

The Internet of Things Reference Model



The Internet of Things (IoT) is a global industry movement that brings together people, process, data, and things to make networked connections more relevant and valuable than ever before. Today, more than 99 percent of things in the world are not connected. By 2020, it is estimated that 4.5 billion new people and 37 billion new things will have joined the Internet. In the near future, the growth and convergence of information, people, and things on the Internet will create unprecedented opportunity for countries, industries, and individuals.

What You Will Learn

There is much discussion surrounding the IoT with multiple ways of describing it. This paper provides a framework for understanding, discussing, and developing solutions for the IoT. It includes:

- Establishing a reference model for the IoT
- Defining standard terminology for each level
- Describing the functionality of, and interactions between each level

The Internet of Things is Everywhere

The IoT is increasing connections between people and things and volumes of data generated on a scale that once was unimaginable. An explosion of devices and new applications combined with the reach and power of the Internet enables new types of intelligent interactions between those things. In fact, devices already outnumber human beings on the planet by a ratio of 1.5 to 1.

The impact of the IoT on business spans industries and vertical markets. For example, a one percent reduction in capital expenditures from IoT-related efficiencies could save the oil and gas industry US\$90 billion over 15 years (Source: GE report, “Industrial Internet: Pushing the Boundaries of Minds and Machines”, 2012). Organizations are finding that they need to create an expanded, adaptable infrastructure that can keep pace with evolving network, compute, application, and data management demands. They must be able to address growing demand from inside and outside of their organizations. And they must be able to secure this more complex, interconnected infrastructure.

For providers of networking equipment and applications, the IoT potentially represents a huge market opportunity. Several trends have emerged over the past several years that are working together to shape the emerging IoT market:

- Rapid growth of data and analytics capabilities enabled by cloud computing
- Rapid growth in smart mobile devices
- Increasing interconnectivity between industrial, operational, and smart mobile devices
- Convergence of industrial and enterprise networks that enable applications such as video surveillance, smart meters, asset tracking, fleet management, digital health monitoring, and a host of other next-generation connected services

Various research sources suggest the world is adopting digital infrastructure five times faster than it adopted electricity and telephony. Industry estimates place the number of smart devices at 50 billion by 2020. Already today, there are approximately 46 million smart utility meters in the U.S. alone, generating 1.1 billion data points (0.5TB) every day.

It is clear that the number of network connections and amount of traffic generated by the IoT will soon dwarf the number of connections and amount of traffic generated by the Internet today. As an example, for every 30 minutes of flight, the typical commercial jet airliner generates 10TB of data. With more than 25,000 flights every day, petabytes of data are being generated daily just from commercial airliners. The amount and types of traffic traversing the Internet will change in several ways:

- More devices attached to the Internet will operate independently of human interaction
- Smart devices will sometimes have limited processing capabilities, which changes interactions between the device and the network
- Amounts of data generated are unlike anything seen to date

A New Reference Model is Needed

Network, compute, application, and data management architectures that are IoT-ready require a different communication and processing model. Today, there is not a standard way of understanding or describing these models for the IoT. As a result, the lines are blurred between IoT devices and systems and non-IoT devices and systems. The fact is, not every network is an IoT network. Nor does it need to be. And not every application is an IoT application. In general, when data is generated under the control of machines or equipment and sent across a network, it is probably an IoT system. But in the IoT, even generalizations may be inappropriate. There may be many, many exceptions.

This paper sets forth an IoT Reference Model. Its purpose is to provide clear definitions and descriptions that can be applied accurately to elements and functions of IoT systems and applications. This reference model:

- Simplifies: It helps break down complex systems so that each part is more understandable.
- Clarifies: It provides additional information to precisely identify levels of the IoT and to establish common terminology.
- Identifies: It identifies where specific types of processing is optimized across different parts of the system.
- Standardizes: It provides a first step in enabling vendors to create IoT products that work with each other.
- Organizes: It makes the IoT real and approachable, instead of simply conceptual.

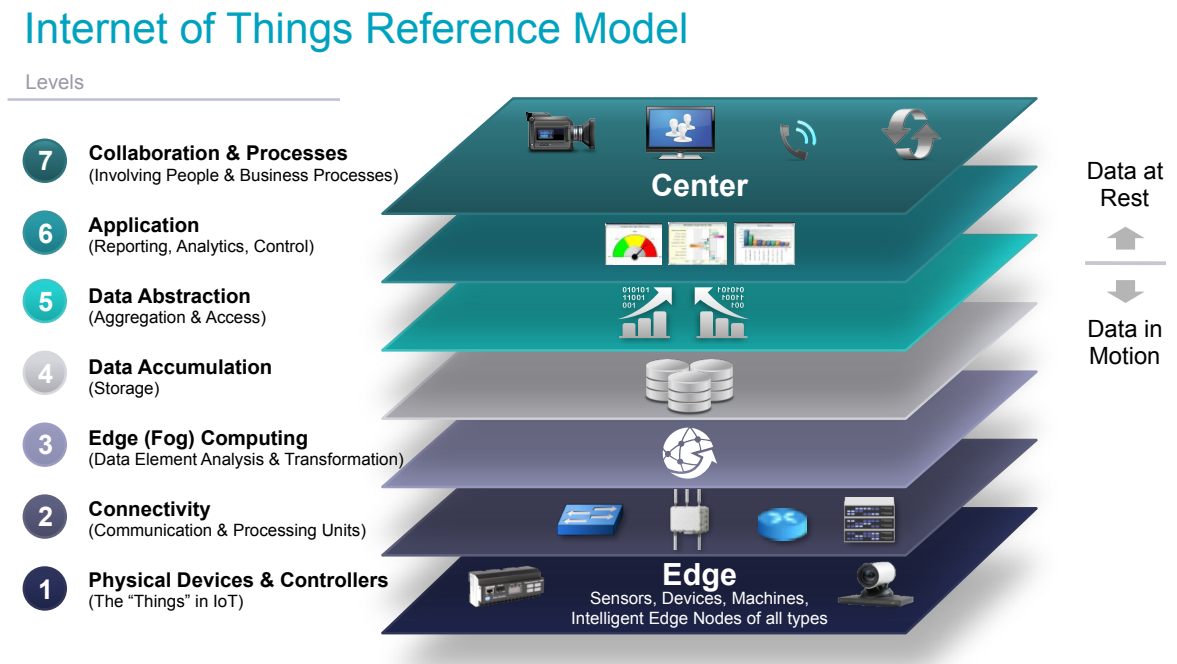
A Comprehensive, Multilevel Model for IoT

In an IoT system, data is generated by multiple kinds of devices, processed in different ways, transmitted to different locations, and acted upon by applications. The proposed IoT reference model is comprised of seven levels. Each level is defined with terminology that can be standardized to create a globally accepted frame of reference.

The IoT Reference Model does not restrict the scope or locality of its components. For example, from a physical perspective, every element could reside in a single rack of equipment or it could be distributed across the world. The IoT Reference Model also allows the processing occurring at each level to range from trivial to complex, depending on the situation. The model describes how tasks at each level should be handled to maintain simplicity, allow high scalability, and ensure supportability. Finally, the model defines the functions required for an IoT system to be complete.

Figure 1 illustrates the IoT Reference model and its levels. It is important to note that in the IoT, data flows in both directions. In a control pattern, control information flows from the top of the model (level 7) to the bottom (level 1). In a monitoring pattern, the flow of information is the reverse. In most systems, the flow will be bidirectional.

Figure 1. The IoT Reference Model



Level 1: Physical Devices and Controllers

The IoT Reference Model starts with Level 1: physical devices and controllers that might control multiple devices. These are the "things" in the IoT, and they include a wide range of endpoint devices that send and receive information. Today, the list of devices is already extensive. It will become almost unlimited as more equipment is added to the IoT over time.

Devices are diverse, and there are no rules about size, location, form factor, or origin. Some devices will be the size of a silicon chip. Some will be as large as vehicles. The IoT must support the entire range. Dozens or hundreds of equipment manufacturers will produce IoT devices. To simplify compatibility and support

manufacturability, the IoT Reference Model generally describes the level of processing needed from Level 1 devices. Figure 2 describes basic capabilities for a device.

Figure 2. Level 1 Physical Devices and Controllers

Internet of Things Reference Model

1 Physical Devices & Device Controllers (The “Things” in IoT)

IoT “devices” are capable of:

- Analog to digital conversion, as required
- Generating data
- Being queried / controlled over-the-net



Level 2: Connectivity

Communications and connectivity are concentrated in one level—Level 2. The most important function of Level 2 is reliable, timely information transmission. This includes transmissions:

- Between devices (Level 1) and the network
- Across networks (east-west)
- Between the network (Level 2) and low-level information processing occurring at Level 3

Traditional data communication networks have multiple functions, as evidenced by the International Organization for Standardization (ISO) 7-layer reference model. However, a complete IoT system contains many levels in addition to the communications network.

One objective of the IoT Reference Model is for communications and processing to be executed by existing networks. The IoT Reference Model does not require or indicate creation of a different network—it relies on existing networks. However, some legacy devices aren’t IP-enabled, which will require introducing communication gateways. Other devices will require proprietary controllers to serve the communication function. However, over time, standardization will increase. As Level 1 devices proliferate, the ways in which they interact with Level 2 connectivity equipment may change. Regardless of the details, Level 1 devices communicate through the IoT system by interacting with Level 2 connectivity equipment, as shown in Figure 3.

Figure 3. Level 2 Connectivity

Internet of Things Reference Model

2

Connectivity (Communication & Processing Units)

Level 2 functionality focuses on East-West communications

Connectivity includes:

- Communicating with and between the Level 1 devices
- Reliable delivery across the network(s)
- Implementation of various protocols
- Switching and routing
- Translation between protocols
- Security at the network level
- (Self Learning) Networking Analytics



Level 3: Edge (Fog) Computing

The functions of Level 3 are driven by the need to convert network data flows into information that is suitable for storage and higher level processing at Level 4 (data accumulation). This means that Level 3 activities focus on high-volume data analysis and transformation. For example, a Level 1 sensor device might generate data samples multiple times per second, 24 hours a day, 365 days a year. A basic tenet of the IoT Reference Model is that the most intelligent system initiates information processing as early and as close to the edge of the network as possible. This is sometimes referred to as fog computing. Level 3 is where this occurs.

Given that data is usually submitted to the connectivity level (Level 2) networking equipment by devices in small units, Level 3 processing is performed on a packet-by-packet basis. This processing is limited, because there is only awareness of data units—not “sessions” or “transactions.” Level 3 processing can encompass many examples, such as:

- Evaluation: Evaluating data for criteria as to whether it should be processed at a higher level
- Formatting: Reformatting data for consistent higher-level processing
- Expanding/decoding: Handling cryptic data with additional context (such as the origin)
- Distillation/reduction: Reducing and/or summarizing data to minimize the impact of data and traffic on the network and higher-level processing systems
- Assessment: Determining whether data represents a threshold or alert; this could include redirecting data to additional destinations

Figs 4 and 5 illustrate the functionality of Level 3 data element.

Figure 4. Level 3 Edge (Fog) Computing

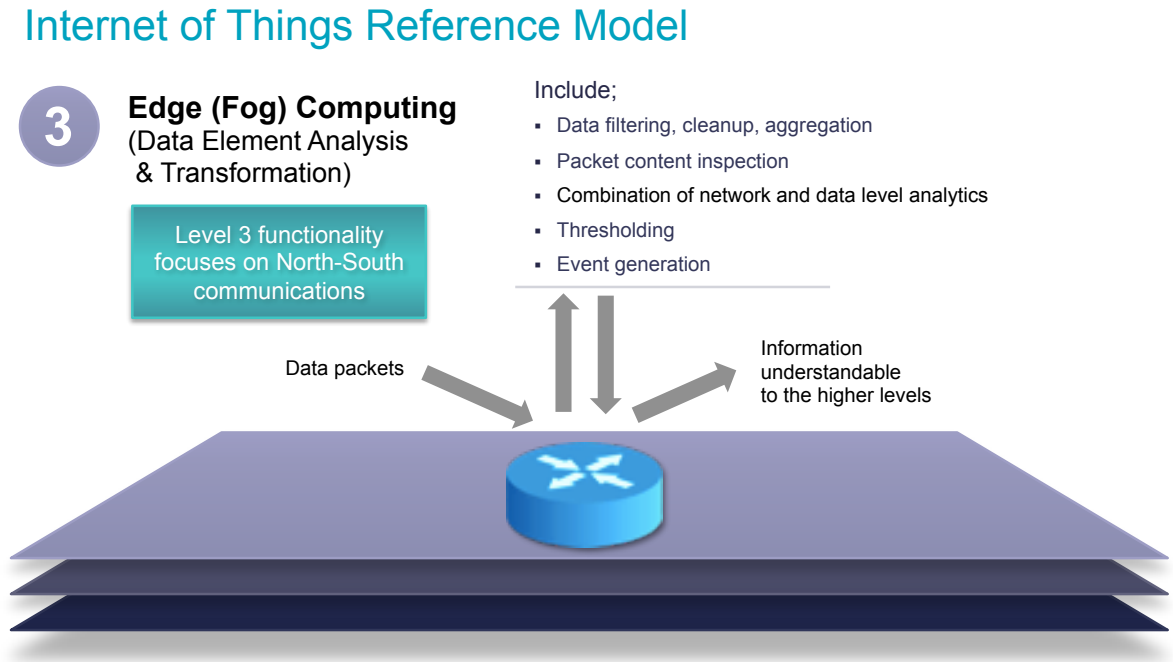
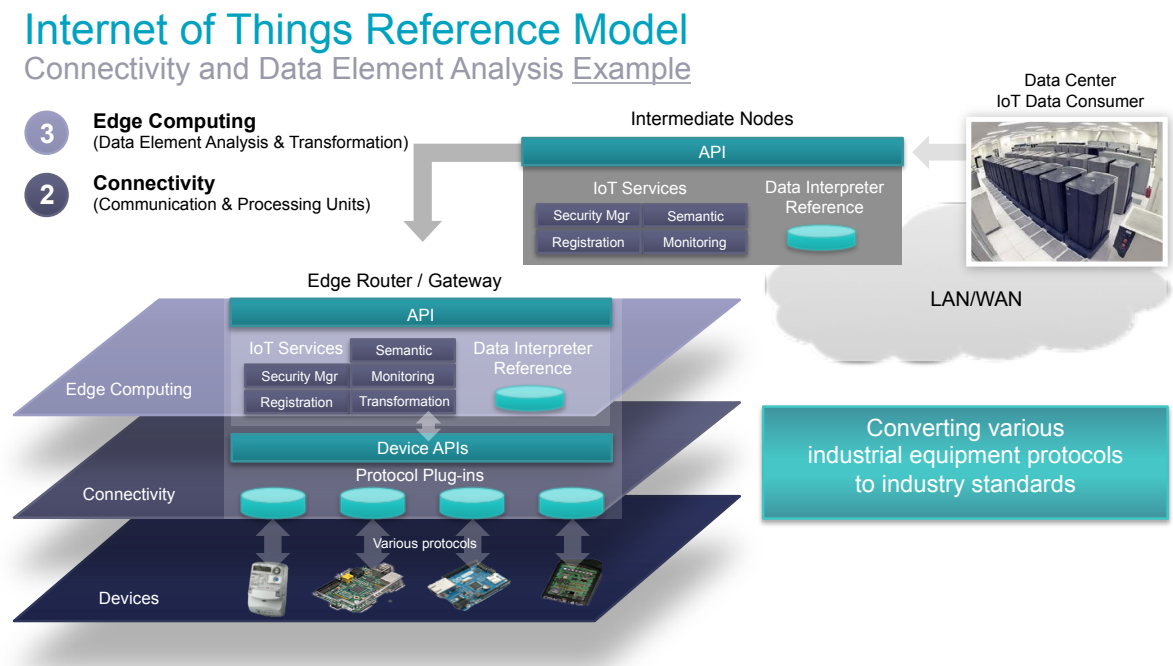


Figure 5. Level 2 and 3 Connectivity and Data Element Analysis Example



Level 4: Data Accumulation

Networking systems are built to reliably move data. The data is “in motion.” Prior to Level 4, data is moving through the network at the rate and organization determined by the devices generating the data. The model is event driven. As defined earlier, Level 1 devices do not include computing capabilities themselves. However, some computational activities could occur at Level 2, such as protocol translation or application of network security policy. Additional compute tasks can be performed at Level 3, such as packet inspection. Driving computational tasks as close to the edge of the IoT as possible, with heterogeneous systems distributed across multiple management domains represents an example of fog computing. Fog computing and fog services will be a distinguishing characteristic of the IoT.

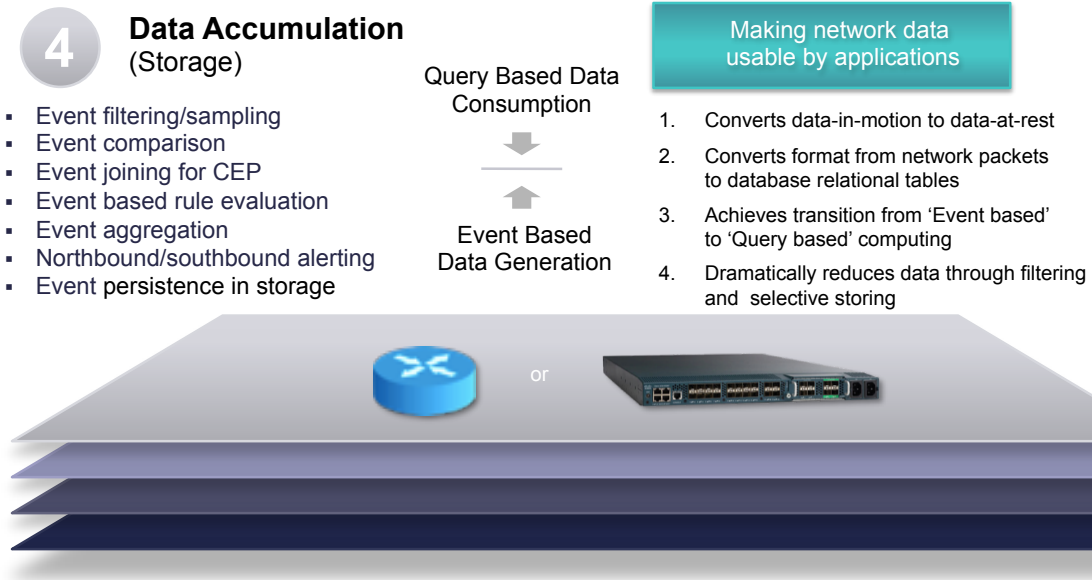
Most applications cannot, or do not need to, process data at network wire speed. Applications typically assume that data is “at rest”—or unchanging—in memory or on disk. At Level 4, Data Accumulation, data in motion is converted to data at rest. Level 4 determines:

- If data is of interest to higher levels: If so, Level 4 processing is the first level that is configured to serve the specific needs of a higher level.
- If data must be persisted: Should data be kept on disk in a non-volatile state or accumulated in memory for short-term use?
- The type of storage needed: Does persistency require a file system, big data system, or relational database?
- If data is organized properly: Is the data appropriately organized for the required storage system?
- If data must be recombined or recomputed: Data might be combined, recomputed, or aggregated with previously stored information, some of which may have come from non-IoT sources.

As Level 4 captures data and puts it at rest, it is now usable by applications on a non-real-time basis. Applications access the data when necessary. In short, Level 4 converts event-based data to query-based processing. This is a crucial step in bridging the differences between the real-time networking world and the non-real-time application world. Figure 6 summarizes the activities that occur at Level 4.

Figure 6. Level 4 Data Accumulation

Internet of Things Reference Model



Level 5: Data Abstraction

IoT systems will need to scale to a corporate—or even global—level and will require multiple storage systems to accommodate IoT device data and data from traditional enterprise ERP, HRMS, CRM, and other systems. The data abstraction functions of Level 5 are focused on rendering data and its storage in ways that enable developing simpler, performance-enhanced applications.

With multiple devices generating data, there are many reasons why this data may not land in the same data storage:

- There might be too much data to put in one place.
- Moving data into a database might consume too much processing power, so that retrieving it must be separated from the data generation process. This is done today with online transaction processing (OLTP) databases and data warehouses.
- Devices might be geographically separated, and processing is optimized locally.
- Levels 3 and 4 might separate “continuous streams of raw data” from “data that represents an event.” Data storage for streaming data may be a big data system, such as Hadoop. Storage for event data may be a relational database management system (RDBMS) with faster query times.
- Different kinds of data processing might be required. For example, in-store processing will focus on different things than across-all-stores summary processing.

For these reasons, the data abstraction level must process many different things. These include:

- Reconciling multiple data formats from different sources
- Assuring consistent semantics of data across sources
- Confirming that data is complete to the higher-level application

- Consolidating data into one place (with ETL, ELT, or data replication) or providing access to multiple data stores through data virtualization
- Protecting data with appropriate authentication and authorization
- Normalizing or denormalizing and indexing data to provide fast application access

Figure 7 illustrates the functions of the data abstraction level.

Figure 7. Level 5 Data Abstraction

Internet of Things Reference Model

5

Data Abstraction (Aggregation & Access)

Abstracting the data
interface for applications



Information Integration

1. Creates schemas and views of data in the manner that applications want
2. Combines data from multiple sources, simplifying the application
3. Filtering, selecting, projecting, and reformatting the data to serve the client applications
4. Reconciles differences in data shape, format, semantics, access protocol, and security

Level 6: Application

Level 6 is the application level, where information interpretation occurs. Software at this level interacts with Level 5 and data at rest, so it does not have to operate at network speeds.

The IoT Reference Model does not strictly define an application. Applications vary based on vertical markets, the nature of device data, and business needs. For example, some applications will focus on monitoring device data. Some will focus on controlling devices. Some will combine device and non-device data. Monitoring and control applications represent many different application models, programming patterns, and software stacks, leading to discussions of operating systems, mobility, application servers, hypervisors, multi-threading, multi-tenancy, etc. These topics are beyond the scope of the IoT Reference Model discussion. Suffice it to say that application complexity will vary widely.

Examples include:

- Mission-critical business applications, such as generalized ERP or specialized industry solutions
- Mobile applications that handle simple interactions
- Business intelligence reports, where the application is the BI server
- Analytic applications that interpret data for business decisions

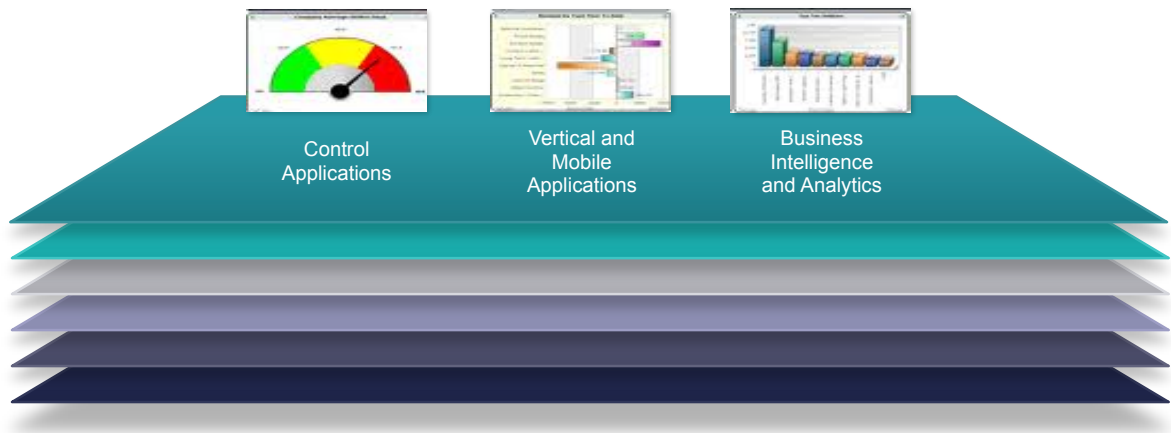
- System management/control center applications that control the IoT system itself and don't act on the data produced by it

If Levels 1-5 are architected properly, the amount of work required by Level 6 will be reduced. If Level 6 is designed properly, users will be able to do their jobs better. Figure 8 depicts Level 6.

Figure 8. Level 6 Application

Internet of Things Reference Model

6 Application (Reporting, Analytics, Control)



Level 7: Collaboration and Processes

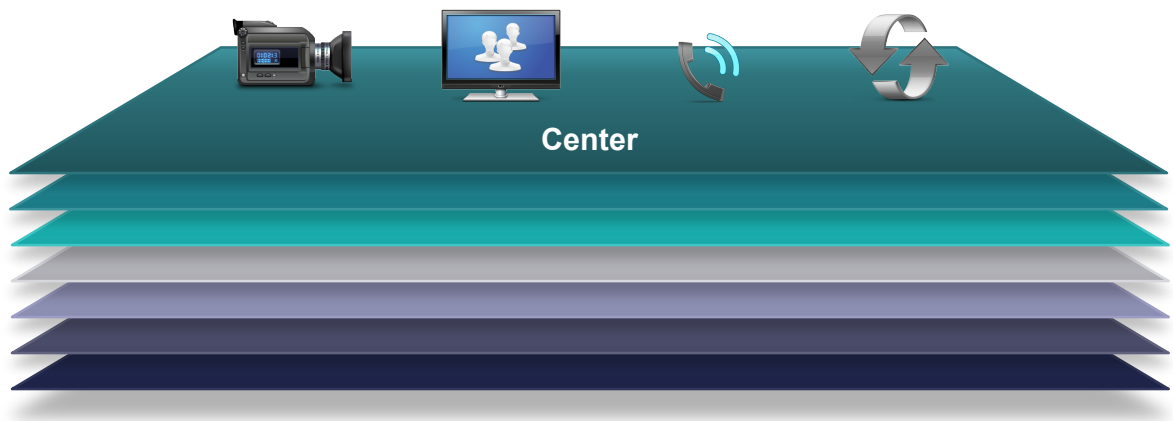
One of the main distinctions between the Internet of Things (IoT) and loT is that loT includes people and processes. This difference becomes particularly clear at Level 7: Collaboration and Processes. The IoT system, and the information it creates, is of little value unless it yields action, which often requires people and processes. Applications execute business logic to empower people. People use applications and associated data for their specific needs. Often, multiple people use the same application for a range of different purposes. So the objective is not the application—it is to empower people to do their work better. Applications (Level 6) give business people the right data, at the right time, so they can do the right thing.

But frequently, the action needed requires more than one person. People must be able to communicate and collaborate, sometimes using the traditional Internet, to make the IoT useful. Communication and collaboration often requires multiple steps. And it usually transcends multiple applications. This is why Level 7, as shown in Figure 9, represents a higher level than a single application.

Figure 9. Level 7 Collaboration and Processes

Internet of Things Reference Model

7 Collaboration & Processes (Involving people and business processes)



Security in the IoT

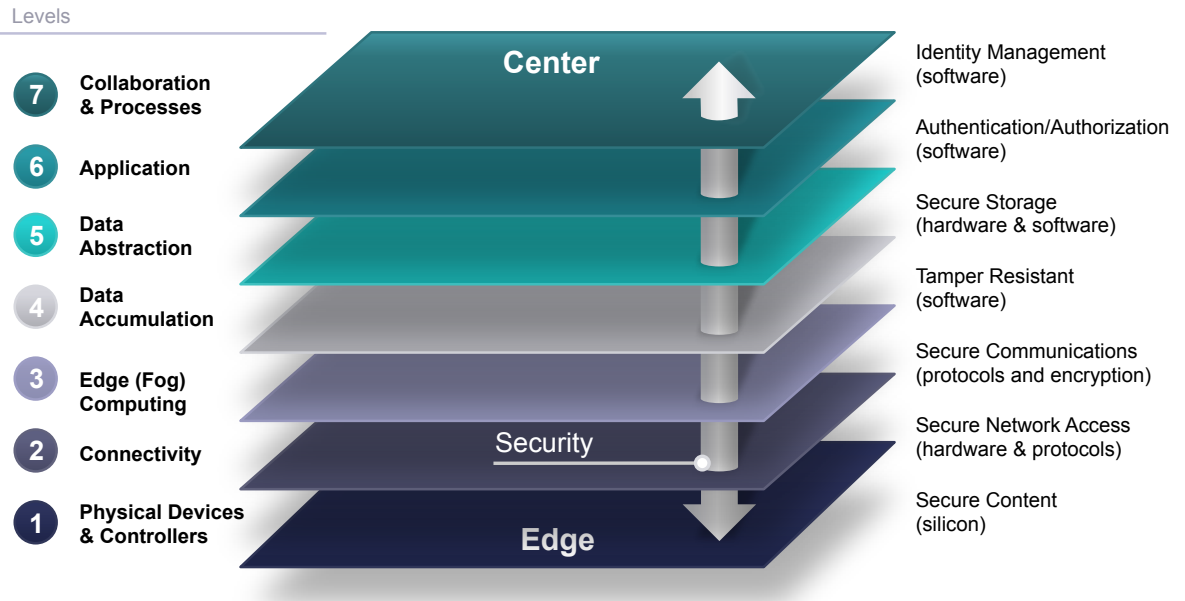
Discussions of security for each level and for the movement of data between levels could fill a multitude of papers. For the purpose of the IoT Reference Model, security measures must:

- Secure each device or system
- Provide security for all processes at each level
- Secure movement and communication between each level, whether north- or south-bound

As shown in Figure 10, security must pervade the entire model.

Figure 10. Pervasive security throughout the IoT Reference Model

Internet of Things Reference Model: Security



Summary

The Internet of Everything (IoT) Reference Model is a decisive first step toward standardizing the concept and terminology surrounding the IoT. From physical devices and controllers at Level 1 to the collaboration and processes at Level 7, the IoT Reference Model sets out the functionalities required and concerns that must be addressed before the industry can realize the value of the IoT. With the goal of enabling the IoT, this reference model provides a baseline for understanding its requirements and its potential.



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)