

Between Cloud adoption, new technologies (IoT devices, machine agents, APIs, etc.), worldwide changes in consumer behaviors, and the overall shift of pace in innovation and deployment, the traditional IT model has come to the transdisciplinary impasse we now call *digital transformation*. At the root of it, of course, is data. More specifically, ever-Bigger, ever-Faster, more complex, more diverse, sometimes-incomplete data that challenges traditional performance and service management strategies and tactics... in more than one way.

To streamline IT operations, observational data (system monitoring, logs, etc.) needs to be aggregated with engagement data (events, tickets, and incidents), then distilled into streaming insight that can yield continuous improvements.

Continuous improvements, however, cannot happen without Automation, which, in turn, cannot happen without Machine Learning and Artificial Intelligence. Both of which, of course, cannot happen without successful and specific algorithmic models—and algorithmic models can only perform as well as the rule application, pattern recognition, and data quality they've been trained on.

But how do you build and train an algo?

Usually, algorithms learn from structured data- but first, ground truth for the data needs to be established with human annotators. This covers all the important cases in each of the datasets, and accuracy standards need to be baselined on each specific use case. As the model moves between sets and gets calibrated through rules training, it needs to also be tested, continuously, in production—all while keeping a close eye on performance indicators in a contained, scalable, and self-feeding process known as a *flywheel*.

The seamless integration of the training flywheel with a business' core functions is imperative, as it directly influences how straightforward the path is from data source

to automation—and that’s where AIOps comes into play. A clear and still nascent paradigm shift in the way we handle digital transformation in IT Operations, **AIOps integrates and applies ML and fast analytics to data from IT Operations tools and devices, in order to automatically detect and address issues in real time.**

AIOps platforms ultimately enhance alarm filtering, anomaly monitoring, and automated repairs, addressing some of the challenges related to the ever increasing **volume, velocity, variety, and veracity of data**, by:

- Extending the capacity of real-time analytics beyond manual reporting through streaming performance monitoring
- Expanding reach beyond the traditional infrastructure model through inglobation of new environments (public cloud, private cloud, hybrid cloud, SaaS, mobile, third party services)
- Allowing for data processing and analytics at the edge
- Redistributing the responsibility of IT and speeding up decision making.

Many companies produce and/or ingest Peta-, Exa-, even Zettabytes of data, which—depending on their needs, processing technology, and limitations—can be used in effective, near-real time. After insight has been extracted and action—taken, the largest bulk of data is often discarded for lack of space and relevance.

Repurposing high-quality, already-cleaned, already-sorted and filtered data to serve as training sets for ML and AI initiatives can be a difficult, intensive endeavour. But with such a plentiful resource at hand, a company should look no further—because we simply cannot build, test, and move AI and ML algorithms from experiment to core business without the aid of training.

The ML/AI journey to build the right frameworks, create the right pipelines, and implement the right tools is thus a story of training—and many business leaders are turning to AIOps guide their path.

[Download](#) our **AIOps for Business Leaders** whitepaper to learn more about:

- Ways to scale your AI and power your tech stack with AIOps training data
- Ways to build models that move seamlessly from experiment, to program, to core product with a ML/Big Data flywheel
- Key data-labeling resources for AIOps.