AN ENTERPRISE, END-TO-END, OPEN SOURCE ARCHITECTURE FOR IoT

A joint solution from Cloudera, Eurotech, and Red Hat

TECHNOLOGY OVERVIEW

BENEFITS

- Cut costs and reduce risk and complexity associated with managing millions of connected devices in large-scale disparate environments
- Optimize data stream transformation, aggregation, filtering, and routing
- Enable machine learning to create deep business insights and actionable intelligence
- Streamline cloud-native application development, delivery, and agile integration
- Secure the environment and data end-to-end
- Enable agile, cost-effective business innovation while reducing operational complexity
- Easily scale from proof-of-concept, to pilot, to full production

The Internet of Things (IoT) represents revenue opportunities, operational efficiencies, and the emergence of new products and services enabled by digital transformation. This transformation requires designing, planning, and executing on a coordinated and collaborative level across functions, operations, departments, and business units.

IoT creates new challenges for enterprises, related to both technology and business. From a technology perspective, the IoT market is moving quickly. There are many evolving players and standards. Early adopter organizations that adopted a proprietary IoT platform now find themselves tied to limited functionality, locked into a particular vendor, and rethinking their choices. Many organizations now seek open source alternatives, recognizing the value of the open source community as a hub for innovation and continuous development for IoT. And many organizations now realize that no single provider can completely address the end-to-end challenges IoT presents.

However, it is a complex challenge to manage multiple vendors' solutions, incorporate various open source projects, validate that they work together, integrate them to provide the right functionality, and ensure future enhancement compatibility. That is why Red Hat, Eurotech, and Cloudera have combined their strengths and integrated their technologies to deliver the first end-to-end, open source IoT architecture that addresses enterprise IoT needs. Red Hat, Eurotech, and Cloudera are making IoT easier for organizations by providing a validated, modular, flexible architecture built to be open, interoperable, and cost effective.

ADDRESSING IoT MARKET CHALLENGES

This end-to-end, open source architecture for IoT:

- Connects and manages millions of distributed IoT devices and gateways with added security
- Simplifies data flow management with intelligence and analytics at the edge.
- Provides a comprehensive, centralized advanced analytics and data management platform with the ability to build or refine machine learning models and push these to the edge.
- Enables modern, cloud-native application development, delivery, and agile integration.

Whether you are designing a complete IoT system or developing individual components of an intelligent solution, this end-to-end architecture can help you simplify development and integration tasks, save time, and reduce costs. The architecture provides the components and foundation needed for an end-to-end IoT solution, but with the benefits of open source innovation and interoperability. Its modular nature allows you to swap out system components over time so you can keep pace with advances in technology while protecting previous investments.

KEY COMPONENTS

The components of the IoT architecture deliver the capability to manage connected "things," control and manage the flow of data from device to the cloud, analyze data for insights and machine learning, and develop, deploy, and integrate applications.

FEATURES

- Enterprise-ready
- Modular, secure, end-to-end architecture
- Validated, integrated, and tested
- Deployment flexibility
- End-to-end analytics
- End-to-end security

The key components of the architecture are:

- **Connected "things"** that generate device data and require management, a secure connection, and seamless protocol translation.
- **Intelligent IoT edge stack** to support data ingestion and control and enable analytics at the edge.
- **IoT integration hub** to manage disparate devices and control the operational flow of data directly to enterprise applications for input, or to a data management platform for analysis.
- **Data management and analytics platform** for IoT data processing, persistent storage, analytics, and machine learning to enable deep business insights and actionable intelligence.
- **Application development, delivery, and integration environment** to create cloud-native applications with DevOps disciplines, and integrate to distributed IoT and traditional applications.

KEY FEATURES

Enterprise-ready, open, and interoperable, the architecture is validated, integrated, and tested by Red Hat, Eurotech, and Cloudera. The architecture has pre-integrated security and manageability across devices, access, authentication, and applications, as well as data that is in-motion and at-rest. Its modular nature enables choice, protects your existing technology investments, and provides the flexibility to build out an IoT environment on-premise or in public, private, and hybrid cloud environments.

End-to-end analytics are realized through the use of the integrated components of the architecture. Business rules and advanced analytical models can be deployed both at the edge and within the core platform, enabling decisions based on historic data and real-time device data.

KEY FUNCTIONALITY

As seen in figure 1, the architecture enables bi-directional communication with devices via intelligent-edge IoT gateways. Data is routed through the IoT/edge Integration hub for application integration within the application environment and for aggregation into the data management platform for deep analysis and machine learning. Data can be flexibly processed throughout the architecture based on use case requirements, including the capability to apply machine learning models and advanced analytics at the edge.

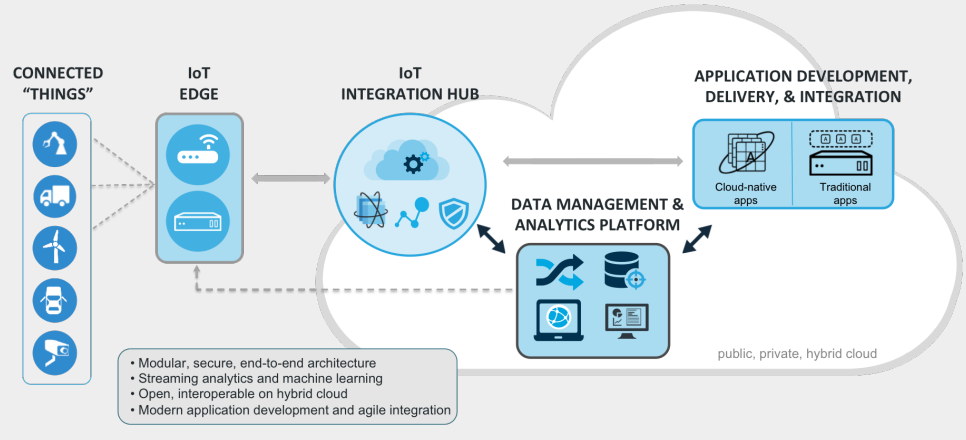

Figure 1. End-to-end, open source IoT architecture to capture, process, and analyze data from connected devices, as well as ship machine learning models and intelligence back to the edge

FEATURES

- Edge processing and analytics
- Data integration, routing, device management
- Advanced analytics and machine learning
- Application development, deployment, integration

This offering provides a production-ready foundational architecture upon which you can layer your own business logic, data, and applications. You can accelerate time to market and reduce development costs by focusing on creating business value and competitive differentiation instead of building and managing IoT infrastructure. The solution allows you to streamline application life-cycle management across the entire intelligent system and easily scale from proof-of-concept, to pilot, to full production.

The features and functions of each module in the architecture are outlined below.

IoT EDGE STACK

FEATURE	FUNCTION
Device connectivity	Connect devices to the cloud using MQ telemetry transport (MQTT), a lightweight publish-subscribe communications protocol designed to tolerate intermittent connections and to minimize bandwidth consumption.
Remote management	Manage devices, administrators, and settings from a browser-based console.
Data transformation	Convert legacy or proprietary data payloads to standards-based protocols (slow data).
Intelligent routing	Provides continued connectivity and processing resources. Route data to different back-end locations based on priority, nature, and network efficiency.
Business logic	Implement business rules and field applications, and execute control logic in near real time (fast data).
Real-time decisions	Automated decisions at the edge based on results from machine learning analytics
Machine learning execution	Machine learning predictive model markup language (PMML) model executed at the edge

IoT INTEGRATION HUB

FEATURE	FUNCTION
Integration services	Interface with back-end business applications and other cloud services and systems using open application programming interfaces (APIs).
Device registry and management	Perform remote operations on connected devices. Configure operating parameters. Execute operating system commands. Manage applications and services running on devices.
Access control	Control access to the cloud platform using user-based authentication or secure sockets layer (SSL).
Event management	Orchestrate events, alerts, and status checks.
Device provisioning	Automatic, secure device on-boarding procedure that remotely configures a single or a large number of newly deployed devices.
Container application platform	Scalable container-based environment to operate IoT Hub services and distributed applications.

ABOUT CLOUDERA

Cloudera delivers the modern platform for machine learning and advanced analytics built on the latest open source technologies. The world's leading organizations trust Cloudera to help solve their most challenging business problems with Cloudera Enterprise, the fastest, easiest and most secure data platform available for the modern world. Cloudera customers efficiently capture, store, process and analyze vast amounts of data, empowering them to use advanced analytics and machine learning to drive business decisions quickly, flexibly and at lower cost than has been possible before. Visit www.cloudera.com

ABOUT EUROTECH

Eurotech is a global company that designs, creates and delivers full Internet of Things solutions, including services, software, and hardware, to leading systems integrators and enterprises large and small. With Eurotech solutions, clients have access to the latest open source and standardized software stacks, flexible and rugged multiservice IoT gateways, and sophisticated sensors to bring actionable data from the field into business operations. For more details, visit www.eurotech.com

ABOUT RED HAT

Red Hat, the world's leading provider of open source software solutions, delivers reliable and high-performing cloud, Linux®, middleware, storage, and virtualization technologies that help you collect, communicate, transform, store, and act on critical data generated by the Internet of Things (IoT). Red Hat offers a single, extendable, and secure foundation to support the end-to-end life cycle of IoT solutions—from development to production. www.redhat.com/iot

DATA MANAGEMENT AND ANALYTICS PLATFORM

FEATURE	FUNCTION
Real-time data ingest	Ingest data from multiple data sources, in batch and real time.
Data variety management	Handle all types of data sources, multiple data formats, structures, and schemas.
Real-time analytics	Enable real-time data processing on streaming data using in-memory processing engines.
Machine learning capabilities	Out-of the-box machine learning libraries to easily build and iterate on predictive models.
Data science for the enterprise	Self-service data science environment.
Diverse advanced analytical tools	Analytics engines, including search and SQL analytics, with tools to suit diverse needs.

APPLICATION DEVELOPMENT, DELIVERY, AND INTEGRATION

FEATURE	FUNCTION
Application development and management	Build cloud-native applications utilizing a container-based application platform that supports developer self-service provisioning and DevOps principles.
Agile integration	Container-based middleware and API management enable distributed integration across IoT environment, cloud-native and traditional applications.
Polyglot, multilanguage support	Developers can use various languages, frameworks, and databases, all on the same platform with ease.
Automation	Streamline application deployment, infrastructure configuration, and cross-silo operations to automate end-to-end management.
Scalability	Applications can easily scale to thousands of instances across hundreds of nodes in a matter of seconds.
Container portability	Standardized Linux container model ensures that applications are fully supported and portable across private, public, and hybrid cloud environments.

IoT is transforming the way enterprises do business, providing significant opportunities for solution providers and technology suppliers but also creating considerable design, implementation, and management challenges for system developers and operators.

This end-to-end, open source IoT framework can help you efficiently deliver and manage highly secure, reliable, and scalable IoT solutions. Building on Red Hat's experience as the world's largest open source company, Eurotech's capabilities and experience in operational technology environments, and Cloudera's skills as a leading data analytics and data management company, this IoT framework lets you harness the power of community innovation, contain costs, and avoid lock-in.

To learn more about the architecture, contact iotquestions@redhat.com

Copyright © 2018 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.