Cisco Kinetic Data Control Module

Enforcing policy and getting the right data
to the right applications at the right time

Table of Contents

Executive Summary

The Internet of Things (IoT) has dramatically increased the volume and variety of data being produced, opening the door to a wave of new possibilities. Companies can use that data to grow revenue, increase efficiency, enhance customer experience, decrease downtime, reduce costs, and in countless other ways.

The key is being able to unlock data from it source, process it to make is usable, and programmatically and securely move the right data to the right application to put it to work.

Cisco Kinetic is the cornerstone of the Cisco IoT portfolio. With this platform, Cisco is not only fulfilling the need for IoT technology and products, but is also revolutionizing our understanding of networking. Instead of simply providing connectivity, the network will host computation, improving the value of data as it is transported. The network is now smarter. Instead of viewing “the cloud” as a destination, modern IoT networks will pass data through clouds as well as to clouds. Instead of simply enabling connectivity, the Kinetic platform will control the distribution of data, providing far reaching distribution while maintaining and restricting access to data. In numerous ways, Cisco is advancing the state of the art in both IoT and networking.

This white paper describes the Data Control Module (DCM) in the Cisco Kinetic platform, which provides a framework for the distribution of IoT data using cloud technology. It provides secure transport and control of data distribution, enabling companies to get maximum value from their IoT data.

Unlock data from its source, make it usable, and securely move the right data to the right application to put it to work

The key capabilities of DCM include:

- Policy-based data distribution to or through a cloud, or multiple clouds
- Retention of control by the data owner to limit distribution of, and access to, their data
- Enforcement of ownership, privacy, and security across the networks and through the clouds throughout the entire data communications path
- The ability to transform or subset the data for different targeted data consumers
- Multiple mechanisms for integrating with heterogeneous application systems, in the cloud and in a data center
- Seamless integration and connectivity with the Cisco Kinetic Edge & Fog Module (EFM) for on premise data processing and device connectivity
- Synergy with the Cisco Kinetic Gateway Management Module (GMM) for remote deployment and configuration of IoT gateways
- The ability to create and maintain synchronization between a physical device and a software representation of the device's state (device mirroring)
- Integration with cloud-based application development systems, such as the Microsoft Azure IoT Hub

In summary, this technology enables IoT systems to scale indefinitely across multiple networks and through multiple clouds, enabling the vision of unlimited IoT.

Introduction

IoT is a combination of data generating devices, communications, and data processing that effectively leverages machine generated information for business advantage, including analytics and integration with existing IT systems.

Leverage Cisco's unmatched networking capabilities, and scale to meet all your needs

While a lot of attention has been given to device connectivity and the beginning of the data path from the hardware, full benefits cannot be achieved until the data can be leveraged by all interested consumers. In many cases, the consuming applications, displays, and data storage systems are remote from the device. Networking is implicit in these cases. Increasingly, cloud computing is also a requirement. The Kinetic Data Control Module (DCM) is designed to meet the unique requirements of cloud computing for IoT data distribution.

The Cisco Kinetic IoT platform consists of three modules:

- **Gateway Management Module** (GMM) – provision, monitor and manage gateways at scale
- **Edge & Fog Processing Management Module** (EFM) – enable computing on distributed nodes of the network
- **Data Control Module** (DCM) – enforce policy and get the right data to the right apps at the right time

Cisco Kinetic is designed to:

- Provide full connectivity to a wide array of devices
- Leverage Cisco's unmatched networking capabilities
- Scale to meet all future needs

The DCM Module is a companion to the Edge & Fog Module (EFM) and the Gateway Management Module (GMM) to collectively fulfill the complete set of requirements.

The Complete IoT System

IoT systems are unique in that they span the entire spectrum of computing alternatives. The devices are typically located in different places, requiring a distributed system architecture. Data from the devices sometimes needs near real-time processing, requiring computation close to the source of the data. Yet the target recipients of the data may be remotely located, with applications residing in private and/or public clouds.

It is impractical, and definitely unwise, to attempt a single computation system to meet these varied needs. The Kinetic Data Control Module focuses on the wide area distribution of the data. Modern cloud technology is heavily leveraged, as is Cisco's industry leading networking technology.

The Kinetic platform is the infrastructure for Cisco® IoT solutions, used by Cisco Advanced Services for custom implementations, and is available to Cisco partners and third parties as an offering to enable them to convert their IoT concepts into reality in record time.

Cloud-based data distribution gives you unlimited scalability

Scalability

A cloud based approach to data distribution provides many advantages, including elasticity and unlimited scalability. There is no better way to scale a system than by leveraging a cloud.

The DCM software leverages the underlying cloud system's ability to easily scale as needed. Further, the system is configurable to provide parallelism as required. The system is unique, in its ability to encompass multiple clouds, including public and private clouds by different providers, while retaining the integrity of the services described in this paper.

DCM has been designed and tested to handle tens of thousands of messages per second, where the messages can be up to several thousands of bytes each, from thousands of devices.

Aside from the issues of horizontal scalability (multiple data generators and receivers), there is sometimes the issue of scale and reliability of a given data path between one data generator and one data consumer. In these cases, data is delivered as quickly as possible, but when the consumer cannot ingest data as fast as it is being generated, the system will queue the data, providing a buffer between the parties, enabling them to operate asynchronously without any loss of functionality or data.

In summary, the DCM system can handle the most demanding IoT needs, up to thousands of devices and applications, with unlimited communications between them.

Effectively control IoT data distribution to ensure the validity and integrity of data

Control

Data owners have several legitimate concerns about the treatment of their proprietary and important data, once it passes outside their firewalls. Many of the historic protection mechanisms of their data centers are lost as the data moves to the cloud. Dealing with these concerns about "controlling data distribution" is a major component of the Data Control Module.

Issues include:

- **Data governance**: DCM provides a system whereby the original owner of the data retains complete control over who has access to the data, and how access is obtained.
- **Data providence**: The origin of the data is known and cannot be spoofed. Using unique device registration and credentials, there is never any doubt as to whether the source of the data is valid.
- **Data integrity**: The data is protected, across the network, into the cloud, and all the way to its destination. Integrity is retained throughout the entire path using network security, tunneling, and encryption.
- **Data sovereignty**: In some cases, distribution of the data beyond sovereign borders is prohibited. The system provides controls to assure that data remains within the acceptable geographic area.

- **Unwanted date exposure**: The owner of the data can specify which portions of the data can be shared. The “rules” can be arbitrarily sophisticated. Once the data owner specifies who has access to the data, no other party can override that protection.
- **Undesired data recipients**: Recipients must be authenticated to access the system. Further, they must be authorized to access select data.
- **Unanticipated derivative information**: In some cases, there is a hesitancy to share data for fear that it may provide unwanted insights. For example, data about the efficiency of a factory machine may lead to insight about a company’s manufacturing productivity. The system provides a number of ways to limit the possible risks, including:
  - Obscuring the source of the data
  - Anonymizing multiple sources of the data
  - Masking portions of the data
  - Limiting access in multiple ways, as described above

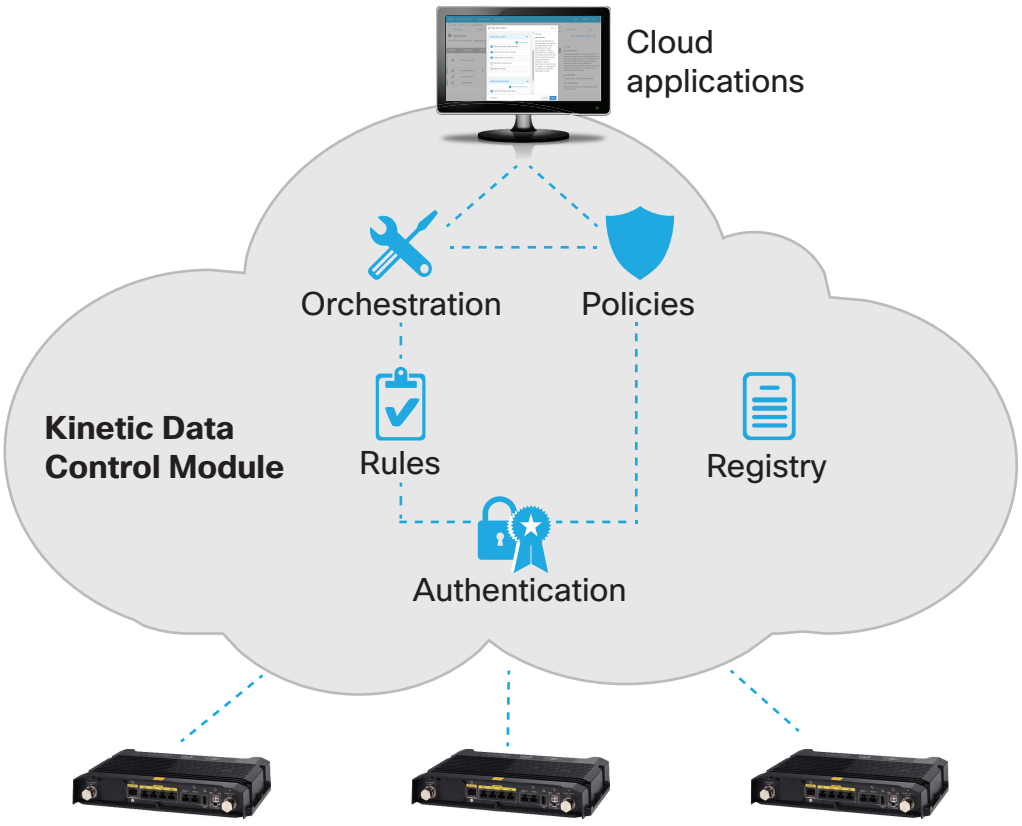

Policies

A fundamental mechanism for implementing the security and control of the IoT data is the establishment of policies. These policies can be implemented at different levels. Examples include:

Create policies to transform, filter, or anonymize data for specific use cases or audiences

- **Data field**: The telemetry from a device may contain some data that should be distributed and shared, and other data that should not, for any of a number of reasons. Processing can easily be established to control only disclosure of data that is appropriate to one or some of the recipients. Further, the individual data fields can be:
  - Transformed for easier consumption (such as adherence to a database schema)
  - Transformed to a different data type, for combining with other data
  - Removed, to reduce the data clutter or to avoid unwanted disclosure
  - Obscured or anonymized for privacy concerns
- **Data structure**: The structure or tuple, of an instance of data provided by an IoT source, can have policies applied to it prior to communications with a consumer. It can be:
  - Anonymized with regards to the originator
  - Portions of it can be masked, so only desired data is communicated
  - Blocked entirely, based on inspection of its content
- **Recipients**: Policies can be established such that only authorized recipients can receive the data. Note that these policies are controlled by the originator of the data, not the receivers.
- **Routing**: Policies can determine the routing of data to other systems or instances of DCM (such as in another cloud), based on the content of the data, or the originator of the data.

The desired policies are sometimes complex and include multiple “rules” and/or logic. As such, they can be specified using a pseudo-English scripting language. This approach was created to enable a system more complex than is feasible with a visual programming tool, but more straightforward and foolproof than a programming language.

Following are two examples of what can be done with the capability to set data policies:

Policy Description	DCM Rule

When the brake temperature exceeds 100 degrees for 300 ms, then convert temperature and pressure to metric units and send an alert	When brake.temp > 100 for (300 ms) then {
	brakes.temp_fahrenheit = brakes.temp*1.8 + 32;
	brakes.pressure_psi = brakes.pressure * .14503;
	SEND TO “starlord” TOPIC “BRAKES_Overheating_Alert”
	JSON brakes.temp_fahrenheit, brakes.pressure_psi; }

Calculate sliding averages for a 3 second window, and send an alert if the average temperature is higher than 100 degrees for 3 seconds, but not more often than once every 10 seconds	TS_UPDATE temp WITH msg(engine.temp) WINDOW (3 sec); WHEN ts_avg(temp) > 100 for (1 s) THEN throttle(10 sec) SEND TO “starlord” TOPIC “EMAIL/TEMP_TO_HIGH” JSON ts_avg(temp);

Orchestration

The capabilities listed above can be combined in any number of ways to meet your exact needs. This is a “policy-based data routing network” inside a cloud, or transcending multiple cloud domains. The domains can be different instances of a single vendor’s cloud offerings, or different public and private clouds.

Interoperability with Microsoft

DCM is a system to distribute data in a secure controlled manner. Other parties, including Microsoft, are focused on providing capabilities and tools to create IoT applications. The two systems are complementary.

- Cisco's networking, hardware, protocols, and the Kinetic platform, create a complete data path for IoT data. The Cisco Kinetic system, through the EFM and DCM modules, handles all aspects of data from the devices to the cloud.
- The Microsoft Azure IoT system provides cloud-based application enablement for systems that are recipients of IoT data.

DCM's data path to the Microsoft Azure cloud includes interoperability with the Azure IoT Hub. Users will be able to combine the best technologies from these two leading vendors, each of which will specialize in the portion of the system in which they excel.

The following diagram shows the comprehensiveness of the Cisco Kinetic platform, including:

- Hardware and software for the edge nodes (routers and switches)
- Fog nodes (Cisco UCS computers)
- Cisco Kinetic Edge & Fog firmware and software
- Policy-driven data distribution through the cloud (Data Control Module)
- Connecting to the Azure IoT Hub, which supplies application enablement services to cloud-based applications (such as analytics)

Leverage the best technologies to optimize data distribution and capture value in the cloud

CISCO

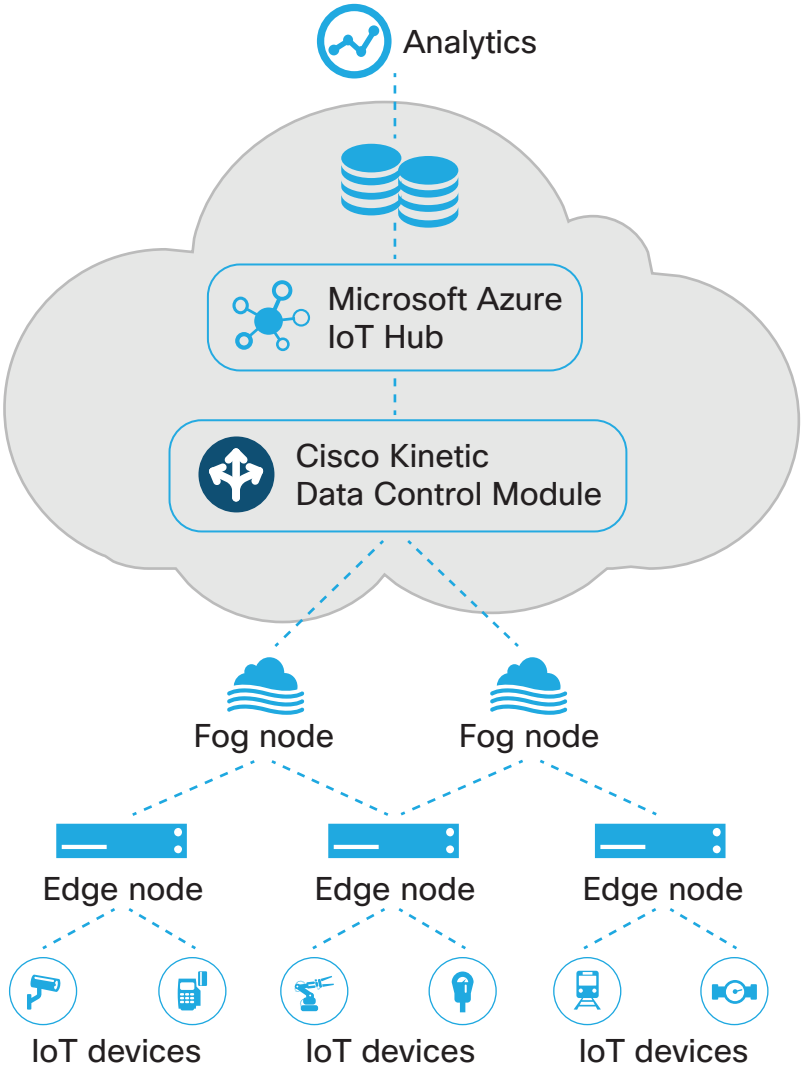

Cisco Kinetic enables you to build comprehensive IoT systems specific to your exact needs. As a result, you can address tomorrow's demands for data driven business systems that were previously neither feasible nor cost-effective.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

WP-CK-DCM-03012018