

MCP SERVER

NO CODE

CLOUD HOSTED

Azure Service Bus Queue MCP for AI Agents

Build scalable background workers and reliable message processing systems

The Azure Service Bus Queue MCP lets your AI agents safely connect to a single, specific message queue. It pulls pending tasks and confirms when those tasks are finished, ensuring reliable background processing without ever accessing other system data.

F Quality Score 3.6/100

message-queue

event-driven

task-processing

asynchronous-messaging

cloud-messaging

worker-patterns



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.

Azure Service Bus Queue MCP

2 tools available

Cloud-hosted on Vinkius

Building scalable systems requires handling tasks asynchronously. This MCP gives your agent one surgical capability: pulling jobs from a designated Azure Service Bus Queue and confirming completion. Because access is strictly scoped to this single queue, you can safely build highly robust workers that process millions of queued items without risk.

Your AI client handles the entire lifecycle. It first pulls a message, which temporarily locks it away from other processes. Once your agent finishes processing the payload, it uses a confirmation mechanism to permanently delete the task. This reliable pattern is essential for any background worker consuming event-driven data. Connecting this MCP via Vinkius lets you integrate this robust messaging capability into any existing workflow or development pipeline.

Core Capabilities

01 — Retrieve pending tasks from a specific queue

Your agent pulls one message from the Azure Service Bus Queue, locking it so that no other worker can see or process it until you acknowledge completion.

02 — Confirm and delete processed messages

After processing a task, your agent confirms its success using the unique ID and lock token provided when the message was pulled, permanently removing it from the queue.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/azure-service-bus-queue — connect your AI agent in three steps.

- 01 Your AI client initiates a pull request, retrieving a single pending task that is temporarily hidden from other workers.
- 02 The agent processes the data contained in the message body. This process uses the lock token to ensure the task remains unavailable until you explicitly confirm completion.
- 03 Upon successful processing, the agent sends an acknowledgment with the unique ID and lock token, deleting the task permanently from the queue.

The bottom line is that your AI client manages the entire lifecycle: pull a message, process it safely under a lock, and then confirm deletion in one transaction.

Built For

This MCP is built for backend developers, DevOps engineers, and system architects. You're the person who spends days building reliable systems that handle massive volumes of data without dropping a single task or message. If your application needs to process tasks in the background reliably, this is what you need.

Backend Developer

Uses this MCP to build microservices workers that consume events from queues, ensuring data integrity and reliable task completion.

DevOps Engineer

Configures background job processing pipelines using the guaranteed pull/acknowledge pattern for automated system maintenance or large-scale batch operations.

System Architect

Designs resilient, event-driven architectures by implementing message consumption patterns that prevent data loss across multiple worker nodes.

What Changes When You Connect

-
- 01 Guaranteed data integrity: The pull/acknowledge pattern ensures that a task is only deleted after your agent confirms successful processing.

 - 02 Absolute scoping: Your AI client is locked to one specific queue, eliminating the risk of accidental interaction with unrelated workloads.

 - 03 Reliable worker setup: Instantly turn your AI into an asynchronous background processor capable of handling high volumes of queued tasks.

 - 04 Controlled execution: You manage message visibility using standard Peek-Lock mechanisms, preventing race conditions between multiple agents.

 - 05 Clean implementation: The process uses direct confirmation (`acknowledge_message`) to guarantee permanent removal from the queue.
-

Real-World Applications

Processing high-volume user signups

A company needs to handle thousands of new user account creation events. Instead of overwhelming a database, they use their agent to pull messages from the queue, process the signup data (like sending welcome emails), and then `acknowledge_message` once complete.

Running nightly batch report generation

The system needs to trigger complex reports for different departments. The agent pulls a task containing department ID X, runs the reporting logic, and uses `acknowledge_message` when the report is ready for download.

Handling IoT sensor data streams

A factory receives continuous telemetry from thousands of sensors. The AI client repeatedly calls `pull_message` to grab batches of readings, validates them against thresholds, and then `acknowledges_message` if the batch was successfully logged.

Managing workflow state transitions

When an external system signals that a document is ready for review, the agent pulls the message payload. After running OCR or metadata extraction, it uses `acknowledge_message` to signal the next stage in the business process.

Patterns to Avoid

Assuming messages are processed instantly

✗ AVOID

Just trying to read a message and pretending it's done. This leaves the task visible, potentially causing other workers to fail or duplicate work.

✓ INSTEAD

Always follow the full cycle: Use `pull_message` to lock the job first. Process it. Only then use `acknowledge_message` with both the ID and token to confirm deletion.

Using global queue access

✗ AVOID

Connecting an agent that can see all queues in the system. This is a massive security risk if the worker only needs to handle one specific type of job.

✓ INSTEAD

Use this MCP because its scope is surgically limited, guaranteeing your AI client only interacts with the intended Service Bus Queue.

Forgetting the lock token

✗ AVOID

Calling `acknowledge_message` without providing the unique `lockToken`. The system rejects the request, and you're left with an unconfirmed task that might expire or stall.

✓ INSTEAD

The `pull_message` call returns both the `messageId` and the `lockToken`; you must pass both of these to `acknowledge_message`.

The Right Fit

Use this MCP if your core business logic relies on reliably consuming messages from a dedicated, single-purpose queue. The process must be: 1) Pull (lock), 2) Process, 3) Acknowledge (delete). Don't use it if you need to read or list the contents of multiple queues; this MCP is strictly scoped. If your goal is simply storage access or database writes, a dedicated database connector is

better. However, if your process is purely event-driven and requires guaranteed cleanup after processing, then this MCP provides exactly the secure boundary you need.

Azure Service Bus Queue MCP for AI Agents: Reliable Task Consumption

Today, system tasks often involve manual steps. A developer might have to check a dashboard, see an incoming job ID, manually copy the details, and then run a separate script or endpoint call just to confirm that the task was finished. This process is slow, prone to human error, and requires multiple points of failure.

With this MCP, your agent handles it all in one go. Your agent pulls the pending message using `pull_message`, executes the complex logic, and then sends a single confirmation call via `acknowledge_message`. You get reliable task completion without any manual copy-pasting or external script orchestration.

Azure Service Bus Queue MCP for AI Agents: Secure Messaging Boundaries

The biggest headache in large systems is over-permissioning. Traditional methods often give agents broad access, meaning a compromised worker could potentially read or delete data from unrelated queues and services.

This MCP fixes that by enforcing absolute containment. It's locked down to one queue only. Your agent can process tasks efficiently without ever seeing the scope of other workloads in your cloud environment.

Azure Service Bus Queue: 2 Tools for Message Queue Operations

Use these tools to pull pending tasks from the queue and confirm message deletion after processing.

#	TOOL	DESCRIPTION
01	<code>acknowledge_message</code>	This tool permanently deletes a message from the queue once you provide its unique ID and lock token.
02	<code>pull_message</code>	This tool retrieves one pending task from the configured Azure Service Bus Queue, temporarily locking it until you process it.

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 pull the next available task from the system queue.



Task Retrieval Report

✓ Success: 1 message retrieved.

- **Message ID:** msg-8a3f2c
- **Lock Token:** lck-b9e4d5
- **Payload Summary:** A request to summarize the quarterly sales report for the APAC region.

Please process this, and remember to acknowledge it when you're done!

U I finished summarizing the APAC report (msg-8a3f2c) with token lck-b9e4d5.



Message Confirmation

The message msg-8a3f2c was successfully acknowledged.

- Status: Permanently deleted from Service Bus Queue.

Impact: The task is complete and removed from the active workload.
Next steps: Your agent can now pull another pending task.

U What tasks are waiting? I need to confirm if anything is in the queue.



Queue Status Check

The queue currently holds **42** pending messages. The oldest message was queued at 2026-10-05T10:15:00Z. Processing capacity looks healthy.

Frequently Asked Questions

01 How do I ensure my AI agent doesn't lose a task or message?

You guarantee reliability by using the pull/acknowledge cycle. The agent pulls the message, which locks it away from others. Only when you confirm completion does `acknowledge_message` permanently delete it.

02 Is this MCP safe for connecting to multiple different services?

Yes, it's very secure because its scope is limited to a single queue. This means your agent can only work on the specific tasks defined in that one Service Bus Queue.

03 What if my processing takes longer than the lock timeout?

If your job exceeds the configured lock duration, the message will automatically reappear in the queue. You must ensure your agent's processing time is less than or equal to the lock window.

04 Can I read messages from other queues besides this one?

No. This MCP has absolute containment. It is strictly limited to interacting with the single, designated Service Bus Queue and cannot access any other workloads in your Azure environment.

05 Does using Azure Service Bus Queue MCP require complex coding knowledge?

Not necessarily. Your AI agent handles the workflow logic—pulling, processing, and acknowledging—through natural conversation, making it accessible for sophisticated automation without deep code changes.

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 `"azure-service-bus-queue": {
 "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

Azure Service Bus 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 Azure Service Bus 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	Azure Service Bus Queue MCP
Server ID	019e386a-7aa9-7042-9358-0f5e287efed5
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/azure-service-bus-queue.