Using Metrics in the Industrial IoT Data Value Chain to Drive Trustworthiness

Authors:
Jacques Durand
Director of Engineering and Standards
Fujitsu
jdurand@us.fujitsu.com

Frederick Hirsch
Standards Manager
Fujitsu
frederick.hirsch@us.fujitsu.com

Jim Morrish
Head of Strategy & Partnerships
Nokia WING, Nokia
jim.morrish@nokia.com

INTRODUCTION

Confidence that an Industrial Internet of Things (IIoT) system will operate according to expectations is based on assurance that several aspects of the system are under control: security of its data and of its equipment, safety for people and assets, reliability of operations and subsystems, resilience of these in case of hardship, and privacy concerns for all kinds of personal data handled.

This set of properties – security, safety, reliability, resilience, privacy – has been identified in the Industrial Internet Consortium (IIC) Industrial Internet Vocabulary Technical Report[1], IIC Industrial Internet Reference Architecture[2] and ISO-IEC/JTC1/SC41 [3] as defining the *trustworthiness* of a system. While these properties were looked at separately (mostly) in the past, their grouping in a single concept – *trustworthiness* – makes more sense in IIoT systems because of their increased interdependency. This is due to the intricacy of the digital and the physical, the level of automation and the extent to which people and processes depend on it, the overall complexity of these systems and the increased digitization and volume of data generated.

IIoT systems include a value chain related to the use of data. This data value chain starts with the production and collection of data from assets and their environment in the physical world. That data is then contextualized and processed to become intelligence and knowledge. This knowledge, in turn, is analyzed in a business context to translate into decisions and actions that can be used to improve or create new business operations or products in order to drive business value. Metrics defined at various stages of the data value chain can also be used to establish and manage the assurance of trustworthiness, providing confidence in the IIoT system. This is especially important in IIoT systems that can include many distributed components. This data value chain sequence is illustrated in Figure 1.

Assurance of trustworthiness is established based on the IIoT data value chain. This assurance reflects measures that either prevent adverse results from occurring or anticipate and mitigate potential effects when issues occur. Regardless, the data value chain and metrics are essential to maintain confidence in the trustworthiness of the IIoT system.

Trustworthiness metrics are an important concept supporting the management of trustworthiness in an IIoT system. The definition of metrics may aid in

[1] The Industrial Internet of Things Volume G8: Vocabulary, IIC:PUB:G8:V2.1:PB:20180822, Version 2.1, August 2018, IIC. https://www.iiconsortium.org/pdf/IIC_Vocab_Technical_Report_2.1.pdf

[2] The Industrial Internet of Things Vol G1: Reference Architecture (IIRA), IIC:PUB:G1:V1.80:20170104, Industrial Internet Consortium, January 2017, https://www.iiconsortium.org/IIRA.htm

[3] Information technology – Internet of Things Reference Architecture (IoT RA), ISO/IEC DIS30141:20170628, ISO/IEC JTC 1/WG 3, June 2017.

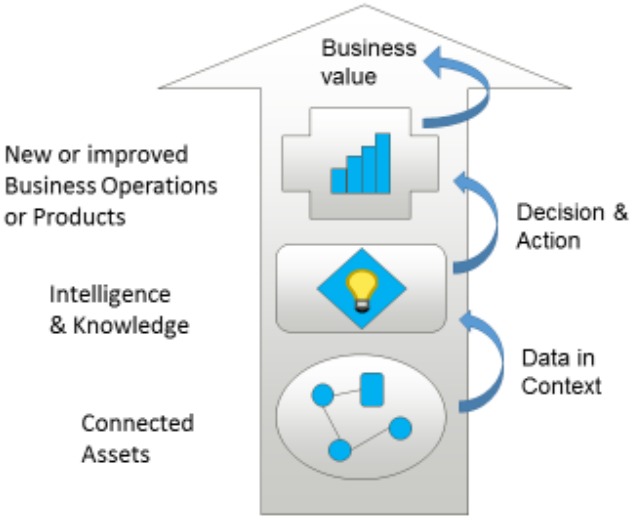

Figure 1: The Value Chain of Industrial IoT

understanding the key considerations in a system and also aid in the analysis and design of a system, especially if historical metrics data is used to inform decisions. Clearly metrics can also be used during the operation of a system to maintain visibility into its operation and to help ensure that trustworthiness and other business and operational targets are met.

Good metrics will typically find many uses. For example, metrics on storage service reliability will help to:

- Clarify the expected service level and quality with any service providers. Metrics also support contractual enforcement such as assessment of penalties in case of failure to fulfill SLA (Service Level Agreement) or SLO (Service Level Obligation) targets, and support the ability to compare providers.
- Evaluate how well the service performs. This in turn allows for precise feedback to providers.
- Understand the nature of shortcomings and failures of a system component so that these can be mitigated within the system, or negotiated with the provider.

A real world example of this third point is a system where edge devices periodically invoke the cloud storage service directly every minute, while having the capacity to handle a backlog of only up to ten minutes of data stream. In this scenario, it is important to prevent downtimes of nine minutes or more of the data storage service. An adequate metric measuring the duration of downtimes – not just the uptime average – will be the basis for negotiating service quality in order to minimize data loss.

This paper describes how operational metrics data may be combined with business and risk management information to support a better understanding of trustworthiness, enabling investments made in trustworthiness to be better managed.

MANAGING TRUSTWORTHINESS

Trustworthiness in the business context means that a satisfactory level of confidence can be established in any system component (be that a sensor, a machine or a factory). Confidence can be established in what it claims to be, whether it fulfills its ascribed tasks, has appropriate performance, and will not endanger people, the environment, partners or the organization due to any issues relating to security, safety, reliability, resilience and privacy.

Tradeoffs and decisions need to be made among business and functional requirements as well as design decisions and

risk mitigations related to trustworthiness. Care must be taken not to overinvest in trustworthiness mechanisms, however, since this can be detrimental to overall system performance, utility and cost. Connected bathroom scales provide a good everyday example of this kind of dynamic: There are many things that could be done to improve the accuracy of the scales (for example controlling for environmental humidity, adjusting for altitude, ensuring a perfectly level surface and so on), but it just doesn't make business sense to invest this much in the pursuit of accuracy.

Figure 2 provides some illustrative examples of the consequences of over-specifying an IIoT solution in each of the five identified aspects of trustworthiness:

strategy and business results. The investments to be made can relate to the "business as usual," such as increasing production, expanding the business and so on. Investments can also be used to address trustworthiness risks either through direct investment with the aim of reducing the frequency of harmful events or by indirect methods such as purchasing insurance (to mitigate the consequences of any harmful event). Making the right investments in trustworthiness requires a detailed understanding of the IIoT system in question and also that appropriate decisions are taken regarding the needs and tradeoffs among functional, non-functional and trustworthiness requirements.

Trustworthiness aspect	Potential consequences of over-specification
Security	Increased costs and reduced usability
Safety	Reduced solution flexibility, compromised accuracy of data generated by an IIoT solution
Reliability	Increased testing costs, increased predicted maintenance costs for additional replicated components, increased complexity, greater capital costs
Resilience	Increased capital, commissioning and maintenance costs and reduced flexibility to maintain larger margin for error in processes and systems
Privacy	Unnecessarily cumbersome processes associated with registering and managing user accounts and resolving technical issues.

Figure 2: Example consequences of over specifying trustworthiness aspects

Managers of a company have choices they can make with regards to the investments they make in most business-related assets. These real options have consequences to the organizational business continuity, health,

Compliance with any prevailing regulations is a baseline trustworthiness consideration, but there may be benefits from going beyond minimum compliance (as discussed in the "Trustworthiness from a Business Perspective" section of this article). Figure 3 illustrates how the current state of trustworthiness aspect may be positioned with respect to a minimum compliance level versus a target state determined by business considerations. This is a sample illustration – actual data will depend on the specifics of the business situation.

Figure 3: Kiviat diagram illustrating Minimum, Current and Target states of Trustworthiness Aspects for an IIoT system

Decisions on investing in real trustworthiness options can be based on a structured analysis of scenarios based on the risks and consequences. Such analysis works best for potentially high-frequency events since the probabilities can be quantified and used in conjunction with an analytical model of the business to review scenarios and make decisions. Such an approach is less well suited to addressing low frequency and high impact events (for example flooding risks in certain locations where floods are a rare and unexpected event). Thus investment decisions should not be based solely on quantitative analysis but also should include judgments and investments based on an understanding of high impact events. Care must be taken that both the data and the analysis used to make decisions are appropriate and that the confidence in the data quality and analysis is appropriate to the concerns. The analysis should take into account operational goals and their corresponding metrics, financial and other targets, risk metrics and trustworthiness considerations.

Risks of various types, including security risks, safety hazards, natural events and privacy risks (among others) can be mitigated through organizational changes (e.g., training staff) as well as technology deployment (for example, deploying identity management processes). Other traditional risk management approaches may also be used, such as purchasing insurance, that effectively transfers risk. Risks may also be accepted as a necessary component of an overall business, but this should only be done if the consequences of those risks (and the company's risk attitude) are well understood. Mitigating risks through the use of technology or organizational changes can

be complicated due to the variety of concerns as well as the number of approaches that can be taken to mitigate risks. For this reason a structured and systematic approach is helpful.

USING METRICS TO ASSESS AND CONTROL TRUSTWORTHINESS

Trustworthiness metrics associated with operational components provide insight into the operation of those components and enable control over trustworthiness aspects, if the metrics are defined correctly. For example metrics related to the Reliability trustworthiness aspect could include:

- Variability of end-to-end data latency from source to storage. Keeping such variability low is desirable as many application only provide quality output when latency is well controlled and within limits. This clearly depends on many factors (potentially including device caching and configuration settings, network latency, and storage service availability).
- Elapsed Time between detection of stress conditions and dynamic scalability operations to restore overall performance expectations.

Trustworthiness metrics are often designed to be shared by a broad class of systems, defining a way to adhere to regulations or industry-defined standards and assessments. This is the case of readiness metrics such as scorecards derived from maturity models.

When it comes to managing the operation of a particular solution, often the performance metrics that are used relate to that specific solution only. These metrics are developed locally by operation managers and service providers and are directly useful in managing the solution during operations. While these metrics are not necessarily shared across systems, a standard representation – and ideally a standard definition – is useful to compare them as well as to compose them. Trustworthiness metrics for an entire system can be derived, incorporating consideration of the trustworthiness of its constituent components and sub-services.

Trustworthiness aspects may contribute – or conflict with – each other. Part of managing trustworthiness in a solution is to define and control these interdependencies. These interdependencies may vary from one system to the other, and sometimes may impact each other within the same system, as illustrated in the following examples:

- **Privacy considerations can impact Security**: Privacy regulations may restrict data replication, prohibit collecting too much data on clients accessing a service, or make strong requirements about disposing of data. In some cases these restrictions may adversely impact the security of the service by preventing useful data collection or tracing, such as the identification of requests and their origin.
- **Investment in Privacy may contribute to Security**: in other cases the opposite is true, as Privacy measures may help reduce data thefts or their consequences.

- **Enhanced Security can reinforce Reliability**: Security of many services in a system will generally contribute directly to Reliability objectives – e.g., by preventing DOS attacks.
- **Reliability management techniques can be detrimental to Security**: A common case of adverse effect on Security – as previously mentioned - is the delaying of security updates and patches in order to preserve the stability of current IT systems.

Uncovering such dependencies between trustworthiness aspects is part of a trustworthiness analysis. Assessing these dependencies, their synergy and, in case of conflicts, the acceptable tradeoffs, will rely on measurements and objectives.

The goal in both design and continued operation of an IIoT solution is to keep a system operating within the bounds determined to be acceptable. These bounds are defined by objectives set for the various trustworthiness aspects relative to some metrics, as well as functional and other business considerations, as illustrated in Figure 4.

Assessing the impact – positive or negative – of one trustworthiness property on another property will rely on metrics and related measurements. The role of metrics goes beyond assessment though since they also help manage the reinforcements and conflicts between trustworthiness aspects by balancing the metric targets.

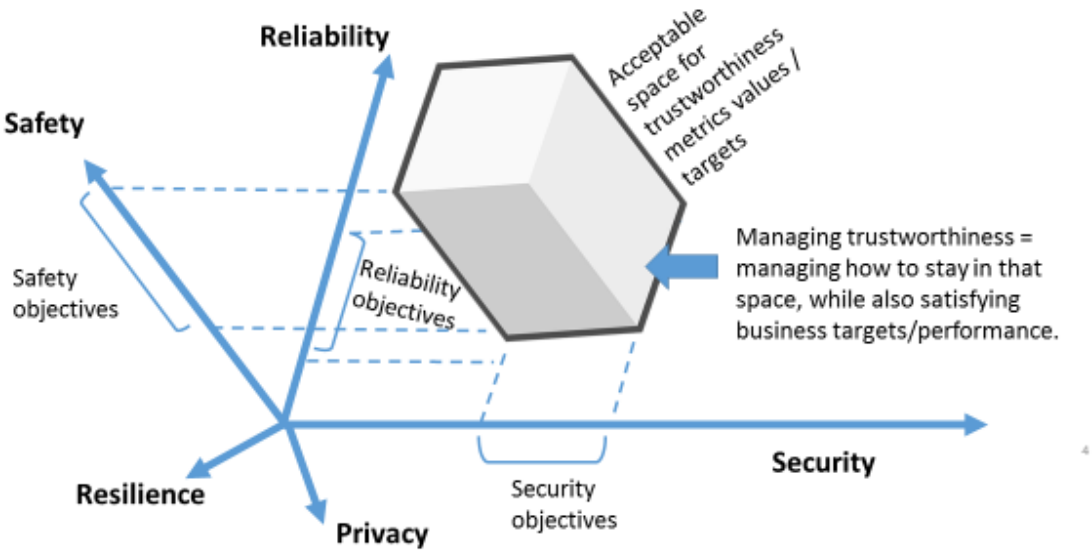

Figure 4: The Trustworthiness Space as Defined by its Metrics

Interdependency Between Trustworthiness and Business Operations

Operational trustworthiness metrics may inform design and analysis as well as assist in keeping an operational system in control,

but are not enough to meet the needs of managing trustworthiness and creating business value. Business context and information must also be considered. Combining appropriate business metrics and information with trustworthiness understanding and metrics will allow an organization to turn data into knowledge that it can act upon to enable a dynamic and successful enterprise.

Organizations are generally concerned with managing risks, both those associated with trustworthiness aspects as well as others (for example product delay or lack of market adoption). A common approach to measuring risks is to calculate the expected value based on probabilities of events as well as the anticipated impact of the event. Leading the business while considering business mission and goals, financial metrics, risk position and trustworthiness considerations will require making complex investment decisions involving many factors. This is complicated since some trustworthiness aspects will support a business metrics and others will detract. When all factors are considered together, the tradeoffs will lead to an "acceptability zone" where all objectives are reachable (e.g., safety and performance). The details depend on the precise definitions as adopted by the business and industry in question.

Trustworthiness properties generally have an impact on operations and business outputs. Consider an IIoT system in a factory that comprises an assembly chain. The resilience of such a system includes the resilience of its assembly chain. The resilience of the assembly chain can be measured by a metric involving the percentage of overtime (OTpercent) for processing of a production lot due to either replacing, repairing or simply reusing a defective machine:

ResACmetric = 100 - OTpercent

where a value of at least 80 is expected for some types of failure.

Consider now a performance metric for this assembly chain:

PerfACmetric = (expected processing time of a production lot)/(actual processing time)

where a value less than 0.9 in average is considered unacceptable.

It is understandable that these two metrics may depend on each other: the more resilient the assembly chain, the greater the certainty that its performance level will be stable, according to these metrics. Figure 5 illustrates a strong dependency between both, in case of hardship:

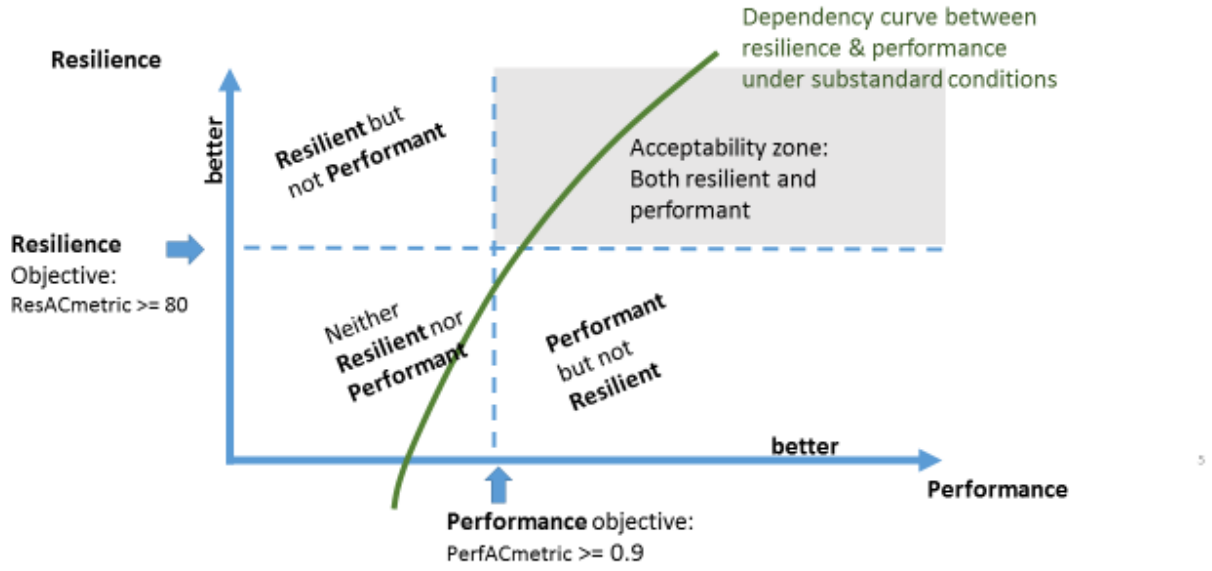

Figure 5: Example of Resilience reinforcing performance

Resilience may take many forms in an IIoT system. The same system that involves the above assembly chain may also use an IT service, say to archive sensor data in a Cloud. On the IT side the resilience of this service may be obtained by clustering a large enough set of servers in different locations under load balancing, thus mitigating server failure. That aspect of resilience will also impact operations in a positive way by improving requests throughput and response time. In this case, achieving the resilience objective by itself does not improve performance, but the means deployed for achieving resilience happen to have a positive impact on performance.

As highlighted earlier in this article, trustworthiness properties may potentially adversely affect business performance. Trustworthiness has investment costs, raising the question of how much the organization is willing to pay in terms of additional time, complexity and operational costs as well as initial investment. There is often a trade-off between safety measures, for example, and operational agility or speed. Security can often have a business impact, in terms of complexity and cost. Of course, the risks that such trustworthiness measures help mitigate or prevent, may have a much larger cost in the long run.

Figure 6 illustrates a negative impact of a safety measure on some business performance indicator. For example, safety may require giving more time to the personnel to manually change machine configurations, often needed in high-mix production. This will require slowing down the assembly chain thus reducing its performance. In this example, the dependency curve shows that increasing safety will decrease performance (e.g., by slowing an assembly chain) and vice-versa.

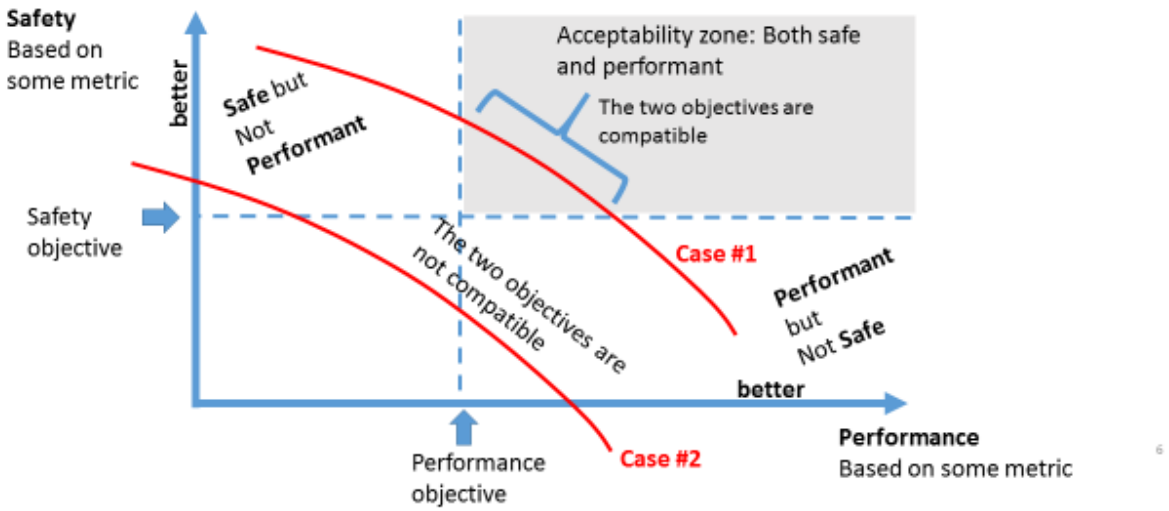

Figure 6: Example Representation of Conflicting Objectives

Combining metrics and context can provide an understanding of the system, its dynamics and the tradeoffs that are made. This can be represented graphically with diagrams showing tradeoffs. These diagrams may oversimplify the situation since they often ignore other variables. In many cases they also just express a correlation, not a causality. Detailed understanding of the factors behind the metrics are needed to understand what is happening.

TRUSTWORTHINESS FROM A BUSINESS PERSPECTIVE

From a business and financial perspective, increased levels of trustworthiness can generate benefits through reduced levels of risk: While 'per user' costs may increase and some processes may become more cumbersome, the risk of trustworthiness-related events reduces so that the overall value of a business increases. In this context, it is worth noting that greater levels of trustworthiness can generate financial-related benefits in many ways, including:

- Reduced levels of compensation payments for outages and other failures
- Avoided payments to regulatory bodies for instances of non-compliance with regulations, or any individual trust-related events
- Increased levels of sales, and revenues per sale, due to stronger brand image
- Lower costs of business insurance
- Lower costs of funding and greater shareholder returns

Clearly any trustworthiness investment decisions that a business might make should be taken with reference to the probability of potential trustworthiness events (such as data loss or worker injury) and also the likely commercial and societal impact of those events.

The high level assessment of the expected consequences of any event should also take into account a range of extenuating and/or mitigating factors that may be relevant to the specific trustworthiness event in question. Such factors include:

- ***Scale of Breach*** – Is the breach very limited in scale, compared to the overall solution or is it complete and fundamental?
- ***Reversibility*** – Can the breach be reversed with a definite cost or will it become an ongoing and open-ended exposure?
- ***Downstream Impacts*** – Is there potential for a trustworthiness event to impact other IIoT solutions (or real world events) that are potentially influenced by any outputs of the IIoT system in question?
- ***Potential Criticality*** – Do specific trustworthiness events potentially have different levels of impact for different user groups?[4]

Efforts should also be made to estimate the financial impact of any potential trustworthiness breach. The expected business risk associated with any specific trustworthiness event is simply a product of the probability of that event and the impact of that event. So these are the two critical inputs to any approach to optimizing trustworthiness within an IIoT system.

Investment in trustworthiness is not, however, a one-sided argument: trustworthiness measures can also themselves generate increased value for a business. This can happen through a number of mechanisms, including:

- ***Brand impact*** – companies can seek to differentiate on the basis of trustworthiness and become recognized as 'more trustworthy' than competitors.
- ***Increased revenues*** – potentially products and services that are underpinned by higher levels of IIoT trustworthiness (or QoS) can be sold for higher unit revenues.
- ***Market access*** – potentially new markets for products and services may become addressable if a company maintains higher levels of trustworthiness.

Accordingly, to appropriately manage trustworthiness within a business it is also necessary to identify how any real options for enhancing trustworthiness might potentially result in an opportunity to enhance value.

It should also be noted that trustworthiness is an evolving concept and trustworthiness measures that are appropriate for the overall context of any IIoT system at any one time may not be suitable at some future time. Figure 7 illustrates a sequence of events whereby the required level of trustworthiness increases twice during the 'operate/maintain' phase of an IIoT solution (potentially due to a change in relevant

[4] For instance information about use of drugs by a rock star is a non-story, whereas for a politician it could be career ending.

regulations).[5] On one of these occasions Target and Current levels of trustworthiness have to be increased to meet Minimum (compliance) levels.

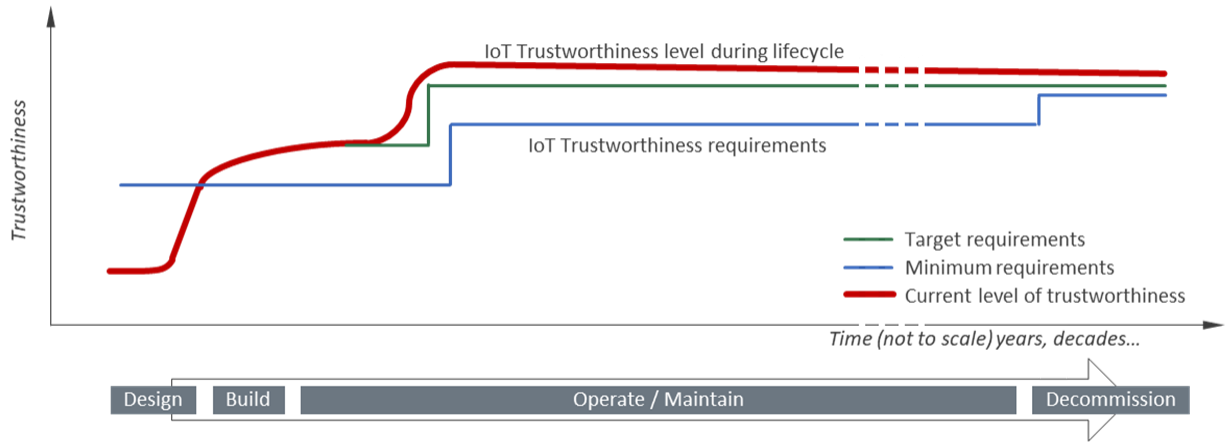

Figure 7: Trustworthiness over time

Accordingly, the trustworthiness of any specific IIoT system needs to be revisited and trustworthiness measures reassessed periodically. A prudent management team will ensure that trustworthiness measures are reviewed as follows:

- ***Periodically*** – on a regular basis, potentially quarterly or annually, depending on the criticality of the IIoT solution in question and the overall levels of risk exposure
- ***Reactively*** – when the trustworthiness environment changes, potentially through the introduction of new regulations (for instance the introduction of GDPR) or changes in the trustworthiness associated with any upstream processes and/or data sources (for instance, the use of a new supplier to provide weather data inputs to a system)
- ***On request*** – potentially when downstream uses of data (or other outputs of the IIoT system in question) change and in response to requests from relevant stakeholders

A company's approach to maintaining trustworthiness and updating any associated analyses and impact assessments should be the subject of a documented and managed security policy.

CONCLUSION

This paper has introduced the idea of a data value chain and the use of metrics in the Industrial Internet of Things to provide assurance of trustworthiness. Data from operational metrics can be used to inform design decisions as well as be used to monitor and take action to keep an

[5] From "IoT Trustworthiness is a Journey and NOT a Project" in Sept 2018 Journal of Innovation, https://www.iiconsortium.org/news/joi-articles/2018-Sept-JoI-IoT-Trustworthiness-is-a-Journey_IGnPower.pdf

operational system within the required trustworthiness region. These operational metrics can be combined with business metrics and risk management analysis to create knowledge that can be used to understand design tradeoffs and enable decisions regarding technical and organizational approaches to achieve both trustworthiness and other business objectives. This is especially important in complex IIoT systems with many components and complex interactions.

We are continuing work at the IIC on developing this approach including the metrics and the analysis approaches to make trade-offs. We are seeking input and use cases to extend this work and invite participation.

➢ Return to IIC Journal of Innovation landing page for more articles and past editions.