VueForge®

Seven principles for achieving Security and Privacy in a world of Machine-Driven Big Data

altran

1 Summary

The increasing prevalence of IoT (Internet of Things) and industrial Big Data technology and solutions has been accompanied by an increasing concern about security. Examples include November 2015 reports suggesting attacks on medical devices such as pacemakers and insulin pumps may represent the biggest cyber security threat for 2016 [1], and the October 2015 TalkTalk data breach in which 155,000 customer personal details were stolen [2]. Addressing security is essential for the successful exploitation of the huge value opportunity from IoT and industrial Big Data.

> "Addressing security is essential for the successful exploitation of the huge value opportunity from IoT and industrial Big Data."

This white paper first summarises the reasons why security is important and the specific security challenges to be overcome, then focuses on how to address IoT and industrial Big Data security by describing the key principles to be addressed. The seven principles are as follows:

1. KNOW YOUR ENEMIES
Understand the security risks posed by a deployment and form a comprehensive policy to deal with them.

2. TAKE SECURITY TO THE EDGE
Address security from end devices to central services, and from initialisation to disposal.

3. KNOW WHAT YOU'RE TALKING TO
Understand the identities, roles and authorisations of people and equipment. Address the provisioning of new identities, maintenance, change of ownership, and withdrawal of trust.

4. CREATE A STRONG NETWORK
Ensure communications are resilient and resistant to attack.

5. DON'T TRUST IT, WATCH IT
Monitor behaviour for signs of attack, don't rely on fixed defences. Use SIEM (Security Information and Event Management) techniques including advanced analytics.

6. BUILD IT RIGHT
Minimise the vulnerabilities exposed to an attacker. Use security-oriented architecture, separation of security domains, highly-assured software and hardware components, and generate assurance evidence during development.

7. BASE ON FIRM FOUNDATIONS
Use trustworthy services for communications, computation, storage and management.

About the author: Dr David Jackson CEng MIET FBCS is Global Technical Director for Intelligent Systems Security within Altran, and lead architect for VueForge®: Altran's framework for Machine-Driven Big Data solutions.

Contents

2 Context: the Emergence of Machine-Driven Big Data

Machine-Driven Big Data (MDBD) represents the combination of two major technological trends – the Internet of Things (pervasive sensing and connectivity) and the Big Data revolution (sophisticated analytics carried out on data sets exhibiting high volume, velocity and variety using standardised IT infrastructure and services, including the Cloud).

MDBD is offering new capabilities and experiences in both personal and professional contexts. For example:

⟶ CONNECTED VEHICLES AND EQUIPMENT can achieve higher performance and lower cost through predictive maintenance and continuous monitoring;

⟶ PRODUCT DEVELOPERS can gain insight into the day-to-day use of their products and target development accordingly;

⟶ OPERATORS OF COMPLEX INFRASTRUCTURE (such as power distribution or transport) can gain real-time visibility of the status of their assets, and the well-being of their staff.

Security will be an important part of these capabilities and experiences – a manufacturer does not want to have details of process or workload available to the world, nor does a private driver want their movements to be publically distributed.

The emergence of MDBD poses particular challenges associated with security, privacy and data sovereignty. For example:

⟶ CONNECTED DEVICES (Such as baby monitors, smart TVs [3], and even refrigerators) capture detailed personal data about individuals and their activities;

⟶ ANALYTICS CAN TAKE DATA FROM MULTIPLE SOURCES and can derive significant personal data from apparently loosely related inputs [4];

⟶ CLOUD COMPUTING INFRASTRUCTURES can relocate data (or establish responsibilities) between organisations and data centres that may be in different legal jurisdictions, with controversial consequences [5].

Maintaining the security of data, together with privacy of individuals and compliance with legal and corporate obligations, becomes increasingly difficult as systems expand and become more interconnected. The success of MDBD relies on trustworthy management of data throughout the end-to-end value chain – from machines in the field to all relevant applications in the data centre or the cloud.

3 Why is Security Important?

Security is important because information is an asset to be protected, and consumers, industry and legislation are demanding this protection.

3.1. INFORMATION AS AN ASSET

"Security attacks have the potential to cause significant economic, physical or environmental damage."

Addressing security and privacy starts with the assets – the data and information – at risk.
A traditional security analysis will consider the types of data held and the environment in which they are held. For example, a 'lone worker assistance' system might hold worker identity, expected work schedule and the identities of clients to be visited.

This traditional approach becomes less effective with the introduction of big data analytics because of the ability of analytic systems to detect patterns and correlations between disparate data sets: the value of a single set of information cannot be understood in isolation. It has become possible to violate a person's privacy by estimating their private characteristics with high accuracy without any single uniquely identifiable characteristic being required [4]. This application of aggregated intelligence has long been understood and used in the defence sector; now advertising agencies and individuals have gained the same capability.

The assets we need to protect in an MDBD deployment may be much more than just information – physical assets and people may be placed at risk by attacks because of the close links between MDBD and operational systems which have the potential to cause significant economic, physical or environmental damage.

MDBD IMPACTS REAL-WORLD OPERATIONS ACROSS ALL INDUSTRIES

3.2. CONSUMER AND INDUSTRY EXPECTATIONS

The information contained in a system has value to its owner, to third parties related to its owner, and to potential attackers. The consequences of security failure can be life-changing for an individual. For example, the Cifas fraud prevention service identifies over 100,000 cases of identity fraud per year in the UK; each involving some loss of time or money to the victim or to the institutions involved [6]. High profile attacks on organisations can threaten their very existence [7].

In considering the impact of security failures desired properties can be identified in terms of information assets and the qualities that they exhibit. A commonly used classification ("CIA") identifies:

⟶ CONFIDENTIALITY : information is not visible to unintended parties;

⟶ INTEGRITY : information is maintained without corruption or unauthorised change, and;

⟶ AVAILABILITY : information is accessible to the intended users in the timescales and format they expect.

While appropriate for traditional environments such as companies and governmental organisations with well-defined information assets, for MDBD a more broad-ranging set of principles is required to address the more complex network of relationships and responsibilities.

One industry collaboration [11] produced a set of principles including, for example:

⟶ Consumers should own their own data;
⟶ Not all data is equally sensitive;
⟶ Consumers must have confidence in how their data is used, stored, and transported.

Similar principles can be considered a 'customer centric' view of high-level security and privacy requirements. A satisfactory response to these needs is essential to realise the promise of MDBD.

Note also that in many cases, IoT and MDBD deployments are subject to stringent constraints arising from the safety, environmental and commercial environment in which they work – the Things in the Internet of Things may be machines critical to the efficient operation of the activities they support, and need to be protected from interference and information leakage. In environments such as nuclear energy generation, air traffic control, or defence facilities, it may ultimately be determined that no risk mitigation can achieve the necessary level of security assurance, and instead resort to solutions with minimal connectivity.

3.3. LEGAL AND REGULATORY EXPECTATIONS

MDBD security and privacy must also take into account that many systems will operate, or at least be constructed, in an international environment, meaning that those involved will have to comply with a range of regulatory expectations.

These will arise at a number of levels:

⟶ TRANS-NATIONAL AGREEMENTS (eg Directives and Regulations within the EU, international treaties);

⟶ NATIONAL LAW (potentially relating to a number of jurisdictions depending on the location of the parties involved and the location in which data is held), and;

⟶ REGULATORY REQUIREMENTS IN ADDITION TO LEGISLATION, for example specific controls on data management in healthcare (HIPAA in USA, corresponding regulation in France), or requirements that support safety or environmental concerns.

These many constraints can conflict at many levels. For example:

⟶ Under US law, THE MAIN RIGHTS to data lie with the organisation that collects it; in European law, the subject of the data has rights over its use and management.

⟶ CONFIDENTIALITY may be a duty in many cases, but many jurisdictions control the extent to which measures such as encryption can be applied or exported.

⟶ Many jurisdictions enforce REQUIREMENTS FOR DISCLOSING INFORMATION for law-enforcement, a need in direct conflict with the desire to protect confidentiality & privacy.

This is also a rapidly evolving field. Recent events include:

⟶ THE EUROPEAN COURT OF JUSTICE striking down the 'safe harbour' agreement under which much data on EU nationals could be handled without restriction in the USA [5].

⟶ MICROSOFT IS CURRENTLY APPEALING A RULING that it must hand over data held in Ireland in response to a subpoena served in the US. The outcome of this appeal could set a wide-ranging precedent about access to data held outside the US by US-controlled corporations [8].

⟶ CALLS FROM PROSECUTORS FOR LIMITATION OF 'FULL DISK ENCRYPTION' ON SMARTPHONES [9], followed by controversy and a statement that the US authorities will not seek built-in 'back doors' in security products [10].

Such activities occur against a background of a continued stream of security breaches in the news.

4 Challenges to be Overcome

A number of challenges constrain available approaches to achieving security and privacy in MDBD deployments. Some of these are systemic – they apply to the deployment and the problem as a whole - while others relate to the specific technical domains involved:

⟶ PERVASIVE CONNECTED OBJECTS (IOT);
⟶ CONVENTIONAL INTERNET NETWORKING (THE I IN IOT), AND;
⟶ MASSIVE-SCALE DATA TECHNOLOGIES FOR STORAGE, PROCESSING AND ANALYSIS.

Each of these aspects gives rise to constraints on how the security & privacy objectives may be achieved.

4.1. CHALLENGES FROM SYSTEMIC ISSUES

> "Market success of technologies and products is usually driven by what they enable, not by what they restrict or prevent. technologies and products is usually driven by what they enable, not by what they restrict or prevent."
>
> Ralph Langner

The issues that arise from the nature of an MDBD deployment as a whole include:

⟶ SECURITY is necessarily an emergent property of a whole system – if the security of any part of a system is weak, the system as a whole is vulnerable.

⟶ THE NUMBER OF DISTINCT ELEMENTS AND TYPES OF INTERACTIONS can give rise to significant complexity – few tools (and few people) are able to address this complexity.

⟶ THE ENVIRONMENT in which a system operates and evolves is constantly changing, both in terms of internal technologies and particularly in terms of the external threat landscape.

⟶ THE EXPECTATIONS OF THIRD PARTIES (including the regulatory and legal aspects described above) are complex and potentially conflicting. For example, what set of security requirements does an internationally-distributed system with access to medical device data and financial information have to meet?

⟶ ULTIMATELY, FEW SYSTEMS ARE PROCURED SPECIFICALLY FOR THEIR SECURITY PROPERTIES – the owners and users are seeking new functionality, not new restrictions.

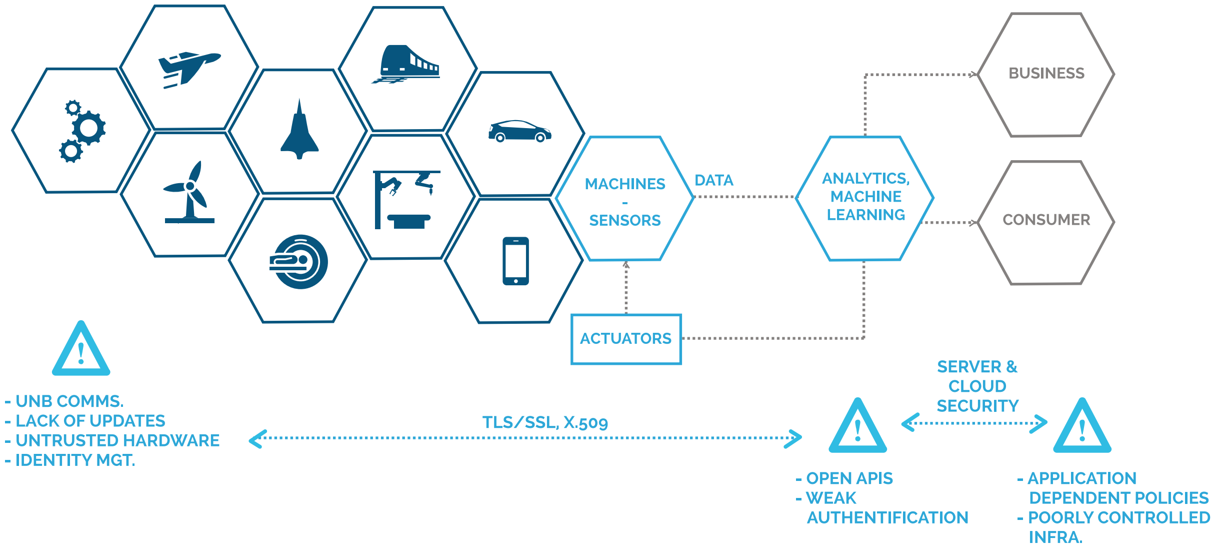

THE MDBD VALUE CHAIN NEEDS TO BE PROTECTED AT ALL STAGES, FROM MACHINES TO USERS

4.2. CHALLENGES FROM PERVASIVE CONNECTED OBJECTS

There are specific technical obstacles raised by the nature of a pervasive, connected system:

> "Imagine the challenge from the engineer's point of view. The IoT is asking an embedded designer with no security experience to protect a $1 MCU against all threats of the Internet"
>
> Mike Muller, CTO ARM®

→ **RESOURCES OF IOT NETWORKS AND DEVICES ARE OFTEN LIMITED** – for example in bandwidth, in computational capacity, in memory - which then limits the extent to which sophisticated cryptography or complex protocols can be used to address security;

→ **IOT NETWORKS** will often incorporate gateways, intermediaries, or differences in communication method that limit the ability to provide end-to-end security;

→ **THERE MAY NOT BE AN OBVIOUS (OR ACCESSIBLE) CENTRAL AUTHORITY** that can be relied on to manage a centralised security architecture;

→ **PROVISION OF ALL IOT** devices with necessary identity and key material may not be easy;

→ **MAINTENANCE ACTIVITIES** may be restricted; update and patching may not be feasible;

→ **NETWORK CONNECTIONS** may be transient and dynamic, for example as a mobile device disconnects and reconnects perhaps through different carriers.

Furthermore, many connected sensors or actuators will be physically exposed to attack, by virtue of being deployed in public or unsupervised spaces, or by using publicly accessible networks (including wireless networks).

4.3. CHALLENGES FROM INTERNET NETWORKING

Inherent in the "Internet of Things" is the use of Internet communication technologies. In spite of the widely reported risks of online security breaches, these technologies form a model which supports millions of adequately-secure financial transactions per day.

The major technical features of this model are:

⟶ EDGE DEVICES (PCs, tablets & smartphones) which are assumed to be secure. These devices are typically not themselves authenticated – it is the user, not the device, whose identity is checked by the application.

⟶ SERVERS ARE AUTHENTICATED by a public key cryptography system (PKI).

⟶ DEVICES AND SERVERS ARE PROTECTED from general network access by restrictions on traffic, for example using firewalls.

⟶ DATA IN TRANSIT IS ENCRYPTED using the TLS/SSL (Transport Layer Security / Secure Sockets Layer) family of protocols.

⟶ SERVICES AND SERVERS ARE IDENTIFIED using the Domain Name System (DNS), preferably in its secured form (DNSSEC).

⟶ HOSTING AND CONNECTIVITY IS PROVIDED BY INTERCONNECTED SUBNETWORKS MANAGED by many individual operators. Peering and transit arrangements between these networks govern the actual delivery of data.

There are a number of limitations in applying this model directly to MDBD (many related to the challenges of pervasive connected objects):

⟶ NECESSARY SECURITY CONSTRAINTS AND POLICIES are expressed at the level of complete applications or ecosystems but the necessary controls must be implemented on individual components.

⟶ DEVICES WITH LIMITED MEMORY AND COMPUTING RESOURCE (or electrical power) may not be able to implement the communication and cryptography necessary to maintain good security practices.

⟶ COMMUNICATION WITH IOT/M2M (Machine to Machine) devices is often limited in bandwidth (eg UNB radio) and availability or reliability; such devices cannot make use of mechanisms that may require substantial volumes of security-related data to be transmitted.

These limitations are sometimes a matter of implementation priorities but are in many cases fundamental consequences of the nature of the devices & data involved.

4.4. CHALLENGES FROM MASSIVE-SCALE DATA TECHNOLOGIES

The creation and operation of large-scale data processing infrastructure, even without the problems of pervasive sensing & communication, sets challenges.

> "Traditional security mechanisms, which are tailored to securing small-scale static (as opposed to streaming) data, are inadequate."
>
> CSA BiG Data Top Ten [16]

Typical challenges include:

→ **PRESERVING INDEPENDENCE BETWEEN APPLICATIONS AND DATA OWNERS** can be difficult in shared cloud and managed hosting facilities with multiple tenants and multiple management structures.

→ **THE SECURITY AND MANAGEMENT PRACTICES FOR NEW BIG DATA TOOLS ARE NOT MATURE** – IT security policies and mechanisms derived by traditional RDBMS structures may not be sufficient.

→ **MAINTAINING SECURE ACCESS TO SECURITY LOGS** and event reports will be difficult in large and potentially distributed infrastructures.

→ **ENSURING DATA INTEGRITY** through input validation and filtering will be difficult in large distributed systems.

Good examples of practice are provided by organisations such as the Cloud Security Alliance (CSA, [16]).

5 Addressing the Challenges

The previous sections described key elements of the MDBD security problem and constraints on the solution. This section describes seven key principles required for achieving adequate privacy and security in MDBD deployments.

These principles are:

1. KNOW YOUR ENEMIES
Understand the security risks posed by a deployment and form a comprehensive policy to deal with them

2. TAKE SECURITY TO THE EDGE
Address security from end devices to central services, and from initialisation to disposal

3. KNOW WHAT YOU'RE TALKING TO
Understand the identities, roles and authorisations of people and equipment

4. CREATE A STRONG NETWORK
Ensure communications are resilient and resistant to attack

5. DON'T TRUST IT, WATCH IT
Monitor behaviour for signs of attack, don't rely on fixed defences

6. BUILD IT RIGHT
Minimise the vulnerabilities exposed to an attacker

7. BASE ON FIRM FOUNDATIONS
Use trustworthy services for communications, computation, storage and management

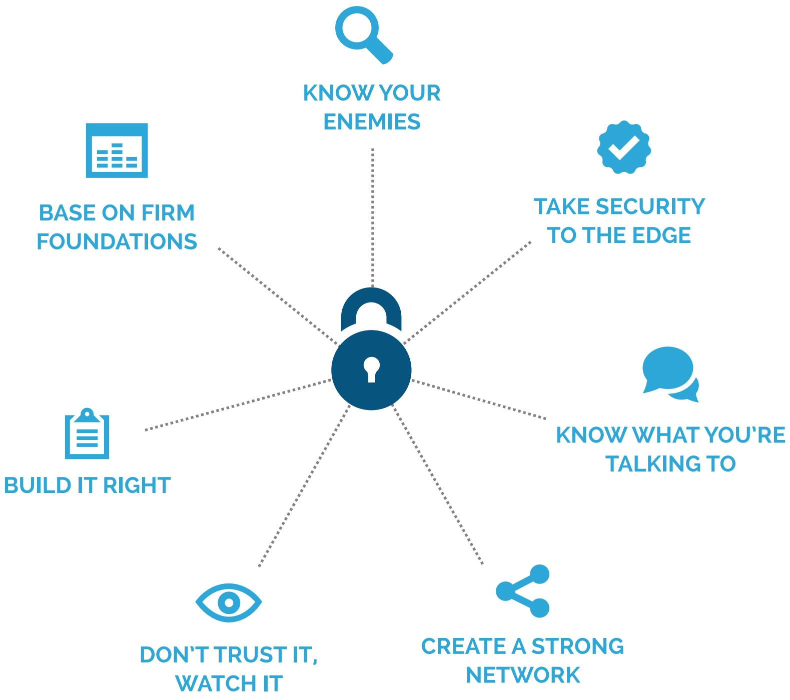

ACHIEVING ACCEPTABLE PRIVACY AND SECURITY REQUIRES ADDRESSING SEVEN KEY PRINCIPLES

5.1. KNOW YOUR ENEMIES

Understand the security risks posed by a deployment and form a comprehensive policy to deal with them.

> "I consider ignorance the biggest danger to the human race's progress."
>
> Sir Hossein Yassaie

To achieve security and privacy we must set reasonable goals in the form of a Security Policy which are achievable and which deliver the necessary levels of confidence and trust to all the stakeholders involved. This Security Policy must address all the stakeholders in a deployment, and all aspects of development and operation – end-to-end and cradle-to-grave.

The most fundamental question to address is to determine the value of security in a given application – how much security should be applied? This will generally be determined by the risk presented by the system and the risk appetite of those responsible for its procurement and operation. Many different models [eg 12, 13, 14] have been proposed for assessment of security risk. They share common elements in combining assessments of the value of an asset, the vulnerability of the systems which contain or protect it, and the attractiveness of the asset to attackers to provide a measure (typically qualitative) of the overall risk.

Risk management is challenging in the context of MDBD because of the number, range, and potential changes in time of stakeholders involved – any of the parameters of risk identified in the previous paragraph may change when a new service connects to an existing data source, or when new devices are added to an installation. In particular, policies and controls established for a typical IT environment may not be sufficient when the IT is connected to operational equipment.

The Security Policy captures the ability and desire of an organisation to accept security risk, and the measures it takes to manage that risk. The key activities of any risk management scheme (see diagram) are applicable to security, and should be adopted.

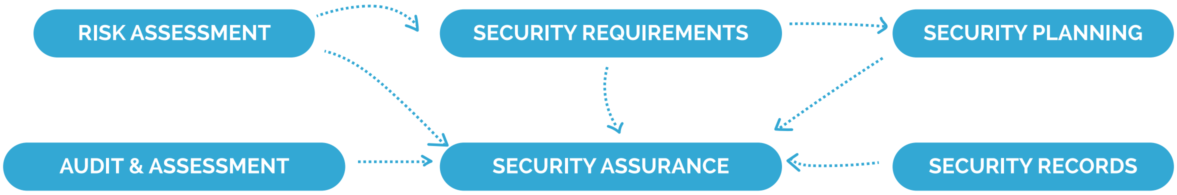

RISK MANAGEMENT ACTIVITIES ARE ESSENTIAL TO "**KNOW YOUR ENEMIES**" AND PROVIDE SECURITY ASSURANCE

The contents of the Security Policy will determine the types and level of the security controls to be applied. These controls typically combine a number of elements:

- ⟶ FUNCTIONAL REQUIREMENTS, OR SECURITY CONTROLS - features of the system which are required to maintain security;
- ⟶ ASSURANCE REQUIREMENTS - steps taken to ensure that the components and construction of the system is adequate (See '**Build it right**' below), and;
- ⟶ SECURITY OPERATIONS - activities undertaken to ensure the system maintains its security in operation: monitoring and responding to attacks, enforcing process controls through audit, and managing change to a system and its environment.

The security controls should cover AAA - Authentication, Authorisation and Accounting (or administration).

5.2. TAKE SECURITY TO THE EDGE

Address security from end devices to central services, and from initialisation to disposal.

> "Partial strength produces total weakness."
>
> Sir Robert Seppings

A security implementation needs to be comprehensive if it is to be secure. The constraints of MDBD deployments described above pose particular challenges in ensuring that security is maintained across the whole physical and logical scope of a deployment (embedded sensors or actuators to servers to client desktops) and across the whole lifespan (which may last decades, and involve replacement of every component through time).

Appropriately trustworthy software environments can be implemented across a networked deployment provided that the hardware is sufficiently robust and has the necessary performance. Achieving this on embedded end nodes will require solutions which are lightweight (so as to support devices with little memory, processing power, or communications bandwidth) and low in defects (as defects in security functions are potential vulnerabilities).

In some cases, efficient implementation of standard PKI technology (ie TLS/SSL) will be appropriate; in other cases custom protocols for specific applications may be sufficient (eg shared secret cryptography for local networks). Hardware support for Roots of Trust will increasingly be required to support authentication, encryption, identity management and the establishment of trusted execution environments.

Maintaining security through time requires consideration of the construction of the system and its components (see '**Build it right**' below) and also addressing deployment, failure scenarios, maintenance and disposal. Ensuring secure initialisation is a particular challenge: the machines connected to an MDBD deployment will typically be distributed geographically or across an organisation, and thus exposed to disruption and interference, and will generally need to be updated or replaced within the lifetime of a system.

Provisioning processes which establish equipment identities and relationships (see '**Know what you're talking to**' below) require a trustworthy environment to be established at "power on", and maintained in the face of interference throughout operation.

Currently there are semiconductor designs (such as ARM's TrustZone and Imagination's Omnishield) capable of providing the hardware functionality needed to **Take security to the edge**, but standards to govern their deployment and the policies and software stacks needed to exploit them are not yet widely adopted.

5.3. KNOW WHAT YOU'RE TALKING TO

Understand the identities, roles and authorisations of people and equipment. Address the provisioning of new identities, maintenance, change of ownership, and withdrawal of trust.

Implementing any practical security control requires establishing identity – of people or roles, of machines and of services. PKI (Public Key Infrastructure) systems provide the basic mechanisms for describing and communicating identity in the digital domain, but exploiting these to define a solution for an MDBD application is not trivial.

We can reduce the threat of untrusted devices on our network by authenticating them and the software that they execute. However, authentication of the device will depend on a certificate installed in a controlled manner. As the machines will be physically exposed to attack (and probably communicating over wireless) there needs to be a degree of protection available on the device, perhaps achieved by hardware.

Even the concept of "identity" has questions associated with it – what happens when a device is sold or transferred? When an intermediate (eg network) supplier changes? Or even simply if a data connection is broken as a result of roaming and a new IP address is allocated? Questions of ownership are also fundamental: who owns & operates the policies that link device behaviour to trustworthy management of data? For example, we can establish a secure certificate store on a microprocessor designed by one company, fabricated by a second and installed in a controller by a third for a car built by a fourth and sold through a fifth – but who owns that certificate?! Answers to these questions will emerge over time as industry standards and practices develop, but in the meantime early adopters of MDBD technology need to define solutions that meet their business needs today.

5.4. CREATE A STRONG NETWORK

Ensure communications are resilient and resistant to attack.

Impairing the capacity of a system to respond is an element of many attacks, typically as a DoS (Denial of Service) attack; this includes potential disruption of communication and routing mechanisms.

To maintain security and privacy an appropriate level of resistance (to overload, but also to interception and misdirection – "hacking") is required of communications infrastructure.

Achieving consistent network security is a challenge in industrial applications because of the varied networks and protocols used. Security solutions that are applicable to TCP/IP connections may not be practical across CAN (Controller Area Network) or Modbus industrial networks. (The accessibility of CAN is a factor in the much-publicised JEEP attack [17].) Even where a single standard is used, connectivity may be provided through several suppliers (eg fixed vs mobile carriers).

THE INTERNET OF THINGS

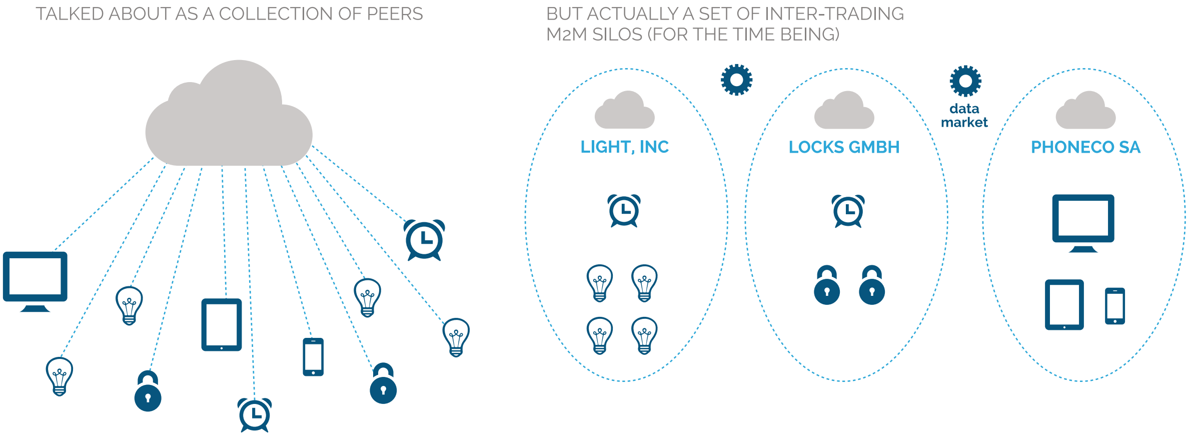
TALKED ABOUT AS A COLLECTION OF PEERS
BUT ACTUALLY A SET OF INTER-TRADING M2M SILOS (FOR THE TIME BEING)
LIGHT, INC
LOCKS GMBH
data market
PHONECO SA

INDUSTRIAL INTERNET APPLICATIONS WILL BE CREATED BE CONNECTING DIVERSE NETWORKS

5.5. DON'T TRUST IT, WATCH IT

Monitor behaviour for signs of attack, don't rely on fixed defences. Use SIEM (Security Information and Event Management) techniques including advanced analytics.

Due to their open and expansive nature, MDBD deployments must not expect to be able to prevent all security-related events or incidents but must have means of identifying them and responding (or allowing management action in response). Existing IT solutions for security monitoring and audit, including Security Information and Event Monitoring (SIEM) technologies may be applicable, but risk being swamped by the volume and variability of the devices connected in support of MDBD.

The incorporation of advanced analytics into SIEM systems will facilitate this – the monitoring and status information from a deployment is itself data that can be analysed, used to build models, and searched for anomalous behaviour. Security actions such as discarding data or disabling devices can be triggered by the analysis.

5.6. BUILD IT RIGHT

Minimise the vulnerabilities exposed to an attacker. Use security-oriented architecture, separation of security domains, highly-assured software and hardware components, and generate assurance evidence during development.

No security architecture or policy will achieve its goals if the fundamental components of the implementation fail to implement their expected functions or leave vulnerabilities open to exploitation by attackers. There are a number of technical measures which should be used:

⟶ **SYSTEMS SHOULD BE DESIGNED, DEPLOYED AND MAINTAINED IN ACCORDANCE WITH ESTABLISHED BEST PRACTICES FOR SECURE IMPLEMENTATION** – patterns of secure development and operation should be shared and implemented (see [15] for one example guide).

⟶ **IN PARTICULAR, THE INTERFACES (HARDWARE OR SOFTWARE) USED TO CONSTRUCT AND ACCESS MDBD SERVICES SHOULD BE SECURE** (demand authorisation) and should be used securely (not abused). Ensuring that APIs have appropriate authentication controls is a technological problem; preventing people abusing APIs to avoid such controls is more a question of training, economy and business relationships.

⟶ **USE OF SECURE PLATFORMS AND MIDDLEWARE.** The co-ordinating and integrating software (typically referred to as an "IoT Platform" or "M2M Platform") must be secure. It must support the security functions necessary to support the system's policies (see '**Know your enemies**') – for example encryption and authentication – but should expose a minimum of new vulnerabilities.

⟶ **USE SECURE SOFTWARE DEVELOPMENT PRACTICES.** Any software which is deployed in an MDBD environment should be developed to appropriate security standards (for example using the Common Criteria). The development lifecycle should include relevant review and verification and tools should be used to detect errors and generate assurance evidence. Security-focussed verification should be used to identify vulnerabilities. Development staff should be trained and supported in adopting secure practices.

⟶ **SEPARATION OF SECURITY DOMAINS.** To minimise the impact of attacks or failures, use of defence-in-depth (separate layers of protection arranged so that a breach in one does not result in a widespread attack) or security based on controls between zones should be considered. Such separation is increasingly achieved by virtualised infrastructure, ie hypervisors and software defined networks negotiating between containers, virtual machines and VPNs. Such technologies are becoming available for the IoT domains (eg with Altran's new open platform for electrical and electronic architectures [18]) but their implementation and adoption is only beginning.

5.7. BASE ON FIRM FOUNDATIONS

Use trustworthy services for communications, computation, storage and management.

"**Build it right**" addresses the technical aspects of vulnerability avoidance. There are also procedural and organisational measures that should be taken. These are particularly applicable to dependencies such as hosting or PaaS services supporting centralised storage and processing which are likely to be bought as services rather than built for a specific deployment:

→ SECURE CONFIGURATION AND DEPLOYMENT OF SOFTWARE AND INFRASTRUCTURE. All tools and IT/ communications infrastructure must be deployed in a manner which respects the security and privacy requirements of the application. They must support necessary authentication, authorisation and administration functions and not needlessly present further vulnerabilities. Within the MDBD domain, more technology may be available (the whole world of data centres and IT security) but few policies yet exist on how to protect & manage the data – and the security features of standard tools such as Hadoop are as yet weak. This is also a domain where national and international regulation on privacy, sovereignty, and data protection becomes key – who owns what and how is it controlled? How can permission be withdrawn?

→ ASSURANCE IN THE SUPPLY CHAIN, ESPECIALLY FOR SOFTWARE. Because no single entity will provide or operate a complete MDBD solution, a supply chain is inevitably introduced. Although many of the operators in such a supply chain (eg IT providers) may be expected to understand and respect security and privacy concerns, not all providers will have experience of the implications of failure in a MDBD context. Suppliers will have to be encouraged to provide relevant assurance of the fitness for purpose of their products (eg including penetration testing and 3rd party audit) and in some cases assistance will have to be given in understanding the application domain (the 'thingness' of the Things in the IoT).

→ TRUSTWORTHY HOSTING SERVICES. Ultimately all MDBD deployments will depend on communication, storage, and computation infrastructure. Without assurance of the trustworthiness of these services, little else can be assured. Note that the many factors determining trust relate not to the information technology, but to the business organisation and environment (can a provider underwrite an SLA) and to the political and regulatory context (can data from one jurisdiction be held and adequately protected in another?)

In a MDBD deployment, the performance and security of IT infrastructure is crucial because of the strong links to operational systems.

6 Example: MDBD Architecture

This paper concludes with an example describing how the key principles described above may be applied to a simple example: enabling an individual's access (mediated by a personal mobile device such as a smartphone) to a facility (industrial or domestic) protected by an on-line managed access control system. The key challenge is to identify a means by which the access control functions and personal devices may be linked securely.

The access control system is likely to be a typical "machine to machine" application combining remote devices (eg locks) local gateways (eg per building) and a central administrative function as shown:

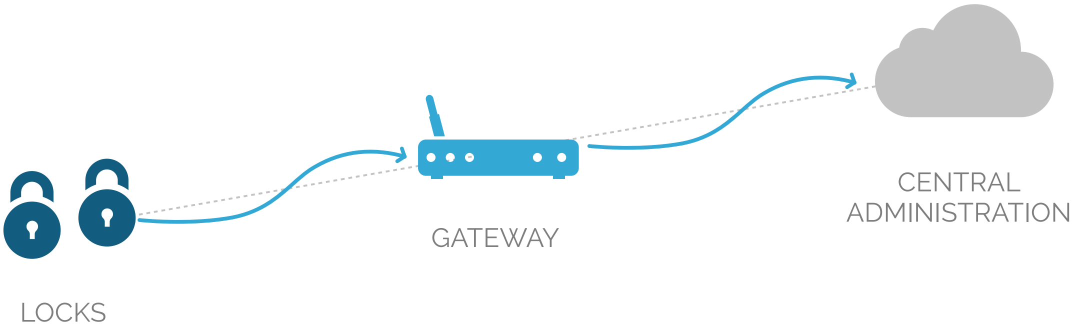

EXAMPLE ACCESS CONTROL SYSTEM: COMBINATION OF REMOTE DEVICES, LOCAL GATEWAYS AND CENTRAL ADMIN

The principles may be applied to this stand-alone system as follows:

⟶ **KNOW YOUR ENEMIES.** The primary risk of an access control system is granting unintended access. Evaluating the severity of this risk and defining an appropriate security policy will depend on the facility being secured, the assets within it and the consequences of unintended access. Consideration will be needed of failure modes, emergency access and maintenance.

⟶ **TAKE SECURITY TO THE EDGE.** Control messages between the end devices (the locks) and the central authority must be protected, particularly from spoofing or modification. A significant risk lies in the 'final metres' between the gateway and the lock if these links are poorly secured. Therefore one solution is to base the locks on microcontrollers able to support full IP connection to allow secure communications from the cloud to the lock.

⟶ **KNOW WHAT YOU'RE TALKING TO.** Mutual authentication will be specified between locks, gateways and central services. Initial provisioning is beyond the scope of this simple example.

⟶ **CREATE A STRONG NETWORK.** Both local- and wide-area network (WAN) connections will be encrypted. Fall-back WAN connections (eg via a cellular modem in the gateway) might be provided in case of telco network failure.

⟶ **DON'T TRUST IT, WATCH IT.** Monitoring of all remote facilities is assumed to be a key function of the central management service.

⟶ BUILD IT RIGHT. Key components must be developed to an appropriate level of assurance (which may depend on the asset being protected by the lock!). Interfaces and APIs must be protected.

⟶ BASE ON FIRM FOUNDATIONS. Employ Trustworthy Central Services Hosting and management of the central management service must be assured to an appropriate level.

Now consider an alternative but similar architecture (albeit including the cellular network rather than a simple local hub) applied to the personal portable devices. These personal devices are likely to be powerful enough to host full IP stacks and PKI infrastructure but will still be centrally managed by a corporate owner, a telecom network operator or by the manufacturer's cloud services (eg as provided by Apple or Google).

When considering the use of the personal device as an authentication mechanism to the security system, there is a choice of how the authority is communicated and information exchanged in order to negotiate access, as shown in the following:

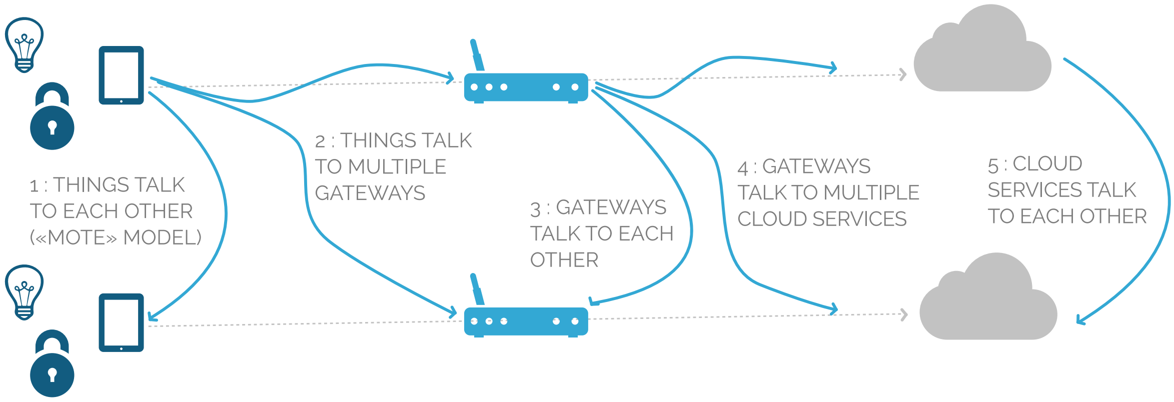
1 : THINGS TALK TO EACH OTHER («MOTE» MODEL)
2 : THINGS TALK TO MULTIPLE GATEWAYS
3 : GATEWAYS TALK TO EACH OTHER
4 : GATEWAYS TALK TO MULTIPLE CLOUD SERVICES
5 : CLOUD SERVICES TALK TO EACH OTHER

USING THE PERSONAL DEVICE FOR AUTHENTICATION PROVIDES FIVE ALTERNATIVES FOR COMMUNICATION

These alternatives will have different characteristics:

⟶ COMMUNICATION MIGHT BE ENTIRELY LOCAL (eg via Bluetooth or NFC, case (1) in the diagam) – this is probably appropriate for the actual unlock requests, but rather restrictive if it entails individually authorising personal devices with locks. This approach would, however, minimise the number of additional components (gateways, cellular networks) and APIs that require security measures to be put in place.

⟶ EASIER MANAGEMENT might be achieved by using a system that authenticated a device with a site-level gateway that was trusted by individual locks (cases (2) or (3)).

⟶ THE EASIEST SCHEME to manage operationally might be to allow an authentication agreement to be arranged between the respective managing services – allowing the security company to negotiate authentication details with the organisation that manages the mobile devices (cases (4) or (5)).

"This example shows that the principles described in this paper provide a strong contribution to ensuring end-to-end security."

The principles can be used to inform the choices between these alternatives, for example:

→ KNOW YOUR ENEMIES. Knowledge of security risk will determine whether agreements can be set up to publish APIs at cloud level which allow such cross-organisational authentication, and to determine key decisions (eg who is liable for a failure?)

→ KNOW WHAT YOU'RE TALKING TO. Managing identity becomes the key function of the system – how are mobile device identities to be checked, and how are they correlated to access control system roles. Ultimately this is likely to be a procedural, not a technical Issue.

→ CREATE A STRONG NETWORK. The access control system may be provided with a fall-back WAN link, but can the mobile device be authenticated if it loses its own WAN link?

→ DON'T TRUST THE COMBINED SYSTEM, WATCH IT. What monitoring is carried out by each party? Do they share SIEM data and processes?

→ BUILD IT RIGHT. Actual control of physical access still resides with the local devices – the lock must have a trusted execution environment that allows it to be trusted by the access control cloud server, and the mobile device must not allow impersonation by an attacker. In domestic applications, conventional consumer-grade security might be sufficient to allow a phone to take this role.
In more sensitive environments, a virtualisation mechanism allowing a secure partition to be established on a device might be required.

7 About VueForge®

VueForge® is Altran's industrial end-to-end offer for Machine-Driven Big Data whereby data is harvested, transported and transformed into information, which in turn is converted into services paving the way to new business.
VueForge® is underpinned by a reference architecture, technological accelerators developed in-house, a partner ecosystem, specific domain applications and the cross-industry expertise of Altran, including security.

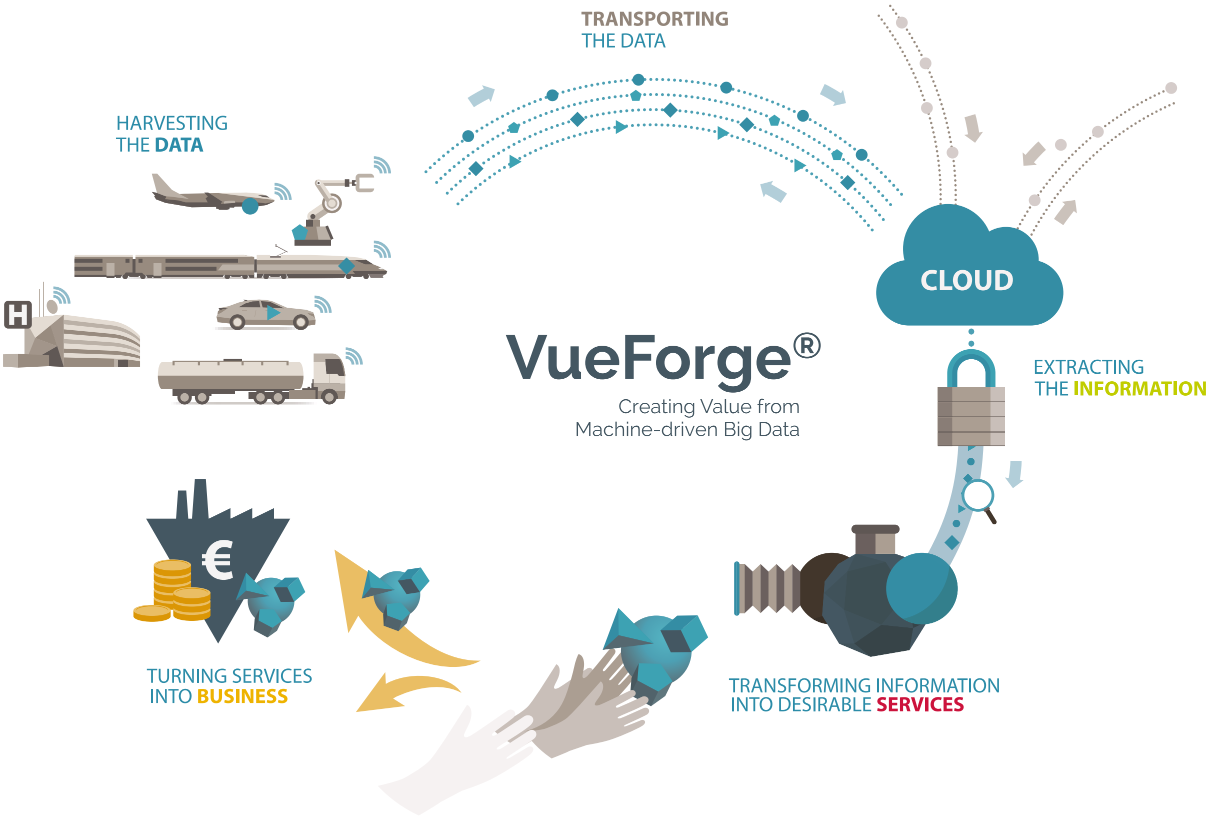

Altran's security expertise includes:

⟶ RISK ASSESSMENT & SECURITY MANAGEMENT

Using a range of industry best practices, standards and methods, including ISO 27001, EBIOS, MEHARI, HMG IS1 and ISO 15408. Also as part of Intelligent Systems / Altran's proprietary STORM™ approach to overall operational risk management.

⟶ AUDIT, ASSURANCE AND TESTING

Evaluation of system and organisational security, including penetration testing, and product evaluation, including assurance development and supplier follow-up.

⟶ SECURITY ARCHITECTURE

which includes technologies, application principles and specific patterns for security such as built-in security filters and policies in the VueForge® Play rapid-prototyping IoT platform, and analytics techniques for identifying attack behaviour.

⟶ SECURITY TECHNOLOGIES

in a variety of domains including thru-wall radar, SPARK tools for software development and assurance in support of MILS, and the VueForge® Moon carrier-grade standards-compliant IoT platform.

⟶ SECURITY IMPLEMENTATION

- server and web application security, implementations of key technologies such as SSL, including the Pico TCP implementation of a complete TCP/IP protocol stack suitable for small devices.

⟶ HIGH-ASSURANCE SECURE SOFTWARE DEVELOPMENT

Intelligent Systems / Altran's proprietary software development techniques are one of few approaches able to provide the assurance required for the highest levels of product approval (Common Criteria EAL6–7).

⟶ AVIATION SECURITY

providing the international Air Transportation community with pragmatic and cost-effective end-to-end security solutions for passengers, personnel, information, goods and infrastructures, protecting against international threats.

ABOUT ALTRAN

As global leader in innovation and high-tech engineering consulting, Altran offers its clients a new way to innovate. The Group develops with or for its clients the products and services of tomorrow. Altran works along with its clients on every link in the value chain of their project, from conception to industrialization. The Group has been providing its expertise for over thirty years to key players in the Aerospace, Automotive, Defence, Energy, Finance, Life Sciences, Railway, and Telecoms sectors, among others. In 2014, the Altran group generated revenues of €1.756bn. With a headcount of nearly 25,000 employees, Altran is present in more than 20 countries.

DOCUMENT REFERENCES

1 Forrester, reported by Motherboard: http://motherboard.vice.com/read/ransomware-is-coming-to-medical-devices

2 Various, including http://www.bbc.co.uk/news/uk-34784980

3 http://www.bbc.co.uk/news/technology-31296188

4 Private traits and attributes are predictable from digital records of human behaviour, Kosinski et al, http://www.pnas.org/content/110/15/5802.full

5 http://www.theguardian.com/technology/2015/oct/06/safe-harbour-european-court-declare-invalid-data-protection

6 Cifas, https://www.cifas.org.uk/fraudscape_latest

7 http://www.bbc.co.uk/news/business-33984017

8 In re Search Warrant, No. 13 Mag. 2814, M9-150, US District Court South New York

9 When Phone Encryption Blocks Justice, http://www.nytimes.com/2015/08/12/opinion/apple-google-when-phone-encryption-blocks-justice.html?_r=0

10 https://www.washingtonpost.com/world/national-security/obama-administration-opts-not-to-force-firms-to-decrypt-data--for-now/2015/10/08/1d6a6012-6dca-11e5-aa5b-f78a98956699_story.html

11 http://www.informationweek.com/government/big-data-analytics/data-protection-in-internet-of-things-era/d/d-id/1204428

12 EBIOS — Expression des Besoins et Identification des Objectifs de Sécurité, http://www.ssi.gouv.fr/guide/ebios-2010-expression-des-besoins-et-identification-des-objectifs-de-securite/

13 Mehari: Information risk analysis and management methodology, https://www.clusif.asso.fr/en/production/mehari/

14 HMG IA Standard No. 1, https://www.cesg.gov.uk/guidance/information-risk-management-hmg-ia-standard-numbers-1-2

15 https://www.owasp.org/index.php/OWASP_Guide_Project

16 https://cloudsecurityalliance.org/download/top-ten-big-data-security-and-privacy-challenges/

17 http://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/

18 http://www.altran.com/fileadmin/medias/1.altran.com/Finance/2015/PR_ALT_02122015_JLR-f.pdf

All trademarks are acknowledged.

VueForge®

Seven principles for achieving Security and Privacy in a world of Machine-Driven Big Data

http://intelligent-systems.altran.com/core-offers/vueforge.html