Accenture Labs

FIND YOUR EDGE

Bringing analytics to data at the edge of IoT

With Internet of Things (IoT) deployments becoming more complex and businesses more data-driven, there's a big challenge: how to combine scalability with manageability? It's a challenge that's exacerbated when dealing with the large volumes of data generated by things, and with network limitations of bandwidth, latency and connectivity. Bringing the data to the analytics is no longer sufficient. It's time to bring the analytics to the data.

Pushing analytics to the edge of IoT brings its own challenge: equipping the edge to act as well as possible in the absence of network connectivity or when compute resources are scarce. The way forward? A flexible, knowledge-driven framework that captures the appropriate business logic needed to push analytics to the edge of IoT—where data is high-fidelity and decisions are more reliable and more real time.

Today's IoT solutions rely on centralized, platform-based solutions to collect, store, and analyze sensor data from devices at the network edge. Platform solutions built on a cloud-centric paradigm require reliable, low-latency, high-bandwidth network connectivity—a requirement that's clearly unsuitable for implementations where industrial assets are located in remote geographic regions with limited, unreliable and expensive network connectivity, or when data volumes are immense.

This scenario makes a compelling case for pushing more computation out to the edge of IoT. Advanced digital businesses are looking beyond simple storage and rules-based processing of sensor data at the edge. Their goal? Applying advanced machine learning and AI-based analytics to take advantage of the high-fidelity data available from devices at the edge, and then taking immediate action without the round-trip time to the cloud—bringing the analytics to the data.

COLLECT

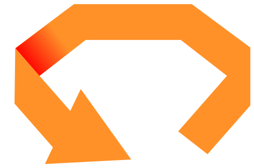

STORE

ANALYZE

The edge introduces the challenge of extending the role of the platform. For *edge analytics*, a central cloud-based platform is still critical to manage and develop the analytics applications and models based on populations of devices, which when deployed at the edge are customized for the specific instance and scenario.

To help enable this coordination of cloud and edge, the organization needs domain expertise around analytics applications, devices, and field engineering. They need to combine traditional analytics with AI that can understand and adapt to the dynamic conditions in the field. Collectively, operationalizing how these domains work together in a centrally managed manner through to fluid deployment, operation, and monitoring from the cloud to the edge is critical for any practical IoT solution.

Operationalize Edge Analytics

1 **DEVICE EXPERT**

2 **APPLICATION EXPERT**

3 **FIELD ENGINEER**

CLOUD IS **MOVING TO THE FOG**

We distinguish edge analytics from edge computing in the need to run a variety of machine learning, predictive and prescriptive models that require a more powerful compute environment with specific CPU, memory, and storage requirements. Today, not every edge device is suitable for providing such an environment to train, run, and retrain these models.

For example, traditional mobile phones can process rule-based logic, but are not suitable for retraining deep learning models. Or historians that specialize in data collection and its processing at the edge, but are not suitable for deploying generalized analytics and machine learning.

Instead, environments suitable for edge analytics occur at the fog layer and work to connect the dumb or legacy devices at the furthest ends of the edge. Figure 1 shows the various layers of an IoT deployment topology where the fog layer offers suitable environments for edge analytics.

In the fog layer, edge gateways typically serve as the primary connectivity points in the field for various IoT devices and industrial equipment, including legacy and passive sensors. Their small footprint means they can be deployed more closely to the IoT devices in the physical environment. They offer a basic level of compute, storage, and network capabilities suitable for analytics.

Additionally, fog can be a small server configuration typically comprised of one or more appliances as an extension from the cloud. The compute, storage and network capabilities available here can be up to an order of magnitude higher than that of a gateway. This configuration bridges the gap between the limited compute resources provided on edge gateways, and the seemingly unlimited and highly scalable compute resources available in the cloud. It is suitable for analytics that require collecting, storing and processing data from a larger number of IoT devices than the gateway.

Figure 1: Layers of an IoT solution

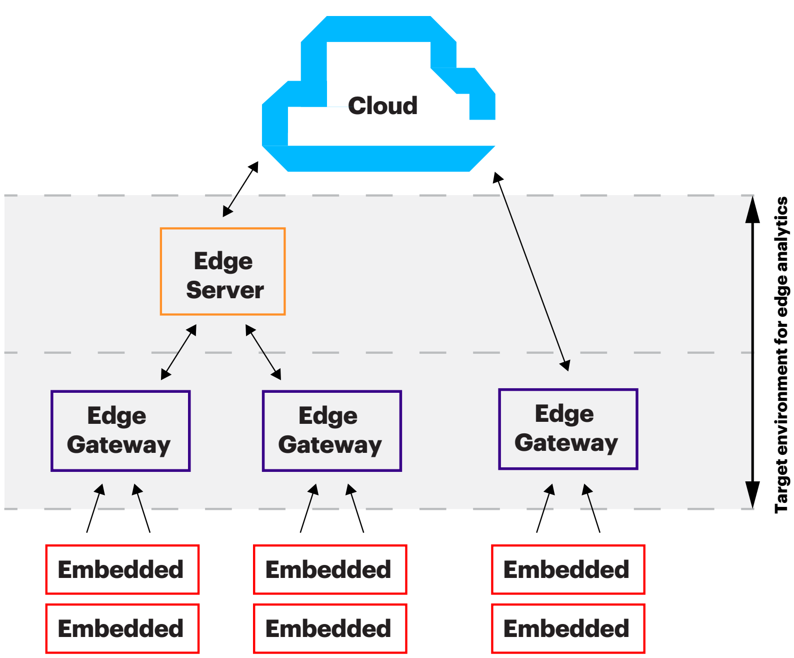

Cloud Layer
Centralized enterprise data centers with scalable and reliable IT resources and processes. Cloud is ideal for governing, storing, and processing data requiring global awareness across a population of devices and scenarios.

Edge Server Layer (optional)
A server or collection of servers extend cloud-like capabilities to a single industrial site (eg ship, factory, oil site) for servicing a collection of edge gateways. Characterized by reliable network downstream, and limitednetwork upstream. This layer extends the capability of the traditional hardware data historian.

Edge Gateway Layer
Gateways provide compute, storage and network connectivity for industrial assets and sensors that are physically or wirelessly connected. Often gateways have limited compute and network resources, but have access to highest resolution of data. This layer extends the remote terminalunits (RTUs) or programmable logic controllers (PLCs).

Embedded Layer
Asset-level hardware with domain-specific sensors and purpose-built embedded devices and controllers that measure and control industrial equipment (eg trigger emergency shutdowns).

A key functionality of an edge analytics platform is the way it supports the management of the models and applications lifecycle across the cloud and fog. A cloud platform sets the policy for model creation, training, governance, and deployment that must extend to the edge. When this functionality is enabled seamlessly, with coordination between edge and cloud portions of the platform, it provides an autonomous operation of model lifecycle management against the dynamic factors within the edge environment. Even when not connected, centrally governed models and operations are carried forward.

An edge analytics platform needs to be able to manage the deployment of the models associated with the edge device in the fog today, and then as edge device capability improves to push it out to the device. The platform manages all layers of the cloud, fog, and edge as a fabric to optimize analytics deployment based on data and response requirements.

Specifically, an edge analytics platform will enable a business to 1) centrally develop, train and manage analytics models in the cloud by leveraging data from a global population of edge devices, 2) deploy these models to execute in the fog to take advantage of unfiltered, high-fidelity data and low latency response times, and 3) seamlessly coordinate with the cloud-based platform for the models to adapt to the specific dynamics of the local environment, and to buffer data based on specific application needs according to centrally defined objectives. All these combine to enable the enterprise to provide insights where they are relevant and to drive immediate outcomes.

1 **Centrally develop, train and manage analytics models in the cloud by leveraging data from a global population of edge devices.**

2 **Deploy these models to execute in the fog to take advantage of unfiltered, high-fidelity data and low latency response times.**

3 **Seamlessly coordinate with the cloud-based platform for the models to adapt to the specific dynamics of the local environment, and to buffer data based on specific application needs according to centrally defined objectives.**

GET FLEXIBLE: **AVOIDING THE PITFALLS OF VERTICAL INTEGRATION**

Consider the following scenario. A device expert onboards a new type of sensor. Elsewhere in the business, an application expert develops an analytics application that consumes data generated by that sensor. The application uses models developed by a data scientist. When these steps are complete, a field engineer deploys the new instances of physical sensors across multiple sites.

An ad hoc solution could support this scenario. But it would require explicit coordination among the three experts to help enable a comprehensive solution: knowledge of the sensor capabilities, of the application, of the models, and then working with the field engineer to correctly deploy and configure the edge instances in situ.

Such a tight interdependency runs counter to the principle of separation of concerns and limits the operational scalability and maintainability of the solution in large enterprise settings. There's no reusability and only limited flexibility in streamlining operational processes. It lacks organization governance and business oversight.

Vertically integrated solutions address this problem, up to a point. But they do so by controlling every aspect of the IoT stack—from computing hardware on the edge to the centralized orchestration of components in the cloud. This type of solution locks an organization into a vendor ecosystem, and limits its ability to use existing technology or to deploy a best-of-breed solution.

We see common challenges across industrial IoT businesses including those from manufacturing, transportation, and oil and gas. These businesses have been digitizing and networking their industrial assets for decades. And they have a mature understanding of industrial IoT technologies and tools. Many are bound by regulatory requirements that hamper their adoption of innovation. Over the years, many have accumulated technological "debt", ranging from proprietary solutions, heterogeneous approaches and dated hardware, involving several business units gathering data and running analytics on the edge.

The brownfield nature of these operational environments means any edge analytics framework must be flexible: it will have to support existing operations and provide a seamless path toward ecosystem modernization and the adoption of new technology. Diverse requirements across business units mean the company must support a variety of edge computing hardware, operating systems, data-processing, storage and analytics runtimes and languages.

Any solution must offer a way to balance the variability of resource limits at the edge in a way that best serves the overall needs of the business.

The situation is further complicated by the operations that generate high volumes of data at the terabyte scale. Or with operations in geographically remote locations where broadband network connectivity is either unreliable, or needs cost-prohibitive cellular or satellite coverage.

In today's case, high-value data is usually filtered and down-sampled to lower rates for transmission. Or high-fidelity data faces high transmission costs to cloud for storage and processing—often transmission occurs literally over "the air" via manual extraction flown in by helicopter.

Cloud-based analytics is uniquely suitable for analytics models looking at data across populations. And the edge offers unique access to high volumes of data to customize context-specific models and take immediate action in the field.

High-value data is usually filtered and down-sampled to lower rates for transmission.

MINIMIZING COMPLEXITY: **AN EDGE ANALYTICS FRAMEWORK**

Accenture Labs has created an edge analytics framework that addresses these challenges without re-inventing the wheel. A key goal of this framework is not to require an enterprise to make wholesale changes to its IT and OT environments, or to migrate to a vertically integrated solution.

We want to work with an organization's existing heterogeneous environment with a varied applications, models, and hardware infrastructure. At the same time, we pave a path forward for the organization to take advantage of technology refresh step by step.

Our approach focuses on managing the challenges presented by this heterogeneous ecosystem.

Table 1: Illustrative Use Cases

	Robotics automation in retail warehouses	Oil and gas operations	Smart Transit
Cloud-based enterprise data center	• Order routing to specific warehouses • Closely correlated fulfillment items • Predictive maintenence models • Update optimal routing models	• Global tracking of asset utilization and optimizing resource allocation • Fleet management to track vehicle telematics and combine with weather and traffic conditions • Update anomaly detection and predictive maintenence models from global data	• Expected traffic and load calculations • Positions of all vehicles • Time targets passed down to vehicles • Update video analytics/ routing models
On-premise location-wide fog compute servers	• Real-time location and status of all robots • Place item locations in robot queues for fulfillment	• Site-wide situational awareness by monitoring status and condition of surface and downhole equipment • Utilize site-specific historic data to identify optimal drilling parameters to enhance wellbore quality	• Processes/Models on how to hit time targets • Incoming real-time traffic and load data for local surroundings
Asset-specific edge gateways	• Intelligent routing to specified location • Obstacle avoidance • Run predictive maintenance models and inform gateway of potential issues	• Optimize fluid pressure and chemical mixture composition to enhance oil extraction • Real-time sediment analysis to determine borehole composition • Predicitve and anomaly detection models to reduce eqipment failures, and extend drillbit life	• Video analystic to detect density of crowd at transit stations • Visual indicators to vehicle operators to help them slow down or speed up in order to meet current goal

Hierarchical architecture: distributing analytics across layers

Hierarchical layers across embedded, fog and cloud results in a diverse set of standards and requirements. Each of these layers offers increasing amount of compute, storage, and network capacity suitable for implementing analytics with varying complexity and latency requirements. The edge framework must provide the flexibility for the business to deploy analytics that span across these layers in a loosely coupled way allowing for interchangeability of hardware, and model and application software components as needed.

Our solution allows for decoupling where each component can be replaced. A microservices architecture manages each software component independently with its own defined goals and deployment requirements. Asynchronous messaging between various layers, including connectivity to the sensors, enables these components to communicate with each other using open libraries. Apart from taking advantage of industry-wide standards and practices, using open libraries means the architecture can be extended to support any business- or industry-specific custom protocols.

This hierarchical design supports analytics that reside in a set of fog servers handling all the data across a population of devices in a site, or where gateways at the bottom of the hierarchy handle the analytics for a smaller subset of sensors paired with it (see Table 1 for illustrative use cases).

To ensure interoperability and encourage reuse, the framework must offer a standardized approach to manage the devices and the applications that rely on them.

An abstraction layer: providing a common framework for multiple devices and owners

Enterprises deal with a wide variety of edge computing devices and sensors, many of which are owned and managed by different units within the business. Each type of device or sensor brings a unique set of hardware capabilities, protocols, data formats and interfaces that adds to the overall complexity. To ensure interoperability and encourage reuse, the framework must offer a standardized approach to manage the devices and the applications that rely on them.

Figure 2: Components at the edge

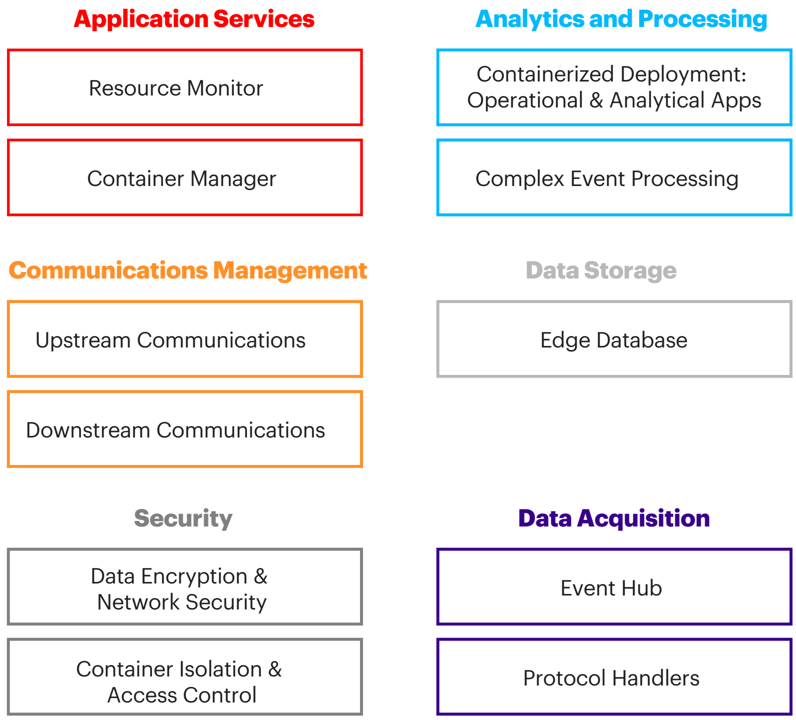

Application Services
Provides the APIs to interface with the platform. Handles control messages and abtraction of incoming data streams for apllications.

Communications Management
Sends processed data northbound to the edge server and cloud tiers. Handles routing of southbound data and control messages.

Analytics and Processing
Ochestrates the provisioning, delployment, and monitoring of containerized edge applicaitons, and user-defined analytics workflows.

Data Storage
Stores raw sensor readings and processed data from edge applications. Stres configuration information for core platform components.

Data Acquisition
Interfaces with embedded layer to acquire and ingest sensor data.

Security
Ecrypts data-at-rest and data-in-motion. Utilizes certificates and container isolation to enforce access control policies.

Our framework provides abstraction over the underlying complexity on the edge devices by utilizing containerization technology (like Docker), along with an asynchronous event-hub (see Figure 2).

Containerization provides a standardized deployment environment for developers to build and package their edge application and analytics models. It helps enable portability to deploy edge applications and models across various edge computing hardware, irrespective of device-specific capabilities, settings and configurations. Similarly, an asynchronous event-hub aided by a protocol translator library abstracts the variation in sensor interfaces, protocols and data formats. It acts as the single interface for all edge applications to communicate with sensors, other edge applications, or other components in the cloud.

This abstraction in the fog is complemented by a cloud-based knowledge graph (see Figure 3) and an intelligent orchestration server. The knowledge graph captures semantic metadata about edge device's hardware capabilities, sensor data format and protocols. The graph structure enables flexibility to relate these capabilities to application and model needs.

When a new device or sensor is onboarded for the first time, these details are added to the knowledge graph via an API. The orchestration layer automates operations by querying the knowledge graph to determine specific device configuration, and orchestrating the deployment and monitoring of containerized components on each device.

Model lifecycle automation meets cross-organizational requirements

In addition to devices, there is a myriad of cross-organizational application and model requirements. These include heterogeneous software and tools with different operating systems, analytics runtimes, platforms, and languages needed to take advantage of modern, higher-performing tools, domain-specialized capability, and reusing legacy code. The edge framework must provide a common enterprise-wide process to govern and ensure that the right mix of models and applications is used, and then to simplify packaging, deployment, monitoring and governance at the edge.

Figure 3: Knowledge graph example

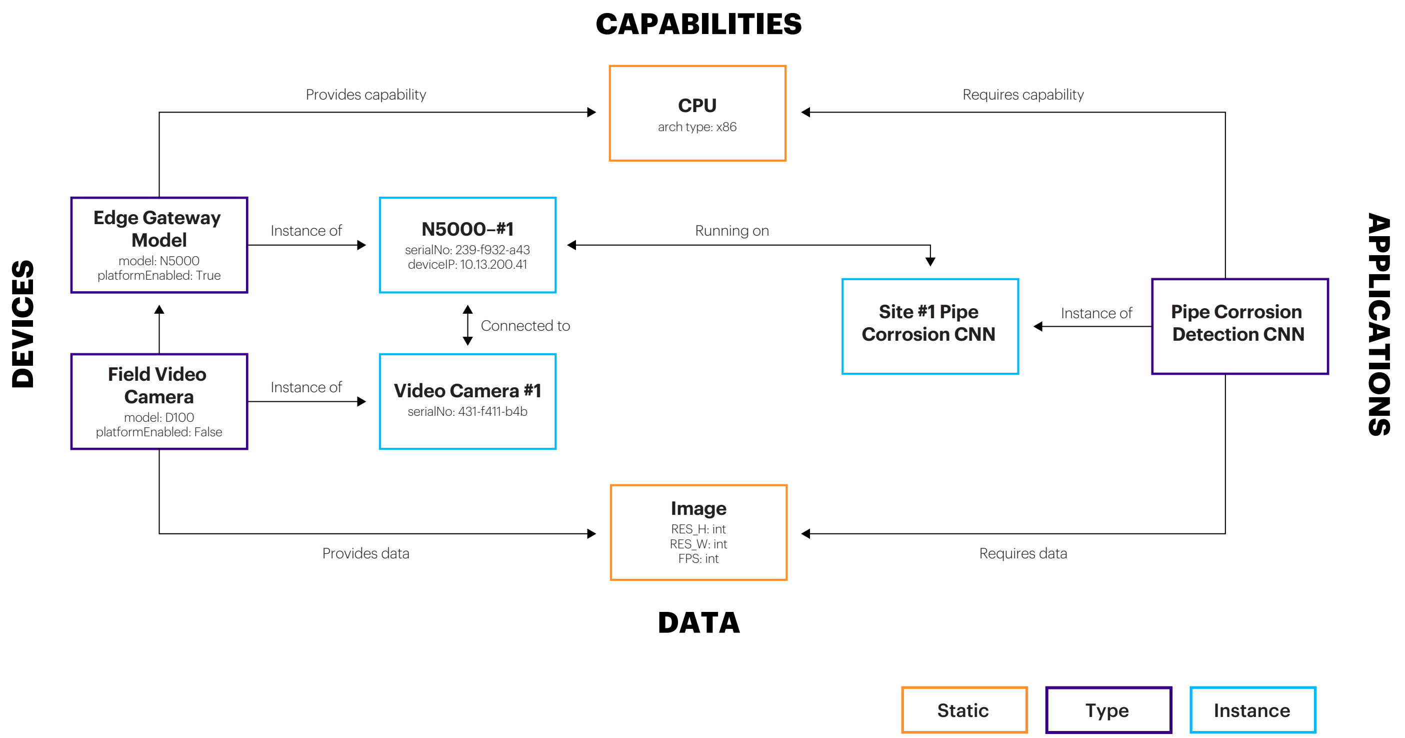

A cloud-based model management framework provides lifecycle management and governance capability to ensure that the right models and applications are used. Since edge deals with limited resources and, sometimes, no network connectivity, it's critical to pass on the logic around priorities so that the edge environment can run as intended.

To do so, our model management capability integrates with the centralized knowledge graph to relate application and models with hardware and data requirements. For example, for training or retraining a deep learning model to function effectively, access to specialized hardware such as a GPU may be needed. The developer simply declares this as a model requirement as part of the onboarding process to capture this requirement in the knowledge graph.

For deployment on an edge device, a cloud-based orchestration service, and an agent on the edge device coordinate to query the knowledge graph and validate model requirements against device capability and configuration. Together they orchestrate the deployment, execution and monitoring of containerized models on the edge devices, thereby enabling various business units to operate independently (see Figure 4).

Figure 4: End-to-end analytics management

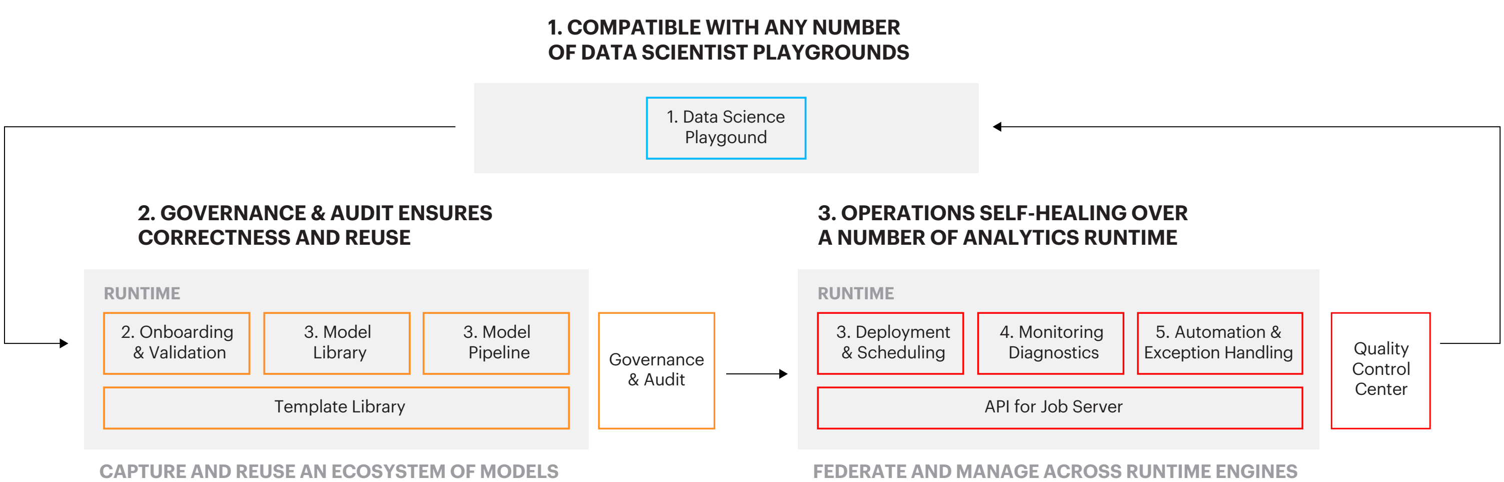

Monitoring and metering: governing usage of limited resources on the edge

Edge environments in industrial settings, such as an offshore oil exploration platform, are often characterized by limited availability of compute, storage and network resources. What's more, today's data-rich Industry X.0 machinery generates terabyte-per-day scale of sensor data that should be analyzed but might not make sense to move to the cloud wholesale. Sharing these resources effectively across edge applications from various business units requires support for provisioning, monitoring and metering of resource utilization on the edge devices. Additionally, the deployed edge applications and models must continue operating through intermittent, low bandwidth, or non-existent internet connectivity.

Again, our knowledge graph supports the discovery and management of edge devices that can support a particular analytics model and application. Once a target device is selected, the orchestration server uses this API to first validate that the device capabilities match the requirements specified by the analytics model. It then coordinates with the orchestration agent on the device to provision resources and instantiate the model on the edge device at the specified priority.

Once the model is deployed, a monitoring agent tracks resource utilization for each deployed application or model. Similarly, an upstream agent prioritizes any data generated by model and controls access to the network resources. This process helps enable the business to specify, allocate and meter edge resources based on the priority in a seamless and frictionless fashion.

ENABLING THE DIGITAL TWIN: **BRIDGING THE EDGE AND THE CLOUD**

Many businesses have digitized their industrial assets and established centralized IoT platforms to collect, process, and analyze data from these assets. They have data-science teams to create a portfolio of analytics applications.

But limiting this to the cloud is insufficient, given the nature of applying machine-learning and deep-learning models to large volumes of data typical of industrial operations. And there is much value to be realized from these investments by extending capability to hi-fidelity data and low-latency compute at the edge.

These technology trends show that data immobility demands implementation of analytics capabilities closer to the data-sources. Another emerging trend in analytics is an increasing utilization of unconventional computing hardware such as GPUs, FPGAs and quantum computing. Both data immobility and a dependence on hardware are exacerbated in IoT environments where there are network limitations and resource constraints. Inevitably, there's a need to extend analytics platform capabilities from the cloud to the IoT "fog" layer (see Figure 5).

Figure 5: A roadmap for IoT

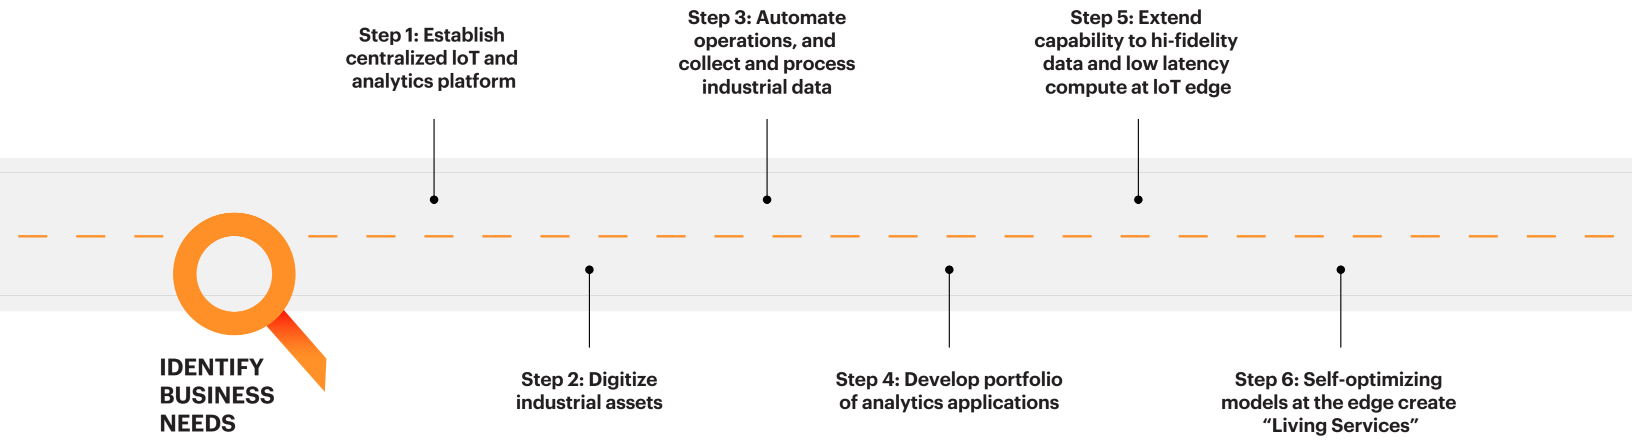

We describe an edge analytics framework designed to do just that. It addresses the key challenges that we know businesses face with their IoT implementations today. Its knowledge-driven approach provides the flexibility they need to easily develop, reuse, and seamlessly deploy applications to edge devices through a manageable and scalable framework.

It addresses the key challenges that we know businesses face with their IoT implementations today.

Its knowledge graph provides an abstraction layer that helps enable the various IoT stakeholders in an enterprise to collaborate—without losing focus on their areas of expertise. It seamlessly tracks and manages the multitude of hardware and software parameters present in existing heterogeneous deployments. And, coupled with interoperability through containerization techniques on the edge, it simplifies the development and deployment of edge analytics applications on the technology of today, while future-proofing the solutions to come.

At a time when new solutions are coming online all the time, our framework also helps enterprises to avoid the pitfalls of vertical integration. Freed from vendor lock-in, they're on the path to flexible computing and analytics at the edge of IoT.

By continuing to deploy and train these models on the edge, it allows for customization of assets at a finer level of granularity. The cloud serves to combine insights from populations of data across devices, and the edge specializes in instances of individual assets, groups, or locales. This paired approach, bridging the governance of the cloud to the edge, leads to a virtuous cycle for creating the self-optimizing models that are essential to digital twins.

Our comprehensive approach—extending an enterprise-wide model management framework across organizations, across capabilities, and across cloud and edge—helps enable the complex models that digital twins require. A twin simultaneously resides on multiple instances of cloud nodes and edge devices and requires coordination to continue learning and improving over time—creating a feedback loop that results in a continuously evolving living service.

Now we can handle edge analytics at scale and tap into high-fidelity data and contextual processing directly at the edge. Rather than bringing the data to analytics in the cloud, it's time to bring analytics to the data at the edge.

CONTACTS

Teresa Tung, Ph.D.
Managing Director, Accenture Labs,
Systems and Platforms R&D
teresa.tung@accenture.com

Jean-Luc Chatelain
Managing Director, CTO, Accenture Digital,
Applied Intelligence
jean-luc.chatelain@accenture.com

Jurgen Weichenberger
Managing Director, Accenture Digital,
Applied Intelligence
jurgen.weichenberger@accenture.com

William Gatehouse
Managing Director, Accenture Digital,
Applied Intelligence
william.r.gatehouse@accenture.com

John Roe
Associate Director, Accenture Digital,
Industry X.0
john.x.roe@accenture.com

Ramoj Paruchuri
Director, Accenture Digital,
Innovation Studios
ramoj.paruchuri@accenture.com

CONTRIBUTORS

Michael Giba

Anuraag Chintalapally

Srinivas Yelisetty

Louis Farfan

ABOUT ACCENTURE

Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. Combining unmatched experience and specialized skills across more than 40 industries and all business functions – underpinned by the world's largest delivery network – Accenture works at the intersection of business and technology to help clients improve their performance and create sustainable value for their stakeholders. With more than 435,000 people serving clients in more than 120 countries, Accenture drives innovation to improve the way the world works and lives. Visit us at www.accenture.com.

ABOUT ACCENTURE LABS

Accenture Labs incubates and prototypes new concepts through applied R&D projects that are expected to have a significant strategic impact on clients' businesses. Our dedicated team of technologists and researchers work with leaders across the company to invest in, incubate and deliver breakthrough ideas and solutions that help our clients create new sources of business advantage.

Accenture Labs is located in seven key research hubs around the world: Silicon Valley, CA; Sophia Antipolis, France; Arlington, Virginia; Beijing, China; Bangalore, India; Herzliya, Israel and Dublin, Ireland. The Labs collaborates extensively with Accenture's network of nearly 400 innovation centers, studios and centers of excellence located in 92 cities and 35 countries globally to deliver cutting-edge research, insights and solutions to clients where they operate and live. For more information, please visit www.accenture.com/labs.