



AN INFLUXDATA CASE STUDY

# NodeSource Builds Custom APM Application in Node.js Powered by InfluxDB



# NodeSource Delivers Enterprise-Level Node.js Monitoring With NISolid And InfluxDB

## Company in brief

NodeSource empowers developers by providing OSS and Enterprise levels of Node.js runtime, tooling and support. NodeSource is the primary distributor of Node.js on Linux with over 34 million downloads. The company aims to solve pain points across the software development lifecycle for developers and DevOps engineers.

## Case overview

NodeSource's flagship product, NISolid Runtime, includes hardened Node.js LTS releases, provides increased visibility into production applications, and enables security monitoring and alerts.

NodeSource needed a database that could handle time series data and allowed users to access and visualize that data in real time. NodeSource uses InfluxDB to power NISolid's analytics and diagnostic capabilities.

## The business challenge

As the primary distributor of Node.js on Linux, NodeSource saw an opportunity to leverage its expertise in Node.js to help users understand what's going on "under the hood" in their Node.js runtime environments. To accomplish this, NodeSource offers a range of products and services intended to help enterprises gain insight about application performance and security.

NISolid, the company's flagship product, exposes data on the internal behavior of node.js runtimes in a web console. NISolid provides deep metrics on node.js runtimes with minimal overhead. It also delivers performance and security insights without any additional instrumentation. NISolid console also offers a bi-directional control mechanism that enables users to control runtime behavior.

For the NISolid console to work, it needed a time series database that could handle all the necessary process data and data aggregation. NodeSource chose InfluxDB as its centralized analytics hub because of its ability to ingest, process, and aggregate data, and to query that same data, all in real time.

## | The technical challenge

NISolid is a drop-in replacement for Node. As such, it serves large installations of nodes, often dealing with hundreds or thousands of processes running simultaneously across different environments. NodeSource needed a way to access metrics for these nodes and processes without a lot of overhead.

Initially, the NodeSource team used etcd as the central database, but it quickly became apparent that etcd, a key-value database, did not allow the kind of access to, and manipulation of data that NISolid needed. So, they decided to go with a time series database instead, quickly settling on InfluxDB.

NodeSource chose InfluxDB, in part, because of its write capability. NISolid provides real-time updates, sampling data every three seconds, which requires low latency and high write volumes. The team looked at external application process management (APM) tools, such as Datadog and New Relic, but they were unable to deliver the required latency. NISolid also needed to be able to query data in real time and present it in multiple different ways. InfluxDB's rich querying capability solved that need as well.

## | The solution

NodeSource built the open-source version of InfluxDB into NISolid so customers automatically get InfluxDB capabilities when they install a new NISolid node. This reduces the amount of configuration that end users need to do to InfluxDB because it works right out of the box.

With the open-source version, NISolid can handle up to 3,000 simultaneous processes. For customers with greater needs, NodeSource offers a version of NISolid that leverages InfluxDB Enterprise, which delivers a 3x improvement, monitoring of up to 9,000 simultaneous processes. The NodeSource team noted that these constraints are due to the console's web interface and not InfluxDB's ability to handle data.

With InfluxDB as the central data hub, the NISolid console lets users visualize data, create filters, set thresholds and alerts based on those criteria, and generate reports on security vulnerabilities within an NPM package or an ecosystem.

“

*“InfluxDB was made for writing, so you can throw a huge amount of data at it, and it can really handle it. It's fast, it's scalable, and it allows for aggregating that [data] in real time.”*

---

**Nathan White**, Principal Solutions Architect, NodeSource

With the open-source version, NISolid can handle up to 3,000 simultaneous processes. For customers with greater needs, NodeSource offers a version of NISolid that leverages InfluxDB Enterprise, which delivers a 3x improvement, monitoring of up to 9,000 simultaneous processes. The NodeSource team noted that these constraints are due to the console's web interface and not InfluxDB's ability to handle data.

With InfluxDB as the central data hub, the NISolid console lets users visualize data, create filters, set thresholds and alerts based on those criteria, and generate reports on security vulnerabilities within an NPM package or an ecosystem.

## Results

NISolid continues to grow and evolve along with InfluxDB. Through this evolutionary process the NodeSource team learned how to fine-tune InfluxDB to optimize performance. For users getting started with InfluxDB, the NodeSource team offers the following tips. InfluxDB logs are very robust, so optimizing log output by turning off HTTP logging and implementing a log rotation can make them more manageable.

Figuring out retention policies for data is vital for making data performant and can affect the ease of scaling in the long term. Understanding cardinality for your use case is also critical because this affects memory usage and storage, which can impact the bottom line.

It is easy to get carried away with InfluxDB's querying capabilities. Queries can eat up critical resources, depending on how they are written and how they access InfluxDB, so spend some time optimizing them.

Finally, be sure to pay attention to schema versioning. A little planning here can save a lot of time and effort in the long run if you need to perform a data migration.

## Benefits

InfluxDB is the backbone that enables some of NISolid’s most powerful features. The time that the NodeSource team spend fine tuning InfluxDB to work with NISolid produced several benefits.

Continuous queries allow users to leave a query open so as new data comes in that fits the query, that data gets pulled in and the query output updates automatically.

NodeSource’s Principal Solutions Architect Nathan White noted, “continuous queries, alone, have been a game changer in terms of being able to provide rich and robust features; and [they are] so much fun to play with.”

InfluxDB’s writing ability makes it possible for NISolid to monitor thousands of processes across an entire ecosystem. “InfluxDB was made for writing, so you can throw a huge amount of data at it, and it can really handle it. It’s fast, it’s scalable, and it allows for aggregating that [data] in real time,” White said.

Another area where InfluxDB proves critical is testing and debugging. The NodeSource team pointed out that once they understood their data shapes and schema they were able to emulate and replay data. NodeSource has done this with a year’s worth of data for 2,000 processes. Having this type of power and volume of data enables NodeSource developers to run through different scenarios, perform capacity testing, and attempt to replicate issues that their larger enterprise customers encounter.

## | What’s next

NodeSource continues to build out features for NISolid. Some roadmap items include functionality around vulnerability exposure, memory leak detection, and async tracing.

## About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit [influxdata.com](https://influxdata.com) and follow us [@InfluxDB](https://twitter.com/InfluxDB).



## Try InfluxDB

Get InfluxDB

Contact us for a personalized demo [influxdata.com/get-influxdb/](https://influxdata.com/get-influxdb/)