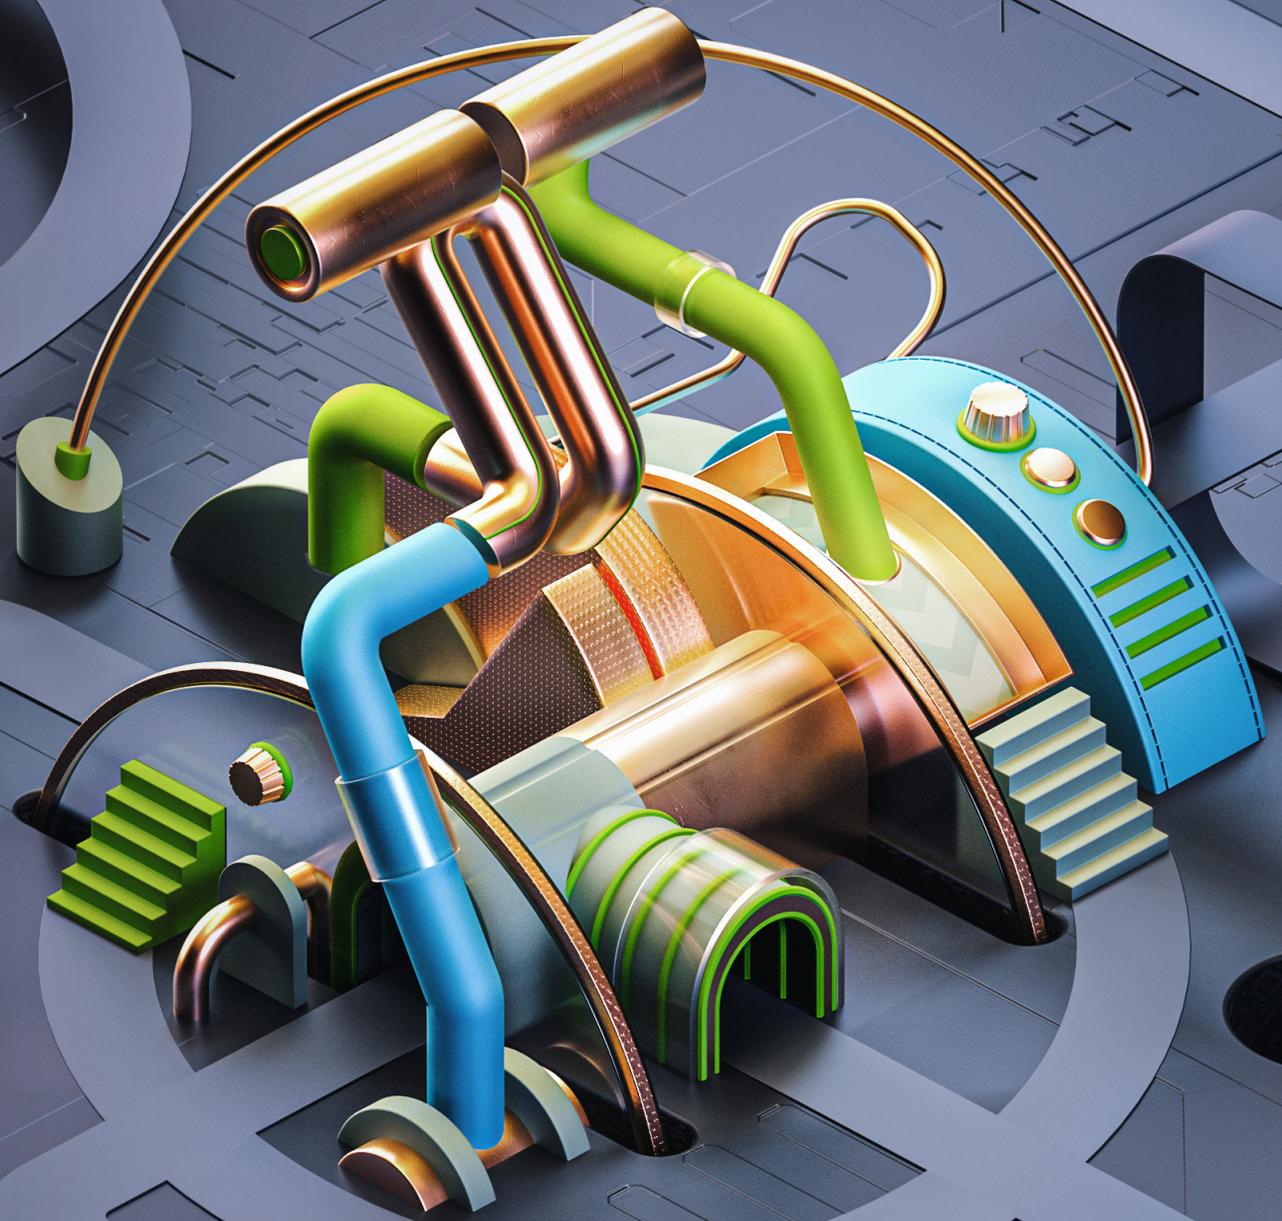


**Deloitte.**  
Insights

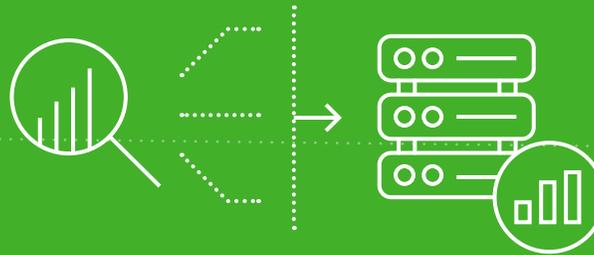
# Tech Trends 2021



# Machine data revolution: Feeding the machine

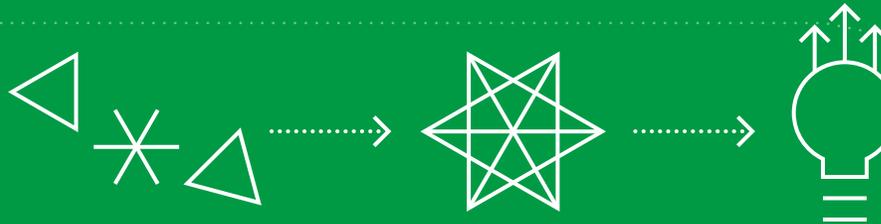
## CAPTURE AND STORE

Search legacy systems and databases for valuable unstructured and nontraditional data. Aggregate and store it in cloud-based offerings optimized for machine learning.



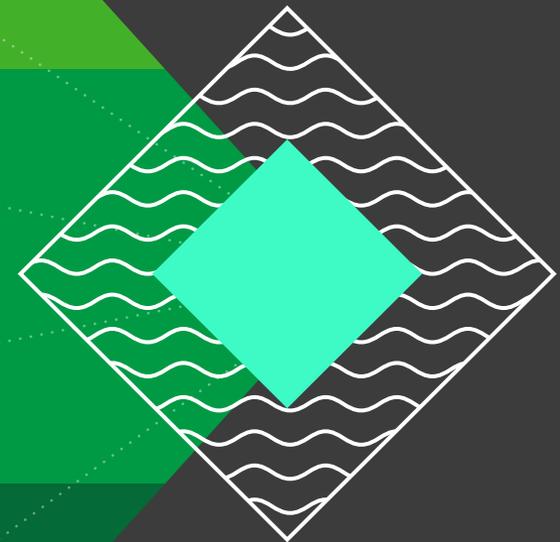
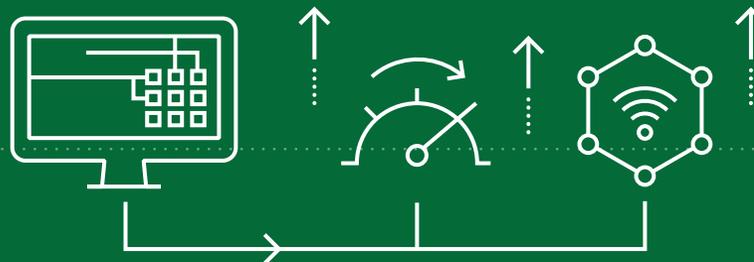
## DISCOVER AND CONNECT

Use cognitive data steward technology to illuminate insights and connections among disparate data.



## AMPLIFY ML CAPABILITIES

Explore opportunities to use edge computing and 5G to boost real-time data capabilities and decrease latency.



## TREND 5

# Machine data revolution: Feeding the machine

## Disrupting the data management value chain for the ML age

With machine learning (ML) poised to augment and in some cases replace human decision-making, chief data officers, data scientists, and CIOs are recognizing that traditional ways of organizing data for human consumption will not suffice in the coming age of artificial intelligence (AI)-based decision-making. This leaves a growing number of future-focused companies with only one path forward: *For their ML strategies to succeed, they will need to fundamentally disrupt the data management value chain from end to end.*

In the next 18 to 24 months, we expect to see companies begin addressing this challenge by reengineering the way they capture, store, and process data. As part of this effort, they will deploy an array of tools and approaches

including advanced data capture and structuring capabilities, analytics to identify connections among random data, and next-generation cloud-based data stores to support complex modeling.

Some companies are already embracing this trend as part of larger AI initiatives. In Deloitte's third annual *State of AI in the Enterprise* survey, when asked to select the top initiative for increasing their competitive advantage from AI, respondents singled out "modernizing our data infrastructure for AI."<sup>1</sup>

For digital nonnatives participating in this trend, the stakes are high. Some of their digital native competitors, largely unburdened by outmoded data models and processing

capabilities, are already monetizing more diverse data, more quickly.<sup>2</sup> Importantly, end users have less and less patience for the kind of latency that legacy systems and data models often deliver. The optimum latency time between click and desired response is less than 50 milliseconds—any longer and users become irritated by the delay and make "executive decisions" themselves.<sup>3</sup>

## Humans, machines, and data

In the coming months, participants in the *machine data revolution* trend will explore opportunities to reengineer their data management value chains to support ML's possibilities. In the arena of data management,

this marks a distinct change of course. For decades, companies have collected, organized, and analyzed data with one goal in mind: helping humans make decisions based on statistical fact rather than hunches and emotion. Humans tend to look at aggregated data characterized by two or three major factors. When faced with more complex data, many humans struggle to process the information presented and to articulate a useful decision. As such, we typically organize data for humans in clean tables and rows, with precise labeling. Machines, by contrast, can assess multiple factors simultaneously and objectively. ML models can extract low levels of statistical significance across massive volumes of structured and unstructured data. They work around the clock and can make clever decisions in real time.

When used in areas in which human decision-making is nonscalable—such as cleaning up raw data<sup>4</sup> or making personalized product

recommendations—ML may only need to make *good enough decisions*, not perfect ones. For example, a retailer would presumably see value in the ability to recommend, in real time, an assortment of products tailored very broadly to thousands of individual online shoppers simultaneously. The products that ML algorithms recommend might not perfectly match each customer's unique tastes, but they might be sufficient, in that moment, to drive a sale. Across an enterprise, each good-enough data-based decision that machines make, rather than humans, drives down the overall cost per decision, which in turn enables companies to extract value from even the lowest-level decisions. Creating an automated pipeline that replaces low-level or nonscalable human decisions with those made by machines brings to mind the promise of Moore's Law. Over time, speed and capability will increase so dramatically that making that data-based decision in the future will cost a fraction of what it does today.

Though approaches can vary by industry, market, and organizational need, many trend participants will likely focus on their reengineering efforts on the following areas:

### *Capture and store*

Chances are, your organization has troves of data that's potentially valuable yet untapped. Some of it is probably traditional enterprise data residing in databases, files, and systems; other troves may be more recent data generated by machines or mobile devices. Still others may be unstructured text, or nontraditional data from video or audio recordings. In all likelihood, this data was previously too hard or too expensive to capture and utilize in a cost-effective way, so it lies fallow. This is a lost opportunity. No one knows which data amid vast stores of raw information might turn out to be predictive or confer some decisioning value down the line, so it is critical to capture all the data you can.

Moreover, you are probably throwing out some data today that, with the right tools and approaches, you can use. Take utility companies, for example. What information do they need to predict power or equipment failures? Traditionally, they may have collected data only on failure. But for predictive purposes, they would also need data on uneventful everyday operations to understand what *normal* looks like. This same idea applies to people visiting your company's website. Do you have website data for both success and failure? In a world where data quality no longer matters as much as it once did, what changes can you make to your current data practices to make them more predictive?

In terms of storage, organizations are becoming less focused on storing clean data that fits neatly into tables, rows, and columns. To feed ML algorithms and advanced analytics tools, many are exploring opportunities to store massive volumes of unstructured data

from IoT, social media, and AI in a variety of modern database technologies, including:

- **Cloud data warehouses.** The cloud-based data warehouse, which a growing array of major and emerging public cloud vendors are offering as a service, aggregates data from disparate sources across an enterprise and makes it available to users for real-time processing and mining. This permissions-based, centralized system eliminates the need for colocated data and data pipelines. In addition to collation and storage capabilities, cloud data warehouses also typically offer search engine tools for querying data and analytics capabilities.<sup>5</sup> This combination of public cloud ease-of-use, the ability to scale up or down as needed, and advanced data processing and analysis tools is fueling considerable growth in the cloud data warehouse market. Prescient & Strategic Intelligence

forecasts the data warehouse-as-a-service market will reach US\$23.8 billion in value by 2030.<sup>6</sup>

- **Feature stores.** In the near future, it will be commonplace for an organization to have hundreds or thousands of data models operating independently of each other, and in parallel. Each of these models will use different feature sets. For example, some require immediate decisions while others do not, thus placing broadly different demands on data and on processing power. Pushing real-time compute uniformly in every model is a waste of computing power. Likewise, some models probably share features while other features may be used exclusively in a single model. How can you manage all of these competing demands across data models? Feature stores provide a mechanism for allocating compute, sharing features, and managing data efficiently,

and at scale, which makes this mechanism integral to driving down decision costs. What's more, by leveraging AI, feature stores may eventually be able to predict demand of certain features based on the types of data being modeled.<sup>7</sup>

- **Time series databases.** The popularity of time series database technologies has grown considerably over the last two years, with good reason.<sup>8</sup> Unlike relational databases that record each change to data as an update, time series databases track and record them—and the specific time they were made—as a unique insert into a dataset. With the explosion of temporal data from IoT and monitoring technologies, among others, both historical and predictive analysis increasingly depends on the ability to query a data value from one point in time and track it continuously, accurately, and efficiently.<sup>9</sup>

- **Graph databases.** Highly interconnected data can be challenging to analyze and use to its fullest potential. Using traditional relational databases in which data are organized in tables, one can identify and manage a limited number of data relationships. But as data grows more voluminous and less structured, the number of relationships and interconnections increases exponentially, thus becoming unmanageable (and unsearchable) in traditional database models. Graph databases are designed specifically to address this challenge by storing not only data but information on each data point's relationships in a native way. With this model, queries about complex relationships among data can be fast, efficient, and more accurate.<sup>10</sup>

With storage costs continuing to fall, aggregating and organizing massive volumes of data is no longer cost-prohibitive.<sup>11</sup> What's

more, modern self-healing, fault-tolerant data architecture typically requires less maintenance, which can reduce administrative and repair costs. Thus, the potential benefit of increasing storage capacity could far outweigh whatever costs you may incur. ML and advanced analytics can discern low levels of statistical significance across a large number of factors, which in turn can provide a significant lift that would be near impossible to achieve using traditional data storage and modeling techniques.

**Organizations are becoming less focused on storing clean data that fits neatly into tables, rows, and columns.**

## Discover and connect

As you begin capturing more data, it will likely include fragmented data generated across different devices, channels, and geographies. How can you connect fragmented data in a way that characterizes an individual customer in an individual context—or reveals an unmet need the marketplace or an internal opportunity for greater efficiency? Unlocking the full value of all data resources, including dark and nontraditional data, can be complex and expensive, particularly in large, established enterprises with hundreds of legacy systems, duplicate data stored around the globe, and inconsistent naming practices. As you start work to build data's future-ready foundation, you will likely face a twofold challenge. First, to make the strongest data-driven decisions, you will need to analyze more than just the obvious data. Indeed, you will need the *nonobvious* data—information that no one knows even exists. Then, even

if you can collect all known and unknown enterprise data, how can you tie these disparate, inconsistently formatted and named data points together in a way that is meaningful? The work of discovering and connecting enterprise data can be formidable and costly. Yet shirking this challenge could cost even more if your company misses out on potentially valuable opportunities.

## To make the strongest data-driven decisions, you will need to analyze more than just the obvious data.

The good news is that ML-powered *cognitive data steward* technologies available today can help accelerate the processes of discovering data and illuminating its insights and connections.<sup>12</sup> Here's how:

- Analytics, semantic models, and cognitive technology can automate manual, costly stewardship activities—thus freeing up data scientists to focus on more advanced analysis.
- Identifying similarities in underlying data systems' code makes it possible for data scientists to use custom data algorithms in multiple data models.
- Finally, by leveraging ML capabilities to automate the processing of master data, cognitive data stewards can help users visualize relationships in data, improve data readiness and quality, and enable greater data management efficiency.

In the near future, expect data steward capabilities to grow with new tools that aid with ingestion, classification, management, and discovery. And as the trend gains momentum in the coming months, look for

data steward deployments to expand further into transactional systems, supply chain ecosystems, and smart factory environments.

## Serving up ML's secret sauce

The ability to process larger volumes of diverse data in real time is the secret sauce of ML-based data decisioning. The faster that big data systems can capture and process data, feed it into ML and analytics platforms, and then serve up insights to users, the more impactful your data investments can be.

To this end, a growing number of organizations are exploring ways to make decisions at data's point of entry into the network rather than sending it first to the core or cloud. Some are building edge computing capabilities that can decrease latency in data systems while also making these systems more reliable and efficient. Edge computing means pushing

compute and processing power away from a centralized point and closer to a network's "edge" or periphery. It does not replace enterprise or cloud-based data centers but helps distribute processing work—including analysis and decisioning—more evenly across a network. Rather than sending raw data back to a cloud or data center, a device operating at the edge generates action independently or sends only already-refined data to the network, in effect storing, processing, analyzing, and reacting locally. Edge computing can be particularly useful when deploying ML algorithms, which require uninterrupted, real-time access to large quantities of recent data.<sup>13</sup>

Advanced connectivity also has an important enabling role to play in real-time decision-making at data's first point of entry. Current generation connectivity, such as 4G/LTE and Wi-Fi, can support some edge computing and real-time data processing needs, but they are limited by bandwidth, latency, and the number

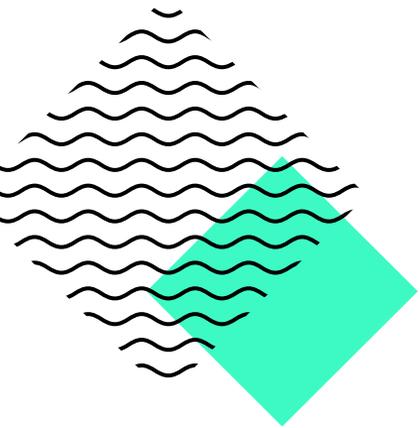
of devices they can effectively manage. 5G can deliver faster speeds and millisecond latency. It can also expand bandwidth capacity to simultaneously manage many more devices per square kilometer.<sup>14</sup>

## The way forward

As we discuss in this report's *ML Ops: Industrialized AI* chapter, ML initiatives are gaining momentum across industries. Indeed, the ML technologies market is currently growing at a rate of 44% annually and is expected to reach US\$8.8 billion in value by 2022.<sup>15</sup> But ML algorithms and platforms will deliver little ROI in companies with outdated data infrastructure and processes that were leading-edge in 2002.

How will you reengineer your data strategies to build a new foundation for your company's future?

# LESSONS FROM THE FRONT LINES



## Adventures in data democracy

Since its founding 144 years ago, AT&T has reinvented itself many times to harness historic disruptive innovations such as the transistor, the communication satellite, and more recently, the solar cell.<sup>16</sup> Today, the global technology, media, and telecommunications giant is reinventing itself again—this time as a pioneer in the use of ML, which it is deploying broadly in areas such as IoT, entertainment, and customer care.<sup>17</sup>

The company is also leveraging ML to reimagine the way it finds, organizes, and uses data. “One of the things we wanted to do was automate some of the routine cleansing and aggregation tasks that data scientists have to perform so they could focus on more sophisticated work,” says Kate Hopkins, vice president of data platforms, AT&T Chief Data Office.<sup>18</sup> Likewise, the company wanted to develop a way to democratize meaningful data, to the extent consistent without privacy, security, and other data use policies, making it more broadly available to qualified personnel across the enterprise. These efforts, Hopkins says, have already borne fruit. New tools have shrunk the time to market required to go

from prototype to full scale production for ML models. These models have had dramatic results, such as blocking 6.5 billion robocalls to customers, deterring fraud in AT&T stores, and making technicians visits to customer homes more efficient.

AT&T started its data transformation journey in 2013 when it began aggregating large volumes of customer and operational data in data lakes. In 2017, the company created a chief data office with the goal of leveraging these rapidly growing data stores for “hyper-automation, artificial intelligence, and machine learning.” The ongoing work of achieving these goals has presented several significant challenges. First, in a company as large as AT&T, it was sometimes difficult to find and access potentially valuable data residing in legacy systems and databases. And even when data scientists eventually found such data, they occasionally struggled to understand it, since it was often labeled inconsistently and offered no discernable context or meaning. Finally, there was a formidable latency challenge across all data systems that, left unaddressed, would stymie the real-time data needs of ML models.

To address these challenges, the chief data office developed the Amp platform. Amp enables a culture of technology

and data-sharing, reusability, and extensibility at AT&T. Pari Pandya, director of technology and project manager for Amp, says that what began a few years ago as an internal online marketplace (aggregating microservices, APIs, chat bots, designs, etc.) for accelerating automation, has evolved into a single, powerful source of data truth for systems and users. Consider this: As data flows through multiple systems and processes, its definitions change. Amp not only finds legacy system data, it uses metadata to ascribe meaning to this data, and provides a clear lineage to help users better understand the data. “It serves as a business intelligence platform that provides not only meaningful data but analytic and visualization tools that empower business teams, strategists, and product developers to leverage data in more advanced ways and share insights through data communities,” Pandya says.<sup>19</sup>

To meet the challenge of latency, AT&T is on a multiyear journey to move some of its

data and tools to the public cloud. Working closely with cyber teams to ensure data and IP security, the company is leveraging the cloud’s ability to scale up compute power as needed. The cloud’s power is helping create the real-time access that ML—as well as enterprise stakeholders and customers—require. Unlimited access to compute on demand through the cloud and the availability of business-ready data is accelerating the journey.

Hopkins notes that AT&T’s data transformation journey has yielded another welcome benefit. “The business units have become much more knowledgeable about data science and are identifying opportunities to use data in new ways. Across the board they’re requesting much more mature and sophisticated data,” she says, adding that “being able to democratize data and make the process transparent across the enterprise can deliver exponential payback.”

---

## Data and IT double-team digital transformation

---

How can a 100-year-old retail organization efficiently and accurately take data from legacy applications that were designed for very specific use cases to accomplish something that those applications were never intended to do?

“Every legacy company faces this challenge,” says Paul Ballew, chief data and analytics officer at the leading Canadian food and pharmacy retailer Loblaw.<sup>20</sup> “You have to bring those data assets together from across your ecosystem in a way that’s scalable, repeatable, and governable, which is no small task.”

Taking an ecosystem approach to data is particularly formidable in a successful retail organization such as Loblaw, which operates 2,400 stores and maintains an expansive

e-commerce presence. “We are a legacy company trying to leverage technologies that digital natives are born with,” Ballew says.

Yet despite its challenges, data represents a unique opportunity on the path to Loblaw’s digital future. And unique opportunities require unique approaches. Like many digital nonnatives, the company is shifting its focus from traditional data management priorities such as storage, curation, and quality to a new, more complex arena in which data analytics and digital solutions drive day-to-day operations. “It requires a different approach to ‘baking the soufflé,’” Ballew says. “We source and mix ingredients differently, and then serve it in new ways to those consuming it.”

Recognizing the critical importance of data in the company’s digital future, Loblaw set up a distinct data organization that works in tandem with IT to drive digital transformation and engage the business.

From a technology standpoint, Loblaw takes a three-layered approach:

- **Data layer.** An array of data management and digital capabilities that make data assets, many from legacy applications, consumable in near-real time to a variety of complex use cases.
- **Analytics and development layer.** A collection of AI/ML and advanced analytics technologies that bring data assets to life and glean insights from structured and unstructured data to support better decisions and more efficient workflows.
- **Solution delivery layer.** Provides tools to systematically integrate decisions and insights into processes and applications, helping meet the organization’s digital strategy goals.

“Once you’ve coordinated these three layers, you have to manage the analytic solutions, refresh cycles, and monitor to ensure strong and consistent governance,” Ballew says, noting that because Loblaw does a variety of things from selling groceries to providing pharmacy and health services, the company has a sliding scale in terms of data sensitivity, privacy concerns, and legal compliance. As such, robust data governance is critical in protecting sensitive data and determining model sensitivity to bias and other factors. “We have to be proactive stewards of customer data and leverage it in a manner that results in providing benefits to them in a transparent manner,” he says.

Data’s ascent as a decision-driving, business-critical asset has redefined the roles of Loblaw’s data and IT teams. “Our work around data and digital helps the business make critically important decisions: who to talk to, how to optimize marketing, run a factory,

or engage customers—the list goes on,” Ballew says. “In terms of data, we helped the organization and leadership understand the art of the possible and the implications—good and not-so-good—of taking a comprehensive approach. The change has been beneficial overall, but it has impacted our entire ecosystem and those working in it.”

And Ballew’s advice to those managing similar change in their own companies? “Seek to understand before you seek to be understood.”

---

## Banking on distributed data architecture

---

ABN AMRO is taking a modern approach to data management. Rather than engineering endless workarounds to accommodate problems with the data pulsing through its systems, the Netherlands-based global bank has developed a feedback mechanism

that enables data scientists to request data quality issues be fixed at the source and focus on turning data into value. “In the past, data scientists would find a problem, fix it, and keep going,” says Santhosh Pillai, chief architect and data management. “Now they can provide feedback to the source where data is mined, and say, ‘do it differently.’ Over time, data quality improves, and data scientists don’t have to spend as much time on cleansing and querying.”<sup>21</sup>

Strengthening governance at the source is just one component of a three-pronged approach the bank is taking to prepare for what Pillai calls “the AI decade”—an era when AI increasingly augments or even replaces human decision-making. The second component focuses on the consumption side, where ABN AMRO has engineered an advanced analytics and AI layer to support business strategies that are evolving rapidly. “In an increasingly digital world, being client-centric means being data-centric,”

Pillai says. “Particularly in the post-COVID era, companies can’t meet face-to-face with clients, so they rely more heavily on data and analytic insights. The analytics capabilities we have in place deliver these insights and unleash the value contained within our data.”

The third component of ABN AMRO’s data transformation effort is a multifaceted data mesh model that moves data anywhere it needs to go within the ecosystem, from source all the way to consumer. This “data supply chain” serves not only as a distribution mechanism but as a timing guarantee mechanism that enables real-time access to meet demand. It also features a self-service “marketplace” where consumers of data—both human and machine—can access high-quality data that is usage-approved and regulatorily compliant.

Like many established organizations, ABN AMRO didn’t originally design its data architecture

to be event-driven—or for current data usage patterns. Today, algorithms and end users read up-to-the-minute data far more frequently than they use it in transactions. Legacy data management models were not designed to respond to constant read queries and real-time updates.

“We solved this challenge by putting each original record in a data store and replicating it,” Pillai says. “On the consumer end, users see replicated data delivered with minimal latency and think they are seeing real-time data generated at the point of consumption. In fact, that data they are reading is coming from another part of the ecosystem.”

Pillai sees great potential in this data replication model, particularly in the area of cloud storage. “Traditionally, technology was designed to optimize data storage. But as we approach the AI decade, I expect to see more companies develop mechanisms for

replicating data that is stored in several clouds and even moving that data between multiple cloud vendors.”

MY TAKE

## Lutz Beck

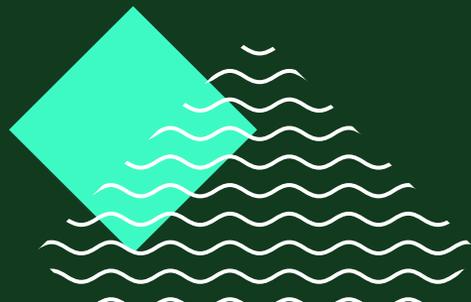
Chief information officer,  
Daimler Trucks North America



Daimler Trucks North America, a leading producer of heavy-duty commercial vehicles, is transforming itself into an intelligent company—one in which data is a key asset.

Whether we are becoming more efficient through automation, creating new services for customers, or making better decisions, by using analytics and other digital technologies to work with data in real time, we can see things in different ways and steer our efforts in new directions.

Consider, for example, our trucks. Each Daimler truck that rolls off the factory floor is a new digital asset. An array of onboard sensors and other technologies monitor vehicle performance continuously, generating data that offers Daimler real-time insights on a truck's health. But rather than simply providing the vehicle owner with a status report, we can now apply analytics to vehicle performance data to predict when a part might fail. In urgent cases, we share this



information with the vehicle owner and initiate a “service now” event, directing the owner to a nearby service facility where mechanics—made aware of the problem and confirmed to have the needed part in stock—can address the issue without delay.

The ability to provide this kind of predictive vehicle performance data presents an opportunity to create an entirely new value-added service. Now, because we can predict and solve a maintenance issue proactively, we can offer our customers an “uptime guarantee”: In this instance, data presents an opportunity for us and our customers.

Of course, thinking of our trucks as digital assets and fully embracing data-driven decision-making represents a cultural shift in a traditional company with deep roots in industrial manufacturing. Traditionally, we designed and built vehicles for sale to customers and, later, provided vehicle service

opportunities. In the intelligent-company model, our relationships with customers deepen after the vehicle purchase. We use technology and data to refine existing products and services to meet unique needs as well as to create new services around the vehicle. This is a completely different way of working—one that we are embracing wholeheartedly.

## Data presents an opportunity for us and our customers.

It is important to note that this new way of working presents challenges as well as opportunities. Data volumes are growing, and in our industry, the pace of this growth will accelerate dramatically with the standardization of automated, connected vehicles, smart traffic management systems, and other digital transportation advances. As such, data governance is more important

than ever. In addition to managing the data itself, there are more and more regulations governing data use, and we must understand which data can be used and which services can be provided while always maintaining regulatory compliance. Likewise, we must clearly understand customer preferences and expectations for the way we use their data, which also affects the way we offer services. Should we offer individual services, or can we bundle them? How do different expectations of privacy in the nations of North America come into play? To meet this complex challenge, we are setting up a data intelligence hub where a chief data officer and data analysts work with our growing data catalog. This team of data experts helps us put in place the governance we need to leverage data to its fullest potential within legal parameters.

We are still on our journey toward becoming a fully data-driven, intelligent company. Right now, the biggest limiting factor remains our

thinking and our behavior, so we must make sure we build a culture in which we learn and think about data holistically. We need to work with the data and look at it in a completely different way. If we do this, within the next three to five years we will reach our goal: Daimler Trucks will have all the data-driven insights necessary to anticipate what our customers and dealers need tomorrow, and the day after that, and on into the future.

# EXECUTIVE PERSPECTIVES



**STRATEGY //** CEOs are rarely involved at the deeper levels of data management. But those whose companies are deploying ML and other data-dependent AI tools may want to spend some time thinking about data strategy. Existing enterprise systems and data pipelines were not designed with AI needs in mind, so CEOs can engage their IT teams to lay the groundwork for this shift. They should hold the CIO or chief data officer accountable to carry out preparations for turning the organization's MLOps vision into a reality. They can also expect ongoing investments, as data is gradually prepared for algorithms first instead of humans. These new investments in data can eventually be paid back through increased profits or lowered cost per decision.



**FINANCE //** As the capabilities of modeling and ML become more robust (as detailed in our [MLOps trend](#)), finance leaders are leaning into data stewardship and even sponsoring data governance roles. With better data comes better modeling, which means CFOs can more accurately produce internal reports, predict quarterly performance, and shape the perceptions of market analysts and shareholders. However, the data needed for good modeling is not always internally owned—consider the example of parts manufacturers that monitor flight statistics to predict aluminum demand. Such combinations of raw, external data with company data can lead to actionable insights amid uncertainty. To enable these insights in the next 18–24 months, CFOs may need to think outside of the box to source new types of data, enable agile model governance, and optimize for modeling outcomes.



**RISK //** In the [Age of With](#),<sup>22</sup> data strategies focus on providing analysis and insights that augment human decision-making. In this construct, algorithms—rather than data scientists and executives—may be consuming raw and unstructured data, but the inherent risks associated with managing large volumes of data remain constant. CROs can play a key role in defining risk parameters around protection, confidentiality, and integrity of large volumes of ML-optimized data. They can also develop separate risk strategies for more sensitive customer data and determine whether and how such data can be used in ongoing MLOps efforts. By delving deeper into these issues, CROs can help their organizations become more deliberate about where to apply AI in the business and, perhaps, explore opportunities to monetize risk data with new products and services.

# ARE YOU READY?



## KEY QUESTIONS

1 How have you reevaluated data governance and data management as you move to increase the use of AI and ML?

2 What mechanisms support accessing data from key legacy systems? Do they meet your current and future needs?

3 Which untapped sources of data could have the greatest impact on your decision-making?

## LEARN MORE



### Data management barriers to AI success

*Learn* about the challenges to preparing data for AI initiatives and where businesses can find support.



### Exploring the future of artificial intelligence

*Read* about the many ways AI can augment standard business processes for better outcomes.



### Seven lessons COVID-19 has taught us about data strategy

*Study* the key lessons learned from governments that have used data to combat the coronavirus pandemic.

# AUTHORS

*Our insights can help you take advantage of emerging trends. If you're looking for fresh ideas to address your challenges, let's talk.*

## **Paul Phillips**

Hux chief technology officer

Deloitte Consulting LLP

[paulphillips@deloitte.com](mailto:paulphillips@deloitte.com)

## **Irfan Saif**

US AI Strategic Growth offering leader

Deloitte & Touche LLP

[isaif@deloitte.com](mailto:isaif@deloitte.com)

## **Sandeep Sharma, PhD**

Deputy chief technology officer

Deloitte Consulting LLP

[sandeepksharma@deloitte.com](mailto:sandeepksharma@deloitte.com)

## **Juan Tello**

Chief data officer

Deloitte Consulting LLP

[jtello@deloitte.com](mailto:jtello@deloitte.com)

## SENIOR CONTRIBUTORS

### **Edward Bowen**

Managing director,

Deloitte & Touche LLP

### **Matt James**

Specialist master,

Deloitte Consulting LLP

### **Carl Gerber**

Managing director,

Deloitte & Touche LLP

### **Adarsh Gosu**

Managing director,

Deloitte Consulting LLP

### **Jason Price**

Managing director,

Deloitte & Touche LLP

### **Piyush Sacheti**

Senior manager,

Deloitte Consulting LLP

# ENDNOTES

1. Beena Ammanath, Susanne Hupfer, and David Jarvis, *Thriving in the era of pervasive AI: Deloitte's State of AI in the Enterprise, 3rd Edition*, Deloitte Insights, July 30, 2020.
2. Suketu Gandhi et al., "Demystifying data monetization," *MIT Sloan Management Review*, November 27, 2018.
3. Samuel Newman, "What is latency and how does it kill fast internet?," *BroadbandDeals.co.uk*, March 30, 2017. A decade ago, Amazon estimated that every 100ms of latency costs the company 1% in sales. See: Yoav Einav, "Amazon found every 100ms of latency cost them 1% in sales," *GigaSpaces*, January 20, 2019.
4. *Deloitte CIO Journal on the Wall Street Journal*, "Tips for transforming data management," August 24, 2016.
5. Sean Michael Kerner, "Top 8 cloud data warehouses," *Datamation*, September 10, 2019.
6. Prescient & Strategic Intelligence, "Data warehouse as a service market," March 16, 2020.
7. Adi Hirschtein, "What are feature stores and why are they critical for scaling data science?," *Towards Data Science*, April 7, 2020.
8. DB-Engines, "DBMS popularity broken down by database model," November 2020.
9. Matt Asay, "Why time series databases are exploding in popularity," *TechRepublic*, June 26, 2019.
10. Serdar Yegulalp, "What is a graph database? A better way to store connected data," *InfoWorld*, March 21, 2018.
11. Robin Harris, "Data storage: Everything you need to know about emerging technologies," *ZDNet*, May 28, 2019.
12. *Deloitte CIO Journal on the Wall Street Journal*, "Tips for transforming data management."
13. Edmond Toutoungi, "Cloud and edge computing explained ... in under 100 words," *Deloitte*, August 27, 2019.
14. Chris Arkenberg et al., *Unbundling the cloud with the intelligent edge: How edge computing, AI, and advanced connectivity are enabling enterprises to become more responsive to a fast-moving world*, *Deloitte Insights*, September 8, 2020.
15. MarketsandMarkets, *Machine learning market by vertical (BFSI, healthcare and life sciences, retail, telecommunication, government and defense, manufacturing, energy and utilities), deployment mode, service, organization size, and region—global forecast to 2022*, September 2017.
16. AT&T, "Corporate profile," accessed November 9, 2020.
17. AT&T, "AT&T's machine learning platform," YouTube, February 1, 2017.
18. Kate Hopkins (vice president of data platforms, AT&T), phone interview with authors, September 8, 2020.
19. Pari Pandya (director of technology and Amp project manager, AT&T), phone interview with authors, September 8, 2020.
20. Paul Ballew (chief data and analytics officer, Loblaw), phone interview with authors, October 6, 2020.
21. Santhosh Pillai (chief architect and data management, ABN AMRO), phone interview with authors, October 16, 2020.
22. Deloitte, *The Age of With: Exploring the future of artificial intelligence*, November 7, 2019.

# Deloitte.

## Insights

Sign up for Deloitte Insights updates at [www.deloitte.com/insights](http://www.deloitte.com/insights).

[www.deloitte.com/us/TechTrends](http://www.deloitte.com/us/TechTrends)



Follow @DeloitteInsight



Follow @DeloitteOnTech

### Deloitte Insights contributors

**Editorial:** Matthew Budman, Blythe Hurley, Abrar Khan, Rupesh Bhat, and Nairita Gangopadhyay

**Creative:** Alexis Werbeck, Dana Kublin, Tristen Click, and Victoria Lee

**Promotion:** Hannah Rapp

**Cover artwork:** Vault49

### About Deloitte Insights

Deloitte Insights publishes original articles, reports and periodicals that provide insights for businesses, the public sector and NGOs. Our goal is to draw upon research and experience from throughout our professional services organization, and that of coauthors in academia and business, to advance the conversation on a broad spectrum of topics of interest to executives and government leaders.

Deloitte Insights is an imprint of Deloitte Development LLC.

### About this publication

This publication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or its and their affiliates are, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your finances or your business. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

None of Deloitte Touche Tohmatsu Limited, its member firms, or its and their respective affiliates shall be responsible for any loss whatsoever sustained by any person who relies on this publication.

### About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the “Deloitte” name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see [www.deloitte.com/about](http://www.deloitte.com/about) to learn more about our global network of member firms.