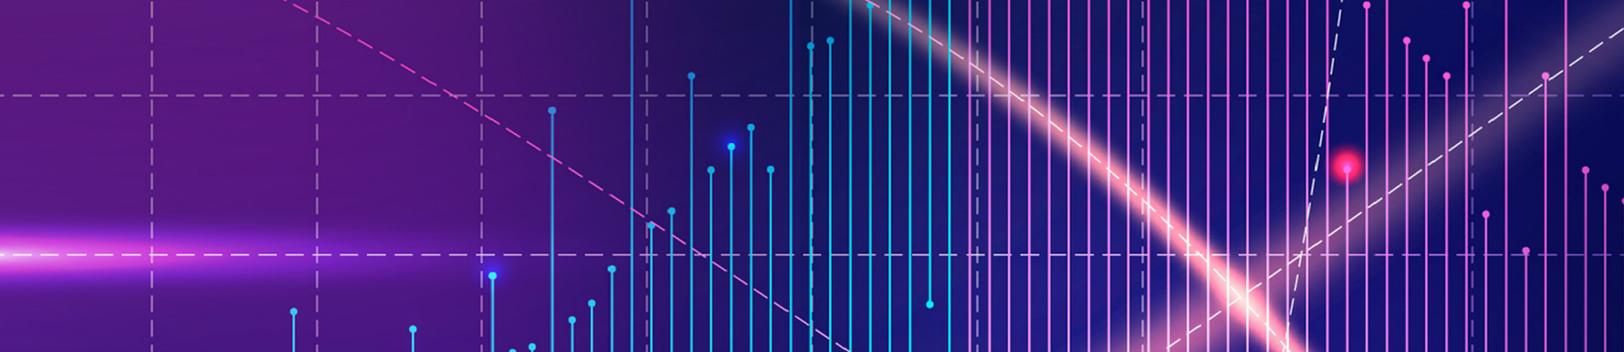


Continuous Intelligence vs. Traditional Streaming Analytics



Special Report: Continuous Intelligence vs. Traditional Streaming Analytics

In the age of digital transformation, it's no longer enough to simply ingest streaming data; organizations need to transition to developing insights that help them understand events in real time. When businesses have such analytics capabilities, they can use that information to decide plans of action.

Beyond Streaming Analytics

Continuous intelligence (CI) produces insights from streaming data to take actions in milliseconds to minutes. It frequently makes use of artificial intelligence (AI) and machine learning (ML) models to perform the real-time analysis and make connections between different events as they are happening.

Such rapid analysis of streaming data requires significant and highly variable amounts of compute and data storage capacity. And in many cases, the various stages of the process (e.g., model training and tuning, data ingestion, and analysis) have distinctly different compute requirements, which can vary over time.

Traditional compute infrastructures for high-performance computing and big data analysis will not do the job. What's needed is a highly scalable infrastructure that supports very dynamic development environments, deployment scenarios, and integration of new technologies.

CI goes beyond streaming analytics, which generally only applies a few filters, transformations, or aggregations to data. For example, a retailer analyzing site visitor click streams might derive information to make the site easier to navigate. By contrast, like the nervous system, CI propels people to act immediately. In the same example, CI might be used to make a personalized recommendation to a customer based on their journey on the site for that day's visit.

To work, CI relies on platforms, architectures, and software that allows organizations to collect, organize, and analyze data, enabling fast actions in response to real-time events. The data can include email messages, clickstreams, social media, data logs, and information from sensors and Internet of Things (IoT) devices.

Another differentiation between CI and streaming analytics is that CI enables event processing by analyzing streams of information about things that are happening now. In other words, CI is a method for processing real-time events and extracting information from event streams as they arrive. The goal of such event processing is to identify meaningful events (such as business opportunities or threats) in real-time situations and respond to them as quickly as possible.

Examples of CI in action include:

- ▶ Conversational chatbots that use natural language processing and text-to-speech capabilities to hold live audio or text-based conversations with customers
- ▶ Customer engagement hubs that process live social media threads and customer clickstreams to interpret what each customer wants and then engage each individual with real-time recommendations based on that analysis
- ▶ A utility that analyzes instantaneous power demands and determines whether to buy more power from a regional grid, tap stored capacity in a battery farm or dam, or offer customers real-time incentives to reduce consumption

In most cases, CI makes use of artificial intelligence (AI) and machine learning (ML) to complement traditional analytics used to investigate event data. Businesses can realize significant benefits from CI when their real-time analytics capabilities are integrated into their operations. In this way, CI can prescribe actions in response to business moments and other events.

Embedding CI into Business Processes

CI run on streaming data lets businesses move away from traditional streaming analytics (here's what's happened in the past). Many companies have already made a move to predictive analytics. Such analytics use regression analysis, modeling, neural network, and simulation algorithms. And they consider multiple independent variables to determine the likelihood of an event happening in the future.



CI offers a way to extend advanced analytics applications into the realm of decision support and decision automation. By processing event-based and streaming data, businesses can understand what's happening in the moment and react rapidly. Running prescriptive analytics, ML models, and AI algorithms against this data can derive actionable information. That information can be used by systems to decide what to do next and automatically take actions on their own.

One way to think about the difference between CI powered by AI versus traditional streaming analytics is to compare a mapping program with something like Google Maps or Waze. Both give people the information they need about getting from point A to point B. However, a mapping program would calculate the shortest distance and recommend that direct route. But Google Maps and Waze use real-time data about traffic plus user-input road situations, historical data about travel times, and more to determine the fastest way, in the moment, to reach a destination.

The difference is the inclusion of situational awareness (what's happening right now) and the delivery of intelligence that has a short lifetime. In this example, asking the same question 10 minutes, an hour later, a day later, or a month later would yield different results. In contrast, a traditional analysis that does not take events data into account would deliver the same route every time. A more advanced application might take average traffic at various times of day to alter the route. But again, CI goes a step further by taking live details into account to provide the best route.

Numerous business processes benefit from the same use of situational awareness and analysis of events-driven data. At the heart of these efforts is the use of AI and ML to make the inferences on the ingested streaming data.

As such, interest in the use of AI in business processes is on the rise. About 90 percent of businesses say they have some plans to implement AI in various parts of their business, including industrial automation, customer relationship management, and inventory planning and logistics, according to a recent industry surveyⁱ. Most believe AI infused into business operations and processes will help them make existing workers more productive and add value to products and services they sell to customers.

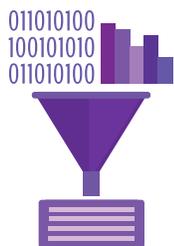
Another industry surveyⁱⁱ found that 59 percent of businesses that are planning or working on AI/ML said that they have AI deployed today. Businesses that are working with AI have, on average, four projects in place, expect to add six more projects in the next 12 months, and another 15 within the next three years. In 2022, those organizations expect to have an average of 35 AI or ML projects in place.

In all cases, the embracement of AI everywhere and enterprise AI is driven by the need to solve sophisticated business problems in highly dynamic environments.

Scaling the AI Ladder

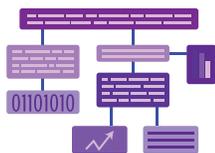
How should businesses transition from traditional streaming analytics to CI based on AI? AI requires machine learning; machine learning requires analytics; and analytics requires the right data and information architecture. That is the heart of a concept called the AI ladder,ⁱⁱⁱ which is loosely defined as the increasing levels of analytic sophistication that lead to, and buttress, a thriving AI environment.

Basically, the road to infusing business processes with AI-based CI starts with these four rungs of the AI ladder:



Collect: Make data simple and accessible.

Businesses can accomplish this by adopting a data and information infrastructure that enables self-service data-prep and ingestion capabilities. Why is this capability needed? Any CI solution or efforts to infuse AI into business processes will not be successful if users have difficulty incorporating new sources of streaming data. AI applications are only going to be as good as the data they use. A suitable application will allow a business to access data no matter where it lives.



Organize: Create a business-ready analytics foundation.

Most business managers do not understand AI. Worse, many organizations lack the data science skills required to implement, develop, and deploy modern analytics solutions. Companies need to create a business-ready data foundation—one that cleanses data and makes it secure and compliant. Organizations also need to have a catalog of their data to know its source, who owns it, and the metadata that maps the data to its business context. Additionally, a trusted analytics foundation can offer the necessary tools, algorithms, and engines to make it easier to conduct advanced analytics. Such an approach can enable businesses to deploy custom CI solutions.



Analyze: Build and scale AI with trust and transparency.

ML applied to streaming data delivers insights that help businesses better understand and cater to their customers, make more strategic business decisions, and optimize company workflows. Businesses need the ability to build and scale AI with trust and transparency. One way to accomplish this is to create or automate AI/ML models to help teams develop insights and make better, smarter decisions.

The analytics used in many AI efforts are black boxes. There is little transparency as to how the algorithms and models work or derive actionable intelligence from the data. At a minimum, this lack of understanding about the analytics can lead to erroneous conclusions. For example, if a healthcare ML model developed to diagnose a particular condition using lab results and a patient's medical history only uses a small segment of the population to train the model, the model's conclusions when run for patients outside that narrow population may not be valid. Another emerging issue is that an ML model might be biased because it uses erroneous assumptions, inherent prejudices, or a limited set of data to train the learning algorithms.

The solution: AI should not be a black box. Business professionals need justification mechanisms that support conclusions. Ensuring transparency might include taking steps such as explaining and documenting why one algorithm or model was selected versus another, and specifying what data was used to train a model and why that data was chosen. In this manner, businesses will be able to trust the recommendations made by AI models.



Infuse: Operationalize AI across the business.

AI offers tremendous value for improving the efficiency and productivity of teams across a business, whether automating routine tasks for sales and service teams or optimizing processes to achieve greater customer lifetime value. Ultimate benefits come when the use of AI executes on its own, without human intervention. With automation, there is no delay waiting for someone to locate the required data, put it into the right format, and run the most suitable algorithms for the problem at hand. As AI becomes part of all business processes, any solution must offer scalable performance that can adjust the processing resources based on the task. In addition to automation, other elements can be used to infuse AI into processes. An example would be the use of natural language processing to extract text from speech or business intelligence algorithms that help a business visualize, analyze, and share data.

How to Get to AI Everywhere

While some businesses are capturing all data in hopes of uncovering some new insights leading to possible actions, others are developing data and analytics architectures that manage processes end-to-end to produce actionable insights.

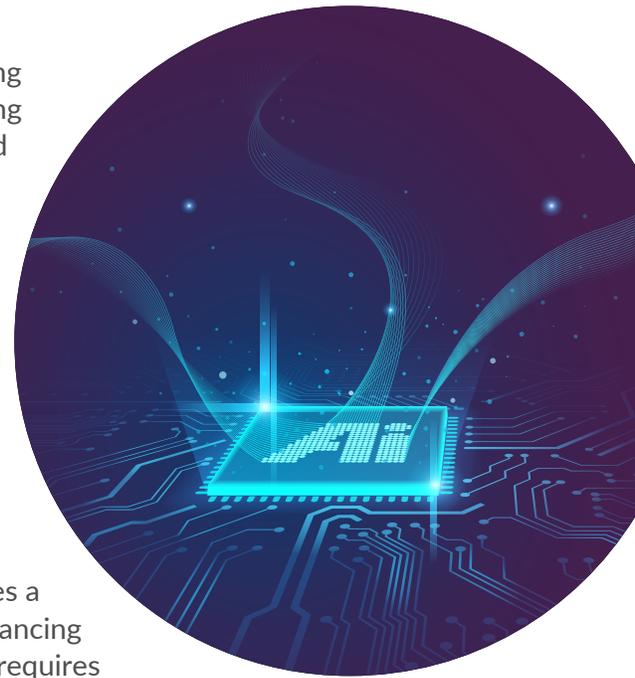
Aggregating and normalizing data is critical for AI success. Not only does the fragmented data from various applications make it impossible to have an effective dataset, but it also creates disjointed or disconnected processes without a full view of the outcome to train algorithms and machine learning models.

Making matters all the more challenging, CI uses flows of data from streaming sources, such as the Internet of Things (IoT), mobile devices, sensors, clickstreams, and transactions. Such data sources have, to date, remained largely unused. CI offers businesses a way to unlock this data to optimize decision making. However, advancing to a state where CI based on AI is infused into business processes requires a special information architecture.

Essentially, old approaches managed data, the derivation of insights, and the actions based on those insights as discrete, siloed efforts. Now, they must be fused into one architecture. That architecture integrates everything into a tight system. Elements of such an architecture must address:

- ✓ **Data acquisition:** Businesses must take into account that data is now generated from numerous external sources and that they can reside in siloes.
- ✓ **Data organization:** Ensure the data being used is clean, trusted, protected, and meets compliance requirements.
- ✓ **Analysis applied where it makes the most sense:** Is the analysis done where the data is generated (e.g., image analysis applied to video streams to guide an autonomous vehicle), or is the data backhauled to a data center or cloud instances for analysis?
- ✓ **Derivation and delivery of insights:** Is the analysis used to complement human activities with just-in-time insights embedded into business processes? Do these insights feed algorithms that analyze data as it streams into the business and automatically take action on the results?

An information architecture that delivers in these areas will provide the needed hybrid data management and unified governance and integration for AI. This architecture should offer great flexibility in where and how things run (on-premises, cloud, hybrid cloud), easily access open source technologies, and make deployment simpler through containers and container orchestrators such as Kubernetes.



Summary

The true value of adopting CI is when it is embedded into decision-making processes. In such cases, CI complements or replaces human decision-making. It can be used in decision support, suggesting to human intermediaries what actions to take next based on situational awareness. Or, CI can be used for decision automation, where the actions are triggered and carried out on their own in response to occurring events.

Achieving this level of success requires that AI be embedded into business processes. And for that to happen mandates that businesses adopt the essential building blocks of the AI ladder. Such an approach will help enable AI everywhere, thus delivering CI capabilities throughout the business.

IBM Cloud Pak for Data

To enable predictive analytics and AI everywhere, IBM developed Cloud Pak™ for Data, a fully integrated data and AI platform that modernizes how businesses collect, organize, and analyze data and infuse AI throughout their organizations. Built on Red Hat® OpenShift® Container Platform, IBM Cloud Pak for Data integrates market-leading IBM Watson® AI technology with IBM Hybrid Data Management Platform, data ops, governance, and real-time streaming analytics technologies. Together, these capabilities provide the information architecture to scale the AI ladder and meet organizations' ever-changing enterprise needs.

Deployable in just hours and easily extendable with a growing array of IBM and third-party services, IBM Cloud Pak for Data runs across any cloud, enabling organizations to more easily integrate their analytics and applications to speed innovation. IBM Cloud Pak for Data lowers the total cost of ownership, accelerates innovation based on open source technologies, and fully supports multi-cloud environments such as Amazon Web Services (AWS), Azure, Google Cloud, IBM Cloud™ and private clouds.

For more information, visit ibm.biz/CloudPak4Data



RTInsights is an independent, expert-driven web resource for senior business and IT enterprise professionals in vertical industries. We help our readers understand how they can transform their businesses to higher-value outcomes and new business models with AI, real-time analytics, and IoT. We provide clarity and direction amid the often confusing array of approaches and vendor solutions. We provide our partners with a unique combination of services and deep domain expertise to improve their product marketing, lead generation, and thought leadership activity.

Copyright © 2020 RTInsights. All rights reserved. All other trademarks are the property of their respective companies. The information contained in this publication has been obtained from sources believed to be reliable. NACG LLC and RTInsights disclaim all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. The information expressed herein is subject to change without notice.

ⁱ <https://www.ifsworld.com/corp/news-and-events/newsroom/2019/11/14/nearly-90-percent-of-companies-have-sights-set-on-artificial-intelligence-investment-ifs-study-reveals/>

ⁱⁱ <https://www.gartner.com/en/newsroom/press-releases/2019-07-15-gartner-survey-reveals-leading-organizations-expect-t>

ⁱⁱⁱ <https://www.ibm.com/blogs/think/2018/02/ibm-ai-ladder/>