Internet of Things: Architectural Framework for eHealth Security

David Lake, Rodolfo Milito, Monique Morrow and Rajesh Vargheese

Cisco Systems, {dlake, romilito, mmorrow, rvarghee}@cisco.com

Received: 26 June, 2013; Accepted: 8 October, 2013

Abstract

The Internet of Things (IoT) holds big promises for Health Care, especially in Proactive Personal eHealth. To fulfil the promises major challenges must be overcome, particularly regarding privacy and security. This paper explores these issues, discusses use case scenarios, and advances a secure architecture framework. We close the paper with a discussion of the state of various standard organizations, in the conviction of the critical role they will play in making eHealth bloom.

Keywords: IoT, IoE, M2M, eHealth Security, Standards.

1 Introduction

During the 1990s, as the Internet grew in the publics' awareness, a number of e-terms emerged to capture new forms of personal and business interactions. "email" brought new possibilities for people to communicate rapidly and share experiences; "e-commerce" enabled new ways to conduct business and financial transactions through the Internet. The introduction of eHealth brought the promise to improve health and the health care system by leveraging Information and Communication Technologies (ICT).

The precise meaning of eHealth varies with the source. There is not a single consensus definition. Some benefits of eHealth extend from established telemedicine systems; others are only practical using a machine-to-machine (M2M) model and assume that patients have access to broadband service.

Journal of ICT, Vol. 3 & 4, 301–328.
doi: 10.13052/jicts2245-800X.133

The World Health Organization [WHO] defines e-Health as:

"E-health is the transfer of health resources and health care by electronic means. It encompasses three main areas:

The delivery of health information, for health professionals and health consumers, through the Internet and telecommunications.

Using the power of IT and e-commerce to improve public health services, e.g. through the education and training of health workers.

The use of e-commerce and e-business practices in health systems management.

E-health provides a new method for using health resources - such as information, money, and medicines - and in time should help to improve efficient use of these resources. The Internet also provides a new medium for information dissemination, and for interaction and collaboration among institutions, health professionals, health providers and the public."[1]

In speaking about eHealth today, we need to understand the relevance of Machine-to-Machine Communication [M2M] and the Internet of Things [IoT].

The term "Internet of Things" was originally associated with applications that involve Radio Frequency Identification (RFID). These make use of so called tags, tiny chips with antennae that start to transmit data when they come in contact with an electromagnetic field. They are passive communication devices, in contrast to active devices, that can transmit because they have access to a power source like a battery.

The term "Machine to Machine communication" (M2M), describes devices that are connected to the Internet, using a variety of fixed and wireless networks and communicate with each other and the wider world. They are active communication devices. The term is slightly misleading in that it seems to assume there is no human in the equation, which quite often there is in one way or another; hence we favour the term IoT[1].

The term eHealth is widely used by academic institutions, professional bodies, and standards organizations. From most of the definitions, two items keep appearing and seem to be the important concepts – health and technology. The definitions include use of the Internet or other electronic media to disseminate health related information or services.

It is in the home and assisted-living environments where many applications for eHealth are expected to flourish. Monitoring systems for the elderly or post trauma patients has gained considerable attention. These new systems come

[1]Lately the more inclusive term Internet of Everything (IoE) has gained traction.

complete with voice and video options. Automated movement monitoring systems allow the identification of falls and notification of medical personnel without any user intervention. Traditional movement monitoring systems are plagued with false alarms. The combination of voice and video allows for verification and a more appropriate response in the case of an alarm.

Another use case is remote monitoring of patients for blood glucose readings, blood pressure, pulse oximetry, or heart monitoring. In the case of blood pressure monitoring, the readings can provide important information to a physician. Furthermore, measurements taken at home, during daily activities, can potentially be of even greater importance to those taken at a doctor's office, since the readings reflect the patient's condition under normal situations.

eHealth brings special characteristics. The monitoring device's environment is a patient; a living and breathing human being. This changes some of the dynamics of the situation. Human interaction with the device means batteries could be changed, problems could be called in to technical support and possibly be resolved over the phone rather than some type of service call. In most cases, the devices on the patient are mobile not static with regard to location.

The environment for monitoring patients has moved from the hospital healthcare services to a patient's context. M2M/IoT eHealth applications enable remote monitoring of patient health and fitness information, the triggering of alarms when critical conditions are detected, and in some cases the remote control of certain medical treatments or parameters.

The rest of the paper is organized as follows. In section 2 we do a deep dive on the role of IoT in eHealth, followed by an architectural framework in section 3. We deal with Security in section 4. Given the importance of security, we pay attention to its different aspects in various subsections (IoT landscape, endpoint devices, networking, cloud-based services, data storage, enterprise, and federated access). Our final section, before the summary and conclusions, is devoted to standards.

2 Internet of Things and eHealth

IoT brings forth a new phase of the Internet evolution that we can characterize as "the Internet meets the physical world". Today's few billions of endpoints will increase in number by several orders of magnitude. This formidable inflation immediately points to an obvious scalability issue. The number game, and consequent scalability issue, may mask deeper issues, though, including

the nature of the endpoints, and the nature of the interactions between the endpoints. We will elaborate on this point later.

While the original Internet connected computers, IoT powered e-Health solutions connect information, people, devices, processes and context to improve outcomes. The intelligent devices that are connected (which were previously passive devices and were not connected) provide a great wealth of information that can be used to make actionable decisions based on algorithms and evidence-based models and can significantly impact how healthcare is delivered and operationalized.

Harnessing the power of Internet of Things for eHealth creates a lot of opportunities to improve outcomes and drive wellness is populations there by reducing the strain points of today's healthcare system. Some of the most promising use cases of connected e-health include preventive health, proactive monitoring, follow-up care and chronic care disease management.

According to BCC research: "Prevention must become a cornerstone of the healthcare system rather than an afterthought. This shift requires a fundamental change in the way individuals perceive and access the system as well as the way care is delivered. The system must support clinical preventive services and community-based wellness approaches at the federal, state, and local levels. With a national culture of wellness, chronic disease and obesity will be better managed and, more importantly, reduced."[2]

The changes, uniqueness, opportunities and complexities of e-health enabled by the power IoT is significant and can be characterised by some of the changes that we anticipate in the ecosystem. They include:

- The number of devices that will come online
- The number of devices that will generate information
- The number of decision making points
- The number of entry points into the system
- The number of types of devices
- The number of types of interactions of devices, applications and processes
- The number of opportunities to leverage the data

Just in the US, the market for preventive healthcare technologies and services by 2014 is projected to increase nearly to $16 billion.

This trend is global in nature. The sensor market is an important component of e-health market. The biosensor market is projected to reach $15 billion by 2015 [3]

According to Cisco's prediction there will be 50 billion devices by the year 2020. While the breakdown for healthcare is unknown at this time, the rate at

which innovations in the bio sensor market is progressing and the increasing use of the data to inform decisions, healthcare will see its fair share of devices in the IoT space.

The most disruptive changes, though, will depend on our ability to organize endpoints in systems that operate as coherent units to deliver applications of interest. Let us discuss this point further.

The IoT space of terminal endpoints decomposes into two major classes. The first class includes the current smart phones, tablets, and laptops. While each one is quite an advanced technological piece, including sensors and cameras, we can ignore their internal complexity and regard them as simple points, providing connectivity to the person who owns them. There is a second, emergent class of complex systems, not decomposable into just the set of sensors and actuators that integrates them. We could draw examples from Smart Grid, Connected Vehicle, Smart Cities, Manufacturing Automation, etc. Most relevant for our purpose, though, is to emphasize that many of the eHealth endpoints fall into this category. There are several critical distinctions between the two classes:

The endpoint-Cloud client-server paradigm dominates the first class. It is essentially a scheme of communication. This paradigm induces a vibrant model of apps development: tens of thousands of individual contributors develop downloadable apps for Apple's IoS and the Android environments, for example. An important fraction of those apps relate to eHealth (monitoring vital signs, tracking physical activity, etc.).

The second class of endpoints is even richer, and more complex. The endpoints are clustered, organized in coherent systems. These systems require networking, computation, and storage resources. The diverse use cases may require low latency and support for mobility, and also be geographically distributed [4]. Most importantly, many include not only sensors but also actuators. Closing the sensor-actuator loop often imposes strict latency requirements. Development apps for these complex endpoints require domain expertise that exceeds the capacity of any individual contributor. For eHealth the actuation part of these systems have security implications that we will discuss later.

The second class of devices that are the source for the generation of the information can be further classified into various categories.

Based on how the devices are connected to the patient, the devices can be classified into implantable, wearable, unconnected, or connected on a need basis.

Based on how the device is connected to the network, the devices can be classified into wired, wireless, non-connected.

Based on the data that the device generates, the devices can be classified into real-time continues (e.g., patient monitoring), discrete data sources (oximeter that generates data at regular intervals), and one-time data source (e.g., MRI scanner).

Not all data will be created equally – for example, a personal monitor designed to track a long-term trend in a medical condition or its treatment may only require to send data to a processing element every few hours of days, and a delay of a few seconds or minutes would be immaterial. At an extreme, the total loss of data for an entire measurement period where that period is a very small fraction of the total collection time would be of little consequence.

In contrast, a device that actively monitors a serious, life-threatening condition that requires specific action to be taken with a given time period or where a single-occurrence is of importance would impose tight requirements on the collection and dissemination of the data. In that case, it would not be acceptable to delay or lose a single packet of data.

Based on how the device is used by a single person or a group of people, it can be classified into dedicated, shared within a limited group, or shareable with a wider population.

The challenge of an eHealth/IoT architecture is to support this wide range of device types in a variety of care needs and settings.

3 Proposed Architectural Framework

The explosion of Proactive Personal eHealth, self-management of health conditions, and the collection of data, will radically change the manner by which health-care is delivered and information is collected.

Already, national and transnational organisations such as the European Commission have identified and begun work on projects to address issues around scalability, security, data collection, and interoperability. The paper "eHealth Action Plan 2012–2020 – Innovative Healthcare for the 21st Century" [5] details a number of these areas very well, calling out interoperability of systems, legal and societal barriers to adoption, and detailing how through support of the eHealth Network, the European Commission aims to research and solve these problems [6].

Of particular note in the Action Plan is the proliferation of mobile health and wellbeing as detailed on page 9 of the report:

> The growth in the mobile health and wellbeing market has been accompanied by a rapid increase in the number of software applications for mobile devices (or 'apps'). Such applications potentially offer information, diagnostic tools, possibilities to 'self-quantify' as well as new modalities of care. They are blurring the distinction between the traditional provision of clinical care by physicians, and the self-administration of care and wellbeing. Network operators, equipment suppliers, software developers and healthcare professionals are all seeking clarity on the roles they could play in the value chain for mobile health.

Coupled with the number of mobile devices "apps" will be specific-use devices, each capable of different levels of security, traffic generation and protocols, each with requirements that mirror their function.

To visualize the architecture framework for IoT enabled e-health, it is very important to understand the lifecycle of the various entities and their interactions. The life cycle of the device data is critical to understand and can be summarized using six C's. They are

1. Connection: the focus for this function is related to how the device is connected to the ecosystem
2. Collection: the focus of this function is related to how data is collected from the sensor. The data can be pushed our pulled from the sensor.
3. Correlation: the focus of this function is related to mapping the data to a context and does correlation to create meaningful and concise data that can be processed and be used to make decisions.
4. Calculation: the focus of this function is to make a decision based on the data that has been filtered and is processed through an algorithm
5. Conclusion: the focus of this function is to take appropriate actions. The action could be to ignore the event or to escalate.
6. Collaboration: the focus of this function is to enable the collaboration between the patient and the care teams.

Architecture for e-health must consider the needs of each step in this life cycle and must address the effective and efficient execution of each function.

The key to e-health architecture is to support an interoperable ecosystem of different types of devices, applications, and backend systems to enable the free flow information for precise and timely decision-making. The

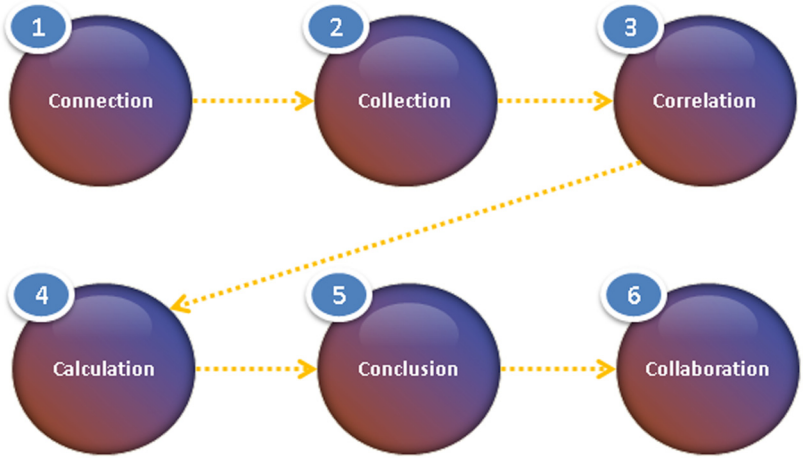

Figure 1 The life cycle of data and processing functions

information service bus enables the communication between the layers and supports multiple protocols.

The device layer must consist of a flexible registry-based model that enables plug-and-play of devices.

Given the number of devices and the information they generate, it is critical that information be filtered. The challenge with filtering is to identify the right information at the right time and eliminate false alarms yet not miss any critical information. The clinical decision support systems are used to process this information to make conclusions and the action that needs to be performed based on the information received from these devices.

The co-relation requires data from multiple systems and hence the architecture must support seamless interoperability between the systems that houses the information. The data includes the real-time data as well as historic data that are stored in the system.

The data flow architecture focuses on the source of the data, the destination the data and path the data. The source of the data is typically the sensor. The data can be either locally cached or is sent to the upstream systems without storing in the sensor. The path taken by the data includes a gateway, which can also cache some of the data and do distributed processing. Intermediate hubs can also store and process the data to filter out or make certain decisions. A distributed rules engine is used to make distributed decisions at the closest point of care. This enables data traffic to be filtered and processed efficiently without having every data being processed by the cloud service.

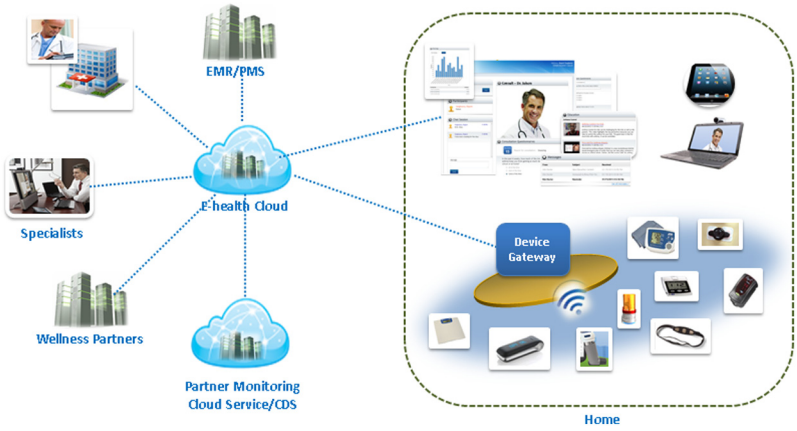

Figure 2 High Level e-Health ecosystem Architecture

The data finally enters the data store in the cloud where it is stored, further processed and archived.

The conclusion could be that a care team member needs to contact the patient to understand further why there is a deviation from the expected readings from the device. This approach is critical to identify problems early in the cycle thereby reducing considerable amount of cost and complexity in dealing with health care issues at the emergency room stage.

Once the conclusion has been made that the care team needs to interact with the patient. Different methods can be used to enable collaboration, which can range from basic text messages to real time video enabled collaboration.

Network architectures must be designed in such a way that these differing, sometime competing requirements may be met.

The specific use and design of each device will impact the choice of underlying access medium and therefore the level and type of security that can be employed. For example, a device that can be powered from a standard household electrical outlet, have no requirements to be available during a power outage will have less constraints on the network media access layer and encryption algorithm than a worn device which is required to be available 24/7, has constraints on size and therefore on the amount of power that can be associated with network access.

However, as pointed out by the eHealth Action Plan, interoperability or at least commonality in application profiles and data will be key to ensuring an ease of exchange of information between these devices or to the consumers of the information – in order to gain an holistic view of an underlying condition,

information is typically gathered from a number of points and by devices with different characteristics, manufacturers, purposes, etc.

Therefore, it is important that a common data-set by employed so that information may be securely consumed by multiple institutions without compromising security.

Looking at the practical implementation of such architecture leads to some requirements that must be met by the components. Generating, moving, and accessing health related information are core activities in any e-Health solution. The associated requirements go beyond those of a traditional data network, and leads naturally to the consideration of a Message Bus, which offers:

- A reliable solution
- Message persistency
- High-performance and scalability
    - Ideally in excess of 100K messages/sec throughput
    - Ability to handle terabytes of messages without performance impact
- Distributed implementation
    - Fault-tolerance with cluster-centric design
- Guaranteed message ordering
- End-End compression support
- Support for online, low-latency communication
- Open interfaces for data connectors

The Apache Kafka distributed messaging system [32] appears as a good fit.

A full assessment of Kafka is beyond the scope of this document. However, in terms of providing a message-bus system appropriate to an e-Health environment, such as distributed system is highly appropriate, offering both peer-to-peer and brokered communications.

For a general discussion of distributed message systems, including Kafka, we refer the reader to the presentation by Max Alexejev [33].

4 Security

4.1 Security in the IoT Landscape

There is a legitimate concern that security vulnerabilities could pose a significant risk to the industry's belief in the ability of M2M/IoT to deliver greater efficiencies and help enterprises optimize the costs of business operations.

More specifically the chances of security breaches increase in direct proportion to the "degree of connectivity" and as more endpoints become connected to the enterprise IT backend systems through public IP networks, the chances of things going wrong, either intentionally or unintentionally, are accelerated.

To evaluate the security architecture for e-health, we break down the architecture into multiple sections and evaluate the security challenges that we have in each of these domains. As depicted in Figure 3, the main domains include endpoint and access, cloud services, partners, and providers.

eHealth applications in an M2M/IoT environment run on a number of components, including sensor devices and actuators, and networking, processing and storage elements. The overall level of security is upper-bounded by the weakest component in this interactive system. Hence, each component, and the overall system must be designed with security in mind.

There are three basic attack vectors, and a corresponding attack surface to each vector. Data is the first attack surface, and the communication channels

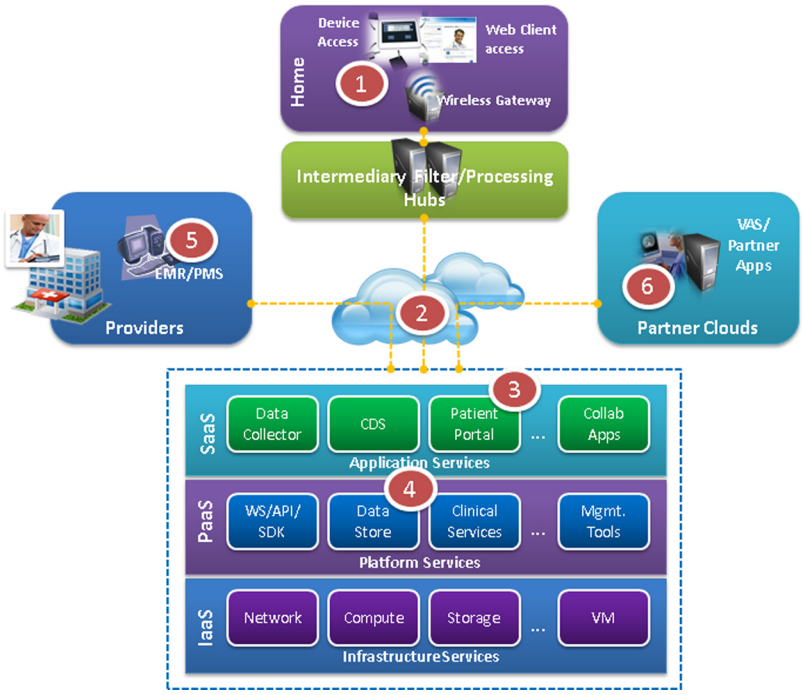

Figure 3 e-health security domain touch points

the second one. M2M/IoT brings forth a third, novel attack surface: physical attacks on or through the medical device.

To deal with the first two issues, one must consider the authentication of the application and/or device, the protection of data, and securing the communication channel, itself. This paper will discuss briefly some of the security issues around device design and use.

A recent editorial [7] welcomes the potential benefits of mobile technology while emphasizing the need to subject the apps to the rigorous standards of evidence-based medicine. Indeed, the smart phone/tablet explosion has spurred a formidable app development activity. A thriving community of tens of thousands of Apple IoS and Android developers has generated a number of interesting apps, a sizable and rich subset of them in the eHealth space. We share with the author the appreciation for the aggregate creativity of this community, which has accelerated bringing apps to the market. We also heartily agree with the need for a strict validation program along the lines suggested by the editorial.

From the security viewpoint we need to go even further, as exemplified by the potential of inflicting physical damage through the compromising of an insulin pump [8]. While developing techniques to prevent hacking of medical devices is beyond the scope of this paper we suggest here a potentially fruitful line of research: a precise definition of the expected interactions for the specific device with the external world, and a clear baseline of the expected behaviour could be the basis for building a tampering-resistant device.

It is important to note that the term "securing the communication channel" is quite broad and should be about ensuring the confidentiality, integrity, and reliability of data sent over telecommunication networks in a connected ecosystem. Security in the M2M/IoT model is not only about ensuring the proper access to the right entities at the right time but also about creating a secure architecture.

The next section will take the reference architecture and a simple version of an eHealth use case to examine the first two aforementioned threat vectors and recommended mitigations.

Consider a patient with a blood pressure monitoring device, which takes a blood pressure reading every 15 minutes and the device itself or another local device, which has a collector function stores the readings. Once a day the device or collector and the medical facility's application server communicate with each other to transfer the reading to the server. The collector function could be in the M2M gateway. The readings may be summarized or some other

data manipulation technique performed, and then the medical staff reviews the results.

This simple example illustrates several areas for security to address.

On the device and collector:

- Secure Boot of the device for platform integrity check and boot loader authenticity
- Secure Storage of the secret keys. The storage should be physically tamper resistant and access control protected
- Secure Storage of the data
- Device identification must be a unique identifier within the eHealth context

For the communication channel:

- Mutual Authentication between the eHealth device and the application server and/or the network
- Data Integrity to protect the data from any alteration during the communication session
- Data Confidentiality uses encryption and decryption of data between the secure device and the application server and/or network during data exchanges

In the Ecosystem:

- Key Management of the secret keys in the eHealth
- Cryptographic Support of cryptographic protocols, such as AES and optionally PKI

There are several things to consider in this example.

Does the M2M monitoring device and collector have proper security characteristics?

For the infrastructure components, M2M communication starts with connected devices. Devices have a certain level of processing power that is then used for monitoring or reporting on specific events or conditions. There is a possibility of the firmware being compromised on the device itself. Device OEMs normally rely on their manufacturing partners to develop and configure these elements. These OEMs have to ensure that their partners adhere to strict security policies themselves. A device that goes rogue or does not perform as expected under certain conditions is equally dangerous (if not more) as a device that is open to easy access and manipulation by unauthorized personnel.

Is the local environment considered secure? If not, then all communications from the device should be protected.

The environment in which the monitoring device operates is very important. If the device is only used in a home, or a medical facility, then one might be able to assume that is a trusted environment. In a trusted environment, the communication between the device and the gateway device may not need to be encrypted. However, in any environment that is not considered trusted, then the communication needs to be encrypted for data conditionality.

Currently a prominent insulin pump vendor provides wireless connectivity between their insulin pump and an USB device plugged into the patient's PC. One might consider this a safe and secure environment. However, the communication protocol is proprietary and not secured. Recently a researcher demonstrated that this insulin pump could be hacked into and controlled if one is in the close proximity. This allowed the hacker to instruct the pump to perform all manner of commands, even dispensing the complete reservoir of insulin to the user without the user's knowledge. This could have fatal consequences.

The second issue and the biggest challenge from the device side is that a lot of M2M/IoT devices do not have enough capability to do the encryption on the device. This is a currently a big issue in M2M/IoT. Many current devices and sensors have a small amount of on-board memory and a microprocessor that is simply unable to handle standard security protocols (such as AES 128 and others). Designs are improving on that front, and there are products coming that are capable of 16-bit processing. The other issue on the device side is regarding the amount of battery power it takes to do all the algorithmic computations for data encryption. This is especially relevant in cases where the endpoints communicate directly with the enterprise IT backend systems without having a local aggregation unit. Some remote deployments are expected to run for several years on battery power without any human intervention (to replace the battery for example) and computationally intensive processes could put a significant strain on the batteries.

The algorithms require a minimal amount of power, occupy a very small footprint in memory and handle computational transactions extremely fast. With these types of advancements it may make it easier to meet the security requirements for integrity and confidentiality.

What protocol does the monitoring device support for communication?

The communication protocol that a device uses is not always IP. The Continua Health Alliance supports ZigBee as the preferred Personal Area Network (PAN) protocol for devices. Other supported protocols by the industry are Z-Wave, ANT and others. Therefore, there is a need for an additional device acting as a protocol gateway between ZigBee and IP networks. This

device is often used for the collection of data, and is usually connected to main power. The protocol situation is similar to the situation 20 years ago with network protocols. There is work by various groups to help promote the use of IP in devices.

The wireless network provides some degree of security. Unfortunately, the current scheme of encryption over GSM/GPRS networks is not totally secure, although it does take a high degree of technical sophistication to break in these ciphers. A5 encryption has been considered broken for some time. Certainly there are significantly more security mechanisms in 3G and 4G, but currently that is a very small percent of the M2M/IoT market.

A rogue base station software, such as OpenBTS and OpenBSC can be used to launch a man-in-the-middle (MITM) attack. It has been demonstrated that with the use of a rogue base station and a patched cell phone, it is possible to get into a vendor's private network. The cell phone builds a bridge to the vendor's network.

One method for added security is to use an encryption mechanism that is layered - 128-bit AES, and then the A-5 encryption on the GPRS channel. The perpetrator would not be able to break the AES encryption. However, most M2M/IoT device manufacturers have not implemented AES, as the present generation of sensor type of devices that are not that powerful from a security perspective.

4.2 Endpoint Devices Security

The management of long-term and chronic conditions will therefore require a number of worn and/or embedded devices, constrained by their specific purposes in their communication and information exchange.

As discussed in the book "Body Sensor Networks" "Body Sensor Networks" [9] network and media access systems will differ between each of these, as will their security capabilities and needs.

It is envisaged that a common device will act as the collector and orchestrator of this information operating at the heart of an autonomous, but connected system. Rather than imposing a single architecture and security scheme on every device, it makes more sense to choose security attributes that are appropriate to each device, wrapping these to a central policy from the collector as part of an autonomic [10] computing domain.

Assuming that identity of the end device can be managed – many of these wearable devices have identity as part of their composition – managing the identity of the collection device and association with the sensors will be

important. Establishing a trust relationship between the collector and the service consumer is vital – technologies such as the Trusted Platform Module[11] as used in compute systems the world over can be embedded in these collection agents.

It is likely that a number of parallel collection agents may be required – for example, for two family members with different conditions or a single person with multiple information sources or where hard separation between data sets is required.

It is suggested that with enhancements to the home-gateway function at the end of most cable, DSL or other Internet connections, the use of virtual machines running in "slices" could provide the ability for a device manufacturer to deploy software images that provide access to use-specific devices. This has the benefit of moving the intensive, power-draining compute functions away from the sensor. Device specific security algorithms and protocols can be employed on the sensor, with translation to systems more appropriate for wide-area network connectivity on the mediating gateway.

Many home-gateway devices also support multiple RF connections – Wi-Fi, Bluetooth, cellular back-up services – run ARM or x86 processors and can support applications today. With the commoditization and rapid price decreases in home-gateway devices, providing this functionality at the edge becomes economically attractive and architecturally beneficial to eHealth.

An example of the MyHeart eHealth architecture in which many of these aspects of sensor usage, security and multiple applications has been piloted is discussed in the referenced paper [12].

According to NIST, the threat profile for handheld devices is a superset of the profile for desktop computers[13]due to the size, portability and availability of wireless interfaces and associated services. The security threats to mobile devices include, but not limited to theft, unauthorized access, malware, spam and electronic tracking. Malware attacks can result in spoofing, data interception, data theft, backdoor access, unauthorized network access, service abuse, and impact the availability and integrity of the data.

The architecture should enable controls and policies that can not only prevent many of these attacks, but also limit the damage in case of a breach. Policies that enforce rules such as authentication, strong passwords, and password changes at periodic interval, disabling services that are not required are critical in ensuring the security.

The use of prevention and detection software in the endpoint devices is critical. These include Data encryption capabilities, firewall, antivirus, intrusion detection, anti-spam, remote diagnostic and auditing software.

If a device is stolen the ability to disable services, lock, wipe out sensitive data is important to limit the damage and is typically performed by MDM solutions, which are important components of the mobile architecture.

The gateway devices can be either hardware or software-based. With the advent of mobile devices, the gateway software is able to use many of the key features that are available in the mobile devices to transfer the device data from the medical device to the final destination. This also provides mobility and a single multipurpose device for the user as compared to a hardware-based gateway.

4.3 Network Security

The e-health network security architecture involves multiple layers of prevention, detection and response controls as the network spans through different types of networks. These include wireless, wired, enterprise, private and public networks.

The mobile security reference architecture published by the department of Homeland security [14] calls out that the devices that use Wi-Fi and cellular network communications are more accessible and exposed than hardwired devices.

The wireless network architecture must consider the protection from various network based threats such as data interception over air, data interception over the network, manipulation of data in transit, connection to untrusted service, jamming [15] and flooding.

The architecture must enable the tuning of quality of service, which can vary based on the devices and the functionality that is used. For example a pulse oximeter generates text data, which is largely transparent to the delays in the network. In contrast stethoscope audio streaming can be extremely sensitive to network delay.

4.4 Cloud-based Application Access Security

e-health catering to multiple actors and patients constitutes one of the most important and largest segments of Health Care. With the patient base being used to accessing consumer-focused applications from anywhere, cloud models that support anywhere and seamless access are important to ensure user adoption of e-health. The architecture must support different models of on boarding, which includes managed-care as well as self-service based models. This requires integration with multiple systems to enable seamless flow of information that is required to on board the device. All these features opens

Figure 4 e-health security building blocks stack

up security challenges and the architecture must provide protection against web-based threats such as phishing, drive by downloads, exploitation of vulnerable browsers to get access to applications and data.

Authentication, Authorization, and Accounting (AAA)[16] are basic pillars of secure mechanisms to enable secure access to resources in web applications. Secure policies play a key role in making the access control methods effective. For example, a hacker using dictionary-based attempts could access a system featuring AAA capabilities if a user uses a weak password. The architecture must enforce not only secure controls but also limit failed attempts to access the system using lockout schemes.

The security building blocks can be stacked in multiple layers from the physical security aspects, internal application and data protection to the secure interface access controls. The building blocks for application security are shown below.

Open Web Application Security Project (OWASP) publishes the top ten web security flaws at its website. The 2013 list includes the following [17]:

1. A1 Injection
2. A2 Broken Authentication and Session Management
3. A3 Cross-Site Scripting (XSS)
4. A4 Insecure Direct Object References

5. A5 Security Misconfiguration
6. A6 Sensitive Data Exposure
7. A7 Missing Function Level Access Control
8. A8 Cross-Site Request Forgery (CSRF)
9. A9 Using Components with Known Vulnerabilities
10. A10 Unvalidated Redirects and Forwards

While the overall architecture can protect from security threats, it is extremely critical for web application to ensure above-mentioned security vulnerabilities does not exist. A proper vulnerability patching and software upgrade policy must be adhered to ensure that security risks are mitigated on a on going basis.

4.5 Data Storage Security

The architecture that enables security for data at rest must take an expanded view outside of securing the physical storage, since there are multiple dependencies that can result in weak points, which can be used as an entry point to access the data. A compromised application can be used to access the data, or a backup disk can be used to get access to data. The key considerations must include data encryptions (at application level, file level, disk level), access rights (physical, application, user), context based access and alternate path access (backup data disk)

The federal information processing standards publication 140–2 [18] lists the security requirements that need to be satisfied by a cryptographic module utilized within a security system protecting sensitive information and defines four qualitative levels of security.

Level 1: Lowest level of security and allows the software and firmware components of a cryptographic module to be executed on a general purpose computing system.

Level 2: Enhances the physical security mechanisms by adding the requirement for tamper evidence. This level requires role-based authentication and authorization of an operator to assume specific role and functions.

Level 3: Enhances security to prevent intruder from gaining access to critical security parameters using an identity based authentication mechanisms.

Level 4: Highest level of security and provides a complete envelope of protection around the cryptographic module with the intent of detecting and responding to unauthorized attempts to physical access.

4.6 Enterprise Application Access Security

To enable applications to be accessed from anywhere, more and more e-health applications are hosted on the cloud. This eliminates boundaries for access, at the same time creates security challenges. However many enterprise applications exists within the enterprise and is used to store sensitive information. The sensitive personal health information data storage must be clearly separated from the external web access.

The architecture can involve multiple considerations including DMZ with double firewall protection, reverse proxy to separate the access boundaries for enterprise applications and external applications.

Public Web servers enforce access to content using authentication and authorization schemes. These include basic authentication, Digest authentication, SSL/TLS based server and client authentication.

NIST guidelines for securing public web servers calls out some of the weakness of authentication schemes including SSL/TLS: "Several limitations are inherent with SSL/TLS. Packets are encrypted at the TCP layer, so IP layer information is not encrypted. Although this protects the Web data being transmitted, a person monitoring an SSL/TLS session can determine both the sender and receiver via the unencrypted IP address information. In addition, SSL/TLS only protects data while it is being transmitted, not when it is stored at either endpoint. Thus, the data is still vulnerable while in storage (e.g., a credit card database) unless additional safeguards are taken at the endpoints."[19] Hence, the architecture for e-health must account for multilayer security and follow the path of the data and ensure data is secured while stored, in transit and when handoffs happen.

4.7 Federated Secure Access to Partner Cloud Services

E-health applications use services and capabilities from different sources. This requires secure and seamless access to services provided by other partner clouds. Federated identity allows identities to be shared securely between applications both within and across organizational boundaries.

Protocols such as Security Assertion Markup Language[20](SAML) and WS-Federation, OAuth and OpenID Connect are commonly used standards for identity federation.

The architecture must support a federated identity management model using standard protocols to enable secure and seamless single sign-on into these services. The access to the services is commonly done through web

services, and hence the architecture must consider the security aspects for public exposed web services.

5 Implications to Standards

The standards landscape for eHealth-M2M-IoT and security is rather nascent. It includes IEEE for wireless, ZigBee Alliance, ITU-T for M2M e-Health Service Layer, Continua Alliance for Use Case profiles and best practices, NIST[21], as well as diverse government initiatives.

The Focus Group on the M2M service layer (FG M2M) was created in TSAG meeting (January 2012) and started in April 2012 [22]. It is expected to conclude its work in December 2013.

M2M (Machine to machine communication) covers very wide area and several standardization activities have already commenced its study in SG13, SG16 and oneM2M. In order to avoid duplicate works with them, FG M2M will focus initially on services and applications for e-health. The specific tasks of the Focus Group are to:

- Perform a "gap analysis" for vertical market M2M service layer needs, initially focusing on applications and services for the health-care market.
- Identify a minimum common set of M2M service layer requirements and capabilities, initially focusing on e-health applications and services.
- Study whether existing APIs and protocols satisfy the above requirements and capabilities to support a common M2M service layer between M2M applications and telecom networks.
- Draft technical reports describing and addressing the gaps and identifying future standardization work for ITU-T in the field of the M2M service layer.
- Support global harmonization and consolidation by inputting final deliverables to the parent Study Group and other relevant Study Groups as appropriate.
- Develop a living list of SDOs, forums and consortia dealing with M2M service layer APIs and protocols, including information concerning their activities and documents in the context of a common M2M service layer platform.

The following figure depicts the Reference Architecture used for the Continua Alliance:

Continua Alliance [23] is focused on establishing industry standards and security for connected health technologies such as smart phones, gateways

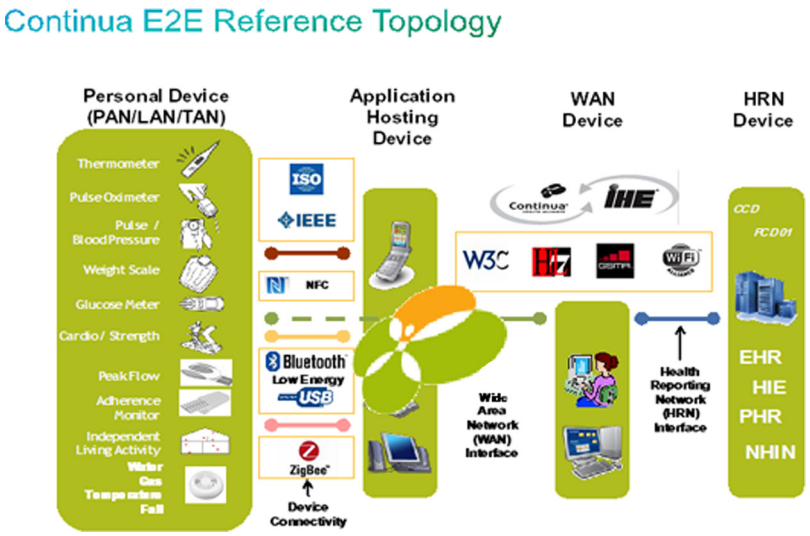

Figure 5 Continua E2E Reference Topology

and remote monitoring devices. Its activities include a certification and brand support program, events and collaborations to support technology and clinical innovation, as well as outreach to employers, payers, governments and care providers.

Multiple Standards exist in healthcare that is used in e-health applications and interactions between humans, devices, processes and applications. These standards can be classified in multiple classes including data standards, message standards, document standards, process standards. They can be syntax based, semantics based, relationship based, purpose based and classification based.

Some of the common data standards include

- ICD (International Classification of Disease) - International standard codes for diagnoses
- CPT (Current Procedural Terminology) - Standard for coding medical procedures
- LOINC (Logical Observation Identifiers Names and Codes) - Standard for Laboratory and clinical observations
- SNOMED CT (Systematized Nomenclature Of Medicine) - Hierarchical Healthcare Terminology
- NDC (National Drug Codes) - FDA's numbering system for medications

The International Classification of Diseases (ICD) is the standard diagnostic tool for epidemiology, health management and clinical purposes. This includes the analysis of the general health situation of population groups. It is used to monitor the incidence and prevalence of diseases and other health problems. ICD codes exist for diseases, signs and symptoms, abnormal findings, complaints, social circumstances, and external causes of injury or diseases. US Healthcare system is in the process of converting from ICD-9 to ICD-10.

IEEE 11073-20601-2008 [24] Standard addresses the need for an openly defined, independent standard for converting the information profile of personal health devices into an interoperable transmission format so the information can be exchanged to and from personal Tele-health devices and compute engines.

HL7 [25](Health Level System 7) provides a framework (and related standards) for the exchange, integration, sharing, and retrieval of electronic health information.

HL7 CDA [26] - Clinical Document Architecture is a XML based markup standard intended to specify the encoding, structure and semantics of clinical documents for exchange. It is based on the HL7 Reference Information Model (RIM) and the HL7 Version 3 Data Types. The purpose is enable exchange of clinical information. It can include multimedia content.

The CCR [27] document standard is used to allow timely and focused transmission of information to other health professionals involved in the patient's care. The CCR data set contains a summary of the patient's health status including problems, medications, allergies, and basic information about health insurance, care documentation, and the patient's care plan

The Continuity of Care Document [28](CCD) is an HL7 CDA implementation of the Continuity of Care Record (CCR). A CCR can be converted to CCD, but not vice versa.

DICOM [29] (Digital Imaging and Communications in Medicine) is a standard for handling, storing, printing, and transmitting information in medical imaging. It includes a file format definition and a network communications protocol. DICOM files can be exchanged between two entities that are capable of receiving image and patient data in DICOM format.

IHE [30] (Integrating the Health Enterprise) is an initiative by healthcare professionals and industry to improve the way computer systems in healthcare share information. IHE defines Integration Profiles, which describe a clinical information need or workflow scenario and document how to use established standards to accomplish it.

Cross-Enterprise Document Sharing [31](XDS) is focused on providing a standards-based specification for managing the sharing of documents between any healthcare enterprise, ranging from a private physician office to a clinic to an acute care in-patient facility and personal health record systems. This is managed through federated document repositories and a document registry to create a longitudinal record of information about a patient within a given clinical affinity domain.

Related to the standards landscape for eHealth-M2M-IoT Security are the various country regulations that define policies for security and privacy specific to health. It is important to monitor these regulations when developing and implementing architecture pertinent to patient safety.

6 Summary and Conclusions

The coming of age of eHealth is intrinsically linked to the successful deployment of a secure and privacy-preserving M2M/IoT infrastructure. The authors have proposed an architecture and framework that support the development and providing of solutions. The authors have further identified core standards and industry bodies where eHealth-M2M-IoT standardization is in progress. While comprehensive, the list is not exhaustive. In closing, we emphasize that security and privacy for eHealth in the emerging IoT landscape offers serious challenges as well as exciting opportunities to the industry.

References

1. http://www.who.int/trade/glossary/story021/en/
2. http://www.bccresearch.com/market-research/healthcare/preventive-healthcare-technologies-hlc070a.html
3. http://www.prweb.com/releases/biosensors/medical_biosensors/prweb8067456.htm
4. http://conferences.sigcomm.org/sigcomm/2012/paper/mcc/p13.pdf
5. “Health Action Plan 2012-2020: Innovative Healthcare for the 21st Century” - https://ec.europa.eu/digital-agenda/en/news/ehealth-action-plan-2012-2020-innovative-healthcare-21st-century
6. http://ec.europa.eu/health/ehealth/policy/network/index_en.htm
7. C. Perera, “The Evolution of e-Health – Mobilte Technology and mHealth”, http://articles.journalmtm.com/1.1.1-2%20Perera.pdf
8. http://www.foxnews.com/tech/2011/08/04/insulin-pumps-vulnerable-to-hacking/

9. "Body Sensor Networks"–Yang, Guang–Zhong (Ed) ISBN 978-1-84628-484-7
10. "The Vision of Autonomic Computing", http://ieeexplore.ieee.org/stamp/stamp.jsp?to=&arnumber=1160055
11. Trusted Computing Group – http://www.trustedcomputinggroup.org/
12. "An Architecture for Secure e-Health Systems" - http://www.tsb.upv.es/eventos/workshophealthcare/documentos/C2-2.pdf
13. W. Jansen, K. Scarfone, Guidelines on Cell phone and PDA security, NIST 800-124
14. https://cio.gov/wp-content/uploads/downloads/2013/05/Mobile-Security-Reference-Architecture.pdf
15. http://en.wikipedia.org/wiki/Wireless_signal_jammer
16. http://en.wikipedia.org/wiki/AAA_protocol
17. https://www.owasp.org/index.php/Top_10_2013-Top_10
18. FIPS 140-2 http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf
19. NIST Publication 800-44: Guidelines for securing public web servers
20. http://en.wikipedia.org/wiki/Security_Assertion_Markup_Language
21. http://csrc.nist.gov/publications/nistpubs/800-66-Rev1/SP-800-66-Revision1.pdf
22. http://www.itu.int/en/ITU-T/focusgroups/m2m/Pages/default.aspx
23. http://www.continuaalliance.org/
24. http://en.wikipedia.org/wiki/ISO/IEEE_11073_Personal_Health_Data_%28PHD%29_Standards
25. http://www.hl7.org/implement/standards/index.cfm?ref=nav
26. http://www.hl7.org/implement/standards/product_brief.cfm?product_id=7
27. http://en.wikipedia.org/wiki/Continuity_of_Care_Record
28. http://en.wikipedia.org/wiki/Continuity_of_Care_Document
29. http://medical.nema.org/standard.html
30. http://www.ihe.net/
31. http://wwww.ihe.net
32. http://kafka.apache.org/index.html
33. http://www.slideshare.net/MaxAlexejev/modern-distributed-messaging-and-rpc

Biographies

David Lake is a consulting engineer in the R&D Group at Cisco. He has more than 20 years of network design and deployment experience, ranging from X.25 and SNA, through the era of multiprotocol routing to IP. He has extensive experience in transporting rich-media technologies across complex enterprise and service provider networks. He is an editor and contributor to the Management and Orchestration (MANO) Working Group with ETSI's Network Function Virtualistion group.

Rodolfo Milito, PhD
Senior Technical LeaderENG Labs
Cisco Systems, Inc Rodolfo Milito, a senior Technical Leader with the ENG Labs of Cisco Systems, is currently engaged in IoT, Big Data and Analytics, and working on a distributed compute, storage, and network platform from the edge to the core of the network nicknamed "Fog Computing".

Rodolfo got his PhD in EE (Control Systems) from UIUC, joined Bell Labs in Holmdel in 1985, and AT&T Labs after the 1996 Lucent spin-off. In 1999 he moved to XStream Logic, a startup in the Silicon Valley, later co-founded ConSentry Networks, and joined Cisco in 2008. Rodolfo's career has straddled research and development in the areas of network design, processor architectures, performance characterization, adaptive control, and algorithms aimed at improving the performance and securing communication networks (load distribution, routing, resource sharing, overload control, and malware detection). He has published extensively, and holds 11 US patents.

Monique Jeanne Morrow CTO Cisco Services
Cisco Systems, Inc. Email: mmorrow@cisco.com

Summary

Monique Morrow holds the title of CTO Cisco Services. Ms. Morrow's focus is in developing strategic technology and business architectures for Cisco customers and partners.

With over 13 years at Cisco, Monique has made significant contributions in a wide range of roles, from Customer Advocacy to Corporate Consulting Engineering. With particular emphasis on the Service Provider segment, her experience includes roles in the field (Asia-Pacific) where she undertook the goal of building a strong technology team, as well as identifying and grooming a successor to assure a smooth transition and continued excellence.

Monique has consistently shown her talent for forward thinking and risk taking in exploring market opportunities for Cisco. She was an early visionary in the realm of MPLS as a technology service enabler, and she was one of the leaders in developing new business opportunities for Cisco in the Service Provider segment, SP NGN.

Monique holds 3 patents, and has an additional nine patent submissions filed with US Patent Office.

Ms. Morrow is the co-author of several books, and has authored numerous articles. She also maintains several technology blogs, and is a major contributor to Cisco's Technology Radar, having achieved Gold Medalist Hall of Fame status for her contributions.

Monique is also very active in industry associations. She is a new member of the Strategic Advisory Board for the School of Computer Science at North Carolina State University.

Monique is particularly passionate about Girls in ICT and has been active at the ITU on this topic - presenting at the EU Parliament in April of 2013 as an advocate for Cisco.

Within the Office of the CTO, first as an individual contributor, and now as CTO, she has built a strong leadership team, and she continues to drive Cisco's globalization and country strategies.

Rajesh Vargheese is the CTO for Cisco Healthcare solutions business unit at Cisco. Rajesh leads the technology team that designs, develops and deploys healthcare solutions and leads the effort in defining the strategy, roadmap and architectures for Cisco Healthcare solutions. Rajesh has 18 years of experience in thought leadership, defining architectures, product and solution development, system integration, marketing and messaging, deploying and problem solving in collaboration, video, cloud, real time applications, big data analytics and security for vertical domains. Rajesh currently focuses on the healthcare vertical solutions and brings together his broad experience in collaboration, infrastructure, workflows, clinical system integrations and cloud architectures in developing and shaping the Cisco Healthcare Solutions. Rajesh works with C levels and IT leadership to understand and provide innovative solutions to customer challenges.