

The Event-Driven Enterprise

Techniques to Rapidly Sense and Effectively Respond to Sudden Change, Risk and Opportunity

Carl Lehmann, Senior Research Analyst

This report defines the event-driven enterprise and explains how such enterprises will develop and modernize applications to create a constant state of execution and performance awareness to quickly sense and effectively respond — reactively, predictively and even autonomously — to sudden change, market dynamics, risks and customer opportunities.

To learn more or to request a demo, visit spglobal.com/marketintelligence.



Table of contents

Executive summary	4
Introduction	4
About this report	5
Key findings	5
The Take	6
The event-driven enterprise defined	6
The value and importance of being event-driven	6
Figure 1: Benefits to being more data-driven	7
Figure 2: Barriers to being more data-driven	8
Figure 3: Approaches to overcome data-driven barriers	9
Enabling the event-driven enterprise	10
Cloud-native application modernization	10
Figure 4: State of application modernization	10
Figure 5: Planned IT spending changes	11
Event-driven architecture	12
Figure 6: An event-driven architecture	13
Event-stream processing	14
Figure 7: Adoption of event-stream processing	14
Use cases of event-driven enterprises	15
Figure 8: Use cases for event-stream processing	15
Intelligent process automation	16
Vendor landscape	17
PaaS providers	17
Low-code application development platforms	17
Digital automation platforms	17
Hybrid integration platforms	18
Vendors offering event-driven technologies	18
Figure 9: Vendors offering enabling event-driven technology	18

Table of contents

Conclusions	22
Further reading	22
Methodology	23
Index of companies	24
About the author	25

Executive summary

Introduction

Our research indicates that enterprises are accelerating their investment in, and pace of, application modernization to realize the benefits of cloud-native technologies. Moreover, they seek to unlock the potential of cloud-native architecture to improve the quality of business decisions and the execution of their digital business operations through the ongoing analysis of real-time data generated from IT and business events.

“Cloud native” describes a set of architectural principles and technologies that decouple applications from infrastructure, and decompose applications into sequences of containers, microservices and serverless functions. It can speed development, deployment and remediation as well as enable rapid adaptation when change is needed. Moreover, workloads of all types can be distributed to the most cost-effective or performant runtime environments required. Enterprises can become more agile in their digital business operations, and developers can be freed to concentrate on more engaging pursuits such as crafting logic and improving the capabilities and quality of mission-critical applications.

The capabilities of cloud native, however, are not enough to fully enable the event-driven enterprise — one that is able to rapidly sense and effectively respond to sudden change, risk and opportunity of any type. Challenges persist. Among them is how to best integrate all the moving parts associated with cloud-native applications (e.g., containers, microservices, functions, code, logic, et al.) across an increasingly distributed hybrid IT landscape of on-premises and edge infrastructure and multicloud environments. Cloud native requires a different approach to integration. This is the role of event-driven architecture.

Event-driven architecture is a software design and communications pattern for building applications composed of decoupled operations that respond to notifications of events. It has existed in various forms for decades, but cloud-native computing gives it new purpose. Event-driven architecture can quickly facilitate the connections between, and interoperability among, highly distributed and diverse infrastructure, cloud services, applications, data, devices and processes of any type that compose today’s modern hybrid IT landscape. It is rapidly becoming a foundational integration environment in support of cloud-native application development and modernization, and a core enabler of the event-driven enterprise.

The collection, correlation and analysis of data from events, and the actions taken in response, are the outcomes of event-stream processing and intelligent process automation. Event-stream processing is a critical capability of event-driven architecture. It actively tracks and analyzes continuous streams of data between producers and consumers of events within and across enterprises so that change, risk and opportunities can be proactively identified and interpreted. Effective responses are then created and developed using cloud-native application development and intelligent process automation platforms to consistently achieve or optimize business outcomes.

Event-driven enterprises are building and modernizing the applications they use to run their digital businesses using a combination of cloud-native and event-driven architecture to maintain a constant state of awareness, and intelligent process automation to take decisive action. They are sharpening their digital business senses and reflexes to be able to sense and effectively respond — reactively, predictively and even autonomously — to sudden changes in execution, performance, market dynamics, business risks and customer opportunities better and faster than before, or at least better and faster than rivals that are not event-driven.

About this report

This report is a 451 Research | S&P Global Market Intelligence, Technology & Business Insight (TBI) report. It examines, in-depth, trends in technology markets that impact the business and IT strategies of enterprises and the product and services strategies of IT vendors. This TBI report is forward-looking and examines the trends we are seeing among enterprises that seek to improve their digital business operations by modernizing their applications to take advantage of cloud-native and event-driven architectures to be able to rapidly respond to business events predicatively or upon occurrence. It is based on information received from dozens of briefings over the last six months with vendors, end-user conversations, industry conferences and quantitative analysis from multiple 451 Research services including Voice of the Connected User Landscape (VoCUL), Voice of the Enterprise (VotE), Market Monitor, the M&A KnowledgeBase and the Datacenter KnowledgeBase.

The purpose of this report is also to trigger curiosity about how business and IT leaders can come together to understand how innovative and evolving technologies can be used to improve the sense and response reflexes of their enterprises. This report does not offer all the answers, but it should spark questions about how to become an event-driven enterprise and what it will take to reap its benefits.

Key findings

- **Event-driven enterprises will be data-driven.** They will secure and exchange data with external parties; create dedicated data engineering functions; establish self-service programs and technologies for data preparation, visualization and analysis; and deploy data fabrics.
- **Enterprises are rapidly moving to cloud-native application development.** The benefits associated with loosely coupling infrastructure, logic and data enable a great deal of development and runtime execution flexibility.
- **Event-driven architecture will be the central nervous system for digital operations.** It can produce, detect, consume and, when supported by the proper development and automation technology, react to events as they occur in real time.
- **Event-stream processing improves the speed and quality of decisions.** It allows decision-makers to harness streams of event data from devices and applications to make data-driven decisions. It plays a critical role in applications that must act quickly upon event data.
- **Intelligent process automation enables decisive action in response to events, and then learns from them.** It is a data-driven approach to visualize, examine, improve and automate business processes and workforce activities. It can also learn by capturing data from execution and enable processes to make recommendations, predictions and decisions.
- **Awareness of the benefits of being event-driven must extend outside the IT organization.** Business leaders must be involved. They need to understand the impact that cloud-native computing, event-driven architecture, event-stream processing and intelligent process automation can have on the performance and execution of individual lines of business, as well as on the enterprise as a whole — enterprises become more agile in their digital business operations, able to rapidly and effectively respond to sudden change, risk and opportunity.

The Take

Event-driven cloud-native applications have significant mindshare among IT professionals, and they are actively pursuing the technologies that enable such applications. The more they research and understand these technologies, the greater their propensity to use them to transform their IT organization. Indeed, this is all well and good. However, it's also important for non-IT business leaders to understand the considerable benefits derived from cloud-native computing and the awareness achieved through event-driven architecture. Organizations can improve upon virtually all of their key performance indicators: Inventory can be maintained at optimal levels, financial transactions can execute without delay, perfect order rates can be improved, production lines can operate unimpeded, claims can be processed and cases resolved quickly, compliance rates can be improved, and error rates can be diminished. These are just a few examples of the types of business outcomes that can be realized when business and IT leaders come together to strengthen the “sense and respond” reflexes of their organizations and invest in the technologies needed to enable the event-driven enterprise.

The event-driven enterprise defined

An event-driven enterprise maintains a constant state of awareness and can quickly sense and effectively respond — reactively, predictively and even autonomously — to any sudden changes in execution, performance, market dynamics, business risks and customer opportunities. Out-of-stock inventory, failed financial transactions, delayed customer orders, unfulfilled claims, unexpected customer demand, production line problems, equipment failure and regulatory non-compliance are but a few examples of events that organizations must respond to and remedy quickly, if not predict and prevent them, before business is negatively impacted or opportunity is lost.

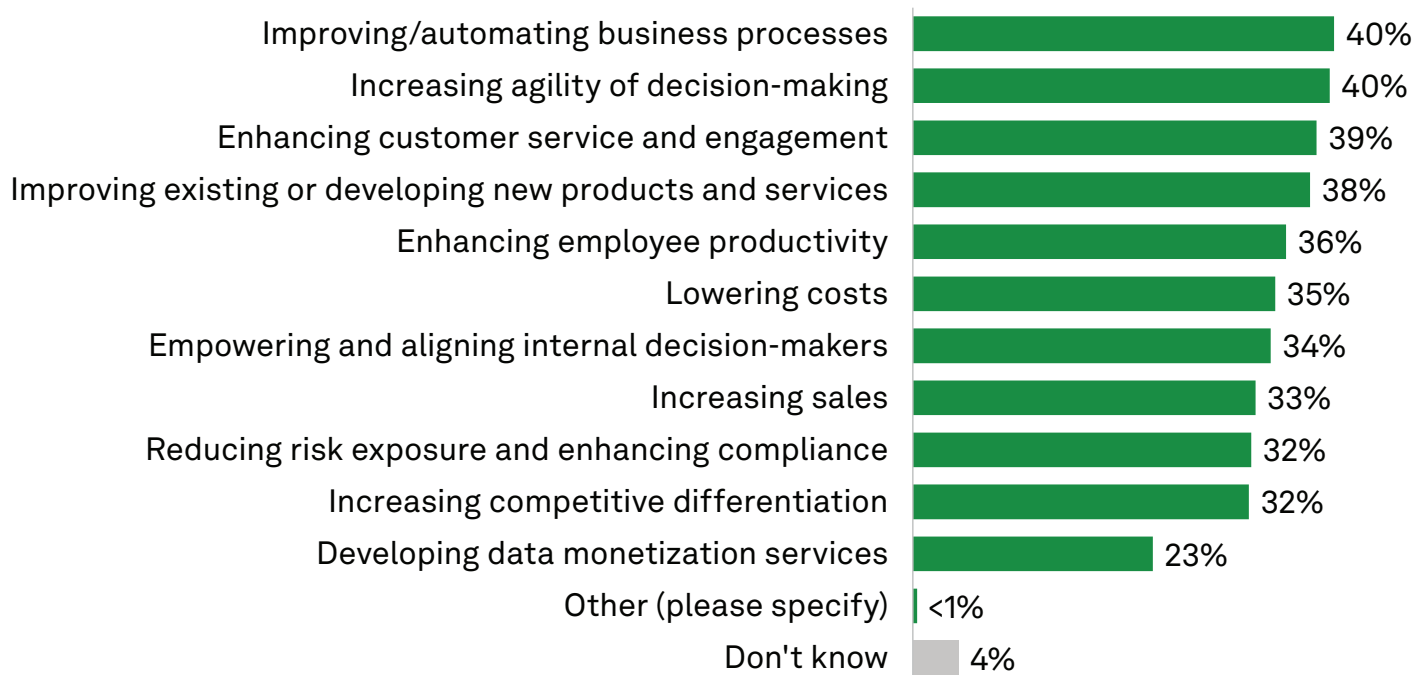
Event-driven enterprises modernize the applications and processes they use to run and manage their digital businesses. They use cloud-native computing and software design patterns that build highly adaptive and portable applications composed of decoupled containers, microservices and serverless functions to perform operations that respond to notifications of current events or predict and prevent undesired potential events. Event-driven enterprises are able to sense, interpret, decide and act accordingly, proactively and in real time, upon events that then rapidly trigger communications between, and actions among, interested and affected services, systems and stakeholders.

The value and importance of being event-driven

All events generate data; therefore, event-driven enterprises are also data-driven. They have a ravenous appetite for accurate, reliable and real-time data about digital operations, execution performance and business outcomes. Being data-driven means that enterprises strive to make as many strategic and tactical decisions as possible based on real-world data rather than on personal experience, anecdotes or “gut feelings.” However, being event-driven goes a step further. It means that enterprises use data to sense change, risk and opportunity; interpret the implications; decide what to do; act (via human intervention or autonomously) within relevant rules, policies and regulations; and track the results and outcomes that then generate feedback data that is codified to enable enterprises to continuously learn about the effectiveness of their sense-and-respond reflexes.

Since data is a fundamental element of event-driven enterprises, we continuously seek information about how enterprises view the benefits of being more data-driven. In a recent VotE survey, we asked 477 business and IT decision-makers about the benefits to their organization of being more data-driven. Figure 1 reveals that improving/automating business processes (40%) and increasing the agility of decision-making (40%) are among the top benefits. Enhancing customer service (39%), improving/developing products and services (38%) and enhancing employee productivity (36%) closely follow.

Figure 1: Benefits to being more data-driven



Q. What are the most significant benefits your organization would expect from being more data-driven?

Please select all that apply.


Base: All respondents (n=477).

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data-Driven Practices 2022.

The benefits of being more data-driven essentially run the gamut of a modern digital business — how it executes, makes decisions, interacts with customers and partners, enhances offerings and improves employee productivity. Being data-driven also reduces risk and increases competitive differentiation; it benefits the entire enterprise, according to our findings.

That said, there are challenges that prevent enterprises from becoming more data-driven. Figure 2 reveals many potential barriers. Data security (30%) and privacy (28%) are perennially among the top concerns of business and IT leaders. Likewise, the availability of skilled human resources (31%) and budget (30%) are common constraints.

Figure 2: Barriers to being more data-driven



Q. And what are the most significant barriers your organization faces in attempting to be more data-driven?

Please select all that apply.

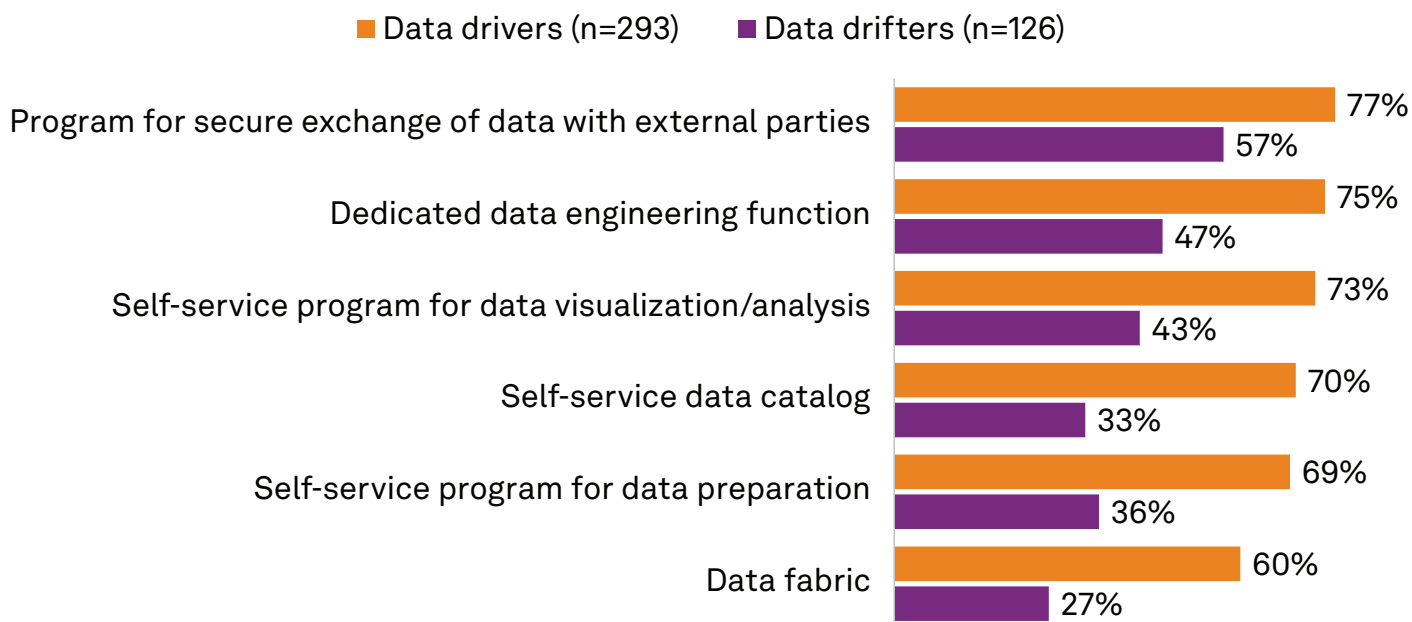
Base: All respondents (n=477).

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data-Driven Practices 2022.

A more important insight from our research, though, is that data quality and consistency (35%) surfaces as the most frequently cited barrier. Other notable obstacles include integration of legacy data (24%), ease of use of relevant tools (23%), maintenance of data architectures (21%), data exchange requirements with external parties (20%), and data navigability and accessibility (20%). Barriers such as these not only impede the benefits of being more data-driven, but they also diminish the ability of enterprises to be event-driven and thwart the ability to craft rapid and accurate sense and response reflexes to business events.

We are also interested in what approaches and technologies organizations use to overcome these barriers. Here, we classify organizations into two camps based on their responses. Organizations that we classify as “data drivers” make nearly all or most decisions based on data, while organizations that make fewer data-driven decisions are categorized as “data drifters.” Figure 3 illustrates that data drivers have a far greater propensity to engage in data quality management programs and invest in data quality management technologies.

Figure 3: Approaches to overcome data-driven barriers



Q. Does your organization currently have any of the following? - Data fabric, Dedicated data engineering function, Program for secure exchange of data with external parties.

Q. And does your organization currently have any of the following self-service platforms or programs? – Self-service data catalog, Self-service data marketplace or exchange, Self-service program for data preparation, Self-service program for data visualization/analysis.

Base: All respondents.

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data-Driven Practices 2022.

Data drivers secure and exchange data with external parties; create dedicated data engineering functions; establish self-service programs and technologies for data preparation, visualization and analysis; and deploy data fabrics — network-based architecture from vendors such as Informatica Inc. and Talend Inc. that addresses the complexity of data management by intelligently integrating and connecting an organization's data sources and making data assets available for consumption.

Based on these findings, we believe that organizations considered data drivers are also likely to be event-driven. However, investing in data quality management alone falls short of enabling the event-driven enterprise. Other evolving technologies and strategic IT initiatives are also necessary to sharpen digital business reflexes. Among them are the drive toward cloud-native application modernization, the reinvigoration of event-driven architecture, the advent of event-stream processing, and the emergence of next-generation cloud-native and event-driven application development and automation platforms.

Enabling the event-driven enterprise

Cloud-native application modernization

In their efforts to create competitive advantage, enterprises continuously pursue a range of IT innovations. A current top priority among enterprises is to modernize their legacy application environment (see Figure 4). Three-quarters of survey respondents say their organization has a formal strategy for application modernization and the effort is underway, or it is planning and researching one. Organizations are doing so to realize the benefits derived from cloud-native applications.

Figure 4: State of application modernization

- **Execution.** We have a formal strategy and are actively modernizing our application environment
- **Evaluation.** We are planning and researching to develop an application modernization strategy
- **Consideration.** We are considering it but have no formal plans
- **No strategy.** We currently have no plans to modernize our application environment



Q. Where does your organization stand regarding its plans to modernize its legacy application environment?

Base: Organizations expressing need for some degree of modernization (n=316).

Source: 451 Research Voice of the Enterprise: Digital Pulse, Application Modernization 2021.

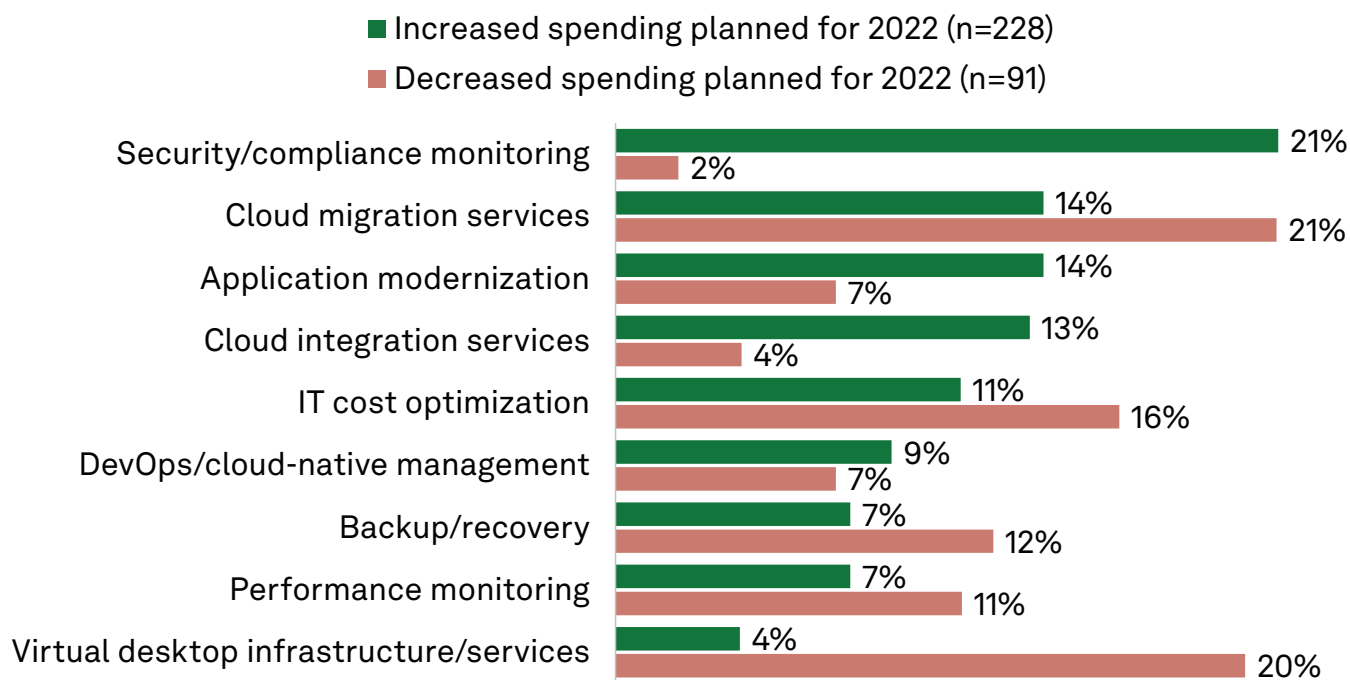
Cloud-native applications are designed, built, delivered and operated differently than traditional/legacy application development efforts. They make use of agile development (rapid software iterations) and DevOps methods (collaborative application and IT operational teams and technologies); cloud resources (virtualized, open source, compute, storage, network); containers (packaged, self-contained code); microservices (discrete logic); serverless functions (code that does not require server provisioning); and application programming interfaces (“APIs” for integration), among other cloud-based services (e.g., databases, analytic tools, backup and recovery services, et al.).

Cloud-native applications are composed of loosely coupled discrete and reusable components that are not hardwired to the underlying IT infrastructure, making them more easily adapted and portable. Logical operations that compose cloud-native applications reside in containers that can include many microservices and serverless functions that are then assembled and orchestrated to perform as a coordinated application experience. Coordinating the application experience requires orchestration technology such as Kubernetes for container orchestration, service mesh technology such as Istio for microservices orchestration, and external integration technology such as APIs when data and logic need to be exchanged among components and outside applications.

Cloud-native applications can be rapidly deployed into production using continuous integration/continuous delivery tools and other repeatable automation and orchestration techniques. They can run in a range of heterogeneous execution venues (e.g., public clouds, managed services, on-premises, hosted private clouds, hybrid clouds and at the network edge). As such, applications can run anywhere and be deployed to the most cost-effective runtime environments and where they are best suited.

Enterprises are steadfastly progressing toward cloud-native application development. The benefits associated with loosely coupling infrastructure, logic and data enable a great deal of development and runtime execution flexibility. Enterprises become more agile in their digital business operations, and developers are freed to concentrate on crafting logic and improving the capabilities and quality of mission-critical applications. In fact, our research shows that the effort to modernize applications is accelerating. Figure 5 illustrates that twice as many respondent organizations will increase their spending on application modernization (14%) as will decrease spending (7%).

Figure 5: Planned IT spending changes



Q. In which of the following do you expect to see the largest increase in spending in 2022, compared to 2021?

Base: Users of each service planning to increase/decrease spending on specific cloud-attached services.

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets & Outlook 2022.

The distributed nature of cloud-native computing affords enterprises flexibility, but it can also create complexity. A great deal of data and logic is distributed across multiple containers, microservices and functions. Container orchestration technologies such as Kubernetes and service mesh technologies such as Istio do a good job coordinating the application experience. However, coordinating an environment of potentially millions of decoupled operations across a diverse hybrid IT landscape also requires a distributed communications and messaging infrastructure that can accommodate data exchange triggered by changes in state, otherwise called events.

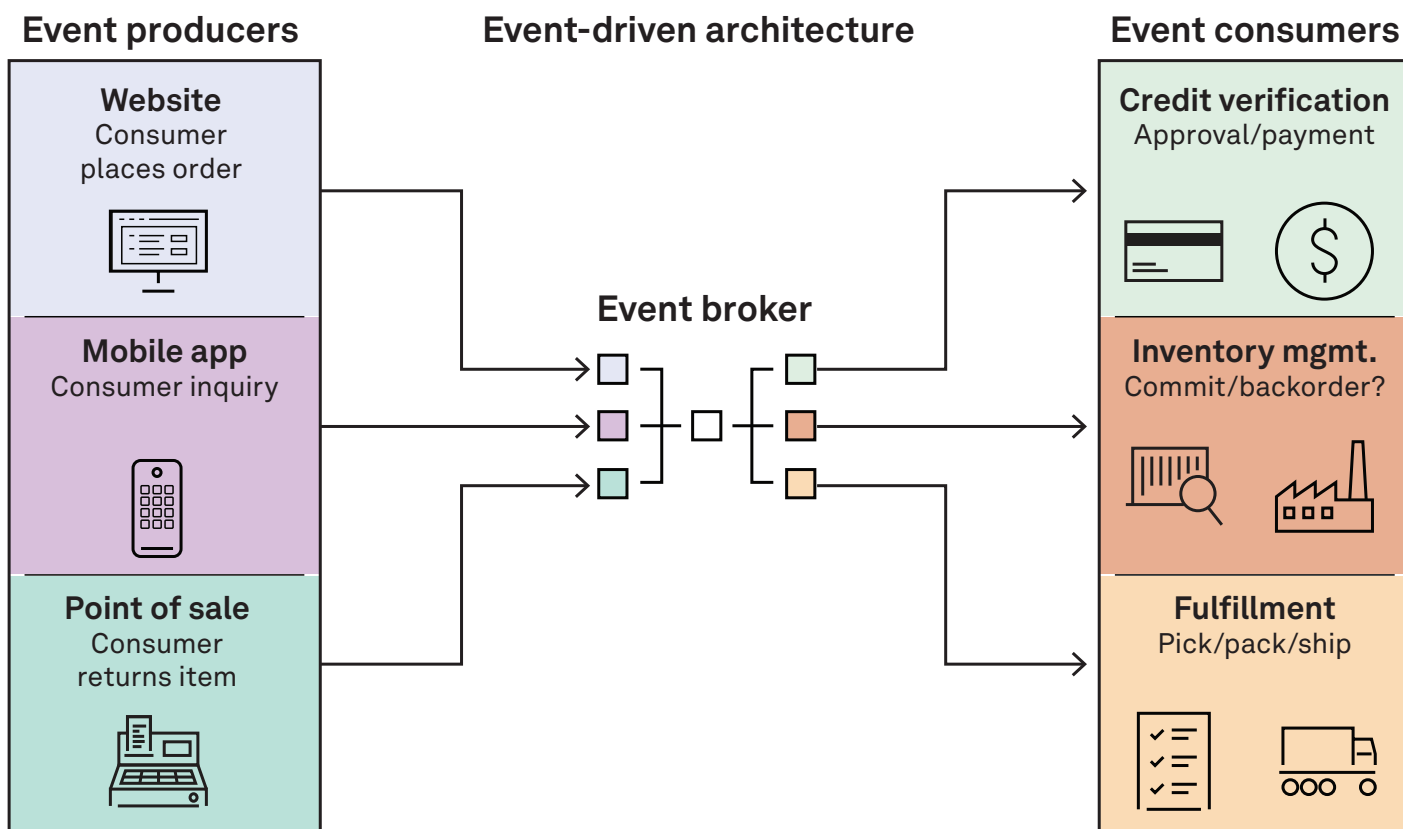
Event-driven architecture

Event-driven architecture (EDA) is a software design pattern for building applications comprising decoupled operations that respond to notifications of events. More specifically, an event is a notification that a change in state has occurred (e.g., for sale/sold, working/not working), and it may include payload data (and metadata) to be acted upon. In an EDA, the publishers of events (called “producers”) do not need to overtly know their subscribers (called “consumers”); the consumers decide which event sources they are interested in. EDA surpasses typical messaging-oriented integration techniques because it can also produce, detect, consume and, when supported by the proper development and automation technology, react to events as they occur in a real time (or near-real-time) and in an asynchronous manner.

An EDA can use a “publish and subscribe” (pub/sub) or an event-stream messaging model. Pub/sub is a messaging technique that keeps track of subscriptions (consumers of events). When an event is published, it sends the event to each subscriber. After an event is received, it typically cannot be replayed, and new subscribers do not see historical events. In event-stream messaging, events are written to a log, strictly ordered within a partition, and durable. Consumers can read from any part of the stream, can join at any time, and can replay events.

The EDA of a retailer illustrated in Figure 6 enables a constant state of awareness, allowing producers to publish or stream data about various types of events, in real time upon occurrence, to all potential consumers. For example, a customer purchase in an e-commerce website produces an event that is consumed by interested services such as credit verification, inventory availability and fulfillment.

Figure 6: An event-driven architecture



Source: 451 Research.

At the heart of most EDAs is an event broker (or cluster of brokers as the architecture requires). Sometimes also referred to as an event router, event hub or event bus, it serves as an active watcher of event producers and pushes events to interested event consumers. They are typically composed of a cluster of multiple event-stream processing nodes for redundancy and can run on-premises or in a cloud. They facilitate the decoupled nature of cloud-native architecture. Distributed services (applications or workloads) do not need to know about each other. Rather, they only need to be aware of the event broker(s). This enables service integration, interoperability and resiliency; if one service fails, the rest can keep running. Developers do not need to write custom code to poll, filter and route events; the event broker will automatically filter and push events to required consumers.

An event broker also acts as a centralized location to audit applications and define policies. These policies can restrict who can publish and subscribe to a router and control which users and resources have permission to access data. It can also enable encryption of events both in transit and at rest. Since all activity in an EDA is push-based, everything happens on demand as the event presents itself in the broker. This enables cost control in cloud-native architecture by eliminating the need for continuous polling, and thus avoids costs associated with network bandwidth consumption and CPU utilization.

Event-stream processing

Event-stream processing (sometimes also known as “data streaming”) is a capability of EDA used to stream the ongoing delivery of event data among producers and consumers of events in real time and asynchronously. It can integrate with and/or be supported by data fabrics that enable access to diverse data sources and improve data quality. Event-stream processing allows decision-makers to harness streams of event data from applications and devices to make decisions. It also plays a critical role in applications that must act quickly upon event data.

Our research indicates that event-stream processing is growing in importance. Figure 7 illustrates the findings from our recent VotE: Data Platforms survey, where over half of the 382 respondents have it in use or are in a proof-of-concept or pilot stage.

Figure 7: Adoption of event-stream processing

- Currently in use
- Currently in a proof-of-concept or pilot stage
- Do not use event-stream processing but plan to within the next 12 months
- Do not use event-stream processing but plan to within the next 13-24 months
- Do not use event-stream processing but plan to within the next 2-3 years
- Not aware of any plans to use event-stream processing within the next 3 years



Q. Event-stream processing technologies allow decision-makers to harness streams of event data from devices and applications and make decisions in real time. Which of the following best describes your organization's event-stream processing adoption plans?

Base: All respondents (n=382).

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Platforms, 2022.

There are several message broker technologies that enable event-stream processing. Among them are ActiveMQ, Amazon Kinesis, Amazon SQS, IBM MQ, Apache Pulsar, RabbitMQ and Red Hat AMQ. However, we have observed growing enterprise adoption of open-source software Apache Kafka in recent years to power event brokers and facilitate an EDA across distributed systems in use by enterprises.

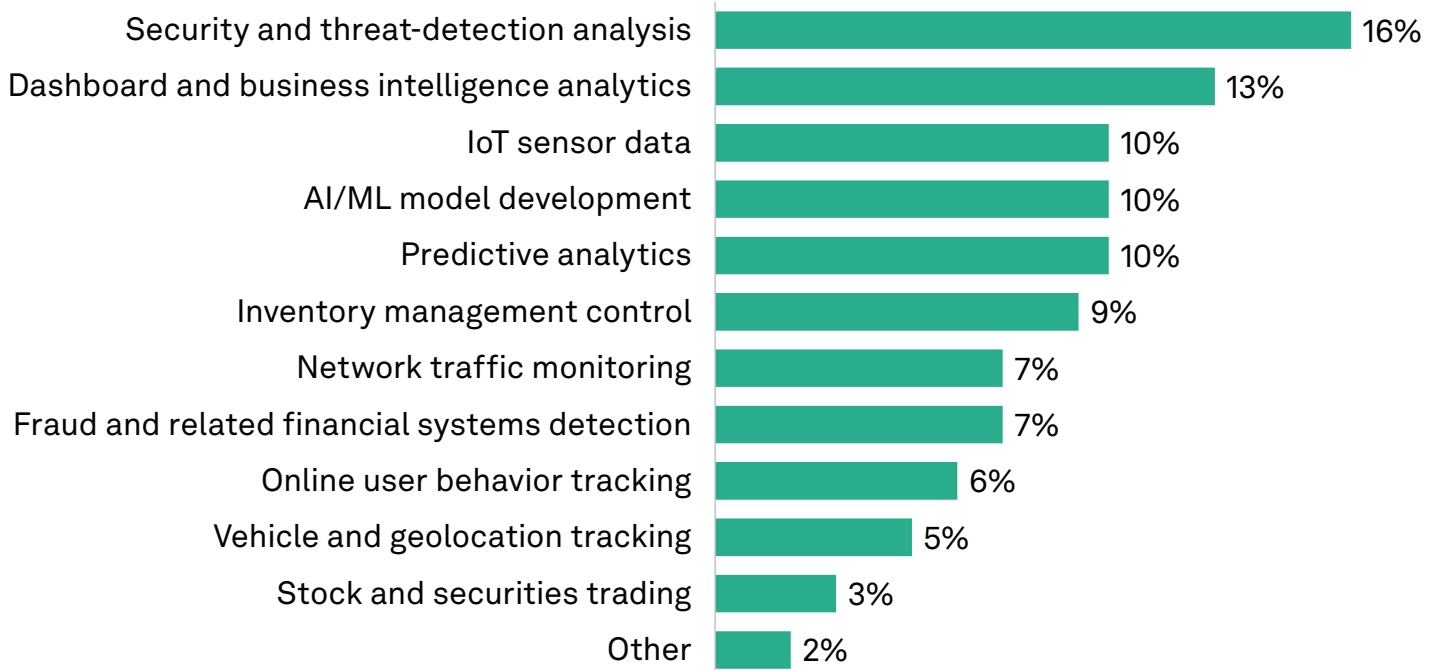
In event-driven IT ecosystems, there can be innumerable events of many types. Kafka provides a means to organize them using ordered logs of events called “topics.” Generally, in Kafka, “topic” refers to a particular heading, category or common name used to store and publish a particular stream of data (e.g., customer purchases, credit verifications, inventory availability, fulfillment, etc.). A producer publishes data to the topics, and a consumer reads that data from the topic by subscribing to it. Topics in Kafka are similar to tables in databases.

This design enables Kafka to be scalable and fault-tolerant. Event data is distributed across multiple clusters and nodes, and it is persistent. Kafka (and all other event brokers) minimizes the need for multiple integrations, and data streams can be consumed in real time or when needed, thereby enabling low latency and high throughput across networks. Indeed, for many companies, Kafka is becoming the de facto digital nervous system for cloud-native computing and a core enabler of event-driven enterprises.

Use cases of event-driven enterprises

In the same survey, we asked respondents to prioritize their top use cases for event-stream processing, as illustrated in Figure 8. All are applications that must act quickly upon event data and represent common use cases pursued by event-driven enterprises.

Figure 8: Use cases for event-stream processing



Q. Which of the following use cases best describes your organization's implementation of event-stream processing technologies?

Base: Respondents whose organization currently has event-stream processing technologies in use, in a proof-of-concept or pilot stage, or plan to use within the next 12 months (n=257).

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Platforms, 2022.

However, it is important to note that event-stream processing is limited in its scope to enable an event-driven enterprise. The term “processing” in “event-stream processing” means to analyze data and use that analysis to make decisions. Event-stream-processing technologies typically lack decision rules engines or AI algorithms to help make decisions, though. Nor are they able to act upon those decisions by enabling humans to build and trigger preconfigured automated workflows, or potentially to trigger automated workflows autonomously. Doing so requires that event-driven enterprises instill in their application design and development strategy the capabilities to enable intelligent process automation.

Intelligent process automation

Intelligent process automation is a data-driven approach intended to visualize, examine, improve and automate business processes and workforce activities. It uses real-world data captured during process execution and as the workforce performs tasks. Analysis reveals opportunities to improve performance, automate, and guide developer and stakeholder actions. Intelligent process automations can also learn by capturing data as they execute to monitor performance and outcomes, giving them the ability to make recommendations, predictions and decisions.

Intelligent process automations are built using cloud-native application development technology and use event-stream processing to continuously listen for and analyze events. Upon occurrence, they interpret their meaning and importance, decide upon a course of action based on statistical probabilities of success and then act by either alerting a human for approval or executing responses autonomously. Results are tracked and measured within relevant rules, policies, compliance and performance metrics, and stored to learn from the response and improve the quality of event-stream processing analysis.

What follows is a description of the application development and integration technologies needed to enable an event-driven enterprise, and a representative list of vendors offering enabling platforms to do so.

Vendor landscape

Event-driven enterprises use a combination of tools from several categories of technology to build cloud-native event-driven, automated, intelligent process applications. Indeed, typical polyglot integrated development environments can be used. However, they place greater a burden on developers to provision and manage the cloud-native infrastructure and services that their applications require for development and runtime.

There are other development platforms and specialty integration technologies that developers can use to improve their productivity and accelerate the pace of application deployment into production. Among them are PaaS offerings, low-code application development platforms (LCADs), digital automation platforms (DAPs) and hybrid integration platforms (HIPs).

PaaS providers

PaaS offerings are cloud services that enable complete application life cycle support, from development to testing, deployment to management, and updating. They free organizations from the complexity and cost of acquiring and maintaining a range of tools and resources such as runtime environments and various types of application services or middleware, operating systems, servers, storage and networking. Developers can concentrate on the quality and capabilities of the code and data relevant to their applications. Representative platforms include Red Hat Inc. OpenShift and Google LLC App Engine.

Low-code application development platforms

LCAD platforms provide life cycle management functionality that supports the design, build, test, deployment and management stages of a software life cycle. They abstract away from the developer the need to entirely use one or more programming languages to code software. Low-code environments use visual models, prepackaged templates and graphical design techniques with drag-and-drop tooling to build software. They typically offer means to craft or insert code (e.g., Java, JavaScript, HTML) as an option to customize logic and user interface (UI) designs. They enable software to be “configured” rather than “coded,” but it can be customized using code. Representative platforms include Mendix Technology BV and OutSystems.

Digital automation platforms

DAPs evolved from earlier business process management suites and generally enable the automation of end-to-end business processes. They include a set of tools and resources structured within a uniform framework to enable developers to rapidly design, develop, deploy, manage and monitor process-, content- and case-oriented applications. Like LCADs, they use low-code/no-code capabilities that include visual models, prepackaged templates and graphical design capabilities with drag-and-drop tooling to compose rather than code software. DAPs combine development and runtime IT environments. Changes can be made on the fly, making them agile DevOps platforms. Representative vendors include Appian Corp. and Pegasystems Inc.

Hybrid integration platforms

HIPs represent the next generation of integration platform-as-a-service (iPaaS) technology that enables data exchange and process interoperability across distributed hybrid IT landscapes. They generally include integration pattern development tools, messaging, data integration and API management capabilities. Some have data quality management tools and use AI to make recommendations to integration developers as to the “next best steps” to take during design.

Vendors in the integration market are just now awakening to the need for event-driven integration. Several, but not many, integration vendors are crafting a new breed of HIP capable of supporting EDA. These platforms will comprise, to varying degrees, a combination of connectors, pub/sub messaging capabilities, event brokers, APIM tooling, and synchronous and asynchronous service and event mesh technologies (i.e., configurable infrastructure for distributing events among decoupled applications). Representative platforms include Confluent and Solace. HIP vendors that also offer data fabric technology are in this category. Representative vendors include Informatica and Talend.

Vendors in each of these markets are also gradually increasing their analytic capabilities and use of AI to support the discovery and make recommendations toward the development of cloud-native event-driven intelligent process automations.

Vendors offering event-driven technologies

Figure 9: Vendors offering enabling event-driven technology

Legend	
PaaS = Platform as a service	DAP = Digital automation platform
LCAD = Low-code application development platform	HIP = Hybrid integration platform

Vendor	Technology pedigrees	Offerings
Amazon Web Services	PaaS	<ul style="list-style-type: none"> – Kinesis – a real-time, fully managed, scalable cloud service for streaming large volumes of data on AWS. – Simple Notification Service (SNS) – a fully managed Pub/Sub service. – Simple Queue Service (SQS) – a fully managed message queuing for microservices, distributed systems, and serverless applications.
	HIP	<ul style="list-style-type: none"> – Managed Streaming for Apache Kafka (MKS) – a managed service for Kafka. – Step Functions – a visual workflow service for distributed applications. – EventBridge – a serverless, fully managed, and scalable event bus that enables integrations between AWS services, SaaS offerings, and applications. – EventBridge Pipes – simplifies the creation of point-to-point integrations between event producers and consumers with data from sources including Amazon SQS, Amazon Kinesis, Amazon DynamoDB, Amazon Managed Streaming Kafka, self-managed Kafka, and Amazon MQ.
Appian	PaaS	<ul style="list-style-type: none"> – Internal Messaging Service – implemented using Apache Kafka and Apache ZooKeeper; it is responsible for relaying data between components of Appian’s architecture.
	DAP	
Boomi	HIP	<ul style="list-style-type: none"> – Partnered with Solace to provide event-streaming and management capabilities. (See Solace table entry for more information.)

Vendor	Technology pedigrees	Offerings
Confluent Cloud	PaaS LCAD HIP	– Confluent Cloud – a fully managed streaming data service based on Apache Kafka.
DataStax	PaaS HIP	– Astra Streaming – scalable event streaming powered by Apache Pulsar delivered as a fully managed service. – Luna Streaming – production-ready distribution of Apache Pulsar.
Google	PaaS HIP	– App Engine – a fully managed, serverless platform for developing and hosting web applications at scale. – Eventarc – offers a standardized way to manage the flow of events between decoupled microservices. – Google Cloud Pub/Sub – distributes events from databases.
Heroku	PaaS	– Heroku – a managed container system with integrated data services and an ecosystem for deploying and running modern apps. – Heroku Developer Experience (DX) – an app-centric approach for software delivery, integrated with popular developer tools and workflows. – Kafka on Heroku – a distributed messaging service for stream processing, optimized for developers.
IBM	PaaS DAP HIP	– IBM Cloud Pak for Integration – provides a comprehensive set of integration tools within a single, unified experience to connect applications and data across any cloud or on-premises environment. Includes IBM MQ (messaging), App Connect (application integration and integration PaaS), API Connect (API lifecycle management), Event Streams (enterprise Kafka) and Aspera (high-speed file transfer). – IBM Cloud Pak for Automation – an automation platform to analyze workflows, design AI-infused apps with low-code tooling, assign tasks to bots and track performance. – Red Hat OpenShift. (See Red Hat table entry for more information.)
Informatica	HIP	– Edge Data Streaming – collects streaming data and delivers it to both real-time and batch processing technologies.
Instaclustr	HIP	– Instaclustr for Apache Kafka – a hosted and fully managed Apache Kafka service. – Instaclustr for Kafka Connect – used to perform streaming integration between Kafka and other systems such as databases, cloud services, search indexes, file systems and key-value stores.
Mendix	LCAD	– Offers a Kafka module; supports Kafka Streams.

Vendor	Technology pedigrees	Offerings
Microsoft	PaaS DAP HIP	<ul style="list-style-type: none"> – Logic Apps – creates workflows and orchestrates business processes to connect hundreds of services in clouds and on-premises. – Service Bus – connects on-premises and cloud-based applications and services to implement secure messaging workflows. – API Management – publishes APIs securely for internal and external developers to use when connecting to back-end systems hosted anywhere. – Event Grid – connects supported Azure and third-party services using a fully managed event-routing service with a publish-subscribe model that simplifies event-based app development. – Azure Functions – simplifies complex orchestration problems with an event-driven serverless compute platform. – Azure Data Factory – visually integrates data sources to construct ETL (extract, transform, load) and ELT (extract, load, transform) processes and accelerate data transformation, using 90-plus pre-built connectors to manage data pipelines and support enterprise workflows.
MuleSoft (Salesforce)	HIP	<ul style="list-style-type: none"> – Mule – the runtime engine of Anypoint Platform, a lightweight Java-based enterprise service bus (ESB) and integration platform that enables application connectivity and data exchange. – Apache Kafka Connector, Mule 3 – ingests streaming data from Kafka and can publish to Kafka. – MuleSoft also partners with Confluent and Solace. (See Confluent and Solace table entries for more information.)
OutSystems	LCAD	<ul style="list-style-type: none"> – Kafka Apache Connector – a component built around Kafka with synchronous functions in OutSystems.
Pegasystems	PaaS LCAD DAP	<ul style="list-style-type: none"> – Pega Event Stream service – service to trigger events and respond to them immediately. For use with Pega Customer Decision Hub.
Qlik	HIP	<ul style="list-style-type: none"> – Data Streaming (CDC) – moves data in real time from source to target, managed through a graphical interface that automates end-to-end replication. It sets up, controls and monitors data pipelines based on the change data capture (CDC) technology.
Red Hat	PaaS HIP	<ul style="list-style-type: none"> – OpenShift – a cloud-based Kubernetes platform that helps developers build applications. It offers automated installation, upgrades and life cycle management throughout the container stack (the operating system, Kubernetes and cluster services, and applications) on any cloud. – Red Hat Integration – a family of products that include Fuse, an ESB; 3Scale API Management; AMQ, a suite of components-based Apache ActiveMQ and Kafka for real-time integration; Runtimes, for developing cloud-native applications; and Change Data Capture. – EventFlow – a lightweight cloud-native distributed microservices framework.

Vendor	Technology pedigrees	Offerings
Salesforce	PaaS	– Lightning – Salesforce’s next-generation PaaS.
	DAP	– Platform Events – connects business processes in Salesforce and external apps through the exchange of real-time event data.
	HIP	– MuleSoft Anypoint Platform – adds event-driven capabilities to the Salesforce platform and uses the Solace platform. – Enterprise Messaging Platform (EMP) – includes event messaging, streaming API and CDC features.
SAP	PaaS	– SAP Cloud Platform Enterprise Messaging – a central event bus.
	DAP	– AMQP Adapter for SAP Cloud Platform integration.
	HIP	– SAP API Business Hub. – SAP Event Mesh, a fully managed cloud service. – SAP also uses and resells Solace. (See Solace table entry for more information.)
SnapLogic	HIP	– Kafka Snap Pack – creates data pipelines for Kafka stream processing without coding.
Snowflake	PaaS	– Snowflake Connector for Kafka (“Kafka connector”) – reads data from one or more Apache Kafka topics and loads the data into a Snowflake table.
Solace	HIP	– PubSub+ Event Broker – a unified event broker technology available as software, purpose-built hardware and as a managed service.
		– PubSub+ Event Portal – provides visibility into Kafka streaming environments and manages Kafka event streams.
Sping.io	PaaS	– Spring Cloud Stream – formed from Spring Integration and Spring Boot as a framework for building scalable event-driven microservices connected with shared messaging systems.
		– Spring Cloud Kafka Streams – adds support for Kafka-specific features such as KStream, KTable and GlobalKTable.
SUSE	PaaS	– Rancher – a software stack for teams adopting containers. It addresses the operational and security challenges of managing multiple Kubernetes clusters across any infrastructure while providing DevOps teams with integrated tools for running containerized workloads.
Swim	HIP	– Swim Continuum – monitors applications, inspects data streams, configures meaningful traits and connects with external systems through a single interface.
Talend	HIP	– Talend Data Streams – a cloud-based self-service web UI designed for streaming data integration.
TIBCO	HIP	– TIBCO StreamBase – a complex event platform for rapidly building and deploying applications that analyze and act on real-time streaming data.
		– TIBCO Streaming – applies learning algorithms to streaming data for real-time insights.
		– TIBCO Cloud Events – detects key events and automates the next best actions to improve business outcomes so organizations can act in real time. – TIBCO Messaging – enables secure, reliable, real-time data distribution.

Vendor	Technology pedigrees	Offerings
TriggerMesh	HIP	<ul style="list-style-type: none"> – TriggerMesh – a cloud-native integration platform for deploying serverless integration that can integrate cloud services, SaaS and legacy applications. – Shaker – an alternative offering to AWS EventBridge; captures, transforms and delivers events from many out-of-the-box and custom event sources in a unified manner.
VMware	PaaS HIP	<ul style="list-style-type: none"> – Tanzu – a DevOps platforms to build, deliver and manage cloud-native apps on any cloud and at the edge. – VMware Event Broker Appliance – enables event-driven automation based on vCenter Server Events. It extends vSphere by triggering custom or prebuilt actions to deliver integrations within datacenters and across public clouds.

Source: 451 Research.

We offer this partial list of vendors to illustrate the range of technologies and platforms that can enable event-driven enterprises. Indeed, the capabilities of these vendors vary greatly, from simple Kafka modules (e.g., Mendix) to more elaborate EDA development and integration environments (e.g., Amazon Web Services Inc., IBM). Many of these vendors are actively discovering what an event-driven enterprise is and are gradually building out capabilities to support such enterprises. Potential users of these platforms should be aware that no single vendor is likely to have all the capabilities needed to enable an event-driven enterprise. However, the vendors presented here do offer the core elements necessary to build cloud-native, event-driven, intelligent process automation applications, so some assembly will be required, which portends opportunities for global systems integrators with deep understanding of both technologies and desired business outcomes.

We believe that as cloud-native application development accelerates within enterprises, more pioneering vendors will enter the market with increasingly comprehensive offerings fueled by organic and inorganic development activities. This will trigger accelerated M&A activity and further development across the IT vendor landscape, as well as accelerate deployments among a broader range of transformative event-driven enterprises.

Conclusions

Event-driven enterprises develop and modernize their applications to create a constant state of execution and performance awareness and are able to quickly sense and respond — reactively, predictively and even autonomously — to sudden change, market dynamics, risks and customer opportunities.

Event-driven enterprises invest in data quality management methods to improve the quality and speed of their decision-making. They modernize their applications to capitalize on the portability, agility and resiliency benefits of cloud-native computing. They deploy event-driven architecture and event-stream processing to quickly sense and respond to events with capabilities to rapidly communicate between and alert interested and affected services and systems. They also support application design and development principles that enable intelligent process automation using next-generation PaaS, low-code application development systems, digital automation platforms and event-driven hybrid integration platforms.

IT organizations within enterprises largely pursue the benefits derived from becoming event-driven because of IT systems' current technical complexity. Most business leaders are unaware of the benefits of being event-driven for the same reason. We believe this situation needs to change, and that it will. Business leaders need to be aware of the use cases noted in Figure 8 of this report, and, more broadly, they need to understand the impact and implications that cloud-native computing, event-driven architecture, event-stream processing and intelligent process automation can have on the performance and execution of individual lines of business, as well as on the enterprise as a whole. Event-driven enterprises, and their constituent lines of business, can respond quickly to change, risk and opportunity, and do so intelligently based on data while learning from execution. This makes event-driven enterprises highly capable, intelligent and competitive digital businesses.

Further reading

[Red Hat intends to automate IT operations with Event-Driven Ansible](#), November 2022

[vFunction looks to reveal tech debt and risk with data-driven approach to application modernization](#), November 2022

[Pega aims to advance intelligent automation platform with Everflow technology](#), August 2022

[ServiceNow provides a single platform for automation, integration with Automation Engine](#), August 2022

[Appian's latest release positions the vendor as a Total Automation platform](#), May 2022

[Recent and projected M&A activity in the converging total automation market](#), March 2022

[Total Integration: The future of hybrid IT interoperability – 2022 research agenda](#), January 2022

[Total Automation: The future of intelligent enterprise execution – 2022 research agenda](#), January 2022

[2022 Trends in Cloud Native](#), December 2021

Methodology

This report is based on information received from dozens of briefings over the last six months with vendors, end-user conversations, industry conferences, and quantitative analysis from multiple 451 Research services including Voice of the Connected User Landscape (VoCUL), Voice of the Enterprise (VotE), Market Monitor, the M&A KnowledgeBase and the Datacenter KnowledgeBase.

S&P Global Market Intelligence provides essential insight into key trends driving digital transformation across the entire technology, media and telecommunications (TMT) landscape. By offering a combination of expert analyst insight and differentiated data, our TMT Research group enables the industry with the information and perspectives they require to make more effective decisions.

Reports such as this offer a holistic perspective on key trends and themes driving the technology, media and telecommunications space over the coming year. These markets evolve quickly, so S&P Global Market Intelligence offers a wide range of research services that provide critical marketplace updates on an ongoing basis. These reports, datasets and perspectives are published frequently, in numerous short- and long-form factors, available on S&P Capital IQ Pro and 451Research.com. Forward-looking M&A analysis and perspectives on strategic acquisitions and the liquidity environment for TMT companies are also updated regularly via these platforms, backed by industry-leading databases such as the 451 Research M&A KnowledgeBase.

Our research is organized into channels that align with the prevailing key issues driving digital transformation across TMT. These channels are: Applied Infrastructure & DevOps; Cloud & Managed Services Transformation; Cloud Native; Customer Experience & Commerce; Data, AI & Analytics; Datacenter Services & Infrastructure; ESG; Fintech; Global Media; Global Mobile; Global Multichannel and Broadband; Information Security; Internet of Things; and Workforce Productivity & Collaboration.

For more information about S&P Global Market Intelligence, please go to: [Vital TMT Insight | S&P Global Market Intelligence \(spglobal.com\)](https://www.spglobal.com/insight).

Index of companies

Amazon Web Services **14, 18, 22**

Appian **17, 18, 23**

Google **17, 19**

IBM **14, 19, 22**

Informatica **9, 18, 19**

Mendix **17, 19, 22**

OutSystems **17, 20**

Pegasystems **17, 20, 23**

Red Hat **14, 17, 19, 20, 23**

Talend **9, 18, 21**

About the author



Carl Lehmann

Senior Research Analyst

Carl Lehmann is a Senior Research Analyst in the Applied Infrastructure & DevOps and Cloud Native research channels at 451 Research, a part of S&P Global Market Intelligence. He leads 451 Research's coverage of process automation and integration in hybrid IT and cloud-native architectures, as well as how hybrid IT affects business strategy and operations. The markets covered in his research include digital automation platforms (including workflow and business process management suites), robotic process automation technology, process discovery and mining technology, and hybrid integration platforms (including integration PaaS and API management).

Prior to joining 451 Research, Carl was a Principal Analyst at BPMMethods, where he advised clients on business strategy and process management. His book, "Strategy and Business Process Management: Techniques for Improving Execution, Adaptability, and Consistency," was published by Taylor and Francis Group in 2012. Carl was also a Senior VP of Strategy and Product Management for a B2B integration firm (now OpenText) and served 10 years as VP of Research for IT advisory firms Gartner and META Group, advising Fortune 500 clients. Carl's career began as a Project Manager for AT&T and a Product Manager for Digital Equipment Corporation (now Hewlett Packard Enterprise). He is a graduate of Boston University School of Management.

About S&P Global Market Intelligence

S&P Global Market Intelligence's Technology, Media and Telecommunications (TMT) Research provides essential insight into the pace and extent of digital transformation across the global TMT landscape. Through the 451 Research and Kagan products, TMT Research offers differentiated insight and data on adoption, innovation and disruption across the telecom, media and technology markets, backed by a global team of industry experts, and delivered via a range of syndicated research, advisory and go-to-market services, and live events.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2023 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.