AI Everywhere:
Modern Predictive Analytics

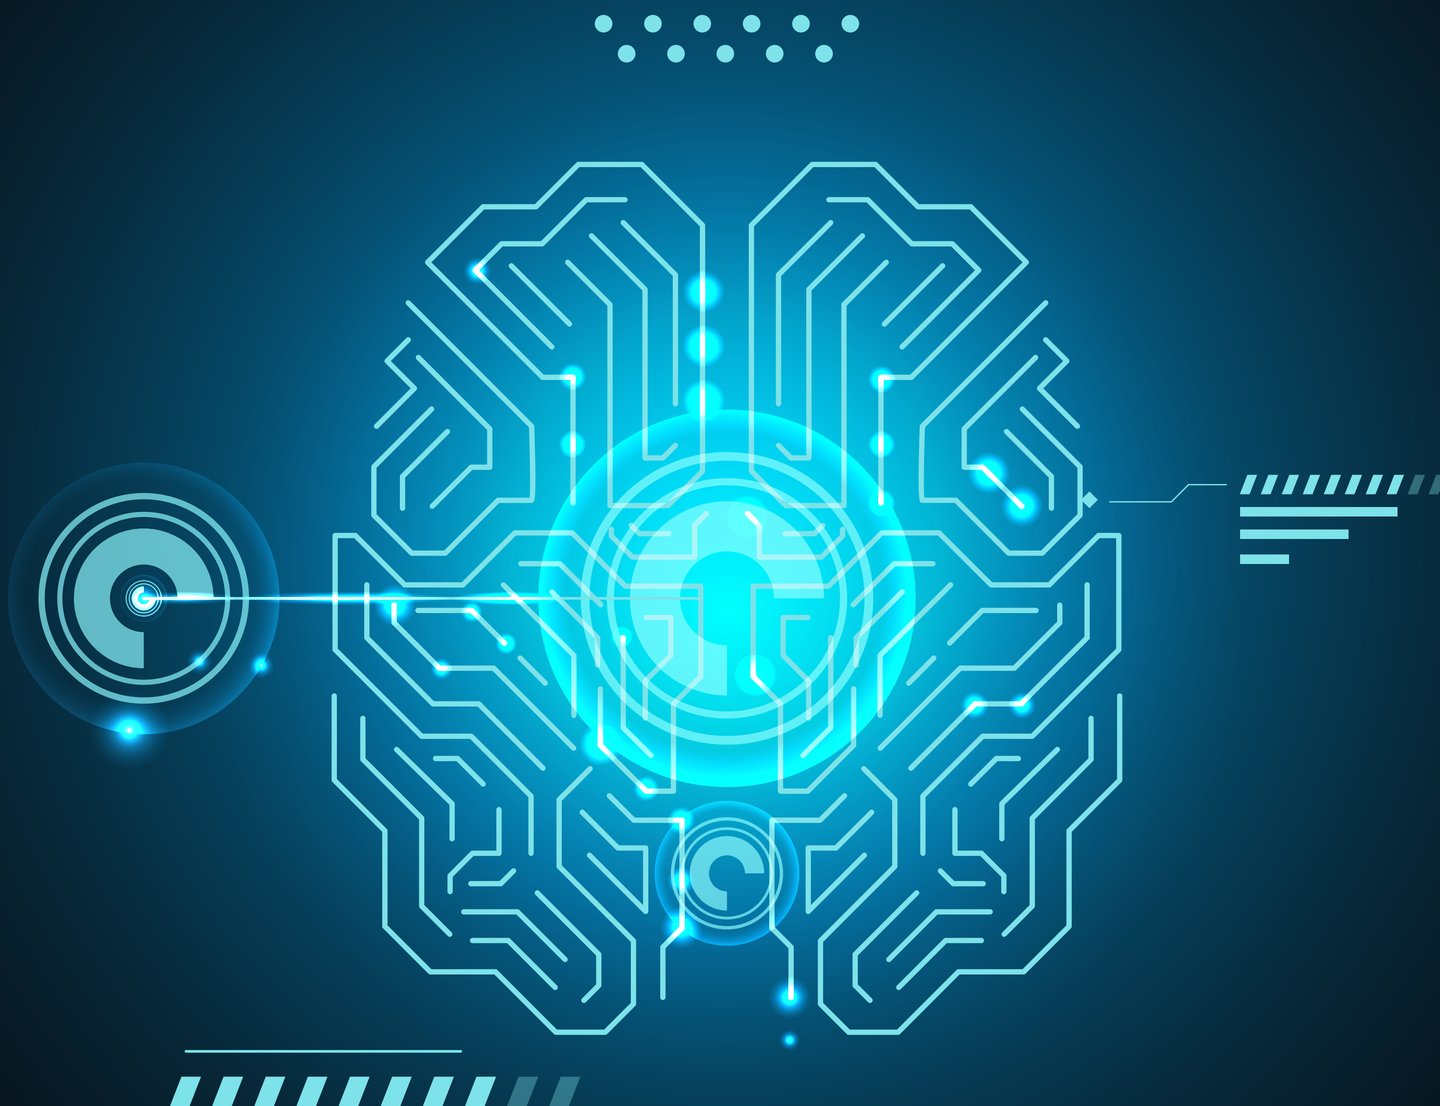

RTInsights.com

IBM

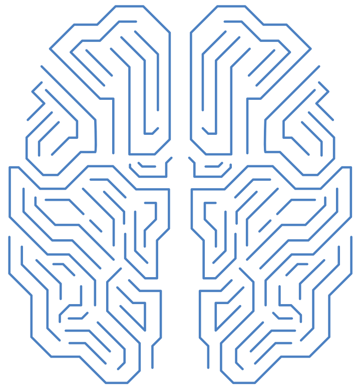

Chapter 1

Introduction

To make faster, smarter decisions, businesses are adopting increasingly sophisticated analytics methods. The transition that is occurring reflects a shift from reporting historical data to making predictions with artificial intelligence (AI). Beyond basic reporting and dynamic dashboards that present business intelligence information, telling companies what has happened or is happening, companies are now using predictive analytics and prescriptive techniques based on AI and machine learning (ML) models that do much more.

The transition from descriptive to predictive analytics and AI is being driven by a need to be less reactive and more proactive. Analytics based on AI and machine learning models simply provide better insights.

The transition from descriptive to predictive analytics and AI is being driven by a need to be less reactive and more proactive. Analytics based on AI and machine learning models simply provide better insights. They build on the benefits of predictive analytics by complementing human decision-making. Often, such analysis is used to recommend one or more courses of action and show the likely outcome of each decision.

For example, a financial services company might use predictive analytics to identify what percentage of customers is most likely to fall behind on loan payments. In contrast, an AI-based analytics approach might provide guidance as to which debt relief approach to use on a specific customer or set of customers.

The real benefits come when such predictive analytics and prescriptive capabilities are embedded into business processes and are used to provide continuous intelligence (CI), allowing for decisions to be based on events that are happening in the moment as well as historical data, enabling actions to be taken in milliseconds to minutes.

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 now and react rapidly. Running prescriptive analytics, ML, and AI algorithms against streaming data can derive actionable information. That information can then be used by systems to decide what to do next and enable the systems to take actions automatically.

Essentially, CI run on streaming data lets businesses move away from traditional descriptive analytics (here's what's happened in the past). CI enhances predictive analytics by applying AI analysis to event streams, allowing businesses to add situational awareness to their decision making. One way to look at this is that predictive analytics might be used to determine what will happen next. A financial transaction looks suspicious, for example. CI can complement this predictive capability using prescriptive analytics or AI rules to determine what to do next in that moment—stop the transaction in real time, preventing loss.

CI for Decision Support and Automation

CI systems learn from new and historical data, using AI to gain insights and act appropriately. CI goes beyond predictive streaming analytics. Like the nervous system, CI propels people to act immediately if necessary. As such, CI can be used for decision support and decision automation. These are applications where time is of the essence.

CI used for decision support helps businesses and people make decisions on how to respond to events that may be rapidly changing. The derived intelligence complements the actions a person or business process would take. So rather than making a decision based on a gut feeling (a change in market conditions, for example), CI would provide the information to pick the right decisions and back up the decision. General application areas for CI in business include risk assessment, research, target marketing, sales acceleration, revenue growth, finding opportunities for funding business growth, and increasing operational efficiencies.

CI used for decision automation takes things to a higher level. Actions are taken without a person involved, based on the derived information. CI for decision automation has many benefits. It offloads work from people; it is a faster, less expensive way to operate than relying on human decisions; it ensures more consistent decisions; and it can guarantee better compliance by following predetermined policies.

The bottom line is that AI-based decision support and decision automation systems have the potential to deliver significant benefits to organizations that need to react in the moment to dynamic situations.

Chapter 2

Architectural and Platform Roadblocks on the Journey to AI

Rather than simply running analytics on data at rest, organizations can achieve greater value by embedding AI capabilities into business processes. However, embedding AI into business processes requires updating information architectures (IAs) to accommodate data in motion and to support continuous intelligence and AI everywhere.

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 by making use of AI.

Unfortunately, there are many obstacles to deploying such solutions. Businesses must deal with data silos, data integrity (quality, compliance, privacy issues, etc.), and vendor lock-in problems. What's needed to enable the required functionality and flexibility is an infrastructure that allows businesses to:

✓ **Continuously analyze data in motion** across multiple sources to deliver actionable insights

✓ **Connect to any data stream** to make predictions and discoveries as data arrives to enhance and improve analytic models and cognitive systems

✓ **Deploy a complete set of streaming analytics applications**—such as natural language processing, geospatial, predictive, and more—to satisfy unique, industry-specific requirements and use cases

If an information architecture is deployed that addresses these issues, businesses will benefit from embedded AI in that it can identify perishable insights—insights that enterprises must act on within a given timeframe or else the opportunity to change the course of business outcomes will pass. The value of such solutions does not stop there, though. They also offer continuous integration and orchestration of business processes—all driven by streaming data from a multitude of data sources.

As a result, embedded AI-based CI solutions need platforms that not only ingest streaming data and perform analytics. They are also execution platforms that can embed code, machine learning models, and rules that businesses can use to detect urgent situations and act immediately.

To embed AI into business processes, the selected information architecture must support a distributed runtime environment and include tools for developing applications that make use of streaming and events-based data.

A suitable solution must offer flexible deployment options. Businesses must be able to run their analytics on-premises and on any cloud, depending on their specific requirements. For example, data privacy issues might require more control over the data and dictate an on-premises approach. Alternatively, the compute requirements to run AI algorithms or train ML models might exceed installed capacity, making a cloud solution a better option. And in other cases, a temporary need for additional processing power might be ideally met using cloud services to complement on-premises capacity.

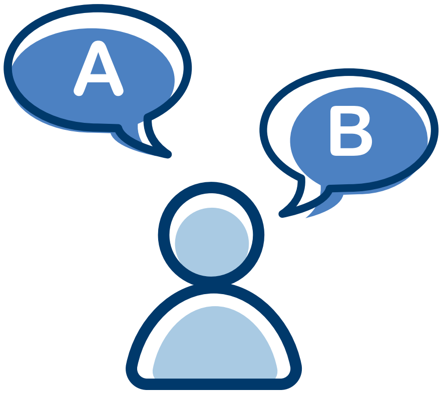

An additional architectural issue to consider is where AI is applied: cloud or edge? Factors that will determine where to run the analytics include efficient use of bandwidth and latency.

Regarding bandwidth, in some cases, it makes no sense to bring all the data from the edge to a central location to run the analysis. A good example of this is the ECG function on an Apple Watch. The watch makes the measurement and then determines if there is an anomaly such as atrial fibrillation (AFib). If so, an alert (rather than all the data) can be sent to a doctor.

In other cases, there simply may not be a way to get the data back to the processing center. For example, a utility might have sensors that monitor transformers in remote locations. If there is no communications channel with enough bandwidth to get all that data back-hauled to a central location for analysis, a better approach would be to do the analytics on the spot and only send a notification if a problem is detected or predictive maintenance is required to prevent downtime.

Even if a communications link is available that delivers the required bandwidth, it still might make sense not to send all the data. In many cases, it is the law of economics. Is it economically feasible to send all that data to one place for analysis? Often, the answer is probably not, when mobile operators charge by the byte.

Latency issues bring other factors into play. As IoT, video, and other data become a more important part of enterprise business operations, the ability to reduce latency in data analytics and processing can make a difference. Latency is particularly important in situations where decisions must be made in milliseconds to seconds. An example is video image analysis to prevent an autonomous vehicle from hitting a pedestrian. If the process is done in the cloud, the video stream must first be sent to the cloud, the AI analysis must then run, and if a pedestrian is detected, instructions must be sent to the vehicle to turn to avoid the person. Realistically, there is not enough time. The AI-based video analysis must be done on the edge, in the vehicle.

Chapter 3

Benefits of Using a Cloud-Native Approach to Develop and Deploy a Modern Data and AI Platform for Predictive Analytics

Events occur in a continuous stream as things happen in the real and digital worlds. By taking advantage of this continuous stream, applications not only may react in real time, but also reason about the future based upon what happened in the past.

Embracing event-driven development for CI applications is foundational to the next generation of digital business applications that embed AI everywhere. Businesses will need to be able to quickly design, develop, deploy, and operate event-driven solutions to have the required agility and speed to innovation to stay competitive.

Increasingly, cloud-native is the architecture of choice to build and deploy AI-embedded CI applications. A cloud-native approach offers benefits to both the business and developers.

How is this achieved? Cloud-native CI applications are a collection of small, independent, and loosely coupled services. At its core, a cloud-native architecture has microservices and containers that use cloud-based platforms as the preferred deployment infrastructure.

Microservices provide the loosely coupled application architecture, which enables deployment in highly distributed patterns. Additionally, microservices support a growing ecosystem of solutions that can complement or extend a cloud platform. Cloud-native platforms with containers and serverless deployment provide the application platform and tools to make use of a microservices architecture.

Using such an architectural approach, businesses can easily move workloads to the most appropriate platform (on-premises, cloud, multi-cloud) at any time. Additional benefits of using a cloud-native architecture for modern applications include:

Faster development and deployment. Time to market is a critical differentiator in today's marketplace. Cloud-native applications using modern DevOps techniques allow businesses to automate many aspects of application development, testing, and deployment. As a result, businesses can quickly create new applications and rapidly deploy them. Thus, they can react to market changes and meet changing customer priorities.

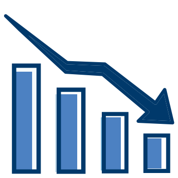

Reduced costs. Cloud-native applications benefit from containerization. Why? Containers make it easy to manage and secure applications independent of the infrastructure that supports them. Increasingly, businesses are using Kubernetes to manage containers and resources in the cloud. When the use of Kubernetes and containers is combined with enhanced cloud-native capabilities such as serverless deployment, businesses can run dynamic workloads and pay-per-use for compute time in milliseconds. This ultimate flexibility in pricing is enabled by cloud-native.

Ease of management. Cloud-native offers ways to make infrastructure management simpler. For example, serverless computing platforms let businesses upload code in the form of functions, and the platform runs those functions without having to manually provision cloud instances, configure networking, or allocate enough storage. Serverless takes care of it all.

Flexibility to incorporate new technologies. Businesses need to keep pace with rapid changes in the field. That may require adding new analytics methods to enhance the capabilities of a CI application. For instance, a CI application to support customer service hub agents might want to incorporate different voice capabilities (e.g., speech-to-text features, and vice versa, using newly available natural language processing routines). A cloud-native architecture would use APIs to easily connect different (and new) analytics solutions offered as microservices.

Extensive use of provisioning automation. Unlike traditional applications where the resources are physically provided, cloud-native applications can automatically provision resources. They facilitate on-demand, self-service provisioning, and automate the release of compute and storage resources when these resources are not being used.

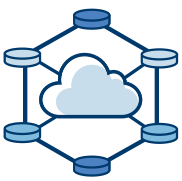

High levels of resiliency. CI applications that make use of AI for real-time analysis of streaming data deliver critical insights. They must be highly available. Application downtime can have serious consequences. An offline fraud detection application will not be able to stop that one bad transaction, or an offline predictive maintenance application will miss the telltale signs of a part failure leading to the shutdown of a production line. Cloud-native applications are inherently resilient to failure as they can automatically handle outages, enable corrective actions, and shift workloads.

Taking these benefits together, it is clear that a cloud-native architecture provides a way for businesses to build new CI applications that incorporate modern analytics. The bottom line is that while event-driven architectures and reactive programming models have been employed in the past, the move to cloud-native architectures with microservices, container-based workloads, and serverless computing is making them more practical and provides many benefits. Such an event-driven architecture offers resiliency, agility, and scalability. The result is that cloud-native CI solutions are both reactive and responsive.

Summary

AI everywhere has the potential to improve operations, enhance customer engagements, and speed the time to insight. The greatest benefits are realized when AI is embedded into business processes. Unfortunately, there are obstacles that can prevent this. The selection and use of the right platform will eliminate obstacles that might limit the benefits of using embedded AI.

Such a platform should:

- Eliminate data silos and provide easy access to data
- Automate and govern data and AI lifecycles
- Avoid vendor lock-in
- Run everywhere
- Embrace open source technologies
- Deploy quickly
- Operationalize AI

By addressing these issues, businesses can quickly adapt to changing customer and market conditions and thus truly benefit from the use of continuous intelligence and AI everywhere.

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 for AI that meets organizations' ever-changing enterprise needs.

Deployable in just hours and easily extendable with a growing array of IBM and third-party microservices, 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 Cloud Pak for Data online.