Your AI gets the facts. Not the keys.

OneIQ is your infrastructure-to-context engine. We lift your infrastructure details, wherever they live, into a governed context cloud, off production, where any AI you choose can reason over them. AI never gets production credentials or a write operation through OneIQ.

FOR SECURITY TEAMS · OCTOBER 2026

Collect
Infrastructure Pulse

Enrich
Context Libraries

Scope
Agent Zones

Serve
MCP Server

AI needs infrastructure context, not production access.

Which hosts reach end of support next year? Where is capacity paid for and not used? What should we know before the licensing renewal? The answers already exist across your platforms, but pulling them together usually takes exports, spreadsheets and a specialist's week.

AI tools can answer questions like these quickly, which is why requests to connect them to infrastructure are reaching security teams. The decision in front of those teams is how to give the AI the information it needs without also giving it access to production, the credentials that come with that access, or the ability to change anything.

OneIQ's Infrastructure-to-Context (I2C) architecture answers that by separating the information from the systems it describes. OneIQ Pulse uses read-only collection to stream your infrastructure's configuration and performance details into the OneIQ context cloud in near real time, so the context AI reasons over stays current with the environment. Agent Zones, which you define, decide which part of that context a given AI conversation can see. Approved AI tools ask questions of that scoped context in plain language, over the Model Context Protocol (MCP).

What this changes

- Production credentials stay separated from AI and under your control. AI clients authenticate to OneIQ with their own, separate credentials.
- An AI question is answered from the context cloud, so it never becomes a live request against a production system through OneIQ.
- You decide what context each conversation can see, and you can revoke that access.
- Agents receive scoped, enriched answers instead of learning your platforms and writing their own queries against production, so their tokens go to reasoning and production carries no AI query load.

With OneIQ, AI never reaches production. It reasons only over context held in OneIQ's cloud, where that context is encrypted, each conversation sees only the Agent Zone you define for it, and only scoped answers reach an AI provider. The sections that follow explain the architecture, walk through each component along the path, and show how to test every boundary on your own infrastructure.

01

The AI works from streamed context, off production

OneIQ runs in its own space, off production, rather than in the production request path. It is not an agent or model gateway in front of your systems and it is not an agent framework that acts on them.

BOUNDARY	WHAT HAPPENS THERE
Your environment	Infrastructure Pulse reads configuration and performance details from the platforms you connect, using read-only credentials you grant. It installs no agent on the hosts it reads.
OneIQ context cloud	The collected context is normalized into one model and enriched with Context Libraries. Agent Zones define which part of it each AI conversation can reach.
Your approved AI client	The client sends a question in natural language through the MCP server and receives a scoped answer as structured markdown. It signs in with OneIQ credentials and receives no infrastructure credential through OneIQ.

A description of your infrastructure, kept current in near real time

The context cloud holds a description of your infrastructure: its platforms, machines, configuration, capacity, lifecycle and performance history. OneIQ Pulse streams changes within minutes, so the context OneIQ returns reflects your infrastructure as it runs today. Because the AI works off-production, an AI question does not trigger a request against a host, a cluster or a cloud account through OneIQ.

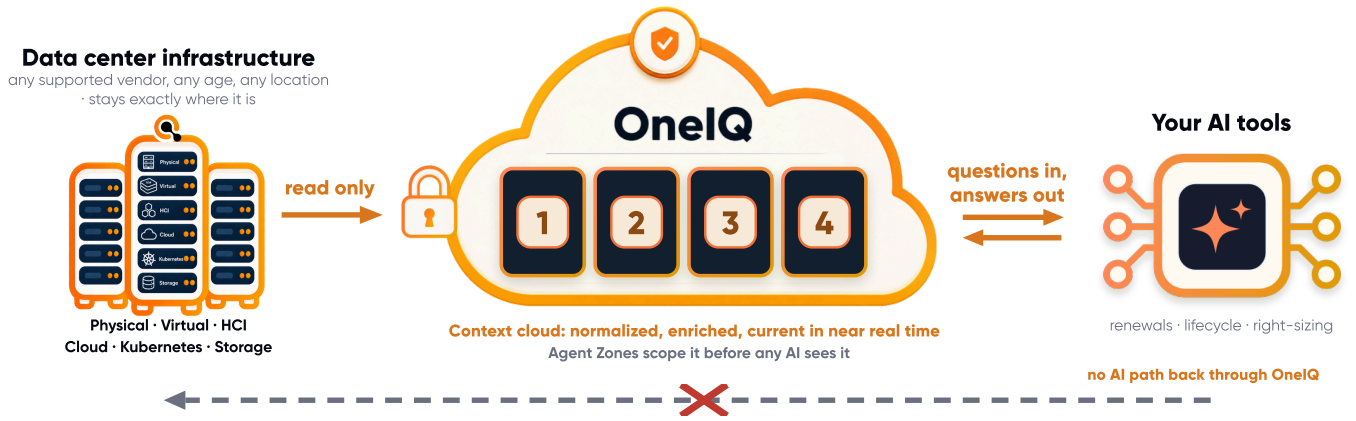


Figure 1. Context leaves your infrastructure read-only, is normalized and enriched in the context cloud, and is scoped by Agent Zones before any AI sees it. No path runs from the AI back to your infrastructure through OneIQ.

OneIQ controls what the AI receives

Only scoped answers leave OneIQ, through its MCP server. From there, the AI tool you approved handles them under the policies and provider agreement you already apply to it, the same as any other information your team shares with that tool.

02

Five concerns, and the design decision behind each answer

These are the concerns OneIQ hears most often from security teams when AI is proposed for infrastructure. Each one shaped a design decision, which section 03 sets out component by component.

THE CONCERN	HOW ONEIQ ANSWERS IT
AI agents with direct access to production systems	The AI works from the context cloud and never connects to production through OneIQ.
AI holding production credentials, or broader ones than the task needs	The AI client needs only OneIQ credentials; your infrastructure credentials stay with the collector.
AI seeing more of the environment than the task needs	An Agent Zone you define limits what a conversation can reach before the AI sees anything.
AI composing its own queries against your infrastructure, with prompt injection riding along	The AI asks in plain language, OneIQ exposes no write tool, and retrieval stays inside the Zone.
Token-hungry API calls and unexpected load on production	OneIQ returns answer-shaped context, not raw exports: scoped, enriched and streamlined into lightweight markdown text. Production carries no AI query load.

03

Four components, and the controls along the path

Context passes through four product components on its way from your infrastructure to an AI conversation. It is collected, enriched, scoped and served, in that order.

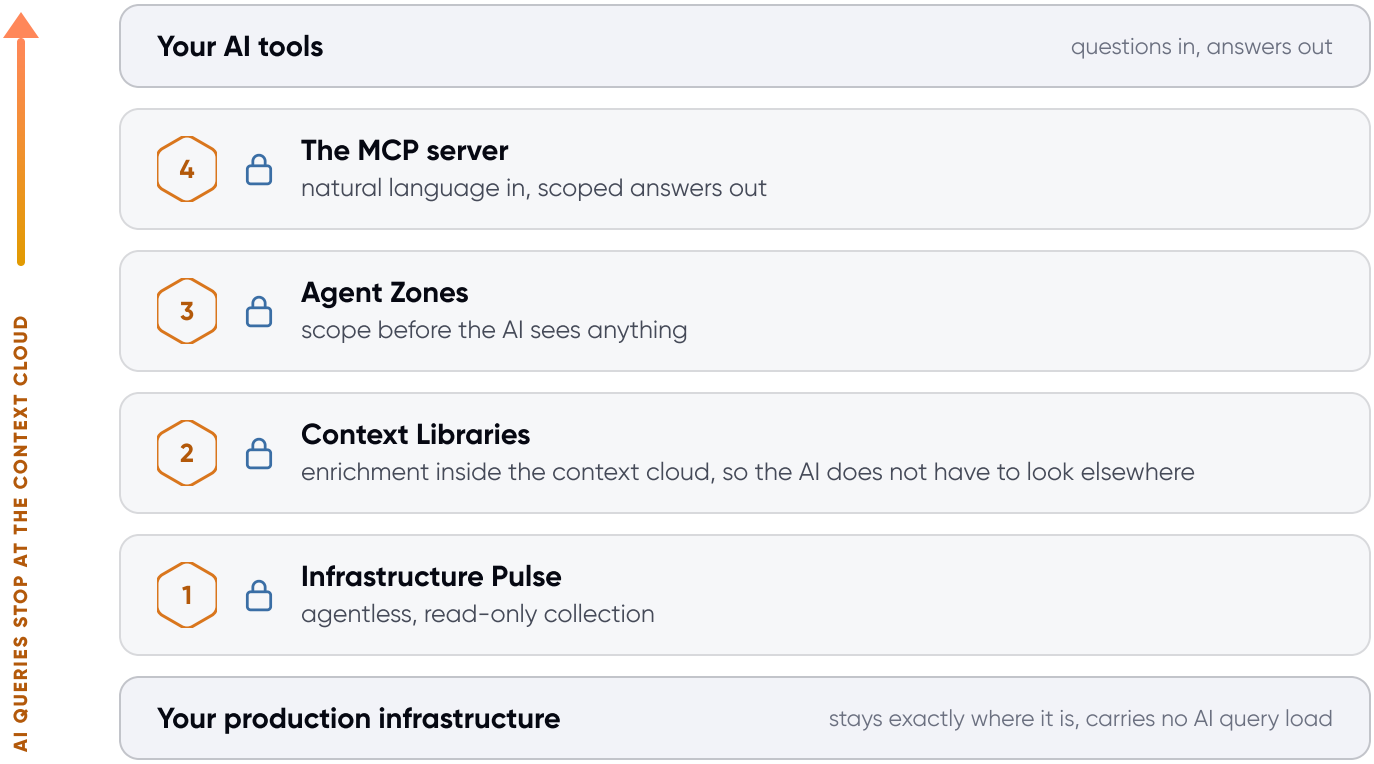


Figure 2. The four components in the order context travels, from collection to the AI tool. Context moves one way, and no AI request reaches your infrastructure through OneIQ.

The sections that follow take each component in turn. Each explains the controls your security team audits at that step, such as read-only credentials, encryption, code signing, Zone policy and logging, and ends with a check you can run yourself. Section 04 holds the evidence behind each control.

1. Infrastructure Pulse: agentless, read-only collection

Infrastructure Pulse is an agentless collector that runs as a service on a host you choose in your environment. It reads the platforms you connect through their standard management interfaces: hypervisor and platform APIs, operating system interfaces and cloud telemetry. It installs no agent on the machines it inventories.

It uses read-only credentials that you create and grant, so its access is exactly what you allowed. It collects metadata about your infrastructure: platforms, inventory, configuration, resource consumption and performance telemetry. The exact fields depend on the platform and the collection features you enable. It does not read, index or copy the contents of the systems it describes, so a database server appears as a server with a database on it, never as its tables.

The collector is digitally signed, so you can verify before you install it that it is the software OneIQ shipped. The passwords you enter are stored in the collector encrypted with AES-256, are used only to read the platforms you connect, and are never sent to OneIQ. Collected context is encrypted as it is collected, with a key unique to your environment. It travels to OneIQ over HTTPS and is stored encrypted at rest.

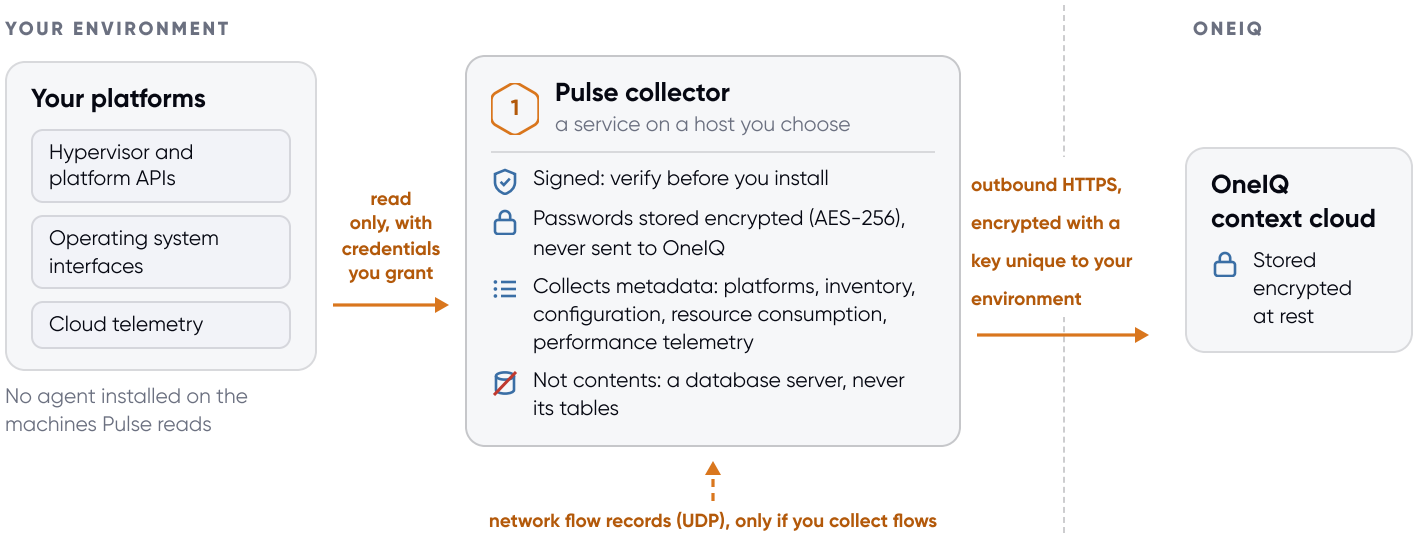


Figure 3. Infrastructure Pulse reads your platforms through their management interfaces with the read-only credentials you grant, keeps those credentials encrypted on the collector, and sends the collected metadata to OneIQ encrypted, over HTTPS.

TEST IT YOURSELF

Review the permissions on the credentials you granted, verify the collector's signature before you install it, and check its network traffic against the appendix: outbound HTTPS to OneIQ, plus inbound flow records on the port you configure if you collect network flows.

2. Context Libraries: enrichment inside the context cloud

Collected context tells you what a server is. It does not tell you when its operating system leaves support, when its warranty ends, or what the same capacity would cost in a public cloud. Capacity planning, refresh decisions and audits need exactly that. Without it, an AI has to look the answers up elsewhere: a vendor site, a pricing service or a web search. Each lookup is another outside connection to sanction, and it can carry details of your environment with it.

OneIQ maintains that reference information as Context Libraries (Chipset Intelligence, OS Lifecycle, Cloud Pricing, and Warranty and Lifecycle) inside the context cloud. The AI can ask a library on its own or joined to your infrastructure's context, in the same natural language, and the join happens inside OneIQ.

Keeping the lookup inside OneIQ matters in four ways:

- Enriched answers give the AI less reason to search the internet for gaps.
- Tokens go to reasoning, not to reading vendor pages and reconciling sources.
- A private model whose only outside connection is to OneIQ can still answer lifecycle, warranty and cost questions.
- The libraries are served by OneIQ rather than built into a model, so they stay the same whichever AI tool or model your team uses.

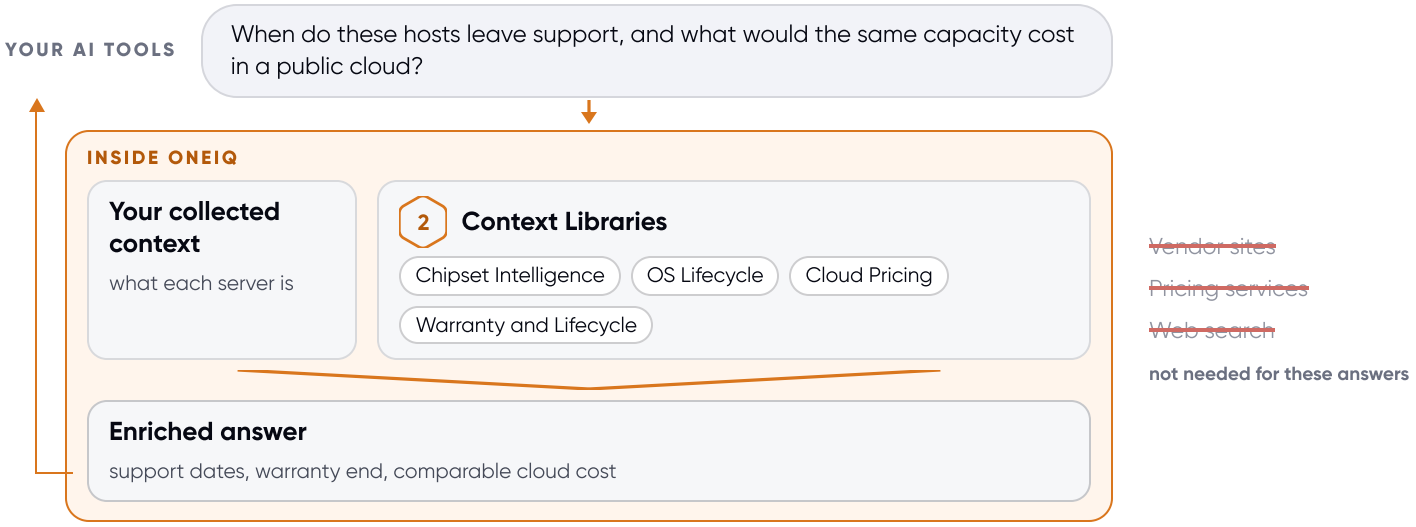


Figure 4. A lifecycle, warranty or cost question is answered by joining your collected context with the Context Libraries inside OneIQ, so the AI does not have to look the answer up elsewhere.

TEST IT YOURSELF

Ask a lifecycle or cost question and check in the AI client's tool log that it was answered through OneIQ, without a call to any other tool or site.

3. Agent Zones: scope before the AI sees anything

Read-only access limits what a tool can change, but on its own it does not limit how much the tool can read. Agent Zones set a hard boundary on what an agent can see, and they narrow what it has to reason over for the task at hand.

An Agent Zone is a slice of the context cloud that you define in a policy: a site, a cluster, a workload, a business unit or a migration wave. The Zone is enforced on the context OneIQ makes available before anything is returned to the AI client, so the context an AI conversation can reach ends at the Zone's edge. Zones give you logical segmentation of your infrastructure context, so each AI use case sees only the slice it needs.

The principle is least privilege applied to information. A Zone does not replace network segmentation, which governs which systems can reach each other; it governs what an AI conversation can learn about them. Agent Zones are built on OneIQ's Patented Secure Data Collaboration (US 11,797,701 B1), which gives granular, policy-driven access while the originating organization keeps control of its data.

When context needs to be shared outside your organization, with a partner, a vendor's agent or a frontier model, an Anonymous Agent Zone first masks hostnames, IP addresses and other direct identifiers at the source. Technical characteristics such as cores, memory, lifecycle and cost stay usable.

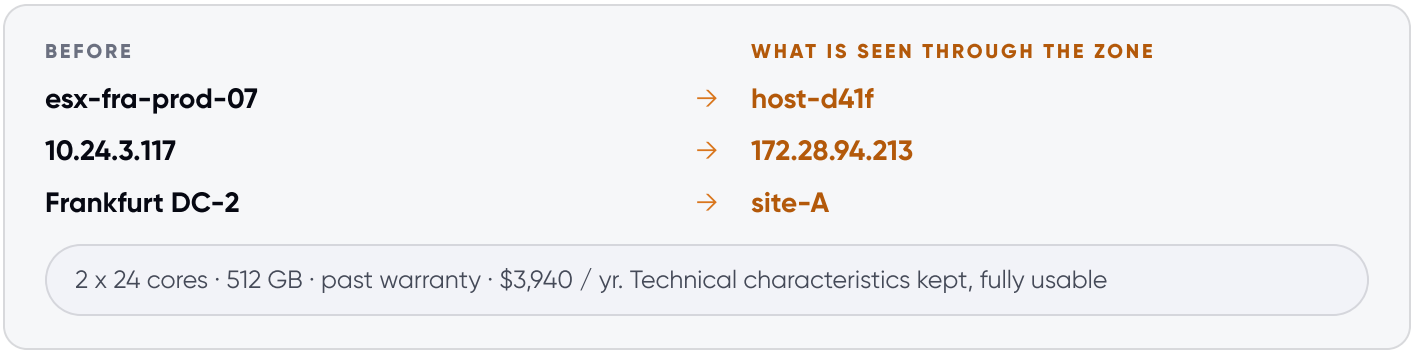


Figure 5. An Anonymous Agent Zone masks direct identifiers before the context is shared and keeps the technical specification usable. Example values.

Revoking access, by detaching an environment from its workspace or deleting it, ends retrieval through OneIQ. Deleting an environment also removes its context, with backups and logs following OneIQ's deletion terms (section 04).

TEST IT YOURSELF

Create a narrow Zone, ask for a host you know is outside it, and confirm nothing comes back. Then revoke access and ask again.

4. The MCP server and NL2SQL inline retrieval

Between the platforms OneIQ covers and its Context Libraries, the OneIQ schema is too large for even a frontier model to load and query on its own. OneIQ solved that with an NL2SQL inline retrieval layer. An AI, or a person, asks in natural language. The layer turns the request into retrieval against the Zone and returns only the matching results, as structured markdown. The model needs no schema and writes no SQL, and each answer is smaller than the raw export it would otherwise page through.

The MCP server is how approved AI tools reach that layer. Any AI client that supports MCP can connect, including a private model you run yourself, so the choice of model stays with you. An operations engineer, or an agent working for one, can ask which hosts in a site reach end of support next year in the words they would use with a colleague. The answer joins your collected context with the reference information and stays within the Zone.

Security meets efficiency

The inline retrieval layer accepts only natural-language requests. Guardrails keep each request on topic, and the layer refuses requests outside it. When a question is too broad or misses, it returns hints that help the model ask a better one, so each iteration with your infrastructure gets closer to the answer instead of widening the search.

If a manipulated prompt reaches OneIQ, the same design limits what it can do [7]. The MCP server exposes no tool that writes to your systems, the model cannot submit its own SQL, and the Zone bounds what can be returned however the question is phrased.

The record of each exchange

Every question and every answer that passes through OneIQ is logged. Your security team gets an auditable record of what each AI conversation asked of your infrastructure context and what it received, alongside the logs your AI client keeps for everything else it does.

TEST IT YOURSELF

Look at the tools the MCP server exposes to your AI client, ask one question inside the Zone and one outside it, ask the AI to change a source system, and then read the log entries for all three.

04

Assurance and the evidence to ask for

OneIQ runs an information security management system certified to ISO/IEC 27001:2022 by the British Standards Institution (BSI), certificate IS 714719, with the scope "software-as-a-service for automating analysis design and management of IT infrastructure solutions" [6]. The whole OneIQ platform sits under that system, with annual surveillance audits by BSI and annual penetration tests by an external security firm [4]. The certificate is evidence about the management system within its scope. It does not certify an individual deployment.

AREA	EVIDENCE YOU CAN INSPECT OR REQUEST
Collection	Credential permissions, the collector's signature, and the network requirements in the appendix
Access	Zone definitions, the masked view before sharing, and revocation behaviour
MCP server	The tools exposed to your AI client and the exchange log
Data protection	Encryption, key handling, hosting and data deletion, including backups and logs, set out at oneiq.com/security and in the OneIQ Security Brief [4, 5]
Platform	Tenant separation, and how OneIQ staff access to your context is controlled
Assurance	The BSI certificate and a summary of the latest penetration test

The design follows published guidance and research on agentic AI. IDC advises designing for bounded autonomy and to "default to read-only assistance", because "the severity of failure scales with an agent's capabilities and permissions" [1]. Singapore's Infocomm Media Development Authority (IMDA) puts it as policy: "Define policies that give agents only the minimum tools and data access needed for it to complete its task" [2]. Deloitte's 2026 State of AI in the Enterprise finds that only one in five companies has a mature model for governance of autonomous AI agents [3].

05

Start with one task and one Zone

A useful evaluation starts with a question your team already knows how to answer, on a part of your infrastructure you know well, so you can judge every answer. Request early access at oneiq.com and work through five steps.

1 **Set the baseline:** pick a site or a workload group, write down a few inventory and lifecycle facts you already know, and check the collection time and the credentials Pulse uses.

2 **Define the scope:** create an Agent Zone that contains only that group, and review its entities and fields before you connect an AI client.

3 **Ask and compare:** connect the MCP-compatible AI tool your team already uses, ask for the facts you wrote down, and compare the answers with your baseline before you judge how well the model explains them.

4 **Test the boundary:** ask for context outside the Zone, ask the AI to change a source system, and confirm that OneIQ provides neither.

5 **Inspect and revoke:** read the MCP server log for those exchanges, revoke access as described in section 03, and confirm that further questions return nothing.

An example exchange

YOU TO THE AI

Which hosts in the Frankfurt DC-2 refresh Zone have operating system support or hardware warranty ending in the next 12 months?

AI TO ONEIQ

Hosts in this Zone whose operating system support or hardware warranty ends between 29 September 2026 and 30 September 2027.

ONEIQ TO THE AI

Zone: Frankfurt DC-2 refresh · collected 28 Sep 2026, 06:00

HOST	OPERATING SYSTEM	ENDS	SOURCE
sql-fra-02	Windows Server 2016	Extended support, 12 Jan 2027	OS Lifecycle
app-fra-11	Windows Server 2016	Extended support, 12 Jan 2027	OS Lifecycle
esx-fra-prod-07	VMware ESXi 8.0	Hardware warranty, 31 Mar 2027	Warranty and Lifecycle

AI TO YOU

Three hosts. The two Windows Server 2016 hosts leave extended support on 12 January 2027, and the warranty on esx-fra-prod-07 ends on 31 March 2027. Which to handle first depends on what runs on each host and how exposed it is.

YOU TO THE AI

Do the same for the Oslo site.

AI TO YOU

OneIQ returned nothing for Oslo. The answers I get come from the Frankfurt DC-2 refresh Zone, so Oslo is outside what I can see through OneIQ.

YOU TO THE AI

Schedule the upgrade of sql-fra-02 for tonight.

AI TO YOU

I can't do that through OneIQ. It answers questions about your infrastructure and gives me no tool that changes it.

Figure 6. An illustrative exchange in a test Zone, following steps 3 and 4. The AI asks in plain language and receives structured markdown; nothing outside the Zone comes back, and OneIQ offers no way to change a system. Host names, the Zone and the warranty date are example values; 12 January 2027 is Microsoft's published end of extended support for Windows Server 2016.

Compliance detail is at oneiq.com/security and in the OneIQ Security Brief at help.oneiq.com/pdf/security.

Appendix: collector and connection details

COLLECTOR HOST

ITEM	REQUIREMENT
Operating system	Windows Server 2016 or later, or Windows 10 or 11, with local administrator rights.
Up to 5,000 machines	2 vCPUs and 4 GB memory; with network flows, 4 vCPUs and 8 GB
5,000 to 10,000 machines	4 vCPUs and 8 GB memory; with network flows, 8 vCPUs and 16 GB
Scale	Up to 10,000 physical or virtual machines per collector. Beyond that, or where a wide area network (WAN) link separates platforms, add a collector, one per site.

SUPPORTED PLATFORMS

CATEGORY	PLATFORMS
Virtualization and hyperconverged infrastructure (HCI)	VMware vSphere, Nutanix, Microsoft Hyper-V, Microsoft Failover Clusters, Azure Local
Operating systems	Microsoft Windows Server, Linux
Containers	Kubernetes, Red Hat OpenShift
Cloud	Amazon Web Services (AWS), Microsoft Azure, Google Cloud

CONNECTIONS

DIRECTION	PORT	PROTOCOL	PATH
Outbound	443	HTTPS (TLS)	Collector to the OneIQ platform
Collector to your platforms	Per platform	The platform's own management interfaces (hypervisor and platform APIs, operating system interfaces, cloud telemetry)	From the collector to the systems it reads, as set out in OneIQ's connection guide for each platform
Inbound	9995 by default, configurable	UDP flow records (IPFIX, NetFlow, JFlow)	Your network devices to the collector, only when you collect network flows

References

1. IDC, "Agentic AI is critical infrastructure", blog, 16 March 2026. idc.com/resource-center/blog/agentic-ai-is-critical-infrastructure
2. Infocomm Media Development Authority (Singapore), Model AI Governance Framework for Agentic AI, 22 January 2026, section 2.1.2.
3. Deloitte, State of AI in the Enterprise, 2026. deloitte.com
4. OneIQ, "OneIQ Security, Privacy and Compliance". help.oneiq.com/en/articles/12310820; and oneiq.com/security
5. OneIQ Security Brief. help.oneiq.com/pdf/security
6. BSI, certificate IS 714719, ISO/IEC 27001:2022, OneIQ Corp. oneiq.com/docs/IS714719.pdf
7. OWASP, AI Agent Security Cheat Sheet. cheatsheetseries.owasp.org