

OPENTELEMETRY FOR HYBRID IT

**Open standards.
Flexible collection.
A stronger data foundation
for Autonomous IT.**





TABLE OF CONTENTS

Your Telemetry Shouldn't Dictate Your Observability Strategy	3
OpenTelemetry is Changing the Observability Architecture	4
Standardize Telemetry Without Standardizing on One Backend	5
The Next Step: OpenTelemetry Metrics in LM Envision	6
What Changes for Infrastructure Teams?	6
From Telemetry Portability to Operational Context	8
Six Ways to Put OpenTelemetry Metrics to Work	9
An Open Architecture for Hybrid Observability	10
The Data Foundation for Autonomous IT	11
Five Questions for Your OpenTelemetry Strategy	12
Open Telemetry. More Choice. Better Operational Outcomes.	12

INTRODUCTION

Your Telemetry Shouldn't Dictate Your Observability Strategy

As enterprises standardize on OpenTelemetry, they gain more control over how telemetry is generated, collected, processed, and transported. Data can be collected from environments that are difficult to reach with traditional monitoring architectures. Higher-frequency metrics can expose infrastructure changes that happen between polling intervals. And telemetry can be routed to the different platforms where teams across a company need it.

That flexibility has often stopped at the edge of infrastructure operations. OpenTelemetry pipelines feed certain tools. The observability platforms that monitor the network, the data center, and the cloud run on their own collection model. The two rarely meet, so teams reconcile them by hand, and the environments that are hardest to poll stay dark.

LogicMonitor is bringing the flexibility of OpenTelemetry Metrics to LM Envision. Now organizations can send standards-based metrics into LogicMonitor, including from the OpenTelemetry Collector pipelines they already run, and put that data to work alongside the infrastructure, network, cloud, and log telemetry their operations teams already rely on.

That solves for a growing observability challenge: collecting the metrics you need without being constrained by a proprietary collector, a fixed collection model, or another monitoring silo.

This also creates a stronger foundation for Autonomous IT. AI can only reason over the data and context available to it. Expanding LM Envision with open, standards-based telemetry gives AI a broader view of the environment and more context for correlating signals, identifying what matters, determining likely causes, and ultimately helping teams move from detection toward increasingly autonomous operations.

The differentiation is not simply that LogicMonitor accepts OpenTelemetry. It is that customers can embrace an open telemetry ecosystem while putting that data to work within the operational context of their hybrid environment.



OpenTelemetry is Changing the Observability Architecture

Today's enterprises operate complex infrastructure environments across data centers, networks, public and private clouds, containers, remote locations, edge environments, and SaaS platforms. Different teams often manage different parts of that environment, often with entirely different monitoring and observability tools. When something breaks, that fragmentation shows up as inconsistent decisions about how to respond.

Historically, collecting telemetry and choosing an observability platform were often part of the same decision. Deploy a vendor's agent or collector. Send telemetry through its pipeline. Analyze the data in its platform.

OpenTelemetry changes that relationship. It provides an open, vendor-neutral framework for generating, collecting, processing, and exporting telemetry. Instead of tying instrumentation and collection to a particular backend, organizations can create a common telemetry layer and decide where that data should go.

The OpenTelemetry Collector is the pipeline that makes this separation practical. Receivers ingest telemetry from supported sources, processors filter, modify, or batch it, and exporters route it to one or more destinations. Those stages let teams apply collection and data-quality controls before metrics, logs, or traces leave the environment.

Because the pipeline is modular, organizations can push telemetry from environments that are difficult to poll, transform it locally, and send the same data to LogicMonitor or other backends. OpenTelemetry therefore separates the mechanics of moving telemetry from the operational question of where teams get the most value from it.



That allows enterprises to separate two important decisions that used to be bundled:

- 1 How should we collect and transport telemetry?
- 2 Where can we get the most operational value from it?

Separating them changes the economics of observability. Collection becomes shared infrastructure that outlasts any single tool choice. Platforms compete on what they help teams understand and do, not on how difficult they are to leave.


Standardize Telemetry Without Standardizing on One Backend

Many enterprises aren't starting their OpenTelemetry journey today.

OpenTelemetry collectors and agents may already be deployed across application, cloud, container, and infrastructure environments. Platform engineering teams may already operate pipelines that collect, process, and route data across

the enterprise. Those investments shouldn't have to be duplicated simply because another team wants to use the telemetry.

An OpenTelemetry architecture can provide a common collection layer while allowing different teams to use the observability platforms that best support their requirements.



A platform engineering team and a network operations team could have two different sets of requirements, while CloudOps and application teams have their own. OpenTelemetry lets the telemetry architecture accommodate those realities instead of reconciling against the differing requirements of each team. LogicMonitor is now part of that open architecture to enable cross-team collaboration.

**Standardize collection
without standardizing on
a single backend.**

The Next Step: OpenTelemetry Metrics in LM Envision

LogicMonitor now supports OpenTelemetry metrics as part of a broader approach to open, hybrid observability.

Organizations can send OpenTelemetry metrics into LogicMonitor over the OpenTelemetry Protocol (OTLP) in two ways:

- From the collectors they already run. Any OpenTelemetry-compatible collector can push metrics to a standard OTLP endpoint in LogicMonitor.
- From a LogicMonitor-provided collector. Teams that are newer to OpenTelemetry can start with the LogicMonitor distribution of the OpenTelemetry Collector.

The two collector paths support different operating preferences. The LogicMonitor distribution is a curated version of the OpenTelemetry Collector built from the contrib project, with a focused receiver set and LogicMonitor-managed setup and upgrades. The community OTel Contrib Collector offers a much larger component catalog and more flexibility for unusual or complex pipelines, with more configuration and lifecycle work owned by the customer.

The choice is therefore about configuration responsibility, not whether OpenTelemetry is supported. Teams can start with the LogicMonitor path or keep the broader community distribution when their pipeline requires a specialized receiver or processor.

For organizations already using OpenTelemetry, the significance is simple: a new collection architecture isn't required to start using those metrics in LogicMonitor. It is a configuration change, not an agent rollout.

Once metrics arrive, teams can put them to work:

- **Discover.** Find available metrics and inspect their attributes, then filter and group them by the systems, services, and conditions that matter.
- **Query.** Use PromQL to filter by attribute, aggregate, and select time ranges for analysis and troubleshooting.
- **Visualize and alert.** Turn queries into dashboards and alerts, so OTel metrics are actively monitored rather than simply stored.

**Keep your instrumentation.
Keep your collectors.
Add LogicMonitor as
a destination.**

The result is more freedom over how telemetry gets into the platform, while keeping the context and workflows operations teams depend on after it arrives.

What Changes for Infrastructure Teams?

Supporting OpenTelemetry metrics expands the kinds of infrastructure monitoring architectures organizations can build.

1 Reuse the OpenTelemetry infrastructure you already have

For enterprises already deploying OpenTelemetry, starting another telemetry collection architecture means more software, configuration, maintenance, security reviews, and operational overhead. Why duplicate infrastructure that's already doing the job?

Existing OpenTelemetry Collector pipelines can send metrics into LogicMonitor, so organizations can build on their current investment rather than starting over. It also gives infrastructure and operations teams a natural way to collaborate with platform engineering organizations that have already standardized on OpenTelemetry.

OpenTelemetry is also extensible when a system does not have a ready-made integration. Organizations can build custom receivers for new platforms. That makes the standard a practical way to extend coverage without waiting for a vendor-built integration.

Extend the value of existing OpenTelemetry investments while reducing unnecessary collection infrastructure.

2 Reach environments that polling can't

Some infrastructure environments place tight restrictions on what software can be installed. Security policy may prohibit remote polling protocols such as SNMP, WMI, or SSH. Credentials may not be allowed to leave a network segment. Organizational boundaries may prevent another monitoring component from being deployed at all.

Polling-based monitoring has no good answer in those cases, because it has to reach into the environment from outside, and strict policy and compliance guidelines prevent that.

OpenTelemetry reverses the direction. Metrics are collected locally and pushed out over OTLP, often by an OpenTelemetry Collector that has already passed security review.

That gives organizations a standards-based path into LogicMonitor for systems that were previously out of reach, without remote polling access and without sharing credentials.

Extend visibility into locked-down and regulated environments where pull-based collection isn't permitted or practical.

3 See changes that happen between polling intervals

Traditional infrastructure monitoring commonly collects metrics at defined intervals. For many metrics, that provides exactly the visibility operations teams need.

But infrastructure doesn't operate according to the monitoring schedule. A resource can spike and recover. Network utilization can shift rapidly. A short-lived performance condition can affect a service and disappear before the next scheduled poll. The monitoring system captures the aftermath without capturing the event itself.

OpenTelemetry metrics provide another option. Where the telemetry source supports it, metrics can be collected at sub-60-second intervals to reveal fast-moving conditions. That complements traditional monitoring rather than replacing it.

Gain granular visibility into short-lived infrastructure behavior that traditional interval-based polling can miss.

**Poll broadly.
Stream selectively.
Gain visibility to what happens in between.**

Higher-frequency telemetry is most valuable where transient events matter. Traditional collection continues to provide broad, efficient monitoring across the rest of the environment.

4 Bring streaming network telemetry into the bigger picture

Networks are an especially important example where streaming and polling data provide value together. Polling remains essential for broad network monitoring, but modern network infrastructure can also generate streaming telemetry that captures operational changes at much higher frequencies.

The two approaches are complementary. Traditional monitoring provides continuous visibility across the network. Streaming telemetry provides greater detail where conditions change quickly. OpenTelemetry offers a standards-based path for bringing supported streaming metrics into LogicMonitor, so they are used alongside the other infrastructure signals teams already monitor rather than sitting in another isolated tool.

Capture fast-moving network conditions and investigate them in the context of the broader hybrid environment.

From Telemetry Portability to Operational Context

OpenTelemetry solves an important problem: it makes telemetry portable. But portable telemetry isn't the same thing as understanding..


A metric can tell you CPU utilization increased. Operations teams still need to know:

- What resource generated it?
- Is the behavior abnormal?
- What else happened at the same time?
- What depends on that resource?
- Is another part of the infrastructure experiencing a related problem?
- What should I investigate next?

That is where the value shifts from collecting telemetry to understanding it, and it is the part of the problem that an open standard, by design, doesn't address.

**OpenTelemetry makes telemetry portable.
LogicMonitor helps make it actionable.**

LogicMonitor's strength has long been broad visibility across hybrid environments — from traditional infrastructure and networks to cloud and other modern technology domains. OpenTelemetry metrics add another telemetry source to that picture. The objective isn't simply to ingest more data. It's to combine useful signals with the context operations teams need to understand what's happening across the environment.




Six Ways to Put OpenTelemetry Metrics to Work

USE CASE	SITUATION	APPROACH	OUTCOME
1 Extend an existing enterprise OTel deployment	Platform engineering already operates OpenTelemetry collectors across the enterprise.	Add LogicMonitor as a destination in the existing pipeline rather than deploying a separate collection architecture.	Infrastructure teams gain access to enterprise telemetry; platform teams keep their established OTel architecture.
2 Observe a locked-down environment	Policy prohibits SNMP, WMI, or SSH polling, and credentials cannot leave the segment. An approved OpenTelemetry Collector is already running.	Push metrics from the approved collector to LogicMonitor over OTLP.	Visibility into systems that were previously dark, with no remote polling access and no shared credentials.
3 Investigate transient infrastructure behavior	Users report intermittent performance problems, but standard metric intervals show nothing obvious.	Collect higher-frequency OpenTelemetry metrics for the infrastructure in question.	Evidence of what happened between polling intervals, not just the aftermath.
4 Capture fast-moving network conditions	Network devices support streaming telemetry, but the team is limited to SNMP polling today.	Bring supported streaming metrics through an OpenTelemetry pipeline into LogicMonitor.	High-frequency network signals investigated alongside the broader infrastructure picture.
5 Bring in services that already speak OTLP	A third-party service, such as a managed DNS provider, can emit latency and health metrics natively over OTLP.	Point that OTLP output at LogicMonitor.	External dependencies appear in the same operational view as the infrastructure they affect.
6 Support multiple observability platforms	Application teams depend on one backend, central IT on another, specialized teams on a third.	Route telemetry from a shared OpenTelemetry pipeline to each destination, including LM Envision.	Collection is modernized now. Tool consolidation happens on the organization's schedule, not as a precondition.

An Open Architecture for Hybrid Observability

OpenTelemetry creates an opportunity to think about observability as a layered architecture rather than a collection of isolated monitoring tools.

The layers below are deliberately ordered. Each one depends on the one beneath it, which is why decisions made at the collection layer end up determining what is possible at the top.



Generate: Infrastructure and applications create the operational signals.

Collect and transport: LogicMonitor's native collection and OpenTelemetry are complementary ways to bring those signals in. One polls broadly. The other accepts what is pushed.

Unify: However telemetry arrives, it lands in one data foundation rather than a separate silo per collection method.

Understand: Context connects signals to the resources, dependencies, and services they affect.

Act: Operations teams, and increasingly AI, use that context to investigate, decide, and take governed action.

**Many paths in.
One data foundation.
Open standards at the
telemetry layer, operational
intelligence above it.**

The Data Foundation for Autonomous IT

The implications of an open telemetry architecture extend well beyond dashboards and alerts.

Autonomous IT is not a hands-off operation. It is an operating model: a closed loop that connects unified telemetry, business-aware context, AI-driven reasoning, and governed execution, with humans in control but no longer stuck in every loop. Every step in that loop inherits the limits of the step before it.

That is why the data foundation matters more than it first appears. AI without context is noise, and context without complete telemetry has gaps. A system can't reason effectively about an infrastructure problem if critical signals are inaccessible, fragmented across disconnected tools, or missing entirely. You can't automate what you cannot see.

Organizations progress toward Autonomous IT in stages, and the first stage is the one OpenTelemetry most directly affects:

Stage 1 Unified visibility

Eliminate blind spots across systems and experience.

Stage 2 Correlated context

Reduce noise and connect signals to service impact.

Stage 3 Guided response

Move from alerts to answers and recommended actions.

Stage 4 Governed automation

Execute trusted actions with approvals and control.

Stage 5 Scaled autonomy

Expand automation safely with continuous validation.

An open telemetry layer strengthens stage one in three ways.

It closes gaps: Locked-down environments, high-velocity network conditions, and OTEL-native workloads are exactly the places where polling-based visibility runs out.

It reduces fragmentation: Telemetry that used to live only in a separate pipeline can sit in the same foundation as the rest of the hybrid estate.

It keeps the foundation current: New infrastructure, including AI workloads, increasingly emits OpenTelemetry by default. A standards-based on-ramp means the data foundation keeps pace without a new integration project for every new source.

The goal isn't to collect everything. It's to make the right operational data available, with enough context to understand what it means. That is the foundation Edwin AI, LogicMonitor's AI operating layer for IT operations, depends on to identify what matters, guide next steps, and support governed automation.

OpenTelemetry expands access to telemetry. Hybrid observability connects it to the environment. AI and governed automation turn that context into action.

Each step depends on the one before it. Autonomy is built from the data foundation up.

Five Questions for Your OpenTelemetry Strategy

As OpenTelemetry becomes part of enterprise observability architecture, infrastructure leaders should look beyond the question, “Do we support OTel?” Ask instead:

1 Where is OpenTelemetry already deployed?

Inventory existing collectors, agents, pipelines, and ownership. You may already have telemetry infrastructure that other operational teams can reuse.

2 Where are we duplicating collection?

Look for environments where multiple monitoring products collect similar telemetry independently. An open collection architecture may offer opportunities to simplify those pipelines.

3 Which signals need higher-frequency visibility?

Not every metric needs to be collected more often. Identify the operational problems where additional granularity would materially improve troubleshooting.

4 Where can we poll?

Locked-down, regulated, or specialized environments are candidates for push-based collection through OpenTelemetry infrastructure that is already approved.

5 What happens after the telemetry arrives?

This may be the most important question. Telemetry collection isn’t the business outcome. Consider how the data contributes to alerting, dashboards, troubleshooting, and service understanding, and whether it lands somewhere AI and automation can use it. Telemetry that arrives without context is a storage cost. Telemetry that arrives in a unified data foundation is the raw material for Autonomous IT.

Open Telemetry. More Choice. Better Operational Outcomes.

The future of enterprise observability isn’t about choosing between open standards and an observability platform. Organizations need both.

OpenTelemetry provides greater control over how telemetry is generated, collected, processed, and transported. LogicMonitor brings that telemetry into a broader view of hybrid IT. With OpenTelemetry Metrics, organizations can reuse existing OTel investments, reach environments that polling can’t, capture higher-frequency infrastructure behavior, take advantage of streaming telemetry, and support architectures that span multiple backends.

Most importantly, organizations gain the freedom to separate the collection decision from the observability decision:

- Collect telemetry using the architecture that works for your environment.
- Choose a platform based on what it helps your teams understand and accomplish.

OpenTelemetry provides the flexibility.

LM Envision provides the operational context.

Together, they create a more open data foundation for Autonomous IT.



Put your OpenTelemetry data to work

Bring OpenTelemetry metrics into LM Envision, and turn it into operational context your teams and AI can use.

**EXPLORE OPENTELEMETRY
WITH LOGICMONITOR**

About LogicMonitor®

LogicMonitor® is the AI-first platform for Autonomous IT, enabling enterprises to operate complex digital systems with greater resilience, efficiency, and confidence. By unifying visibility from user to code across infrastructure, cloud, Internet, and digital experience, LogicMonitor delivers the intelligence required to anticipate issues, eliminate blind spots, and take action automatically. Powered by Edwin AI, LogicMonitor helps IT and business leaders reduce operational toil, protect revenue, and accelerate innovation in an increasingly complex digital world. For more information, visit www.logicmonitor.com and [our blog](#), or follow us on [LinkedIn](#), [X](#), [Facebook](#), and [YouTube](#).

LogicMonitor®, Dexda®, LM Service Insight®, LM®, and the LogicMonitor logo are registered trademarks of LogicMonitor, Inc. in the United States and other countries. Other product and company names mentioned herein may be trademarks of their respective owners. © 2026 LogicMonitor, Inc. All rights reserved. | 10.2026