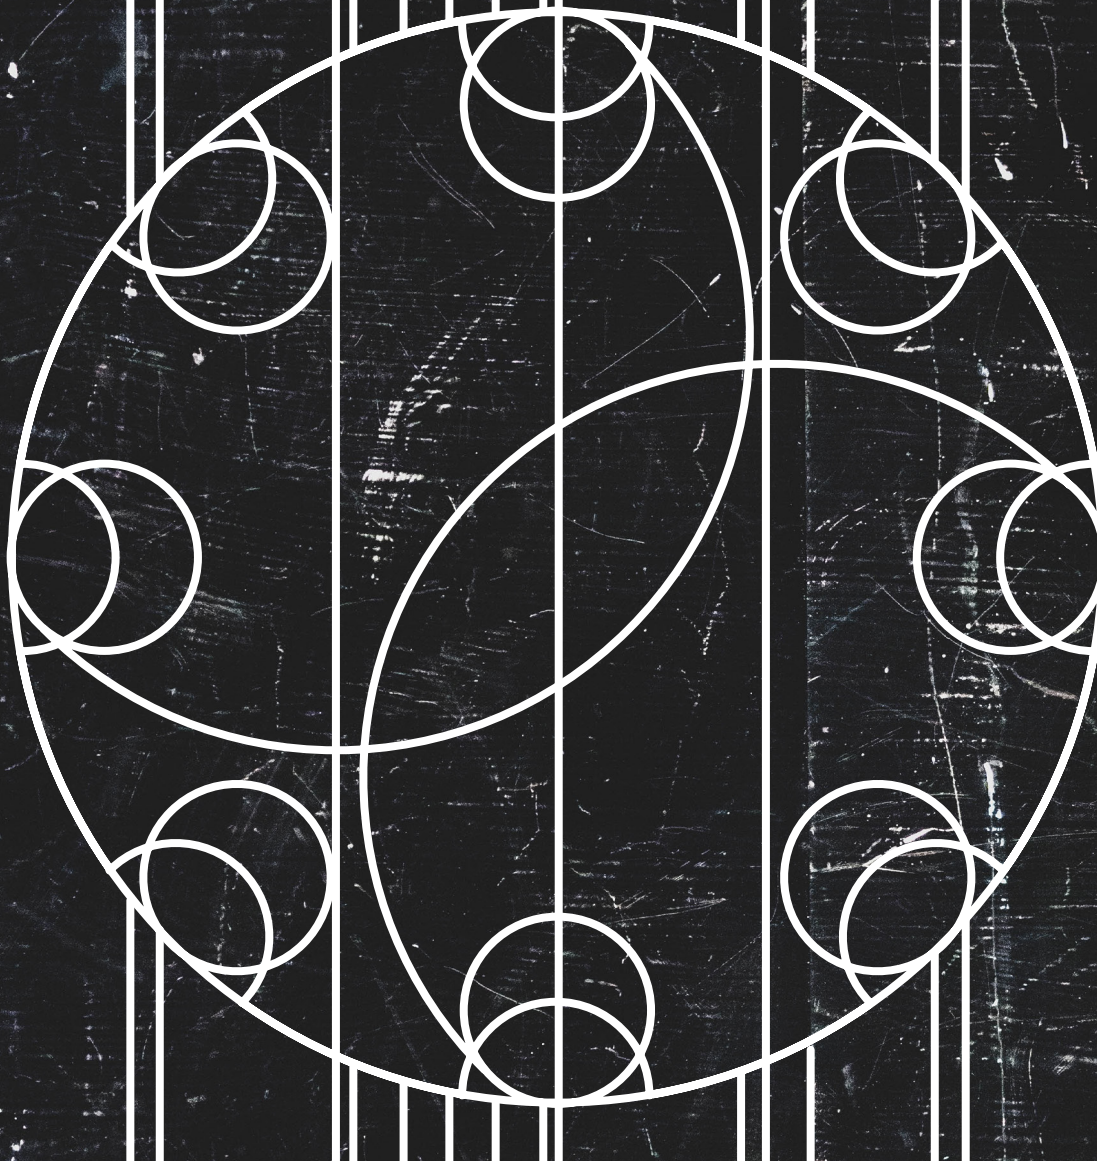


# BEYOND



**Celonis Labs**

**JOURNAL '23**  
**2ND EDITION**

## Beyond A publication from Celonis Labs

---

### Editor

Bill Detwiler,  
Editor, Technical Content & Ecosystem, Celonis

### Celonis Labs

Eugenio Cassiano,  
SVP Strategy & Innovation at Celonis  
Madeleine Just  
Senior Director of Operations at Celonis Labs  
Michael Sodamin  
Innovation Engineer at Celonis Labs

### Authors

Cong Yu  
VP of Engineering, AI and Knowledge at Celonis  
Chris Monkman  
VP Product Intelligence at Celonis  
Marco Tringali  
Senior Project Manager at Celonis Labs  
Claudia Schulze  
Principal Product Manager at Celonis Labs  
Renske Prins  
Head of Products for Sustainability at Celonis  
Andreas Krause  
Distinguished Architect at Celonis Labs  
Wil van der Aalst  
Chief Scientist of Celonis and a full professor at  
RWTH Aachen University  
Jonas Weich  
Application Product Manager at Celonis  
Teodor Fratiloiu  
former Master Thesis student at Celonis  
Jerome Geyer-Klingeberg  
Head of Academic Alliance at Celonis  
Eduardo Goulart Rocha  
former Master Thesis student and now Research  
Software Engineer @PQL team

### Guest Authors

Kyriakos Triantafyllidis  
Head of Growth & Strategy, Center for Advanced  
Manufacturing & Supply Chains at WEF  
Massimo Pezzini  
Independent IT Advisor  
Dimitris Kiritsis (Kyritsis)  
former Professor of ICT for Sustainable  
Manufacturing at EPFL Lausanne, Senior Advisor  
at University of Oslo

### Design & Production

Christian Flaccus,  
CDO, Celonis  
Jan Henkel,  
Senior Art Director, Celonis

### United States - Headquarters

One World Trade Center  
87th Floor New York, NY 10007  
Phone number: +1 212 847 0692

### Germany - Headquarters

Theresienstr. 6  
80333 Munich  
Germany  
Phone number: +49 89 4161596 70

### Twitter @Celonis

**LinkedIn** [linkedin.com/company/celonis](https://www.linkedin.com/company/celonis)  
**Facebook** [facebook.com/Celonis](https://www.facebook.com/Celonis)  
**Xing** [xing.com/company/celonis](https://www.xing.com/company/celonis)

### About Celonis

Celonis enables customers to optimize their  
business processes. Powered by its leading  
process mining technology, Celonis provides

a unique set of capabilities for business  
executives and users to continuously find  
improvement opportunities within and across  
processes, and execute targeted actions to  
rapidly enhance process performance. This  
optimization yields immediate cash impact,  
radically improves customer experience,  
and reduces carbon emissions. Celonis has  
thousands of implementations with global  
customers and is headquartered in Munich,  
Germany and New York City, USA with more  
than 20 offices worldwide.

### Disclaimer and Copyright

The content of this document has been created  
with great care and is solely to promote  
discussion of innovation in the industry. Still we  
cannot guarantee the accuracy, completeness  
or timeliness of the content. The information  
contained herein is not a warranty or promise  
of any kind, including without limitation,  
of performance, future products, features,  
functionality, or otherwise. Celonis SE explicitly  
disclaims all warranties as to the accuracy,  
completeness, or adequacy of the content and  
shall have no liability whatsoever related to this  
information, including for errors, omissions, or  
inadequacies or for the interpretations thereof.

Copyright ©2023 by Celonis SE.

All rights reserved. Celonis and the Celonis  
“droplet” logo are trademarks or registered  
trademarks of Celonis SE in Germany and other  
countries. All other product names or services  
identified throughout this document are  
trademarks or registered trademarks of their  
respective companies.

**Beyond**

# Index

## THE AUGMENTED ENTERPRISE

- 1** Enduring value of Process Intelligence in the age of Generative AI **8-9**
- 2** *GUEST COMMENT* Process Intelligence powers better and faster insights **10-11**
- 3** Collaborative Business Execution Network **12-15**
- 4** *GUEST COMMENT* World Economic Forum **16-19**
- 5** Process-Mining for a sustainable future **20-23**
- 6** The process-centric enterprise IT platform **24-27**
- 7** *GUEST COMMENT* Platform integration benefits **28-31**
- 8** Object-centric process mining: Imagining what's next in business performance **32-35**

## RESEARCH & THESES

- 9** Ai4Pro research project **38-40**
- 10** Revolutionizing customer satisfaction **41**
- 11** Peeking into the decision-making process **42-43**
- 12** Sustainability-Integrated Value Stream Mapping **44**
- 13** Next generation of Conformance Checking in Process Mining **45**
- 14** *GUEST COMMENT* Cognitive Digital Twins **46-47**

---

## FROM THE EDITOR

# “Innovation requires pushing beyond our current understanding.”



**Bill Detwiler**  
Beyond Editor

Bring passionately curious people together around a common objective and amazing things can happen. *Celonis Labs* is that type of group—a collective of innovators, researchers and concept designers that think differently about how the world should work. Their common goal is exploring novel concepts and technologies in pursuit of process excellence.

Why? Because, processes are the fabric of our world. From putting food on our tables and clothes on our backs to letting us travel around the world and down the street, processes make everything happen. They're also our greatest lever for value and our fastest lever for change. As Alexander Rinke, Celonis Co-CEO and Co-Founder, said during his *Celonis World Tour 2023* keynote, “The world is a much better place with better processes in it.”

Beyond, the Celonis Labs journal, is our chance to share the projects and research the team is working on. It's a collection of thoughts, analysis and prototype descriptions from leading experts in process mining and business performance.

In our *inaugural issue published in 2022*, we explored concepts like business collaboration networks, zero-trust sharing and business domain translation. In this Beyond edition, we're taking things to 11. We continue our exploration of collaborative networks and expand our discussions into areas like generative AI / LLM, process-centric IT platforms, software that empowers collaborative action and object-centric process mining. We're pleased to include special guest contributors Kyriakos Triantafyllidis with the World Economic Forum; Massimo Pezzini, a veteran IT industry analyst and research fellow; Eugene Wu, associate professor at Columbia University; and Dimitris Kiritsis (Kyritsis), former Professor of ICT for Sustainable Manufacturing at EPFL Lausanne, Senior Advisor at University of Oslo. There's also a section with abstracts from the latest research and Master Thesis papers on process mining. Now that's cutting edge.

A handwritten signature in black ink that reads "Bill".

# ENTERING THE AGE OF THE AUGMENTED ENTERPRISE



**Eugenio Cassiano**  
SVP Strategy & Innovation  
at Celonis

Last year, we embarked on a remarkable journey by releasing the first edition of *Beyond*, the Celonis Labs Journal. Our goal was to showcase how the concepts and prototypes being developed at Celonis Labs could shape the future of Celonis, process mining as a technology and business performance as a whole.

In our last edition, we gave readers a window into several of our latest initiatives and innovations, such as revealing how organizations can achieve transparency across value chains and how they can leverage unstructured data to enable context awareness. We wanted to push the boundaries beyond our current knowledge. And we did.

But we aren't sitting still. A lot has happened since we published the inaugural edition of *Beyond*. We made process mining more accessible by opening up our platform with the launch of the *Celonis Intelligence API*. We became the first technology vendor to fully embrace *object-centric process mining (OCPM)* with the release

of *Process Sphere™*. Prof. Wil van der Aalst, Chief Scientist at Celonis, contributed an article to this year's edition about OCPM and the benefits of combining it with cutting-edge technologies like generative AI and large language models (LLMs). The excitement and enthusiasm of that first Beyond continues, as we challenge our internal, and for the first time external, contributors to think beyond their day-to-day jobs and explore the possibilities of what's next.


One key area we explore in this edition is the rise of the augmented enterprise. In the last few years, Celonis began to lay the groundwork for how organizations become self-sustaining and business context aware. By bringing humans and AI together, we can

create an assisted human-AI partnership that enhances business performance. A key enabler for this augmented enterprise is the development of a standardized business language, which will act as a bridge between IT and business terminology.

With an emphasis on a standardized business language and an ecosystem of AI-powered applications, Celonis envisions a future where augmented enterprises thrive through collaboration, innovation, and sustainable practices.

Welcome again to the next edition of Beyond where innovations know no bounds.

Eugenio



# 1

## Enduring value of Process Intelligence in the age of Generative AI



**Cong Yu**  
VP of Engineering, AI and Knowledge at Celonis



**Chris Monkman**  
VP Product Intelligence at Celonis

Technologies powered by generative AI and large language model (LLM) will *transform* enterprises in unprecedented ways and enable significant productivity gains.

Many business processes as they exist today will change, for example:

- Manual steps to extract structured information from unstructured data will be automated (thanks to LLM-powered document understanding technologies),
- Understanding correlated events among processes will become more accessible (thanks to advancements in deep learning),
- Executive summaries and ad hoc analysis of data for insights will become more intuitive (thanks to the conversational and generative capabilities of foundational models).

While those processes will undergo significant transformation, the business goals they serve will remain constant—we will continue to optimize objectives such as working capital and labor productivity.

However, foundational LLMs are statistically trained without any notion of factuality and suffer from various

factuality-related challenges, such as hallucination, information staleness and a lack of individualized responses. For example, on hallucination, an LLM can easily produce an incorrect answer confidently for a math problem that a 5 year old would be able to answer correctly. To ground an LLM for math problems, we can provide a calculator and instruct the LLM that, upon seeing a math problem, ask the calculator to answer it. Essentially, the calculator here is the differentiator that separates a math-LLM from the commodity LLMs in the domain of math problems.

Process Intelligence will be key to overcoming LLM limitations like the one above and will thus play an enduring role in the business transformation driven by Generative AI. Built on top of a foundation of process data and knowledge with advanced technologies including process mining, decision intelligence, machine learning, and generative AI models, Process Intelligence offers observability, true differentiation from



commodity large language models, and potentially disruptive technologies for understanding and improving business processes.

First, through highly scalable and low latency process mining technology, Process Intelligence can provide process observability, enabling organizations to effectively track the changes in processes and how those changes are serving the business goals. The more complicated the transformations are, the more critical it is to employ process mining tools to observe the effectiveness of the transformations, and capture potential issues early on.

Second, process data and knowledge is the differentiator that will power LLMs to solve business problems both broadly and for individual enterprises—improving the business processes and achieving the business efficiency goals. Specifically, general process knowledge, as accumulated by industry-leading companies like Celonis, can transform a commodity LLM into one that understands business processes and makes meaningful recommendations for process improvements. Furthermore, enterprise-specific process data and knowledge, integrated in a secure and privacy-preserving way, can tailor an LLM to provide answers that are accurate, up-to-date and based on the customer's data and knowledge.

Third, process discovery and conformance checking, two core techniques of process mining, have striking similarities to language pattern discovery and pattern matching

that recent LLMs excel at, compared to traditional NLP techniques. It is conceivable that, within the next 2-3 years, the field of process mining will have another foundational change, following the move towards object-centric process mining, that leverages LLMs to perform more accurate process discovery, more comprehensive conformance checking, and eventually more proactive business process understanding and improvements. Those changes will bring significantly more power to enterprises in their ability to manage their processes to meet their business needs.

Fourth, Generative AI can drastically enhance the offerings of a broad set of value discovery technologies including decision intelligence and classic machine learning, improving the insights those technologies discover in terms of quality, speeding up the process of discovering those insights, and making them more intuitive. For example, an LLM-powered conversation agent could naturally guide a user through a complicated dimensional breakdown (of the underlying data) discovered by decision intelligence's key driver analysis, narratively highlighting what is interesting about this key driver and what value it can be translated into.

In summary, we at Celonis are excited and committed to drive Generative AI efforts forward to support our customers' AI transformation journey and build AI technologies that are secure, privacy-preserving, trust-worthy, tailored and achieve high return on investment. —

# 2

## Process Intelligence powers better and faster insights and actions




**Eugene Wu**  
Associate Professor at  
Columbia University

In the article “Enduring value of Process Intelligence in the age of Generative AI,” Cong Yu and Chris Monkman outlined four ways which large language models (LLMs) synergize with process intelligence to enhance organizational performance.

---

Eugene Wu is an Associate Professor of Computer Science at Columbia University, specializing in data management, analysis, visualization, and intelligence. This commentary draws on his experience as a visiting researcher at Celonis over the past six months, where he collaborated with the CeloAI team to integrate AI capabilities across the Celonis platform.



---

LLMs will certainly advance the capabilities of existing process mining techniques, but at the same time, LLMs are not a panacea for organizational inefficiency. They need to be carefully designed with an organization's processes in mind, and deployed to address actual process bottlenecks. This all leads to the article's first point about process observability, which particularly resonates with me due to its analogue in software systems.

Modern software systems are notoriously complex: dozens of independently managed services run on tens of thousands of servers, and must seamlessly coordinate to provide efficient, reliable services to end-users. To help manage this complexity, observability tools analyze event logs to reconstruct the system state, and help developers monitor metrics to gauge system health, identify and address bottlenecks and disruptions, and evaluate the impact of performance optimizations on user requests.

Modern organizations are similarly complex: distributed across multiple time zones, hundreds or thousands of individuals need to coordinate to deliver fast, seamless, and reliable services to their customers. Process observability is the necessary

first step to understand the state of the organization's processes. It is then possible to ensure conformance, identify and address major bottlenecks, and evaluate the effects of technological or organizational interventions.

It's rare for a research area to have such a direct relationship to organizational value. Although process intelligence is still relatively new outside of Europe, it already overlaps with many fields of computer science (e.g., machine learning, causal inference, data management, HCI) and has the potential to become a major application area that brings together experts across these disparate fields to solve novel and impactful research problems.

For instance,

- How to scalably and quickly discover processes across trillions of events?
- How to identify practical interventions that improve a process or KPI?
- What is the best modality to interact with process intelligence?
- What is the appropriate data model and analysis paradigm for process intelligence that is analogous to the relational model and analytical dashboards for business intelligence? —

# 3

## Collaborative Business Execution Network

### Improved decision-making and outcomes through better data orchestration



**Marco Tringali**  
Senior Project Manager  
at Celonis Labs



**Claudia Schulze**  
Principal Product Manager  
at Celonis Labs

In the first issue of *Beyond Journal*, we presented the Collaborative Business Execution Network concept. Designed to create a singular, unambiguous source of truth across a complex web of business partners, this concept can simplify and accelerate the way businesses communicate and operate, ultimately improving their business outcomes.



Since then, our team has shaped the concept into a tangible reality, crafted a prototype and solicited feedback from our customers on the solution while also exploring potential use cases.

# FEEDBACK & VALUE PROPOSITION

Our customers recognize that the COVID-19 crisis underscored the fragility of global value chains and their vulnerability to global disruptions, emphasizing the need for strategic resilience.

Thanks to the Celonis platform, our customers are already strengthening and streamlining their operations internally. However, when it comes to collaborative processes with business partners, they encounter significant challenges.

Firstly, they expressed that they face a complexity issue, wrestling with the language barrier of data between partners. The multitude of data formats and standards complicates communication. Agreeing at scale on an unambiguous ontology that is flexible enough to enable orchestration has proven difficult. This problem is not new and it is exacerbated by the complex technological landscape that comprises a value chain composed of thousands of business partners,

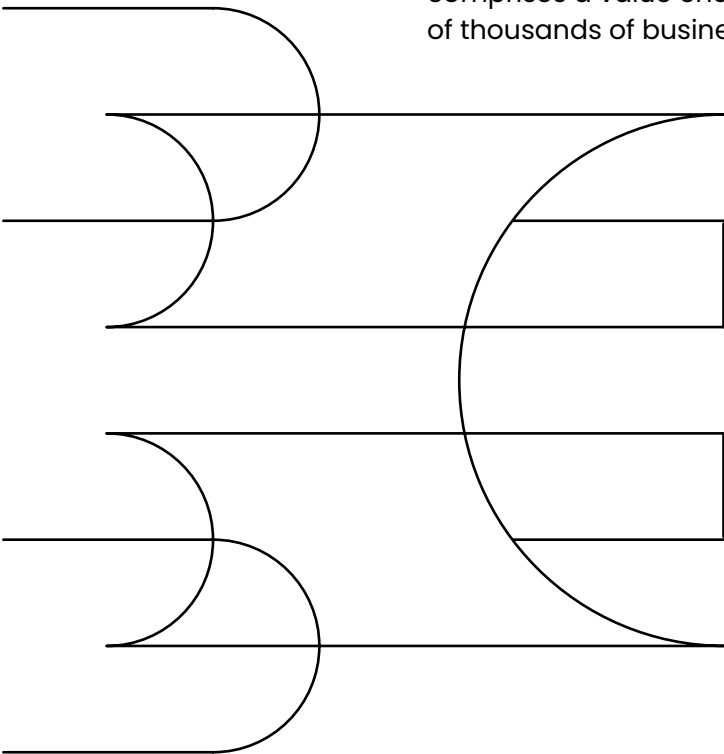
leading existing solutions to leverage customized ontologies that are specific to an organization resulting in a lack of scalability.

Secondly, but none less important, is the issue of data security. Our customers are justifiably cautious about sharing data with partners due to potential data breaches, legal issues, and concerns of revealing strategic insights or proprietary information, all of which could harm their reputation and competitiveness.

These challenges are the primary hurdles in designing a solution capable of effectively exchanging and orchestrating data. The feedback we've garnered from our customers underscores the urgent need for a solution that can handle these challenges, thereby increasing their operational resilience.

The Collaborative Business Execution Network directly addresses these challenges. Rather than attempting to automate complex processes or deep-dive into transactional data, our Network efficiently makes a set of information that is already exchanged between partners available dynamically at the right place and time, enabling orchestration to guide businesses towards their desired outcomes.

This design principle, that implies a low data footprint, also enables the definition of a standardized and simplified taxonomy where each participant can easily agree and understand the meaning of the information shared.



## USE CASES

A wide range of use cases demonstrates the potential of the concept. Picture a world where businesses work on a shared understanding of reality and can act in real time on any disruption that may occur, for example minimizing the risk of stockouts and overproduction. This scenario is just a glimpse of the operational efficiency and transparency the Collaborative Business Execution Network can offer across various industries.

Talking with our customers, they have shared interest in testing use cases across processes such as *Procurement / Order Management* (e.g. When will the materials I ordered arrive?), *Accounts Payable / Accounts Receivable* (e.g. What is the payment status of my invoice?), and recognize that enhanced transparency in business processes provides benefits for all involved parties, such as suppliers and customers.

These use-cases can also bring value in terms of sustainability. By providing a platform for seamless and transparent information sharing, the solution can help

reduce inefficiencies and waste in supply chains. For instance, visibility from suppliers with regards to the status of shipments results in improved inventory management that can help improve working capital by minimizing the need of contingency stock, leading to reduced waste and a smaller carbon footprint. Similarly, more accurate and timely information can lead to better decision-making such as optimizing the capacity utilization of carriers, thus reducing unnecessary transportation and the associated emissions. As such, our solution isn't just about improving business performance; it's also about creating a sustainable future where business growth and environmental responsibility go hand in hand.

The Collaborative Business Execution Network enables organizations to efficiently share, consume, and orchestrate shared data. This leads to actionable information being readily available in the right place, system (such as the Celonis platform or other systems), to the right user, and at the appropriate time, thereby fostering better decision-making and improved outcomes.

## WHAT'S NEXT

Singularly considered, the created transparency offers great value for the business partners engaged in each use case. But reality does not stop with 1, 2 or even 10 business partner relationships in a company. Neither do we - the real magic is added by the scaling effect of a network. Building on the common

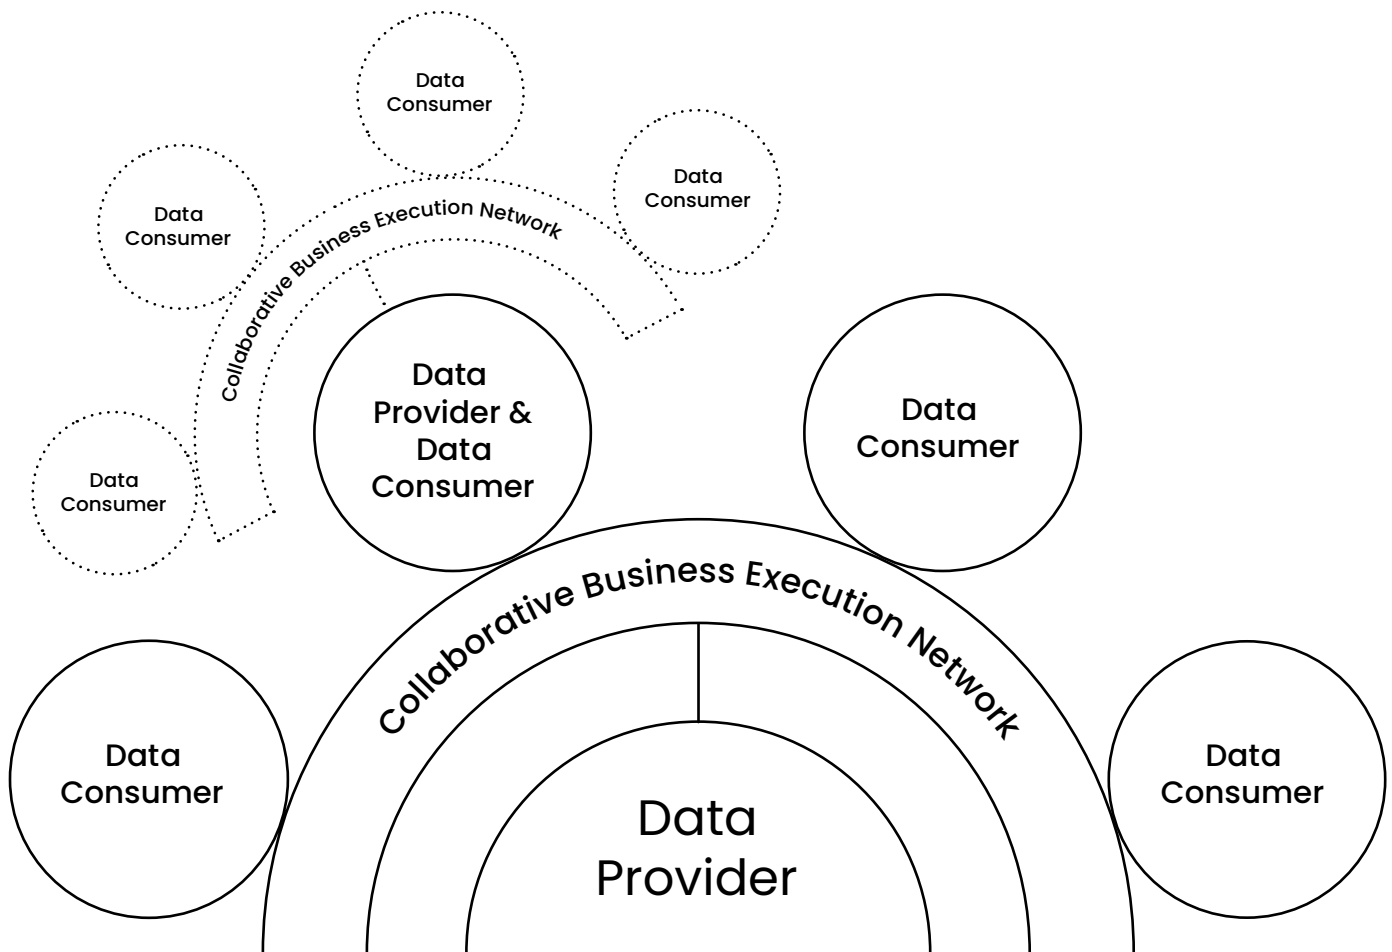
business taxonomy and a standard integration approach leveraging Celonis platform data, Collaborative Business Execution Network offers a low entry barrier to participate from the beginning. To participate you won't have to integrate to each of your partners separately but only once provide and once to consume

information. This will rapidly lead to a critical mass of business partners in the network creating a pull effect for additional ones.

In the past, integrating 3 customers for the sake of providing order updates automatically, meant 3 separate IT projects. With Collaborative Business Execution Network it becomes a matter of hours, owned by business partner managers with low technical knowledge.

Imagine the lever of a platform hosting hundreds of companies for data providers and consolidators. Companies such as Project44, EcoVadis or ParcelLab could integrate with and provide data to new customers in a streamlined and scaled fashion.

Collaborative Business Execution Network is set to become a data marketplace - offering End-to-End value chain collaboration as a service. —



# 4

## World Economic Forum Resilient value chains: Thriving amidst challenges, unleashing endless opportunities



**Kyriakos Triantafyllidis**  
Head of Growth & Strategy,  
Center for Advanced  
Manufacturing & Supply  
Chains at WEF

The recent turmoil in markets is a symptom of global developments driven by three main *mega-trends* - climate change, geopolitical tensions, and emerging technologies - that are fundamentally transforming the business environment.





---

## CLIMATE CHANGE

The climate crisis and widespread environmental degradation are profoundly impacting global value chains through increasing risk to supply continuity, exacerbating resource shortages, and adding greater pressure to act in a carbon-responsible way. For instance, as highlighted in the *World Economic Forum's Global Value Chains Barometer*, large-scale weather events like the US deep freeze in December 2022 have resulted in widespread power outages and road blockages, slowing production and flow of goods.

From pollution and solid waste to water scarcity, desertification and deforestation, the need to address the underlying drivers of environmental degradation is increasingly recognized by companies, governments and civil society. Collective consciousness of these threats has grown significantly during recent years and new measures have been adopted to reduce carbon emissions. However, much more is needed and possible actions include the decarbonization of supply chains, greater support for renewable energy and advancing the circular transformation of industries.

## GEOPOLITICAL TENSIONS


Political shifts, trade tensions, wars and policy uncertainty, expose companies and governments to numerous vulnerabilities – including rising energy prices, global shortages, production scarcity and logistical delays in getting key components. For example, in light of the recent inflation, leading manufacturers are employing *advanced planning tools to review buffers end-to-end and reduce non-critical inventories* to ensure financial availability.

As a result, many governments and industrial companies are looking to adopt more regional, friend-shoring or near-shoring approaches towards globalization and international trade, while attempting to balance the benefits of diversified globalization and concerns for sustainability practices.

## TECHNOLOGICAL ADVANCEMENTS

Ranging from digital platforms, additive manufacturing and artificial intelligence to robotics and the internet of things, emerging technologies are transforming both operating and business models, facilitating the relocation of production closer to consumer demand, enhancing transparency between producers and consumers, and enabling the provision of new and better services. *Generative AI has revealed new opportunities*, promising to transform industrial operations through a variety of applications that can provide information support, convert information to actionable outputs and enhance robotics.

Coupled with a vast increase in connected devices worldwide, these developments are forcing manufacturers to tackle increasing cybercrime and threats. To reap the benefit of emerging technologies in production systems and supply systems, the highest cybersecurity measures are needed.



# Manufacturers are taking a more holistic approach

Against this backdrop, manufacturing leaders are embracing a more holistic strategic perspective of balancing value chain performance with resilience and sustainability to ensure service quality and long-term business

continuity. This requires going beyond *making adjustments in silos to react to a volatile environment*, towards reconfiguring value chains to put resilience and – in light of the climate crisis – sustainability at the core of the business agendas.

---

This reconfiguration will bring a new mindset about global value chains and how they operate. Based on our consultations with senior manufacturing and supply chain executives, we have identified *five trends that will dominate this reshuffling*.

First, in the near future, they will be based around multiple globally connected but highly localized value chains – relying on more local suppliers and customers – rather than one single global value chain.

Second, to manage the resulting complexity and cost implications as well as increase agility, automation and digitalization technologies will become more widely implemented. The judicious use of automation in manufacturing processes can help to offset cost differentials and decrease exposure to labour shortages, which is particularly important when relocating to markets with high labour costs or aging populations. New manufacturing technologies, such as additive manufacturing, can also help with mass customization to better serve customer preferences in local markets. Digitalization can help to build visibility and traceability, ensuring business continuity and alignment between multi-local value

chains. Visibility identifies upstream supply risk, while traceability can help firms quickly react to disruptions and mitigate their impact.

Third, the transition will require equipping the workforce with the right skills and capabilities and attracting new talent in the industrial operations environment.

Fourth, as manufacturers undergo this transition, they will seize the momentum and ingrain sustainability at the core of their more distributed operations.

And fifth, All of this will require careful trade-off management between performance, sustainability and resilience to drive value for the customer.

These five dominating trends are bringing a new mindset about global value chains in which the interconnectedness of globalization will continue to prevail. As leaders reshape their value chains, manufacturers worldwide will begin to feel the impact and ripple effects of the collective rewiring, strengthening the importance of multi-stakeholder collaborative approaches to navigate these shifts and unlock new opportunities. —

# 5

# Process-Mining for a sustainable future



**Renske Prins**  
Head of Products  
for Sustainability  
at Celonis

In the current macroeconomic environment, organizations are forced to strike a carefully considered balance between sustainability goals and other business goals. As businesses grapple with supply chain shortages and rising global costs, some organizations have put sustainability on the backburner. CEOs spoke about economic uncertainty 3 times as often as sustainability in their Q1 earnings calls. But many don't realize that by focusing on sustainability, they can benefit both their all-important *bottom* and *top line* as well as their *green line*, which has faced ample scrutiny in recent years.

Process mining emerges as a technology allowing businesses to optimize their processes, balance

objectives and become more sustainable. The Celonis platform is a pioneering solution in this area. Though originally crafted with the intention of enhancing process performance centered around economic objectives, sustainability has inherently been interwoven into its fabric. Process mining has allowed companies to reduce rework and improve process performance, therewith reducing waste and increasing efficiency. The key is to transform operations in a minimally disruptive way, focusing on processes that lie at the heart of sustainability strategies, and find initiatives to start with that have a positive effect on the *top, bottom, and green line*. This way, businesses can achieve better sustainability and business outcomes simultaneously.

### × top, bottom, green line value

In the realm of business finance, “top line” and “bottom line” are commonly used as barometers for a company’s financial health.

The “top line” is an indication of a company’s gross revenue or sales, often serving as the initial metric of business performance. As the name suggests, it is featured at the top of a company’s income statement, providing an overview of the total revenue generated through various

business activities before any expenses are considered.

On the other hand, the “bottom line” signifies the company’s net income, a more insightful measure of profitability. Derived after accounting for all costs, expenses, and taxes from the top line, the bottom line provides a sharper picture of the financial health and efficiency of a business. It is often the final figure on an income statement, hence the term “bottom line”.<sup>2</sup>

A newer addition to the financial language is the “green line”, an index of a company’s environmental footprint. Unlike its counterparts, the green line moves beyond mere financial performance to assess the environmental sustainability of a company’s operations. It measures the company’s active efforts to reduce environmental harm, such as cutting waste, curbing carbon emissions, conducting due diligence on suppliers, and adopting renewable energy sources.<sup>3</sup>

## THE ROLE OF PROCESS MINING

**Process Mining operationalizes sustainability strategies in three different ways:**

Where cost savings meet sustainability: Optimizing the processes that are at the heart of sustainability strategies

Improving a company’s sustainability performance requires operational data, specific sustainability KPIs like carbon emissions, waste, usage of energy or diversity measures and targeted action. All sustainability metrics are influenced by business processes, particularly supply chain processes: how you source products from responsible and reliable suppliers, how you make products with minimal waste, how you move products efficiently, and how you ship products on-time with minimal emissions. The most straightforward method of applying process mining to make businesses more sustainable is the optimization of processes that have a direct impact on sustainability KPIs. This is an area where cost savings and sustainability

improvements often go hand-in-hand. Reducing the number of order returns not only lowers the number of emissions, it also increases customer satisfaction and cuts shipping costs. Reducing rework in the production process leads to lower scrap rates, decreasing the carbon intensity of the products as well as cutting material cost.

The way to tackle these processes is straightforward and the methodology is very close to the traditional process mining approach: The traditional data model needs to be enhanced with relevant sustainability activities in the event log as well as additional sustainability data points or KPIs, like emission factors or supplier ratings. Based on that, intelligent analyses or even action capabilities can be applied with the Celonis platform. Green line value creation can be tracked and even leveraged for sustainability reporting – in addition, and simultaneously to the traditional process optimization and its value creation.

## 2 When the process hurts the planet: Process deviations as a root cause of high emissions

Process deviations, often prompted by unanticipated circumstances such as rush orders, delivery blocks, or sudden increases in demand, can significantly contribute to high carbon emission levels – and companies are often not able to uncover these relations.

For instance, after delivery blocks caused by master data issues, there is often a rush to fulfill orders to maintain customer agreements resulting in express or expedited transport modes. This leads to a surge in demand for faster and less eco-friendly modes of transport such as airfreight. Air freight accounts for about 90 times as much emissions as sea freight<sup>4</sup>, making it a prominent contributor to corporate carbon emissions. This exemplifies how a seemingly operational hiccup like a delivery block can cascade into a significant environmental impact, amplifying a company's contribution to climate change. Thus, it becomes crucial to understand the true cost of process inefficiencies as well as positive effects that their resolution can have. Unveiling, managing and fixing negative impacts as well as trade-offs, results in more resilient, flexible, and sustainable supply chains that can adapt to fluctuations without compromising environmental responsibilities.

In order to ensure these supply chains, processes need to be enhanced with sustainability KPIs like emission factors, supplier ratings,

and life-cycle analysis results. This allows companies to introduce emission reduction or supplier sustainability as a process objective and evaluate which process deviations influence performance. They can then dive into any hotspots and find the process inefficiencies that cause them. Intelligent process mining algorithms also enable companies to calculate saving or improvement potentials and balance economic and sustainability goals through simulation features. This way it is possible to optimize for sustainability objectives in processes, to transform operations in a minimally disruptive way or to find easy initiatives to start with by harnessing the power of process mining.

## 3 Sustainability excellence: Fixing Inefficiencies in sustainability processes

The last frontier in using process mining for sustainability is optimizing sustainability-specific processes. Processes to implement sustainability strategies are often time and resource-consuming and require a multitude of stakeholders from various departments – thus, the perfect case for process mining. As an example, while complying with the *German Supply Chain Act (German: LkSG)*, companies have to collaborate across several functions in procurement, sustainability, and legal. The regulation requires collecting additional data or connecting various tools, often resulting in manual efforts. Technical solutions that harmonize data sources, enable an intelligent

---

### × LkSG

The German Supply Chain Due Diligence Act, also known as the German Supply Chain Act, came into effect on January 1, 2023. This law regulates corporate responsibility for compliance with human rights in global supply chains. This includes, for example, protection against child labor, the right to fair wages, as well as environmental protection. This benefits people within the supply chains, companies, and consumers alike. With this regulation came the responsibility of companies to do proper due diligence on their own supply chain, which for many companies means sending out rating questionnaires and manually contacting all their suppliers. The European Union is said to follow the principles of the German Supply Chain Act closely as they are developing the European version of this regulation.<sup>5</sup>

analysis of data points, connected activities, and root causes for performance, saving time for the involved teams. Tracking the process, as well as easily unblocking

the process through targeted actions, can make these projects less resource-intensive and regulatory compliance quicker and more proactive.

## CONCLUSION

Infusing sustainability into the DNA of core business operations can initiate a cycle of enhancement, helping companies hit their sustainability targets while staying focused on overall business objectives. Existing data, when used effectively, paves the way for processes to be more efficient and cost-effective. Additional sustainability data sources empower companies to measure the green impact of this process efficiency, resulting in improved sustainability performance.

Process mining experts can easily unlock a new world of possibilities by enhancing data models and applying their existing skills. Aside from cost reductions and revenue increases, they can now use their expertise to balance profitability with sustainability, further increasing the maturity of their

process mining applications and becoming even more strategically relevant. To underline, sustainability should not be an isolated goal but an integral part of the wider business strategy, dovetailing with all other objectives.

Process mining, with its unique capability to identify opportunities for optimization, can help steer businesses towards this ideal balance. Innovative and pragmatic solutions, like those offered by Celonis or built in the Celonis community, assist businesses in taking the first strides towards a more sustainable future without compromising profitability.

Processes are a fast, less capital-intensive way to achieve top, bottom, and green line value simultaneously. Now there is only one question left to answer: what process will you start with? —

---

1 <https://iot-analytics.com/what-ceos-talked-about-in-q1-2023/>

2 <https://www.investopedia.com/ask/answers/difference-between-bottom-line-and-top-line-growth/>


3 <https://www.theneweconomy.com/business/climate-change-managing-the-green-line-to-address-climate-action>

4 <https://www.climatiq.io/data>

5 <https://www.csr-in-deutschland.de/DE/Wirtschaft-Menschenrechte/Gesetz-ueber-die-unternehmerischen-Sorgfaltspflichten-in-Lieferketten/gesetz-ueber-die-unternehmerischen-sorgfaltspflichten-in-lieferketten.html>

# 6

## The process-centric enterprise IT platform Improving business alignment and accelerating digital transformation







**Andreas Krause**  
Distinguished Architect  
at Celonis Labs

IT organizations must evolve from only supplying technology to being digital process providers. Without such a transition, internal IT departments risk becoming obsolete as the work of operating IT infrastructure is commoditized by hyperscale cloud providers.

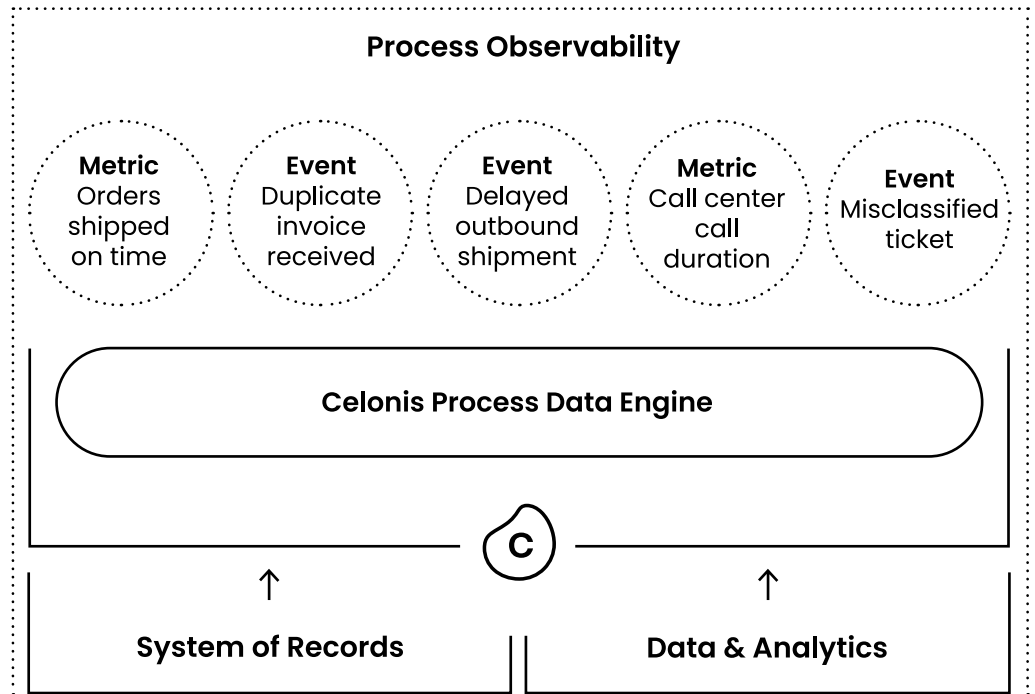
Software however, is now a competitive advantage for many companies. IT can become a core value driver by moving up the stack and delivering comprehensive digital solutions that align with the operational needs of the business. The challenge for IT leaders then becomes how to choose the best solutions to drive value, while managing budget constraints. The answer is a new kind of enterprise IT platform.

A platform is a surface, either physical or virtual, that serves as a stable and supportive foundation for a wide range of activities, objects, or systems. Traditionally, in the realm of IT, this relates to infrastructure and encompasses components like hardware, networking, and essential

software elements such as operating systems and databases. This however, addresses the requirements of an IT department in the pre-cloud and Software as a Service (SaaS) era.

To successfully drive business results, a future enterprise IT platform must combine several crucial elements. First, it must use a mutually understood language, such as business processes, to provide a foundation for effective communication. Second, it must enhance people's visibility and comprehension of business processes by incorporating a means to automatically discover, describe and observe these processes through process mining. Finally, the platform must be able to continuously improve and measure the impact of these processes.

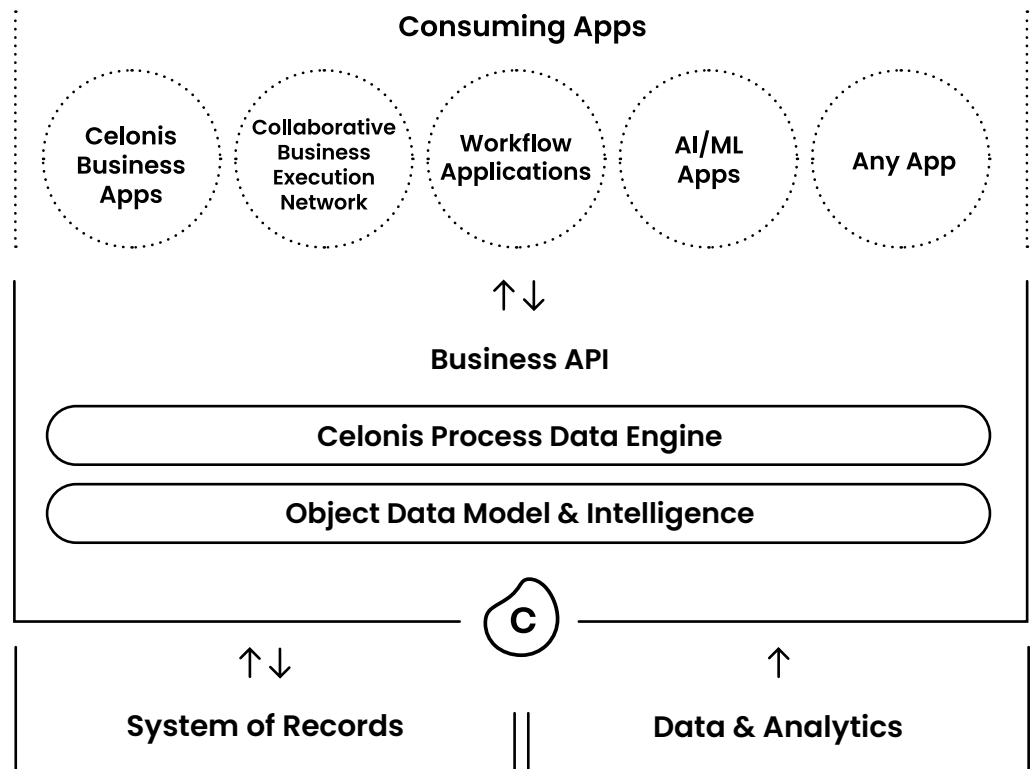
At Celonis, we believe that this platform is implemented as a process-centric *digital twin* of your enterprise. Process Digital Twins constantly observe the state of your processes and make them available through a contextual model.



Process observability, as described in the above visual, sits on top of a data foundation composed of business objects and related events. Data is collected directly from the source systems and represented in a semantically standardized way (i.e. in the form of business objects, business events and their process contexts). Using this data, which is valuable in itself for modeling and describing the current reality of business and IT operations, we can infer and

generate higher order events (e.g. Duplicate Invoice Received), that make observability actionable for humans and software automations. This as well includes business metrics.

Observing processes and identifying opportunities for structural and case-by-case optimization is a good starting point. However, to complete the full cycle, it is essential to take automated actions based on these process-centric insights.



The above visual describes a model in which process observability is consumed by a variety of Apps. These Apps will act upon the insight to improve business processes. For both IT and business organizations we see the following advantages:

- Inspires business users to leverage their innovative and creative potential by creating a common language that allows business users and IT to collaborate more effectively and align around the organization’s strategic goals.
- Common semantic model (set of well-defined Business Objects and Events) is a vendor-neutral starting point for App Modernization
- Protect Legacy Systems like ERP from unintended load spikes
- Celonis managed Performance and Scalability through:

- Elastic cloud-native technology
- Intelligent buffering and caching

Furthermore we believe that standardizing observability across an ecosystem of customers has the potential to enable a collaborative and product centric approach to process improvements.

Here I’ve outlined the ever increasing importance of IT. However, with great power comes great responsibility. Modern IT organizations and their platforms need to be an enabler for more and more technology-centric businesses and their business processes that are responsible for every organization’s value creation. Providing a platform that at the very core is process-centric is key to the evolution of IT organizations and their fate of becoming indispensable for every business. —

# 7

## The process-centric enterprise IT platform delivers integration benefits




**Massimo Pezzini**  
Independent IT Advisor

---

× **A digital integration hub (DIH)** is an architecture meant to support low-latency and high-throughput APIs access to data handled by multiple system-of-record (SoR) sources, by enabling these APIs to access SoR data that are replicated and aggregated in an ad hoc data store, which synchronizes with the original data sources in near real time.

By adopting a process-centric enterprise IT platform, organizations can achieve two important goals. First, it provides the benefits of process observability. Second, by being an extended implementation of the *Digital Integration Hub (DIH)* application architecture, a process-centric IT platform provides an API architecture that is turbocharged in terms of scalability, performance and availability when compared to a classic API architecture.



---

## LIMITATIONS OF THE “CLASSIC” API ARCHITECTURE

Digital operations require low latency and scalable access to data and 24/7 availability. These requirements are hard for organizations to meet given the high fragmentation and partial duplication of key data (customers, products, payments etc.) across multiple “system of record” (SoR) applications such as ERP, CRM, SCM, HCM, WMS etc.

Organizations traditionally use APIs layered on top of an integration infrastructure (for example, an Enterprise Service Bus like Mulesoft Anypoint Platform, SAP Process Integration or Tibco BusinessWork) to enable “digital applications to SoR” integration. This shields the digital layer from the specific protocols, interfaces and data models used by the SoR systems.

This “classic” API architecture is powerful and flexible. For instance it is relatively easy to add new digital applications whilst minimizing the impact on the SoRs.

APIs are not only extremely popular and widely adopted, but also well understood, en-


abled by standards such as OpenAPI and HTTP, and strongly supported by the tools and application vendor community.

However, a classic API architecture has its drawbacks:

- Limited scalability, which may not cope with the high volume of requests generated by digital applications on the SoRs.
- Suboptimal API response time
- Lack of 24/7 operations, impacting digital applications during SoR downtime.
- Complexity in integrating fragmented SoR landscapes, leading to potential fragile and unreliable environments.

In many situations, those challenges are not necessarily showstoppers. However, with the increasing demands of today’s consumers, when it comes, for example, to B2C digital applications, those limitations may lead to catastrophic failures of business initiatives. The digital way is littered with horror stories of ecommerce sites crashing or mobile applications constantly timing out compromising the organization reputation.

The challenge for organizations remains: how can they work around these issues and build an API layer that is scalable, performative and close to 24/7 availability?



---

## HOW THE DIGITAL INTEGRATION HUB TURBOCHARGES APIS

To address the challenges mentioned earlier and implement a more scalable, high-performance, and reliable API layer, a DIH is often the solution. By duplicating and consolidating the necessary data into a single optimized datastore, the DIH ensures low latency, scalability, and 24/7 availability. Additionally, the DIH reorganizes data to align with the organization's digital operations requirements and can incorporate data inferences from external sources.

The DIH architecture brings another significant benefit: it completely decouples the digital applications from the SoRs providing the data. They operate independently, with the DIH data store serving as their common link.

Adding new digital applications that use DIH data doesn't affect the SoR data providers. Similarly, if a legacy SoR is replaced with a modern SaaS alternative, the digital applications remain unaffected as long as the new application feeds the DIH with the same data as the old one. This makes the DIH an excellent enabler for incremental application *modernization strategies*.

Another key benefit of the DIH-based architecture vs. the classic API approach is that the central datastore consolidates data from a variety of sources. This data can be used beyond serving the API layer, allowing organizations to emit events for significant data changes or utilize it for analytical purposes.

---

### × Modernization Strategies

While the classic API approach offers partial decoupling between digital applications and SoRs, also enabling incremental modernization, it demands redesigning of all APIs when transitioning from old to new SoRs. In contrast, the DIH architecture allows seamless SoR replacements without the need for API modifications.

## DIH IMPLEMENTATION CHALLENGES

Despite its conceptual simplicity implementing a DIH in reality has its complexities, to mention the most notable ones:

- Designing the DIH data model requires a deep understanding of the digital business processes it serves.
- Digital processes must work with up-to-date data, which is achieved through real-time or near real-time updates from SoRs using event-driven integration.

- Event-driven integration must include tools to map and transform incoming SoR data to align with the DIH data model.
- DIH implementation poses additional challenges related to data governance and security.

Hence, the entry barrier for DIH is significant, and only financially strong and technically adept organizations have been able to afford implementing DIH-based API layers thus far.

---

# The “process-centric enterprise IT platform” delivers the DIH benefits out-of-the-box

The “process-centric Enterprise IT platform”, by default, has to address many of the DIH implementation challenges as discussed above. For example, it must provide a predefined, yet extensible, data model designed to support a variety of digital processes, a data enrichment (via inference) capability, a predefined set of APIs and events to consume the data and a standardized approach to exchanging data between the DIH and the SoRs.

This prepackaged set of capabilities, coupled with its cloud-based delivery model, will make adoption reasonably easy and fast, therefore not only leading edge, but also mainstream

organizations to take advantage of the DIH architecture and experience its business benefits in terms of business agility, fast innovation, efficiency and observability.

In essence, the process-centric Enterprise IT platform serves a dual purpose. Firstly, it provides process observability, empowering organizations to gain valuable insights into their operations. Secondly, it serves as a turbocharged API platform, offering scalability, performance, and availability advantages, making it an all-encompassing solution for organizations seeking process optimization and digital transformation. —

# 8

## Object-centric process mining Imagining what's next in business performance



**Wil van der Aalst**  
Chief Scientist of Celonis  
and a full professor at RWTH  
Aachen University  
*Latest Whitepaper*

Object-centric process mining (OCPM) significantly increases the speed, value capture, accuracy, and use case potential for process mining. It will also significantly transform many aspects of the discipline, such as conformance checking, predictions and simulations. With this article, I hope to spark the imagination about the advantages and opportunities OCPM offers.


### **OCPM: A REVOLUTION IN PROCESS MINING**

Process mining is one of the few data-driven approaches that is process-centric. Most machine-learning, artificial-intelligence, and data-science techniques tend to focus on a single task, not considering end-to-end processes. Using process mining, we can discover processes models based on event data, check the conformance of processes by comparing a process model with event data, analyze performance by replaying event data on models,

predict the evolution of the whole process or individual process instances, and automatically perform actions to address performance and compliance problems.

Traditional process mining, while powerful and effective, isn't without limitations. Data extraction and transformation can be time-consuming and needs to be repeated. Interactions between process objects (e.g., sales orders,





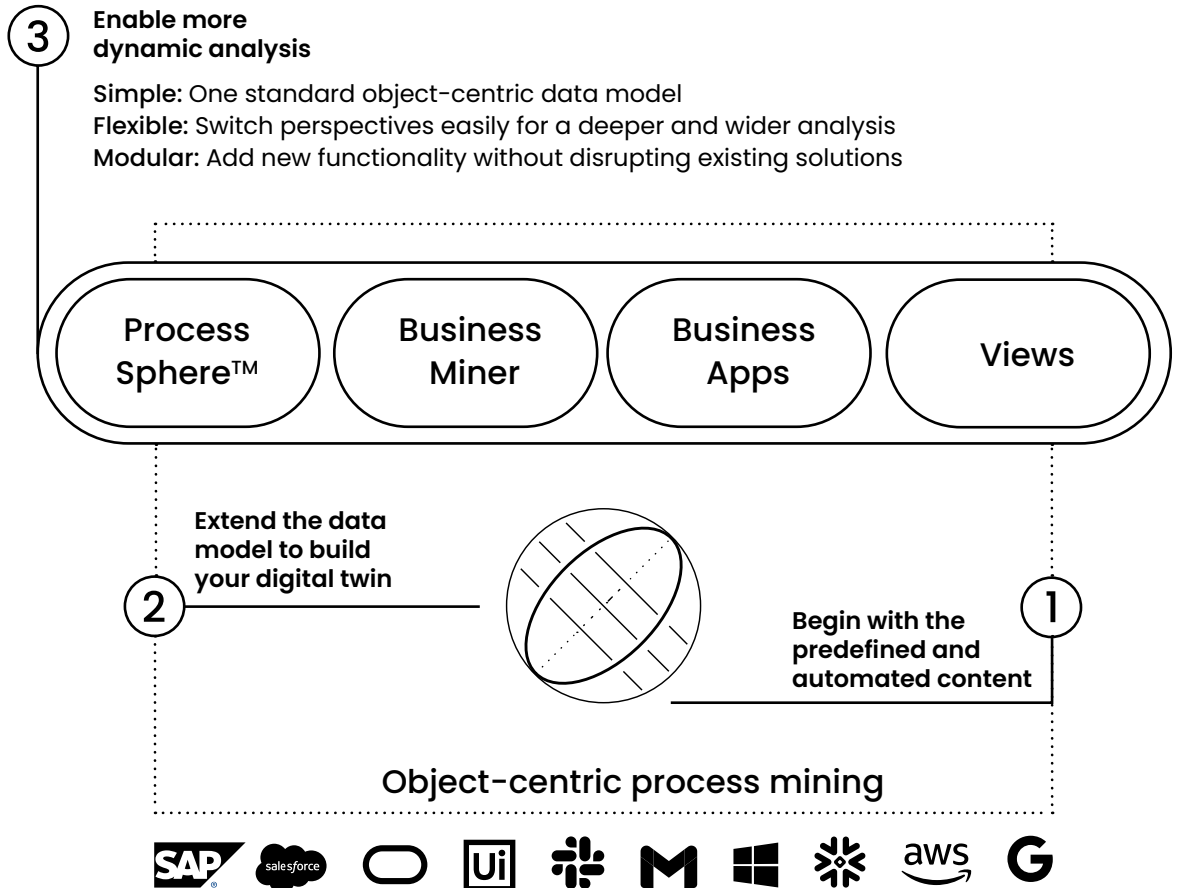
**Download  
Whitepaper**

sales order items, shipments, invoices, etc.) are not captured. And, three-dimensional event data is flattened into two-dimensional event logs causing convergence and divergence problems. Object-centric process mining (OCPM) was developed to overcome these limitations.

OCPM takes a more holistic and comprehensive approach to process analysis and improvement by considering multiple object types and events that involve any number of objects. It can significantly increase the speed, value capture, accuracy, and use case potential for what is already a transformational technology.

Celonis was the first vendor to fully embrace OCPM with *Process Sphere™* and the *Object-Centric Data Model (OCDM)*. These innovations help organizations unravel the true fabric of their operational processes. Instead of using different data models for different analysis viewpoints, we create a single data model that captures all relevant events and objects. Each event may involve any number of objects, and objects may be involved in many events.

The Object-Centric Data Model enables OCPM innovations across the Celonis platform, such as the *Multi-Object Process Explorer* and *End-to-End Lead Times App*.



# OCPM: The next frontier in business performance

OCPM marks a revolution in process mining. To take full advantage of the richer data and process models OCPM offers, we must rethink existing techniques, such as predictive analytics, data-driven simulation,

and comparative process mining. For example, with Celonis Process Sphere, we developed new discovery and conformance-checking techniques. The following are new opportunities with OCPM.

## More meaningful data


Object-centric event data (OCED) used by OCPM, creates opportunities to add semantics, such as specialization and composition, to data. For example, we can create a general payment type as well as types for specific payment activities or objects. We can

also show that one object contains or is composed of other objects. Similarly, we can identify subprocesses. Specialization and composition enable us to think logically about event data and create process models at a higher level of abstraction.

## Org-specific, process aware LLMs

OCED also facilitates the creation of organization-specific *Large Language Models (LLMs)*. GPT-4 and related technologies show the value of bringing a lot of data together in a large model. OCED may contain information about an organization's products, orders, items, suppliers, customers, machines and the like. This information is essential for answering questions and is highly structured. However, uploading organization-specific data like this into

a general-purpose LLM can create confidentiality and accuracy problems. Therefore, many organizations will opt for a proprietary LLM. Combining LLMs and process mining is promising and helps to contextualize process mining results. Moreover, it facilitates interactions using natural language. For example, Celonis *LLM for PQL Generation* translates natural language statements, queries and instructions into PQL (Process Query Language) to query event data.



## More realistic digital twins

OCPM is also a key technology for creating Digital Twins. A *Digital Twin* requires a model that approximates reality and is able to explore different scenarios. It is impossible to develop such models focusing on just one object type. However, faithful data-driven simulation models need to consider the interplay between different

kinds of objects, which is only possible with OCPM.

These examples illustrate that OCPM plays a key role in lifting process analytics to the next level. Although we are just at the beginning, the initial projects using Celonis Process Sphere provide proof points for the advantages mentioned before. —

---

### Additional reading:

- Van der Aalst, W. *Object-Centric Process Mining: Unraveling the Fabric of Real Processes*. *Mathematics* 2023, 11, 2691
- Van der Aalst, W.; Berti, A. *Discovering Object-Centric Petri Nets*. *Fundamenta Informaticae* 2020, 175, 1–40

# RESEARCH & THESES


Welcome to Research and Theses, a new section of Beyond dedicated to showcasing innovative research and emerging concepts.

We're excited to bring you the material here in conjunction with the *Celonis Academic Alliance*, which works alongside talented students and researchers to unlock new opportunities and insights in process mining.

The first four articles focus on the Celonis platform and process mining, exploring how process intelligence can revolutionize business opera-

tions. In the final article, Prof. Dimitris Kiritsis addresses the concept of cognitive digital twins (CDT). Prof. Kiritsis explains the characteristics of a CDT, how it's different from a traditional digital twin and links to several examples of how CDTs are being applied.

Join us on this journey of innovation and exploration as we delve into the world of process mining and its transformative potential. —



# 9

## *Ai4Pro* Research project Streaming Celonis' insights to wearable devices in real-time



**Jonas Weich**  
Application  
Product Manager  
at Celonis

---

× **Ai4Pro**  
is a project to develop an AI- and CPS-based solution that helps small and medium-sized manufacturing companies optimize their production processes in real time. This project was started in 2020 together with the Fraunhofer IIS Institute, Maxsima, AST-X, PASS and Rauschert in 2020 and is funded by the Bavarian State Ministry for Economy, Regional Development and Energy.

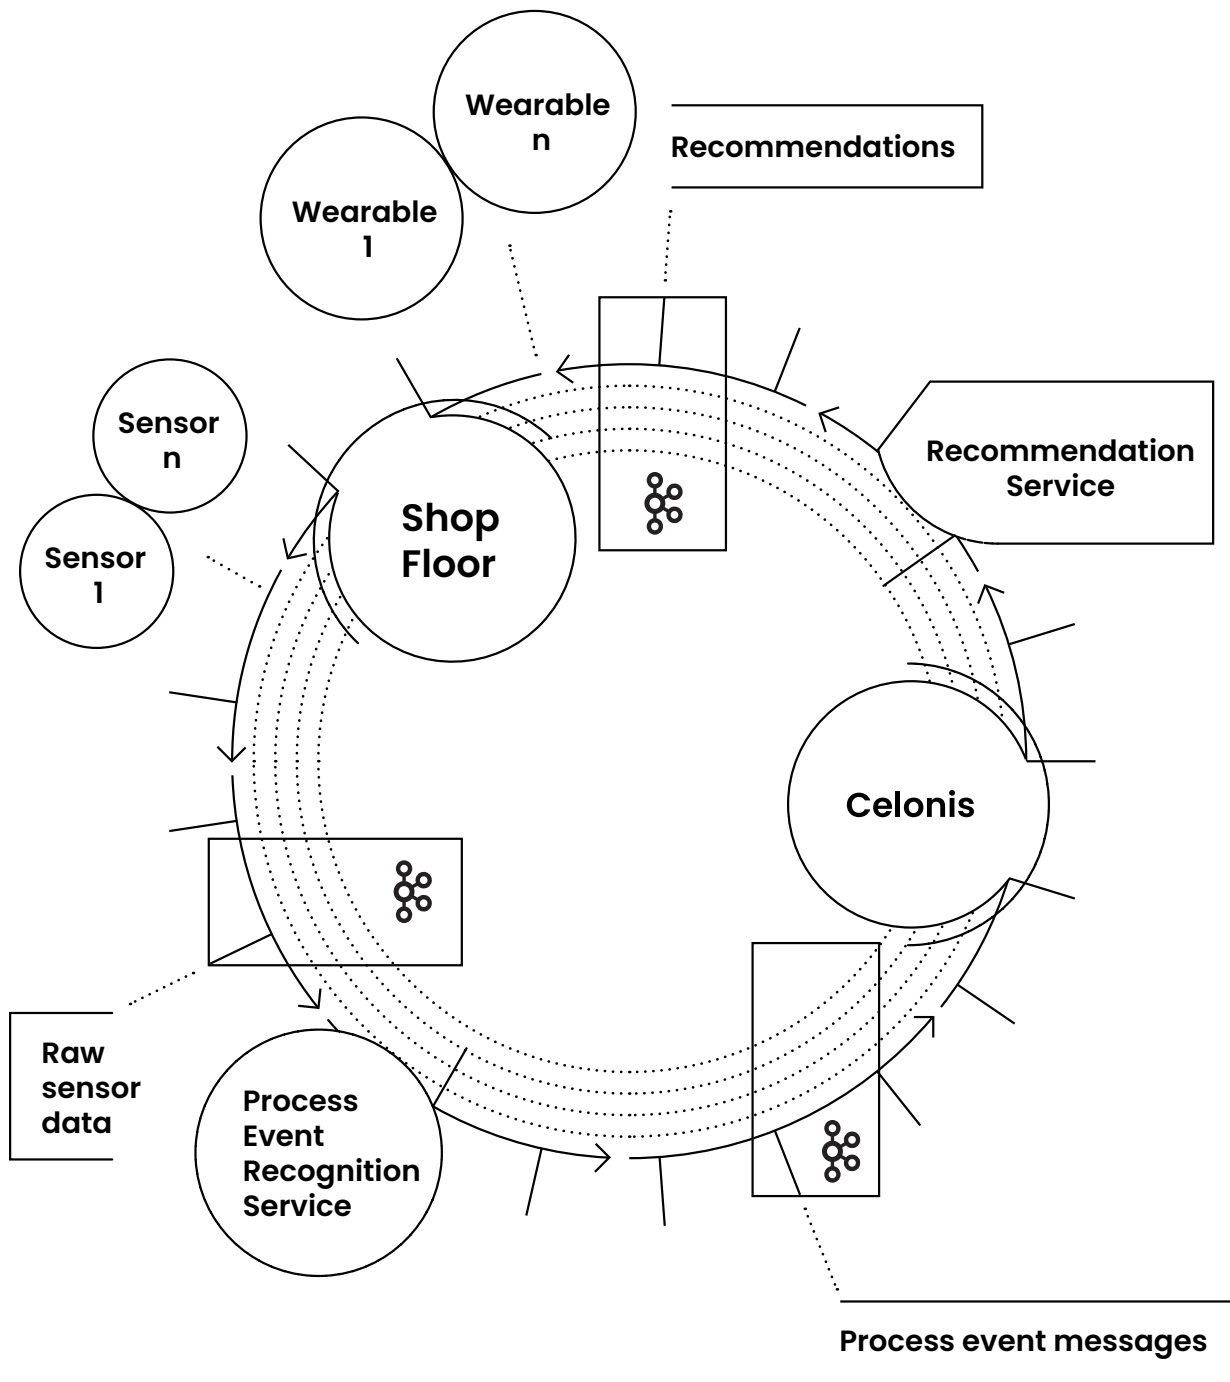
In last year's Beyond Journal, we introduced this visionary research project and talked about how event recognition based on sensor data can bring transparency to a new level, for example knowing exactly where each product component is in the manufacturing process in real-time.

We want to continue to share our insights by focusing on one of our main innovations namely extending the data pipeline with connectivity to wearable devices.

The pipeline extracts process events from raw sensor data. After successfully implementing and deploying this part of the pipeline, we maximized the value of derived insights by sending recommendations to action directly to workers on the manufacturing side. These recommendations are made available via push notifications on wearable devices

(e.g. smartphone, smartwatch), which enables fast action and thus preventing issues in the production process. With this, we are addressing a whole new segment of target persona: Shop floor workers that don't need to have a background in data-driven decision making or enablement on the Celonis platform.

To ensure that this idea is successful, the derived recommendations for production workers need to fulfill three core requirements: Recommendations need to be **1** delivered fast, **2** valuable, and **3** actionable.



# 1

Especially in operational use cases, reducing the time-to-act is a fundamental value driver and thus a competitive advantage. This is why real-time capabilities are crucial. To address this requirement, we implemented an end-to-end streaming architecture based on Apache Kafka and MQTT - starting with sending raw sensor data and ending with the reception of actionable recommendations on the workers' wearable.

With this setup, we are able to deliver recommendations in a matter of minutes after a change in the production environment has been detected by a sensor.

# 2

Instructing workers to use wearables is a burden in the first place and as such, it will only be adopted if both the workers and the management see value in doing so.

To drive value with our recommendations, we implement a Time-of-Event Prediction service. Leveraging the current status of the order, process plans, live context (i.e. queues at machines), and historical data, this event-driven service predicts the time at which upcoming events are about to be completed and thus the arrival of the whole order. This unlocks highly valuable recommendations such as bottleneck avoidance and order prioritization. This eventually leads to less friction in the production process, lower inventory holding cost and higher customer satisfaction.

The workers themselves profit from this direct access to information by saving communicational overhead and acting immediately upon issues, which would otherwise result in additional workload if delayed.

# 3

Workers need to be able to act upon recommendations. In contrast to an analyst persona, the worker persona is not a data-driven role and usually does not have access to devices other than wearables. Thus, we need to be mindful that the messages deliver just as much context as the worker needs to act without any ambiguity.

To give a more tangible example: If a production order is expected to be finished late, we send a message to prioritize this exact production order by skipping queues. There is very little context given as more information would just distract from taking immediate action.

We believe that selectively exposing information from Celonis to other platforms unlocks a new dimension of business management by addressing a wider range of target persona. —



# 10

## Peeking into the decision-making process in the human mind via *DL Models*



**Teodor Fratiloiu**  
former Master Thesis student (TU Munich) at Celonis

---

### × DL Models

Deep learning is a method in artificial intelligence (AI) that teaches computers to process data in a way that is inspired by the human brain. Deep learning models can recognize complex patterns in pictures, text, sounds, and other data to produce accurate insights and predictions.

People? Challenging. What makes investors buy shares in a particular fund? Why do people prefer one team over another in sports?

Decisions are a “black-box” – we can see the outcome, but never the process. What if we told you there might be a quantitative way to peek into that?

We began by training DL classifiers on data from a study of people’s choices when making “bets” – selecting one of multiple options, each described by payoff/probability pairs, none overwhelmingly more enticing than the rest by trivial correlation. We observed these models could indeed replicate real people’s choices. We then constructed “see-through” AI models (decision trees, cascading architectures etc.) to reproduce their behavior. This resulted in a collection of visible weights and biases which can make decisions that align with the original data. Voilà, we have a transparent decision engine.

Next, we analyzed public trade data (buy or sell orders), and fed the now quantified share price sets into our engine. When we compared the outcome to what the majority of traders did during the same timeframe, we found our automated models’ choices correlated to real-life market sentiment. Neat.

Decisions are complicated, tiring, and prone to error. Imagine the value of an automated, fully transparent and accountable decision engine: not only can that increase process consistency, but you might never again have to run an old-school root cause analysis: your logs will have already recorded all the steps that went into those decisions by the time you get around to it. —

# 11

## Revolutionizing customer satisfaction in insurance claims handling

### The fusion of Process Mining and Machine Learning



**Johannes Sasowski**  
former Master  
Thesis student  
(TU Munich)  
at Celonis

One of the most crucial customer facing processes in the insurance business is the handling of a claim. Regulating a claim is a highly complex task, involving feedback loops, information retrieval and compliance. Leveraging advancements in the process mining and machine learning research, can we combine both of these technologies to predict when customers might churn as a direct result of a negative experience in claims handling? And can we derive insights about our process inefficiencies leading to such churn?

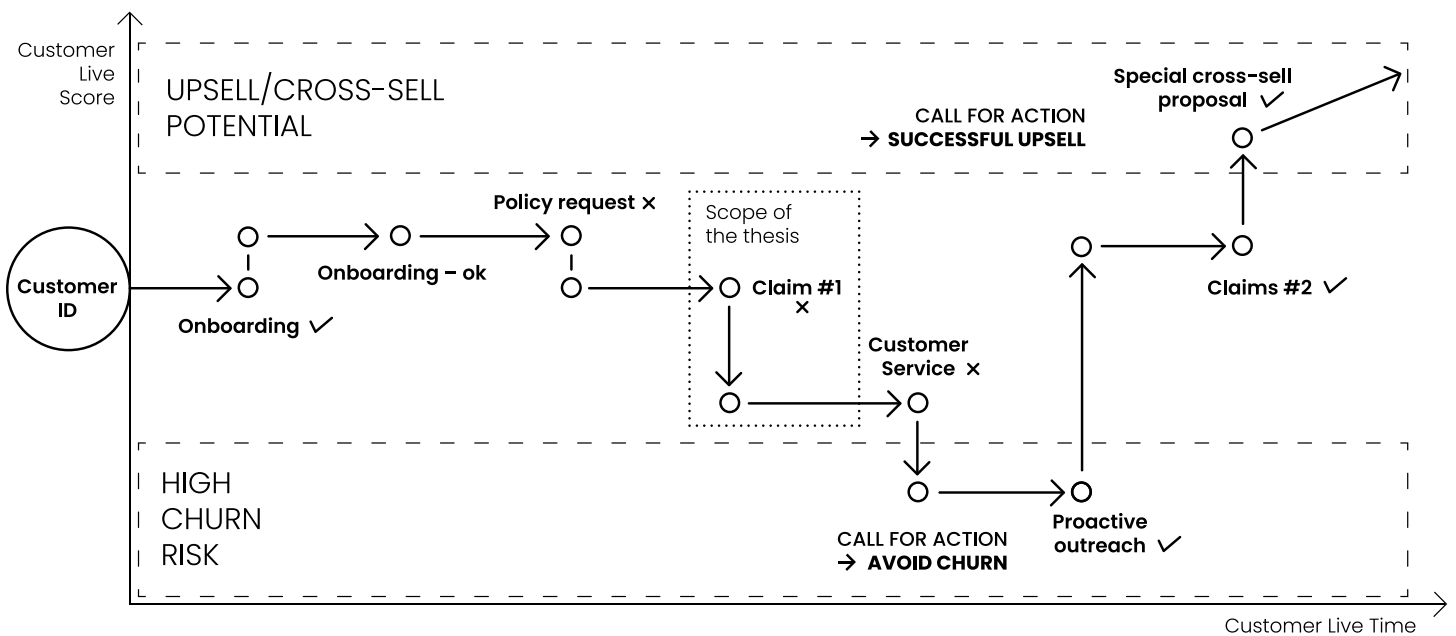
The research project which was conducted together with one of Germany's largest insurance providers aimed at developing the "customer live score model" – a

framework for churn prediction forecasting using process mining combined with machine learning. The introduced pipeline allows for flexible feature selection and automatic identification of an optimal machine learning algorithm by analyzing claims handling data. The project demonstrates promising results, enabling the identification of unsatisfied customers after a claim handling process is concluded and uncovering the underlying reasons behind their dissatisfaction. Meaning that we can not only say if a customer might churn as a direct result of a negative customer experience, but we can also identify the reasons behind it. Furthermore, by aggregating the individual feature attributions for every single churned customer we get the overall picture

of the current factors in our process which are leading to churn, paving the way for strategic initiatives to enhance long-term customer satisfaction.

One notable advantage of the introduced framework is its process-agnostic nature, enabling its application to any business process for binary outcome prediction (suc-

cess or failure). The methodology significantly contributes to ongoing research, showcasing the immense power of process mining and machine learning as decision-support systems that support daily business activities in real-life scenarios. Process Mining can help us in getting more data to open up the possibility for new applications of machine learning use cases. —



Every process interaction gets scored and can be linked to customer ID  
 → Customer Satisfaction Live Score & Intelligent Action Triggers

# 12

## Sustainability- Integrated Value Stream Mapping with Process Mining



**Jonas Weich**  
Application Product  
Manager at Celonis



**Jerome Geyer-Klingeberg**  
Head of Academic  
Alliance at Celonis

In manufacturing, Value Stream Mapping (VSM) is used to improve production processes. The methodology involves a comprehensive analysis of production process steps, from acquiring raw materials to the delivery of the final product. The goal of VSM is to identify and eliminate non-value-adding activities, reduce waste, and streamline the production process to maximize the value delivered to the customer while minimizing costs. Sustainability metrics such as electricity consumption and carbon emissions have recently become an area of waste that is attracting attention in the context of increased awareness of climate change.

In this publicly funded research project, we introduced and implemented an Celonis-native approach for sustainability-integrated value stream mapping with Process Mining. After identifying relevant user stories and personas both in the manufacturing and supply chain area, we developed a digital VSM that

solves crucial shortcomings of the traditional, brown-paper based VSM. Being a manual, one-off approach, the traditional VSM suffers from low data accuracy, subjectivity, cost, and scope limitations. For integrating sustainability metrics into the digital VSM, we incorporated data from power sensors, scrap rates and transportations between plants. Combined with the CO2 percentage in the energy mix, and ERP data on the production process, this enables labeling each production order and thus each material with an individual CO2 footprint and thus unlocks a wide range of use cases - from assessing, reporting, to optimizing sustainability metrics. Apart from sustainability metrics, the developed solution enables users to identify wasteful activities such as overproduction, excess inventory, defects, over-processing, waiting, and unnecessary transportation. By analyzing these areas, manufacturers can create targeted improvements that reduce costs, improve quality, and enhance customer satisfaction through service levels. —

# 13

## Next generation of Conformance Checking in Process Mining



**Eduardo Goulart Rocha**  
former Master Thesis  
student and now  
Research Software  
Engineer @PQL team

with  
**Wil van der Aalst**



Conformance Checking is the field in process mining that compares a normative process model with the actual process behavior recorded in an event log. The main tasks of conformance checking are to measure the degree of deviation between the model and the event log using a single number evaluation metric and to provide useful diagnostics about these deviations to the end user.

State-of-the-art techniques are based on trace alignment, which finds the model execution that most closely matches each variant recorded in the log according to a user-defined measure of similarity. However, trace alignment has two inherent limitations. First, it does not scale for large event logs and process models, making it impractical for large industrial datasets. Second, even if the computation successfully terminates, interpreting its result requires technical understanding on the workings of the algorithm, which hampers its broader adoption.

In my research, I explore alternative methods for conformance checking that overcome these limitations. My most recent paper takes a step in that direction by presenting a method to efficiently measure the conformance for process trees. Process trees are a widely used process modeling notation in process mining. At Celonis, they are used, among other things, as the backbone of our new Process Sphere product. The new method achieves an improvement of over an order of magnitude compared to state-of-the-art techniques. This enables us to monitor the quality of our automatically discovered process models in production and to improve our mining algorithms based on this data. In the future, we plan to use the same technique to provide understandable conformance diagnostics to business users. At Celonis, we believe that alternative, more pragmatic methods for conformance checking are crucial to its widespread adoption in the industry. —

## 14

Cognitive  
Digital Twins

**Dimitris Kiritsis (Kyritsis)**  
former Professor of ICT for  
Sustainable Manufacturing  
at EPFL Lausanne, Senior  
Advisor at University of Oslo

---

× **Cognitive Digital Twin (CDT)**

is a digital representation of a physical system that is augmented with certain cognitive capabilities and support to execute autonomous activities; comprises a set of semantically interlinked digital models related to different lifecycle phases of the physical system including its subsystems and components; and evolves continuously with the physical system across the entire lifecycle

The concept of *Digital Twin* (DT) combines and integrates key enabling technologies of the 4th Industrial revolution, such as all kinds of sensors collecting a big variety of data, Internet of Things, Data Analytics, AI etc. It represents a comprehensive physical and virtual description of a product or a system. After a decade of development, DT has recently started to be applied to many industrial sectors covering specific applications along the lifecycle of a product or a system. The first successful applications of DTs on specific isolated problems, have already revealed the need for the interconnection of DT solutions that operate at different aspects and different phases of the lifecycle of a complex system consisting of products, assets (equipment) and processes with the final goal the optimal and efficient design and operation of such complex systems.

The concept of *Cognitive Digital Twin (CDT)* has been recently proposed which reveals a promising evolution beyond the current

DT paradigm towards a system of interconnected Digital Twins covering linked applications all along the complete lifecycle of a complex system.

While there is no widespread definition of CDT yet, let's explore a few principal features of a CDT:

1. A CDT contains at least the three basic elements of DT, including the physical entity, virtual representation, and the connections between the virtual and physical spaces. The difference is that a CDT may contain multiple DTs, including digital models of its subsystems, and each of them has different status across the entire lifecycle.
2. Cognition capability: As indicated in its name, a CDT should have certain human-like cognition capabilities, such as attention, perception, comprehension, memory, reasoning, prediction, decision-making, problem-solv-

---

ing, reaction and so on. Thus, a CDT is defined to recognise complex and unpredicted behaviours with optimisation strategies dynamically.

3. Full lifecycle management: A CDT should consist of digital models covering different phases across the entire lifecycle of the system, including beginning-of-life (BOL, e.g. design, building, testing), middle-of-life (MOL, e.g. operating, usage, maintenance) and end-of-life (EOL, e.g. disassembly, recycling, re-manufacturing). It should also be capable of integrating and analysing all available data, information and knowledge from different lifecycle phases thus to support aforementioned cognitive activities.
4. Autonomy capability: A CDT should conduct autonomous activities without human assistance or minimum level of human intervention. This capability is partially overlapped with and empowered by the cognition capabilities of a CDT. For example, based on the perception and prediction results, a CDT can autonomously make decisions and react for design, production or operations adaptively.
5. Continuous evolving: A CDT should be able to evolve along the entire system lifecycle. There are three levels of evolution. First, for a single digital model, it updates itself according to the change of relevant data, infor-

mation from the physical system; second, due to the interactions among different digital models contained in the same lifecycle phase, each model evolves dynamically according to the impact of other models; third, due to the feedback from other lifecycle phases, the previous two situations may happen simultaneously, even new models and components will be added.

6. Governance approach: A holistic governance approach of a CDT system is required, integrating three different views: i) business and sustainability; ii) data governance and ii) cognition (AI) models governance.

The concept of Cognitive Digital Twin (CDT) represents a promising advancement in the evolution of Digital Twins (DT). Although this is still in its early stage, the concept of CDT has already been applied in Proof-of Concept projects involving Academia and Industry such as the EU funded project *Qu4lity*. Given the importance and application of DTs in recent years, we anticipate that CDT, as the next evolution, will garner increasing attention from both academia and industry. Furthermore, CDT is poised to become a potent tool in achieving the intelligent manufacturing paradigm.

More details about the concept of CDT can be found in the recently published paper available to download at this [link](#). —

# Celonis Labs

