

MCP SERVER

NO CODE

CLOUD HOSTED

Amazon SQS Queue MCP for AI Agents

Reliably Managing Asynchronous Task Processing in Cloud Environments

Amazon SQS Queue MCP connects your AI agent to a single Amazon Simple Queue Service queue. This allows you to reliably pull tasks and acknowledge their completion, treating the system like a dedicated background worker. It ensures your processes run asynchronously and safely without needing complex AWS permissions.

F Quality Score 3.6/100

message-queue

aws

async-processing

task-queue

background-worker

security-scoping



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.

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.

Amazon SQS Queue MCP

3 tools available

Cloud-hosted on Vinkius

This MCP gives your AI client one specific job: managing message flow from a single SQS queue. Forget juggling global AWS credentials or worrying about scope creep. Your agent gets surgical access, letting it pull tasks through the queue and confirm they're done processing.

Think of it as setting up a highly reliable, dedicated background worker that processes items one by one. Whether you're running image resizing jobs or complex data transformations, your AI can handle the load without ever peeking into other queues—it's strictly contained. This makes it perfect for building scalable, fault-tolerant automation.

If your current setup involves manual polling scripts or managing overly broad cloud permissions, this MCP fixes that. It turns your agent into an asynchronous worker capable of chewing through millions of queued items safely and efficiently. Vinkius hosts this MCP, giving you instant access to professional message queue management for any compatible AI client.

Core Capabilities

01 — Receive Tasks from the Queue

Your agent fetches batches of pending messages from the designated SQS queue.

02 — Dispatch New Tasks to the Queue

You can write new payloads and send them into the queue for later processing by your worker agent.

03 — Acknowledge Message Completion

The agent deletes a message after successful processing, ensuring it's never processed twice.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/amazon-sqs-queue — connect your AI agent in three steps.

- 01 You instruct your AI client to check for pending tasks using the MCP.
- 02 Your agent polls the SQS queue, receiving a batch of messages that require processing.
- 03 After your custom logic completes work on each message, you tell the agent to delete the message from the queue.

The bottom line is: it lets your agent reliably pull tasks in batches, execute code against them, and then confirm deletion so they are permanently removed from the system.

Built For

This MCP is essential for backend engineers, DevOps teams, and data architects who deal with high-volume, asynchronous tasks. If your application relies on background jobs, message queues, or scalable worker processes, you need this. It solves the pain point of brittle infrastructure and overly permissive cloud access.

Backend Engineer

Uses this to implement reliable job processors that consume payloads (like image processing requests) from a queue rather than handling them synchronously.

DevOps Architect

Configures and tests the messaging pipeline, ensuring tasks are processed exactly once and that cleanup operations run correctly after deployment.

Data Engineer

Builds data ingestion pipelines where incoming records (like CSV uploads or sensor readings) must be reliably queued before transformation logic runs.

What Changes When You Connect

- 01 **Guaranteed processing:** The agent handles message deletion using the `delete_message` tool, ensuring that successfully processed tasks are removed from the queue and never re-run.

-
- 02 Safe operations: By strictly scoping access to one single queue, your AI client cannot interfere with other cloud workloads or queues in your AWS environment.

 - 03 High throughput: You can use `receive_messages` to process large batches of data efficiently, letting your agent chew through millions of queued tasks without rate limiting concerns.

 - 04 Full lifecycle control: The combination of `send_message` and `receive_messages` gives you total command over the task flow—from creation to final consumption.

 - 05 Simplified worker setup: It instantly converts your AI into a robust background worker, eliminating complex polling logic from your core application code.
-

Real-World Applications

Processing User-Submitted Media

A user uploads 100 video files. Instead of making the agent wait for all processing to finish, you `send_message` a payload for each file into the queue. Your agent then uses `receive_messages` to pull them and process them in parallel until finished.

Batch Reporting Generation

A nightly job needs to generate reports for 50 departments. You `send_message` a record for each department ID. The agent retrieves them using `receive_messages`, executes the report logic, and `deletes_message` the task only after confirmation.

Handling IoT Sensor Readings

Thousands of sensors stream data points minute by minute. The system `sends_message` each reading payload to the queue. An AI client polls with `receive_messages`, processes the data point, and then `deletes_message` it once stored in the database.

Patterns to Avoid

Using direct API calls for everything

X AVOID

Writing custom Python scripts that directly hit SQS APIs. This code is brittle, hard to maintain, and requires managing complex credential files.

✓ INSTEAD

Use your MCP's structured tools. Your agent handles the `receive_messages` cycle safely through your AI client, abstracting away the raw API calls.

Assuming task completion

X AVOID

Processing a message and then just ignoring it, assuming someone else will clean up or that the system knows it's done. This leads to 'poison pills' in your queue.

✓ INSTEAD

Always confirm deletion. After processing data using `receive_messages`, you must use `delete_message` to formally remove the record.

Over-scoping access

X AVOID

Giving an AI agent broad AWS permissions (like full SQS access). If compromised, it could affect dozens of unrelated systems.

✓ INSTEAD

This MCP limits your agent's scope to one specific queue. It only allows the necessary operations, drastically reducing the blast radius if something goes wrong.

The Right Fit

Use this MCP when you need reliable, decoupled background processing for tasks that don't require an immediate response. If your application works by 'send a request and process it later,' this is the right tool. It excels at managing high-volume asynchronous work streams using `receive_messages`, `send_message`, and `delete_message`.

Don't use this MCP if you need real-time, synchronous communication (like calling an external API immediately). For those cases, a direct HTTP connector or function call is better. Also, don't use it just to 'read' messages; you must complete the cycle by deleting them after processing.

Amazon SQS Queue MCP: Solving Asynchronous Task Processing with Message Queues

Right now, when your application has a long-running task—like generating a massive quarterly report or resizing hundreds of photos—you face friction. You either force the user to wait on a spinning loader until the job is done, or you build complex retry logic using multiple services and permissions.

With this MCP, you simply `send_message` a payload to the queue. Your agent picks it up with `receive_messages`, processes the task in the background, and then confirms deletion. The user gets an immediate 'Job Submitted' confirmation, while the heavy lifting happens reliably behind the scenes.

Amazon SQS Queue MCP: Controlling Message Reliability in AWS Architectures

The pain point is guaranteeing that a task runs *exactly once*. If your current system fails after processing but before acknowledging it, the message might get stuck or require manual cleanup. That's fragile.

This MCP enforces reliability by allowing you to control the entire lifecycle. You use `receive_messages` for consumption and `delete_message` when successful. It provides that critical 'commit' step, making your worker processes robust enough for production-grade scale.

Amazon SQS Queue: 3 Tools for Message Flow Control

Use these tools to send new messages, pull existing tasks, and confirm deletion from your dedicated message queue.

#	TOOL	DESCRIPTION
01	<code>delete_message</code>	Removes a successfully processed message from the queue, completing the cycle.
02	<code>receive_messages</code>	Pulls a batch of messages from the SQS queue for your agent to process.
03	<code>send_message</code>	Sends a new payload or task message into the designated SQS queue.

See It in Action

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

U I need to run a batch of reports for last month's sales data. Send them over.



Report Queue Status

The following 5 report IDs are ready to be processed and added to the queue:

REPORT ID	DEPARTMENT	DATE RANGE
R-4921	Marketing	Jan - Mar
R-3012	Sales	Jan - Mar
R-7789	Operations	Jan - Mar
R-5501	HR	Jan - Mar
R-6110	Finance	Jan - Mar

I've sent them. They are queued and ready for the worker agent to pick up.

U Check if there are any pending video encoding jobs right now.



Queue Poll Results

Found 3 messages waiting:

- Video ID: v-901 (Priority: High)
- Video ID: v-245 (Priority: Medium)
- Video ID: v-888 (Priority: Low)

I have pulled them into the workspace. Once you confirm they've finished encoding, I can delete_message all three from the queue for cleanup.

U The video encoding for v-901 is done. Can you clean up?



Cleanup Confirmation

The message associated with `v-901` has been successfully deleted from the SQS Queue. It will no longer appear in queue polls. The processing cycle is complete.

Frequently Asked Questions

01 How does Amazon SQS Queue MCP help me process large numbers of tasks?

It manages task flow by allowing your agent to pull batches of items with `receive_messages`. This is designed for high throughput, letting you handle millions of queued payloads reliably without overwhelming any single connection or service.

02 Do I have to worry about tasks getting processed multiple times?

No. The core function requires the agent to confirm completion using `delete_message`. This ensures that once a task is successfully handled by your logic, it's permanently removed from the queue.

03 What if I need to send tasks into the queue from my application?

You use the dedicated tool for sending messages (`send_message`). This lets you inject new payloads—like user IDs or file names—into the task stream, waiting for your worker agent to pick them up later.

04 Is this MCP safe regarding cloud permissions?

Yes. The design strictly scopes access to one specific queue. It's like giving your AI a key that only opens one door in your entire AWS environment, preventing unintended damage or overreach.

05 Can I use Amazon SQS Queue MCP for data pipelines?







Absolutely. You can build robust data ingestion pipes where incoming raw data is `sent_message` to the queue, processed by your agent, and then deleted after safe storage.

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>"amazon-sqs-queue": { "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

Amazon SQS Queue 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](https://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 Amazon SQS Queue. 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	Amazon SQS Queue MCP
Server ID	019e3863-ae42-72b4-b12c-bcecfb98412d
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/amazon-sqs-queue.