineering time.

Patterns to Avoid

Relying only on public APIs

✗ AVOID

Manually logging into Kaggle, searching multiple tabs (datasets, forums, notebooks), and trying to cross-reference keywords across disparate pages is too slow for real-time insights.

✓ INSTEAD

Use the agent's search_datasets and search_notebooks tools together. The agent pulls data from both sources simultaneously, giving you a consolidated view of where the technical need intersects with available information.

Ignoring thread context

✗ AVOID

Just pulling random discussion threads doesn't tell you if the problem is solvable or just theoretical. You might spend time on outdated discussions.

✓ INSTEAD

Always filter your searches by checking the most recent activity, and use the agent to get_competition_leaderboard to ensure the pain point relates to a current, active challenge.

Assuming data quality

✗ AVOID

Quickly grabbing a dataset and assuming it’s ready for production without checking its structure or origin.

✓ INSTEAD

First, use list_dataset_files to audit the schema. Then, cross-reference those files with search_models results to see if existing community models were built using that specific data set.

The Right Fit

Use this MCP if your core activity involves deep competitive intelligence, pattern detection, and direct engagement within technical communities like Kaggle. You need an agent to read the room—to know *where* people are asking questions and *what* they are struggling with right now. Don't use it if you simply need to download a file or check basic API status; for that, standard data connectors suffice. However, if your goal is understanding market sentiment based on technical failures (e.g., 'Why is the community frustrated?') or tracking developer trends across multiple content types, this MCP is required. It gives you more than just raw data; it provides a structured map of pain points.

The Pain of Manual Market Research

Today, understanding the pulse of the data science community means jumping between tabs: checking datasets for keywords, scrolling through dozens of discussion forums to find user errors, and manually cross-referencing notebooks to see what solutions people are trying. This process is time-consuming, requires constant context switching, and often leads you straight into information overload.

With this MCP, your agent handles the legwork. You ask it to intercept all discussions related to data preparation failures or model performance issues in a specific competition. Instead of wading through endless pages of text, you get an actionable intelligence report, telling you exactly where to focus your efforts.

Kaggle Market Intelligence: Actionable Community Insights

The MCP eliminates the need for repetitive searches across different types of content. You don't have to search datasets, then switch over to searching notebooks, and then check discussions separately. The agent combines these streams automatically.

Now, you receive a single output detailing both the problem (a user discussing 'missing data') and the potential solution (a trending notebook that uses an alternative technique). You don't just get information; you get strategic direction.

Kaggle Market Intelligence MCP (10 Tools)

These tools give your agent the ability to search, retrieve, and post content across all major sections of the Kaggle platform.

#	TOOL	DESCRIPTION
01	`list_dataset_files`	Shows what files are contained within a specific Kaggle dataset.
02	`create_dataset`	Allows you to upload and share your own synthesized or processed data back with the community on Kaggle.
03	`get_competition_leaderboard`	Retrieves the ranking and scores for participants in an active competition.
04	`get_notebook_status`	Checks if a running Kaggle notebook has finished its execution status.
05	`pull_notebook`	Downloads the actual code written in an existing Kaggle notebook so you can see how it works.
06	`push_notebook`	Shares a new script or analysis back to Kaggle, making it visible to other users.
07	`search_competitions`	Finds active and past competitions on Kaggle based on criteria you provide.
08	`search_datasets`	Searches the entire library of datasets to find specific data points or topics.
09	`search_models`	Finds machine learning model architectures and tracks their creators on Kaggle.
10	`search_notebooks`	Locates code examples, data explorations, or winning strategies written in Jupyter notebooks.

See It in Action

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

U Scan the trending competition for 'error' or 'missing'.

I scanned 'titanic' and found 3 active discussions mentioning 'error'. Discussion #142 asks for a faster data prep alternative. Should I prepare a reply suggesting our Vinkius infrastructure?

U Get an intelligence report on trending datasets.

Here is the intelligence report for trending datasets: 'new-cool-dataset' has gained 5,000 upvotes this week. There are currently 12 open discussions asking for help. I can help you intercept them.

U Reply to discussion #42 suggesting our platform.

I've successfully posted your comment to discussion #42. The community will now see your technical breakdown of the infrastructure solution.

Frequently Asked Questions

01 **How does Kaggle Market Intelligence help with competitive analysis?**

It lets your agent search models and retrieve the competition leaderboard to show you which architectures are currently winning and gaining traction among data scientists. This is key for understanding market adoption.

02 **Can I use this MCP to find specific bugs or errors?**

Yes, you can scan datasets and discussions using the agent to search for keywords like 'error' or 'missing.' The system will report active threads where these technical issues are being discussed.

03 **Does Kaggle Market Intelligence only read data?**

No. You can also actively engage by having your agent post replies to discussions using the push_notebook or other community engagement tools, making you part of the conversation.

04 **Is this MCP better than just using Kaggle's built-in search?**

Yes. The Vinkius integration wraps multiple searches into one command. You don't have to search datasets, then separately search notebooks; the agent combines all that intelligence for you.

05 **What if I want to share my own cleaned data?**

You can use the create_dataset tool to upload your processed findings. This makes your unique dataset available on Kaggle and helps build your profile as a knowledgeable contributor.

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 `"kaggle-market-intelligence": { "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

Kaggle Market Intelligence 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 Kaggle Market Intelligence. 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	June 2026
MCP Server	Kaggle Market Intelligence MCP
Server ID	`019eede0-16b1-710e-ba0a-baa7081f67fa`
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/mcp/kaggle-market-intelligence.