Enabling Digital Transformation with IoT Performance and Properties Measurement

An Industrial Internet Consortium White Paper

Version 1.0

2020-05-07

TABLE OF CONTENTS

The Case for Metrics in DX Solutions 5
Designing Systems for Measurements, Assessments and Evolution 5
Discovering Opportunities for Improvement 6
Seven Value Areas of DX and Their Assessment 7
The Contractual Aspect of Operating IIoT Solutions 9
Different Metrics for Different Phases of a Solution Lifecycle 11

The Role of Metrics in Managing a DX Solution 12
Business Model Validation and Iterative Improvement over a Solution Lifecycle 12
Business model validation & improvement 13
Solution validation & improvement 14
System Characteristics, Trustworthiness and Their Assurance 14
System characteristic assessment over the lifecycle of a system 14
Trustworthiness and its metrics 16
Dependency of some trustworthiness metrics on performance metrics 17
Managing trade-offs between trustworthiness objectives and business objectives 20
Metrics for assisting the design, development and evolution of an IIoT solution 22
The profile of an IIoT solution 22
The Project Explorer solution profiling tool 25
The role of metrics in system interoperability and service composability 28
Service compatibility and interoperability 28
Sharing and reusing metrics 29
The Value of Standardizing Metrics and Indicators 31

Conclusions and Outlook 32

Appendix: A Survey of Existing Works on Metrics in DX Related Areas 32
Quality Metrics for Network Carriers and Mobile Devices 32
A standard-based, large-scale quality management system 32
Metrics and their outcomes 33
Managing service providers 35
DX relevance 35
The Service Measurement Index (SMI) from CSMIC 36
A Metric Framework for evaluating Cloud Services 36
An example of Readiness Metric 38
The Metric Model Standard of ISO-IEC JTC1/SC38 42
Standardizing a Metric Definition: Structure and Rationale 42
An example of Service Availability Metric 44

Bibliography 46
AUTHORS AND LEGAL NOTICE 48

FIGURES

Figure 1: The role of IoT for issue validation and remediation in the Kaizen process 6
Figure 2: The validation and improvements cycle 13
Figure 3: Dependency and compatibility between objectives as measured by metrics 21
Figure 4: The Trustworthiness space as defined by its metrics 22
Figure 5: Some profile dimensions of an IIoT system 24
Figure 6: General structure for assessing an IIoT project profile 27
Figure 7: Metrics Library and its usages 30
Figure 8: The structure of the TL 9000 standard 33
Figure 9: Chart of monthly performance reports 34
Figure 10: TL 9000 supplier executive dashboard 35
Figure 11: Major business service properties to be assessed 36
Figure 12: The metric model according to the ISO/IEC 19086-2:2018 standard 43

TABLES

Table 1: Some indicators used for assessing the real-time profile of a system 28
Table 2: Examples of monthly reported TL 9000 metrics 33
Table 3: A metric for service legal portability according to CSMIC 42
Table 4: The main element of the service availability metric 45
Table 5: The rules of the service availability metric 45
Table 6: The parameters of the service availability metric 46
Table 7: The expressions of the service availability metric 46

Digital Transformation (DX) for industry leverages connected things to transform processes and operations to produce better outcomes. It is a process, an endeavor for more efficiency, new business, operational opportunities and flexibility. The transformation process requires a prompt assessment of what works and what does not. Developing and assessing a solution to support DX is an incremental process. The value of the solution has to be measured and demonstrated once deployed, as does the value of every incremental step toward this solution. Each move must be validated or dismissed as promptly as possible because of the distraction, costs and disruption it can bring. Successful digital transformation requires a culture of measurement.

Industrial DX solutions elicit and exploit new knowledge about the operational context that was difficult or impossible to acquire until now due to the lack of appropriate technology. In the past, improving on practices in the field, industrial processes and product design was based on incomplete, often delayed, data and on the assumptions of experts. New solutions, relying on emerging technologies such as artificial intelligence (AI), digital twins and more generally, industrial IoT (IIoT), bring more insights and rationale to these tasks but have little to build on in terms of past experience.

Whether developed in a brownfield or greenfield environment, it is hard to predict how successful a DX solution will be or can be. We need trials and adjustments throughout the solution's lifecycle, which in turn relies on measures based on metrics and targets.

Today, many solutions in industrial environments are closed and proprietary from end to end. They are, in other words, siloed. Tomorrow they will be interconnected and many of their functions and resources will be shared or contracted out to service providers. This will increase the importance of contracts—whether informal agreements or contractually binding ones. Service quality and performance are an essential part of these contracts. They require continuous monitoring and measurements, which allow for real time assessment and reaction.

This paper investigates the need for measuring various aspects of an industrial DX solution at various stages of its lifecycle and how measurements are essential to manage it. It shows how metrics serve different purposes, supportive of planning, governing and managing a solution. It provides an overview of existing efforts in relevant areas from which DX solutions can learn.

This paper makes abundant references to IIoT technologies and examples, as the frontier between the physical world and the IT world is where a significant amount of data is generated in new ways. These new sources of data play a crucial role in the digitally transformed organization and in evaluating its new processes and solutions. Hence, we pay particular attention to the performance and properties of IIoT systems.

THE CASE FOR METRICS IN DX SOLUTIONS

DESIGNING SYSTEMS FOR MEASUREMENTS, ASSESSMENTS AND EVOLUTION

The business case for transforming industrial, operational or business processes is rarely clear from the start: are efficiency or efficacy gains worth the investment costs and process changes? Will the product enhancement respond to user expectations? How much disruption and cost are entailed by the operational upgrade and is it worth it? There is an investigative nature to deploying emerging technologies for DX and to achieving value, particularly in brownfield conditions. Since the industrial context is always changing (technologies, equipment, practices, products and customer preferences), an optimal process today may not be so tomorrow. Solutions must be able to evolve, adapt and adjust.

The data set that needs be captured to optimize production processes may not be clear up-front, often the plant operator can just tell there is unused capacity and a potential for improvement, or that an error rate seems too high on the assembly chain. Investigating potential issues or improvement opportunities requires both well-rounded data as produced by IoT technologies, and assessment based on this data. Raw data needs to be processed to serve that assessment. Producing meaningful indicators is the first stage of data processing and decision making.

Consider a service availability metric that is based on an availability indicator calculated as a percentage of time the service is available. The availability indicator is already the product of processing several sources of raw data input. Calculating availability time relies on measuring service outage periods. Defining these periods involves, in turn, several sources of data. For on-line services, recording server shutdown times is obviously a major input for this metric. Now, if very degraded quality of service is considered a form of service outage, then some measurement of server response time is also involved. Finally, if there is a policy that says scheduled server maintenance time should not be counted as downtime for calculating service availability, then such periods have to be measured too.

Targets can then be set and decisions made based on such an indicator, such as generating an alert if service availability drops below 98% and a penalty if it drops below 95%. Without an agreement on the precise definitions of what is measured, or at least a clear statement of the measurement modalities, it is difficult to understand, compare and reuse such availability metrics.

A digital transformation solution flexible enough to adjust to changing conditions and to allow for nimble investigations and assessments must be built for continuous monitoring of operations, equipment status and external context. The solution then must be responsive to metric outputs either automatically or manually. Finally, such a solution must allow for managing transparently

its metric definitions and implementations (not just their outputs and targets) as essential decision-making elements.

Discovering Opportunities for Improvement

Process efficiency improvement is a common form of value expected from industrial IoT, especially in areas such as manufacturing and industrial processes. Manufacturing industries are always trying to improve their operations and product quality. Sometimes the potential for improvement is obvious, for example, if the defect rate suddenly jumps. But often the opportunities are not so evident and uncovering them relies on experience and guesswork. Such a discovery process can be assisted with IoT technologies.

In Japan, there is a continuous improvement process called Kaizen (see Figure 1). It illustrates various steps of some variant of the Kaizen process and how IIoT has changed it, according to Fujitsu, a large computer technology company that runs several manufacturing plants.

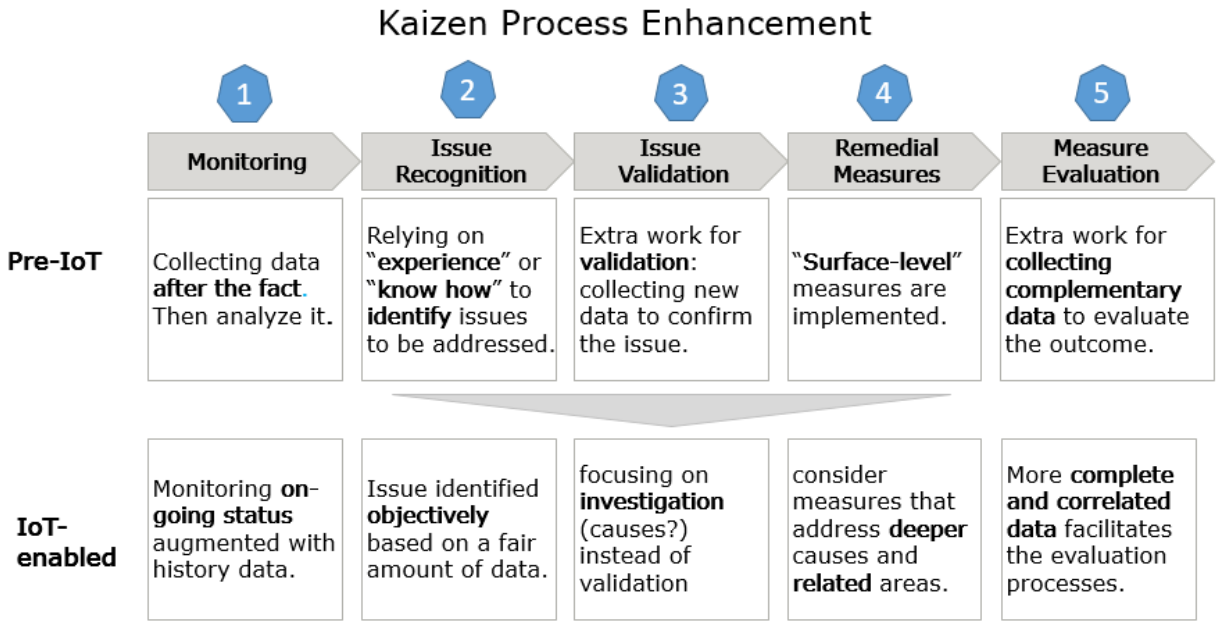

Figure 1: The role of IoT for issue validation and remediation in the Kaizen process

When an issue has been detected, there are broadly two approaches for remediation:

The *shallow resolution approach* employs surface-level counter-measures that would lessen the cost of dealing with defects without treating the root causes of the problem. For example, early detection that a product component is defective avoids the cost of further processing the product, or the cost of shipping the product and having to handle its return.

The *deep resolution approach* is to investigate the causes of the issue with the goal of reducing defects in the first place. For example, it uses manufacturing history data and usage records to

diagnose the cause of the defect. This makes it possible to prevent the problem or to address it earlier. This approach requires some research and takes more time.

In both cases, IIoT enhances this Kaizen process with:

- faster detection of the issues that have good improvement potential (phases 2 and 3),
- better insights on the root causes and how to remediate (phase 4),
- faster assessment of the business or operational impact of a remediation measure and its trade-offs (phase 5),
- better use of personnel expertise and more informed decisions at all phases of the process.

In particular, IIoT helps remove or reduce the cycle time handicap, which is the main impediment to the deep resolution approach.

When improving industrial processes, there is a cost associated with the time it takes to implement the change and to validate its value. These costs are not just about equipment and organizational changes such as training, but about risks, temporary performance loss and various inefficiencies due to disruption. The efficacy of a remediation measure must be evaluated as quickly as possible, with the right indicators. Some of these indicators concern direct operations performance and can be obtained rapidly with proper monitoring (e.g. number of tardy shipments, error rate). Others, usually measuring indirect costs such as disruption of internal operations, other side-effects or financial impact, take more time. These can be assessed informally by experts in the field, and confirmed later by additional monitoring.

SEVEN VALUE AREAS OF DX AND THEIR ASSESSMENT

In the Industrial Internet Consortium (IIC) we have identified seven major business value areas for industrial DX:

Process efficiency: improved agility, speed and reduction in time to market, business process optimization, reduced operation costs, increases in productivity and labor efficiency, enhanced intra-organization collaboration and better integration with greater operational environment and systems.

User experience: improved customer satisfaction, added value for users, better service and customization.

Product quality: reduction of errors and defects, better tracking, measurement and control of quality factors, better consistency in production quality and delivery.

Asset management: better tracking, monitoring and control of physical assets such as machines, tools and other resources or equipment, improved asset utilization, processing, maintenance efficiency (preventive, predictive) and cost efficiency.

Business innovation: new revenue streams with innovative business models and enhancement of existing models. Contributing factors include new services or products, product enhancement, combinations of product and service, opportunities to create services, and faster research, development and engineering processes.

Governance: facilitating strategic decision-making, assessing and assuring compliance to policies and regulations. This also includes informing management strategy to balance dimensions such as product quality, cost, delivery timeliness and environment or regulation impact.

Risk management: identifying, quantifying and managing the risks in business and operations, enabling risk mitigation, monitoring and improving trustworthiness (security, safety, reliability, resilience, privacy) with an understanding of their interdependencies and enabling assurance.

Any DX solution is expected to provide value in one or more of these areas. Cost reduction and revenue increase are not listed as value areas because they are a by-product of any improvement in the value areas above. The value areas represent ways to achieve these goals. Costs and revenues as (financially) quantifiable objectives are better expressed as business-level key performance indicators (KPIs).

Success along any of these value areas has been traditionally assessed by metrics and KPIs, such as:

- order fulfillment time,
- defective products rate,
- inventory turnover,
- number of retained customers,
- hours spent on sales follow up and
- net promoter score.

The success of a DX solution still abides by such metrics and is measured by how quickly a trend can be detected and addressed. In the past, these indicators have been evaluated downstream of operations once all data is recorded and available over a period of time, *a posteriori*. IIoT plays a key role in assessing the value of a DX solution. IIoT technologies generate streams of data that are available in near real-time for a dynamic assessment and faster reactions, such as allowed by real-time analytics or machine learning models.

IIoT metrics are also expected to be used earlier during operations and more dynamically. Consider the product quality value area. Commonly used product quality metrics in this area are:

- the *yield*, as the percentage of product output that meets both quality and compliance standards without the need of re-run or re-work.
- the *scrap rate*, as the percentage of raw materials sent to production that never make it into the finished product and

- *supplier defect rate*, as the percentage of materials or products received from suppliers that do not meet required quality or compliance specifications.

IIoT technologies raise expectations for these metrics, both for their results and for the efficacy of the metric itself:

- The yield is a common target for improvement. Early monitoring and data collection at all stages of the manufacturing process lead to a better understanding of quality factors. In turn this leads to an early quality assessment that makes it possible to avoid further processing of a defective product, apply some remedy, thus improving the yield.
- A more refined metric to assess the quality of raw materials may lead to more efficient use, reducing the scrap rate. Instead of being discarded based on an undiscriminating general quality test, materials of lesser quality may still qualify for the production of less demanding types of product.
- A production system that allows for monitoring and assessing compliance of product parts at the supplier site will prevent the shipment of unqualified parts down the supply-chain, thus reducing the supplier defect rate.

These improvements require metrics that operate on dynamic data during production and as early as possible in a production process.

Other value areas offer similar opportunities when monitored more dynamically. Another area is *risk management***.** The need for metrics here has been well understood for a long time. In IIoT solutions, metrics are instrumental in capturing and controlling the dependencies and trade-offs between operational performance such as productivity or lead time, and trustworthiness factors *[17]*. These dependencies are studied in *[10]* along with ways to manage trade-offs and conflicting goals.

THE CONTRACTUAL ASPECT OF OPERATING IIOT SOLUTIONS

IIoT systems are distributed systems with heterogeneous features and technologies. Some functions and subsystems from device management to data collection and storage, to application-level cloud services are increasingly contracted as a service. Hardware is also increasingly provided as a service, perhaps as an annual subscription plus profit sharing plan.

In a next phase, IIoT solution siloes will integrate horizontally by sharing services and data. For example, road and street traffic data as generated by a smart city IIoT system, including collision incidents rate, will be of interest to the car insurance industry, as well as to car manufacturers. Sharing this data more broadly can be done via a data marketplace, as a service to be contracted

by different parties.[1] Data is increasingly subject to contracts between parties[2] or used as a traded commodity.[3]

Distribution requires agreements at the governance level. As integration of solutions and reusability of subsystems increase, so does the reliance on providers or partners internal or external to the organization. Different managers will be responsible for different sub-systems and functions, adding to the governance divide.

All of the above point at a *flat* governance model for IIoT solutions. In turn, this translates into various forms of agreements or contracts.

Cooperation between the various parties involved in an IIoT solution will take different forms: service level agreements (SLAs), service level objectives (SLOs), service quality objectives(SQOs), joint projects, shared objectives and MoUs. A shared understanding of the conditions for success—and progress towards it—relies on clearly defined metrics and their targets.

Several aspects of an IIoT solution revolve around some form of agreement or contract for which an assessment is needed based on agreed metrics and KPIs. Three major aspects stand out as requiring measurements to be properly controlled:

The business value of an IIoT solution. Value needs to be assessed and measured under different angles (see the value areas above), especially in a brownfield environment. Managing trade-offs requires precise measures of these.

Regulatory and policy requirements. Compliance is subject to audits, certifications and routine monitoring, particularly safety, privacy and security and more generally, trustworthiness.

Third-party services and components. As IIoT systems share components and rely on subsystems or contracted services, the performance of these components and services need to be evaluated and how they affect the whole system understood.

[1] See Interdigital oneTRANSPORT™ initiative, https://www.interdigital.com/videos/onetransport-open-marketplace-for-data#

[2] See the "Data sharing agreement (DSA) framework" ISO/IEC 23751 in the last phases of standardization as of March 2020

[3] See the Data trading alliance in Japan, *https://data-trading.org/en/alliance-outline/*

DIFFERENT METRICS FOR DIFFERENT PHASES OF A SOLUTION LIFECYCLE

Metrics serve different purposes for different stages in the solution lifecycle. We distinguish three kinds of metrics:

Profile metrics help determine the profile of a solution defined in terms of a combination of system properties or parameters: data volumes, data flow patterns, connectivity requirements (latency, reliability, scalability), the kind and quality of physical assets involved, the degree of distribution at the edge, the regulatory environment and more. This profiling provides insights to solution designers. To such profiles one can associate known best practices, architecture design patterns and appropriate technologies.

Readiness metrics assess how prepared a solution is to meet expectations prior to operations. They are based on aspects of the solution such as its functional components and capabilities, its architecture and its administrative processes and governance makeup. These metrics are often qualitative and in the form of manual scorecards. Readiness metrics play a major role when assessing or comparing providers. They also play a role in compliance with regulations and policies, in contracts and in the early phases of a system development. They include:

- customer relations, such as customer service,
- risk management, or financial flexibility,
- business, such as financial assessment indicators (billing structure, financial competitiveness, predictability, flexibility) or usability,
- architectural properties, such as scalability, technical portability (ease of migrating a service or resource to a different provider),
- trustworthiness such as the comprehensiveness and scope of maturity (IIC Security Maturity Model [practitioner's guide]*[19]*) and
- organizational, such as alignment of organizations and the ability to manage change.

Several metrics of the service measurement index (CISMIC) *[4]* mentioned in the next section are readiness metrics.

Performance metrics assess how the solution performs over time and whether it meets expectations. They are often based on quantitative measures and automated. They have quantitative targets. Conventional KPIs and ways to measure services, operations or product quality, rely on performance metrics. Several quality metrics in TL 9000 *[5]* detailed in the next section are performance metrics.

Different metrics may be associated with the same system property, depending on when it needs to be assessed in the lifecycle of a solution. Consider the degree of *availability* of a service. A system may be assessed for its *readiness* to ensure service availability and later for its actual availability *performance* at operation time. The following examples illustrate these two notions of availability and their different metrics:

Readiness metric for service availability: an availability assurance metric is used to produce a scorecard based on a rating. This rating, a qualitative value, is relative to the user's expectations. The rating is determined by the following rules:

- 0 if the provider does not commit to any defined availability for the service.
- 1 if the provider's commitment to availability for the service does not meet the customer's required availability level.
- 2 if the availability window defined for the service meets or exceeds the customer's requirements.

This simple metric is used, for example, to evaluate SLAs and to select a service provider prior to deploying a solution.

Performance metric for service availability: such a metric is used to periodically evaluate the availability of a service. This is a quantitative metric, based for example on service uptime percentage as illustrated in the appendix. Other service availability metrics have been defined based on the ratio of failed requests over total requests.

THE ROLE OF METRICS IN MANAGING A DX SOLUTION

BUSINESS MODEL VALIDATION AND ITERATIVE IMPROVEMENT OVER A SOLUTION LIFECYCLE

As the lifecycle of IIoT solutions is iterative in nature, the IIC has identified two high-level phases that are iterated upon over the evolution of a solution:
*Business model validation and improvemen*t focuses on all aspects of monitoring the financial and strategic KPIs of the solution, measuring overall IIoT maturity and implementing corrective actions if needed.

Solution validation and improvement focuses on monitoring and improving the operational side of the solution from the perspective of functionality, non-functional SLAs and SLOs and other system characteristics, including trustworthiness properties.

Figure 2 illustrates this cycle.

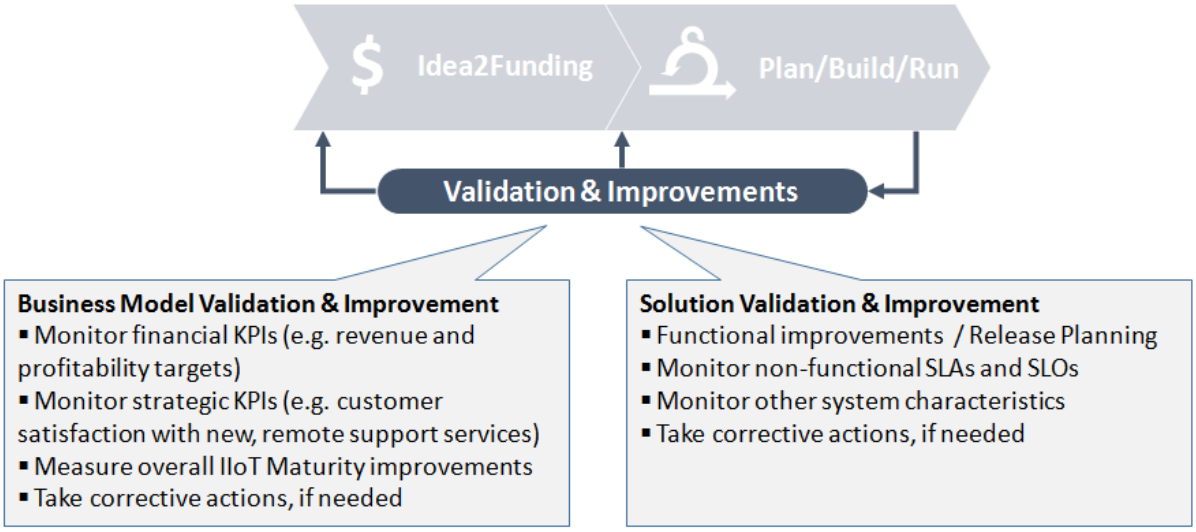

Figure 2: The validation and improvements cycle

In both aspects of the validation and improvements phases shown in Figure 2 (business model and solution), measurements and monitoring play a key role, based on metrics and performance indicators.

BUSINESS MODEL VALIDATION & IMPROVEMENT

Business models associated with digital transformation solutions evolve over time. After the initial rollout, a solution is constantly monitored, and the as-is situation is compared to the initial plan. This usually relies on both financial and strategic KPIs.

Financial KPIs such as revenue, OpEx and profitability need to be closely monitored. Fortunately, in the connected world of the IIoT, this is often much easier than it has traditionally been, and KPIs can be computed in real time. This is especially important for *product as a service* business models.

Strategic KPIs need to be monitored. For example, how satisfied are the customers with the new support services that might make use of remote condition monitoring?

Overall digital transformation maturity should be measured at the portfolio level. Metrics should be defined to evaluate maturity goals and allow projects to compare themselves against them.

These KPIs affect the future evolution of the digital transformation business model and should be considered carefully when prioritizing the requirements backlog. More detailed context is provided in *[12]*.

SOLUTION VALIDATION & IMPROVEMENT

Maintaining and improving both functional performance and adequate levels of non-functional characteristics (such as trustworthiness) depends on how system functions and characteristics are implemented, and on how they are measured and evaluated once the solution is deployed. Metrics need to be defined and based on a set of common norms and measures, agreed upon by all stakeholders. Crucially, metrics should support the overall defined business strategy and business goals, so that managers can be confident that if target metrics are achieved then overall business goals will be delivered.

A consistent set of metrics for both functional and non-functional aspects of an IIoT solution ultimately impact all phases of the solution lifecycle:

- System design relies on profiles (as established by *profile* metrics) that classify IIoT systems based on their quantitative and qualitative requirements: scale, real-time capability, data-intensive aspects, in-and-out data flow characteristics, overall distribution patterns, number and complexity of interaction with physical assets, etc.
- Testing and simulation rely on quantitative measures and tests of functional adequacy. Metrics and targets are essential to analyzing test outcomes, as well as how prepared a system is to handle real conditions when deployed.
- Contracts rely on metrics. SLAs between system end-users and providers, as well as between system administrators and third-party service providers, are based on targets for performance indicators, such as response times, data transfers and QoS, and characteristics such as reliability, timeliness and security (e.g. monitoring of violations).
- Change management is needed as both the operational context and the topology of a system are expected to evolve. Monitoring and testing take place continually using the same metrics involved in previous phases.

SYSTEM CHARACTERISTICS, TRUSTWORTHINESS AND THEIR ASSURANCE

System characteristics are defined in the IIC Industrial Internet Reference Architecture (IIRA, part 2) *[13]* and in the ISO/IEC 30141:2018 Reference architecture *[14]*, as properties of a system such as scalability, manageability, portability and trustworthiness. These properties are subject to assurance procedures both prior to operation and during operation.

SYSTEM CHARACTERISTIC ASSESSMENT OVER THE LIFECYCLE OF A SYSTEM

System characteristics are commonly assessed in three phases in a system lifecycle.

Prior to the deployment or production phase of a system: assessment of the ability of a system (its readiness) to manifest the expected characteristic. Such an assessment relies on readiness metrics which usually apply in two areas:

- *Organizational*: some system characteristics such as those defined in the service measurement index (SMI) of CSMIC *[4]* – financial flexibility, usability and legal portability described in a previous section – rely on organizational procedures. Such procedures in turn involve expert personnel and agencies for their execution as well as for their validation. Auditing and certification are typical validation tools for system characteristics readiness that have organizational requirements. Such requirements are usually stated in regulatory policies and contractual clauses of SLAs. Readiness metrics of qualitative nature are involved here such as the CSMIC SMI qualitative metric for legal portability described earlier.
- *Functional*: what is assessed here are some architectural features in a system that implement expected functions supportive of the system characteristics. For example, assessing the ability of a system to scale will check the presence of scalability enablers such as the capacity of a load balancer, of a cluster of servers and database replicas. The readiness of a system to ensure privacy can be assessed functionally by the presence of a cryptographic components, access to a key management service or a logging service.

Maturity models typically measure readiness both in organizational and functional areas. The IIC IoT Security Maturity Model (SMM) *[18] [19]* considers maturity as appropriate investment according to business need and brings together governance, technology and operations maturity. It combines understanding of process, technological enablement and operations as well as IT and aspects important to OT such as physical security. The IIC IoT Security Maturity Model can consider industry and system scope specifics beyond the general case. Other maturity models include C2M2 for security *[20]* with its *Maturity Indicator Level* (MIL), and the *CERT Resilience* Management *Model* (CERT-RMM) *[16].*

During a testing phase. Does the system manifest the characteristic as expected when under test? Such an assessment relies on performance metrics. Such metrics may be defined on the enablers of the characteristic. For example, the scalability characteristic depends on enablers such as load balancing and a cluster of servers. While a readiness metric is assessing whether a solution architecture exhibits these features, a performance metric is verifying that the load balancing algorithm is doing a good job at allocating workloads under stress testing, e.g. calculating response times for service requests.

Under real operation conditions. Testing conditions only give a partial and approximate rendering of real production conditions. A system may evolve or degrade over time. Does the system manifest the expected characteristic consistently over time in real operational conditions? Such an assessment also relies on performance metrics, and needs to be periodically repeated.

TRUSTWORTHINESS AND ITS METRICS

A set of system characteristics gathered under the term trustworthiness *[17]* is of particular interest to DX solutions that involve IIoT. "*Trustworthiness* is the degree of confidence one has that the system performs as expected. Characteristics include safety, security, privacy, reliability and resilience in the face of environmental disturbances, human errors, system faults and attacks."[1] The five characteristics are defined as trustworthiness characteristics:

- *Safety* ensures that a system operates without unacceptable risk of physical injury or damage to the health of people and indirectly on damage to property or to the environment.
- *Security* protects a system from unintended or unauthorized access, change or destruction. Security concerns equipment, systems and information, ensuring availability, integrity and confidentiality of information.
- *Reliability* describes the ability of a system or component to perform its required functions under stated conditions for a specified period of time. This includes any considerations for physical abrasion, expired software versions, and well-known potential malfunctions that result in frequent maintenance, replacing end-of-life components or software updates. Reliability protects the operation of the system and the system itself, as it is essential for it to be a productive system.
- *Resilience* describes the ability of a system or component to maintain an acceptable level of service in the face of disruption. In contrast to reliability, resilience addresses unexpected and unplanned system statuses that can result from human errors in operation or an environmental event (loss of electric power, earthquake, etc.). The main purpose of resilience is to prevent or at least reduce any serious impact of a disruption to the system by damage or loss of operation.
- *Privacy* protects the right of individuals to control or influence what information related to them may be collected and stored and by whom and to whom that information may be disclosed. Individuals comprise all types of people including customers, guests or employees.

Managing trustworthiness starts with defining criteria for each of its properties in the context of a particular solution and understanding at a high level where to invest resources. In the IIC IoT Security Maturity Model (SMM) *[18] [19]* various domains, subdomains and practices have been identified along with maturity comprehensiveness levels and scope. This provides a means to set a maturity target for each domain, subdomain and ultimately practice with business stakeholders and then determine gaps based on an assessment. Gaps may be addressed by understanding the

[1] Industrial Internet Consortium: Vocabulary, V2.1, August 2018, *https://www.iiconsortium.org/vocab*

associated choices and controls, including the techniques to improve the security or safety level for those practices. Safety and privacy have general assessment guidelines and objectives that are driven by regulation or industry-wide policies and apply to a broad set of systems. These objectives are expressed in terms of security maturity targets for an organization described in in the IIC IoT Security Maturity Model and subsequently as detailed metrics related to corresponding controls and events to be managed.

Metrics can be defined associated with operational business concerns, such as the percentage up-time availability of a service (for its reliability assessment), or the number of stress injuries per month on a machine (for its safety assessment). Trustworthiness objectives translate into targets for these metrics. Several metrics are used to provide a well-rounded assessment of a particular trustworthiness property or practice. Most industries also track security metrics such as the number of detected attack attempts, reporting on the breakdown of those attempts, and categorizing them into successful attacks, incidents, close calls, policy violations and anomalies that have merited investigation. The information from these metrics relates to the evaluation of the associated controls and practice maturity.

For resilience, the CERT Resilience Management Model (CERT-RMM) *[16]* provides a resilience model and basis for a basket of metrics for assessing the resilience of a system. *(See RMM-MUG and the CMU study: https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=30375).*

To be better controlled, these high-level indicators need to be complemented with system-specific metrics that report on actionable components of a system. For example, if safety objectives and measurements can be stated at a high level (such as measured in terms of overall personnel injuries occurring within a month or a year), controlling the factors that influence the output of such metrics will involve finer granularity metrics that are specific to the operational context, such as a metric on the stress injuries caused by a specific machine, or a metric that anticipates likely stress injuries by detecting an increase in error rates or interruptions manifested by personnel on that machine.

DEPENDENCY OF SOME TRUSTWORTHINESS METRICS ON PERFORMANCE METRICS

Safety, security and privacy objectives may generally be stated independently from a specific operational context, even if they need be complemented and measured by system-specific metrics. Reliability and resilience characteristics are more tied to a particular system and to its performance expectations. Consequently, metrics for these are often expressed in terms of expectations for related performance metrics.

Consider a performance metric P_m for a particular assembly chain, based on the actual processing time of an order on the assembly chain, compared with expected time:

Operational performance metric P_m: how good is the performance of this assembly chain in processing a batch of products (production lot)?

Assuming a case where a product unit is expected to be fully processed every 5 minutes, as an example a performance metric can be expressed as the ratio of the total expected time to process the lot of products, over the actual processing time for the lot. Ideally this ratio is 1 (or greater). In case the production takes twice as much time as expected, the ratio is 0.5. Assume a target of 0.9 is given to this ratio, meaning the processing time should not exceed roughly 110% of the target time.

The performance metric formula for this particular assembly chain is:

Performance metric P_m =

(expected time to process a production lot)/ (actual time it took to process the lot)

or:

(planned machines configuration time + (lot size*300 sec)) / (actual time it took to process the lot)

Target: $P_m >= 0.9$ (monthly average)

While the performance metric P_m for this assembly chain measures an *operation* and its effectiveness, a reliability metric for this assembly chain is measuring a *property*, although the latter depends closely on the former:

Reliability metric Rel_m: what *assurance* do we have that the assembly chain will process an order with expected performance (P_m >= 0.9)?

This metric measures the level of performance *assurance* (not the performance itself), in other words the likelihood that the next manufacturing order will be processed within acceptable time, meaning with a performance of 0.9 or more, using the previous P_m performance metric. This likelihood is simply measured here as a percentage of previous orders where the assembly chain showed acceptable performance. Assume that a reliability target of 95% is given to this percentage.

Rel_m = percentage of cases where P_m > 0.9, for the last 100 orders.

Target: Rel_m >= 95%

Similarly, a resilience metric for this assembly chain is measuring a system *property*, although this property, as for reliability, depends closely on the performance metric P_m :

Resilience metric Res_m: what *assurance* do we have that the assembly chain can quickly reconfigure or recover with limited performance degradation under hardship?

What is measured here is the *degree of preserved performance under hardship*, or the *immunity* to hardship. One possible measure for the resilience of this particular assembly chain (which may not be suitable for other systems) is:

Res_m = 1 – (avg (P_m under normal conditions) – avg(P_m under last 10 severe incidents))

The difference between the two P_m averages (in parentheses) may be defined as the degradation in performance under challenging conditions. In an ideal case where there is no performance difference, i.e. the difference is null, the maximum value for the Res_m metric is 1. If performance has been observed to decrease from P_m = 0.9 to 0.5 under hardship, the resilience according to this metric is 0.6 while it would be of 0.8 if degrading from P_m = 0.7 to 0.5, a case where performance was not initially so good. As an example, a target of 0.8 may be assigned for resilience.

Target: Res_m >= 0.8

These resilience and reliability metrics both depend on the same performance metric P_m , yet express distinct and unrelated properties: manufacturing equipment could exhibit very good performance, yet very poor resilience or reliability, or good reliability and poor resilience or vice versa. The metrics in this example are also system specific or ad hoc for this particular system

(here an assembly chain for a specific product) as this will often be expected from IIoT system metrics even if the system is given general trustworthiness objectives such as determined by maturity models.

MANAGING TRADE-OFFS BETWEEN TRUSTWORTHINESS OBJECTIVES AND BUSINESS OBJECTIVES

Different objectives may be associated with the previous value areas of an IIoT solution. These objectives may conflict, or the means to reach one of these may adversely affect the other.

The role of metrics for managing trustworthiness properties and their impact on business and operational objectives has been studied in *[10]*.

Consider the value areas of better risk management and process efficiency. There is a known potential conflict between some trustworthiness objectives such as for security and safety (minimizing risk), and the operational performance of an IIoT system (business process efficiency). For example, increasing the speed of an assembly chain would improve its productivity, but may reduce its safety for the operating personnel. Metrics will assess progress toward objectives in each value area, and also provide insights into how these objectives affect each other.

Consider an assembly chain in a factory. It is subject to two objectives:

- a safety objective, which consists of keeping the level of stress injuries due to operating machinery below some threshold and
- a business objective, which consists of successfully processing a certain amount of production lots over a period of time.

Consider S_m the safety metric calculated as the ratio of actual working days free of stress injuries and accidents for all personnel involved on the assembly chain, over the total of all working days expected from the same personnel on the chain. The maximum value for this metric is 1 (no days off due to injuries).

The business objective has its own performance metric: P_m is roughly the number of products processed over a time period (as defined in the example of a previous section). A known factor of improvement for this metric among others is the speed of the assembly chain.

In addition to measuring progress toward each objective, these metrics will provide insights on dependencies such as how the frequency of stress injuries to personnel correlates with the speed of an assembly chain.

The curves in Figure 3 represent two cases of adverse dependency between the degree of safety of an assembly chain and its performance, as measured by respective metrics S_m and P_m, and how this affects the ability of the system to reach both its safety and performance objectives.

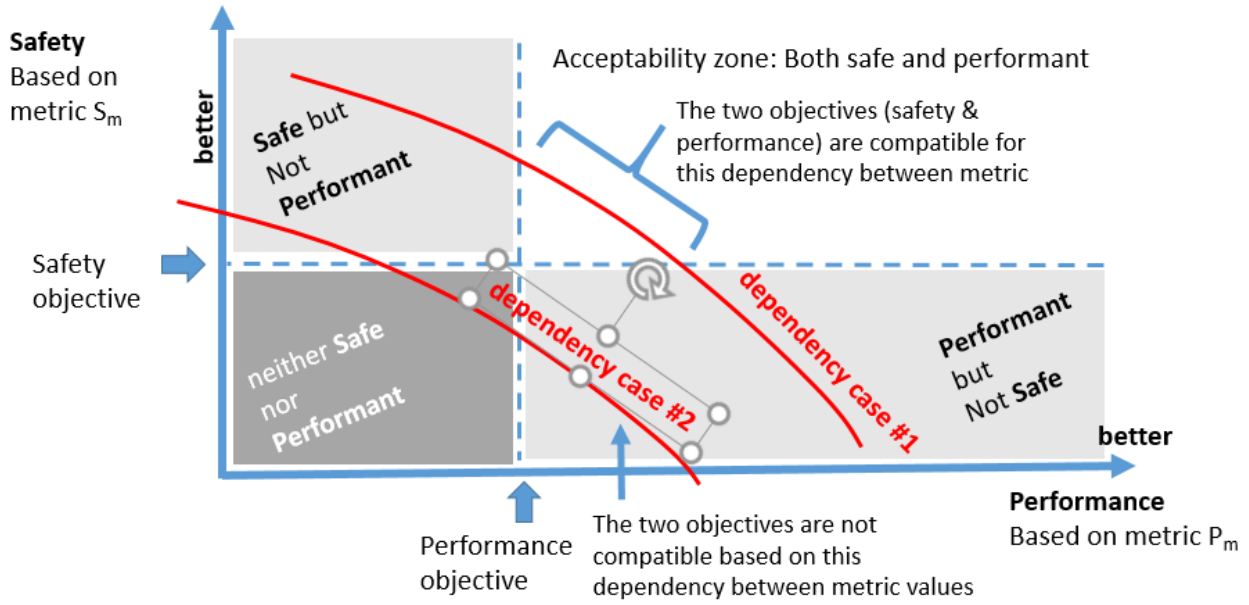

Figure 3: Dependency and compatibility between objectives as measured by metrics

The first dependency case shows a situation where safety and performance as defined by their metrics, adversely affect each other yet in a way that keeps their objectives still compatible. There is an area where both safety and performance objectives (defined as threshold targets for their respective metrics S_m and P_m) can be satisfied. In contrast, the dependency case shows that the two objectives as measured by these metrics are not compatible; reaching the target of one metric will cause failure to reach the target of the other metric.

As various trustworthiness objectives are assigned to IIoT systems along the value areas, progress toward these is measured using a set of appropriate metrics. These metrics go beyond assessing whether trustworthiness objectives are met, they also help assess how these objectives affect each other as well as business performance. Metrics allow for the understanding of hidden dependencies between seemingly unrelated objectives, and for managing trade-offs between them when they conflict.

Figure 4 illustrates the *trustworthiness space* in a case where three trustworthiness properties are of interest: safety, security and reliability. The acceptability zone of this space is where the trustworthiness objectives are satisfied for all properties based on their metrics. Each trustworthiness property in itself will likely have several dimensions (i.e. define a space by itself, such as the aspects of security covering various areas of a solution) but is represented here linearly for simplicity.

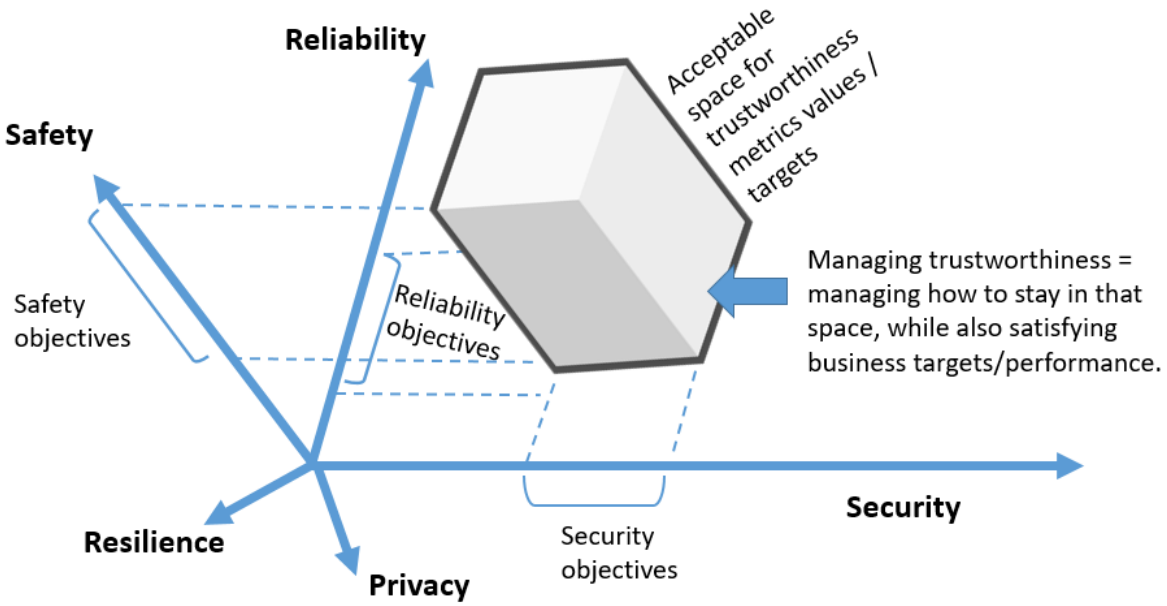

Figure 4: The Trustworthiness space as defined by its metrics

Operating an IIoT solution while keeping it within the trustworthiness acceptability zone requires additional considerations, as trustworthiness properties also affect business performance.

METRICS FOR ASSISTING THE DESIGN, DEVELOPMENT AND EVOLUTION OF AN IIOT SOLUTION

THE PROFILE OF AN IIOT SOLUTION

The diversity of IIoT solutions makes it difficult to identify and capture best practices in design architectures and appropriate technologies. This diversity partly reflects the variety of usage requirements and contexts. Metrics can capture and characterize this diversity in architectures and usage requirements. As requirements and contexts evolve, metrics can assess them periodically.

Metrics used to establish the profile of an IIoT solution are capturing two aspects of this solution:

The *functional aspect* evaluates the functions of the system and their capacity under various perspectives: performance, throughput, data volumes, transfer time, connectivity, quantity of assets, etc.

The *system characteristic aspect* assesses higher-level properties such as trustworthiness properties and architecture-related properties such as scalability, modularity and adaptability.

The functional aspect can be divided into specific areas such as physical assets, data or connectivity. A cluster of indicators may be assigned to each one of these areas to produce a well-rounded profile of the system for this particular area.

An *asset profile* may involve the following indicators:

- quantity of assets to handle at any given time,
- complexity and heterogeneity of assets,
- complexity of assets control (monitoring/tracking, control functions),
- level of compliance/regulatory requirements for deploying asset control,
- maintenance/deployment costs and
- human training and assistance required.

A *data profile* may involve the following indicators:

- importance and complexity of data management,
- data volume generated,
- velocity, at any time,
- variety / heterogeneity,
- long term storage and archival needs and
- level of compliance/regulatory requirements for collecting this data.

A *communication and connectivity profile* may involve the following indicators:

- real-time communication requirements (latency, jitter) for applications,
- dependency on existing networks (Internet, 3G/4G/5G),
- WAN level: required QoS level and bandwidth,
- LAN level: local connectivity complexity and heterogeneity and
- importance of human communication (human-to-human or machine-to-human)

Determining these profiles requires both qualitative and quantitative metrics. Once a profile is defined for a solution, values for the parameters of the profile can be seen as a form of expression for the system requirements, to which known design patterns for this profile can be applied. Figure 5 shows a visualization of these profiles using a starfish representation (Kiviat diagrams).

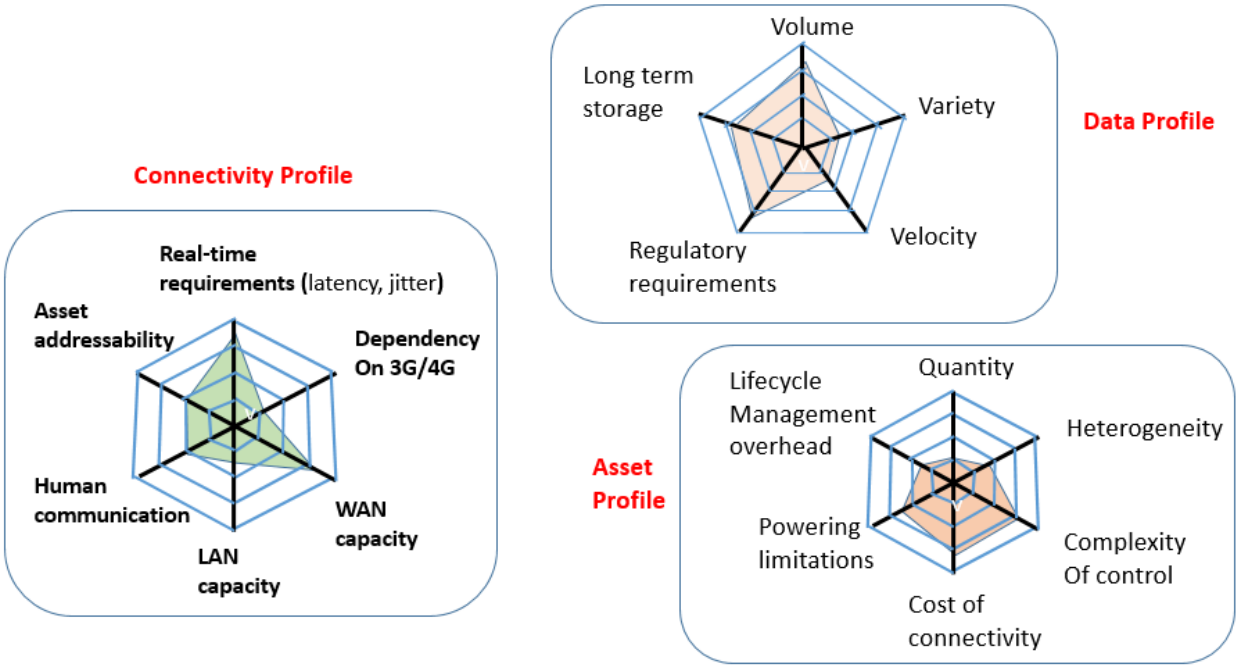

Figure 5: Some profile dimensions of an IIoT system

The profile of an IIoT solution may evolve, particularly in its quantitative properties. The real-time requirements or the network capacity of a connectivity profile may become inadequate over time. The characteristics of a data profile may change, data volumes may increase or the system is deployed in a region with different regulatory requirements. Metrics help capture the requirements and the trends in these areas.

Although every combination of these dimensions and their values is possible in theory, typical profiles have been reported in *[15]* that match specific combinations of indicator values:

Real-time systems that have precise requirements about guaranteed latency and response (often more important than speed). Many systems require low average latency (delivery delay), but real-time systems succeed only if they always respond on time (bounded by maximum latency, or *jitter*).

Data distribution focus systems are those that inform more than 25 data recipients when single data items are changing.
Data collection focus systems transmit significant information generated at the edge to be stored or analyzed in higher-level servers or the cloud. Systems that are restricted to the data collection pattern do not share significant data between devices, but they must move information to a common destination efficiently. A metric of interest here will be the concentration ratio of one-way data flows from a number of edge sources, defined in terms of data aggregation or consolidation. A concentration of more than 100 sources indicates a data collection system.

Such profiles have significant architectural constraints and best practices associated with them. For example, an architecture that can satisfy a human user willing to wait no more than 5 seconds for a web site will never satisfy an industrial control that must respond in 2ms. For a real-time system that is also a distributed system the most important architectural issue is the potential jitter (maximum latency) imposed by a server or broker in the data path.

THE PROJECT EXPLORER SOLUTION PROFILING TOOL

The IIoT Project Explorer[1] assessment tool developed by the IIC evaluates the profile of a solution prior to its design and development. It relies on a set of integrated profile metrics covering the various perspectives of a solution and captures a set of characteristics that can be seen as requirements. This allows for evaluating the resources needed by project managers to develop the solution and for associating best practices and technologies with the solution based on its profile.

Five major perspectives or profile dimensions of an IIoT solution have been identified in this tool:

Project environment captures contextual aspects of the project such as its operational environment and constraints of various nature: budget, timing, regulations and skills set required.

System-wide challenges captures overall system expected properties such as dependability, availability and expected end-to-end connectivity characteristics.

Field assets and devices reflects on the physical assets on the edge and devices with a set of indicators that characterize a type of asset. This dimension may be repeated, as a solution may involve a fleet of different types of assets and devices.

Backend services and their access reflects the properties of the service endpoints in a solution such as data volumes and management constraints, processing and analytics, and type and complexity of the application. These may reside on corporate servers, in the cloud or in other servers such as fog nodes. This dimension may be repeated for various endpoints.

Business model and requirements captures business expectations for the solution, such as the precision and stability of requirements, the targeted level of productization, the regions of deployment and functional complexity.

Each one of these areas groups a small set of dimensions. In turn, each dimension is captured by a set of indicators.

[1] *https://www.iiconsortium.org/project-explorer.htm*

Each indicator is given a range of four values regardless of its quantitative or qualitative aspect. Its value ranges from 1 (low/simple) to 4 (high/complex/challenging)

For example, the *technical skills and experience* indicator is rated as follows:

Indicator	1	2	3	4
Technical skills & experience	Existing team, has done similar project before	Like 1, but geographically distributed	Completely new team, individual team members have little technical experience in relevant area	Like 3, but distributed

Under the *assets and devices* profile dimension, the *number of assets* indicator is mapped to this range as follows:

Indicator	1	2	3	4
Number of Assets	<100	100s or 1000s	10.000s	Millions

While the operating environment (of the asset) qualitative indicator is rated as follows:

Indicator	1	2	3	4
Operating Environment (of the asset)	Indoor	Rough Indoor, e.g. factory	Outdoor, moving (e.g. car in winter)	Critical conditions, e.g. aircraft, space

All the indicators of a particular project dimension (project environment, field assets and devices, etc.) can be rolled up into a summary assessment for this dimension. For example, an assessment of 3 for the *project environment* dimension indicates strong constraints associated with this dimension, requiring management attention and specific expertise, while an assessment of 1 indicates non-significant constraints.

Based on the summary rating of each of the five dimensions a general profile of the IIoT project can be visualized with a Kiviat diagram (Figure 6) and compared with other projects' profiles.

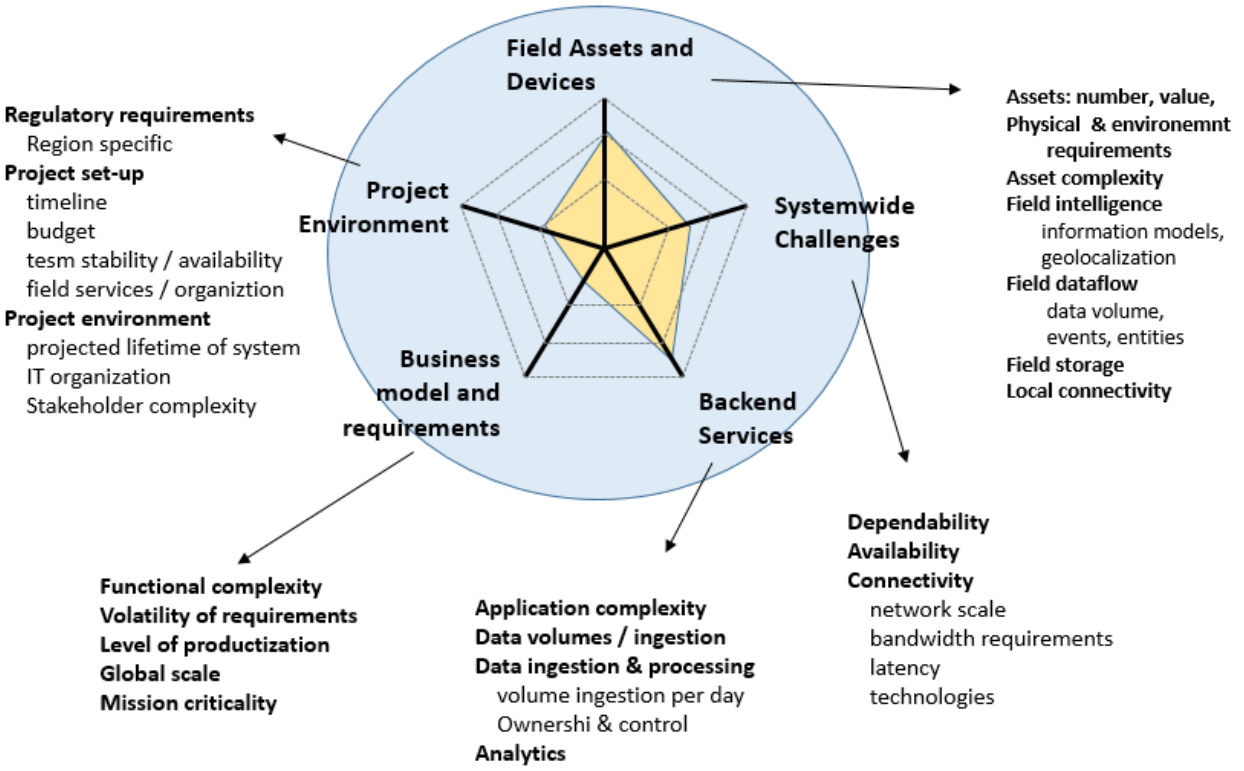

Figure 6: General structure for assessing an IIoT project profile

As a system or its context of operations evolves, it may be necessary to assess if the system still conforms to its profile during operations, or if its current profile is still adequate over time.

As the system evolves, the indicators can be updated. The evolution of a system can then be represented as a succession of such profiles.

Indicators may be grouped across perspectives to provide a well-rounded rendering of a particular aspect of the solution. Table 1 is a gathering of some of the indicators that capture the real-time dimension or requirements of a system.

Indicator	1	2	3	4
System-wide data synchronization requirements	Daily batch synch	synchronization within minutes	synchronization within seconds	synchronization within sub-second
Control loop latency at the edge	>10 ms (e.g. RS 232)	1-10 milli seconds (e.g. WLAN, BlueTooth)	micro seconds (e.g. EtherCAT, Sercos)	nano seconds (e.g. ASIC, FPGA
Service connectivity maximum latency	>10 ms (e.g. RS 232)	1-10 milli seconds (e.g. WLAN, BlueTooth)	micro seconds (e.g. EtherCAT, Sercos)	nano seconds (e.g. ASIC, FPGA)
Global maximum latency	90 Min (LEO, e.g. OrbComm; text messages)	seconds (GPRS)	milli seconds (WAN)	micro seconds (e.g. LAN)

Table 1: Some indicators used for assessing the real-time profile of a system

These indicators may be coupled with performance metrics used to re-evaluate a profile. Performance metrics define the details and protocols to measure an indicator, such as the real-time properties of the system in operation, answering questions such as how will response times be measured in practice? Under which conditions? Based on which events?

THE ROLE OF METRICS IN SYSTEM INTEROPERABILITY AND SERVICE COMPOSABILITY

SERVICE COMPATIBILITY AND INTEROPERABILITY

As DX solutions increasingly rely on contracted services, such as data storage, networking, device management, security, AI model training or entire data platforms such as the transport data marketplace of the OneTRANSPORT™ initiative for smart cities. *[2]* Using third-party services causes a fragmentation of ownership and governance even within an organization. Consider asset or product tracking in a manufacturing plant that was originally designed for managing production. This tracking can be provided as an internal service used by other units or departments, such as shipping/receiving, inventory or equipment maintenance. These services and their quality need to be monitored.

More generally, the evolution toward value networks as opposed to independent, siloed value-chains means that a provider in the larger value network has responsibilities to several consumers, possibly with different quality requirements. Conversely, a service user (either a person or a system) expects a service to be substitutable. This requires more than standardized interfaces and compatible data models and protocols. Such flexibility requires a common understanding in the modes of usage, service quality and monitoring procedures that are supported by a service or a component, which is a prerequisite for an agreement about them. Contracts are key to establishing service objectives and responsibilities between parties and become significant for enabling the integration of services.

Consider a system where edge devices periodically invoke a data service directly to store data generated, say, every five minutes at most. When these devices have the capacity to handle a backlog of only up to ten minutes of data stream, there might be loss of data if the data storage service is down for more than five minutes. A SLA with the storage service provider should then give the latter strong incentives to keep outages shorter than five minute, by including steep penalties when outages exceed this limit.

However, if the service provider cannot commit to this level of service, but operates on a standard SLA common to a class of customers, the solution designer may have to modify the system architecture to mitigate the risk of data loss by adding caching capability on the edge or by using an alternate logging system as backup in case of long-lasting downtime.

SHARING AND REUSING METRICS

Standardized SLAs and metric representations, such as the standard metric model ISO/IEC 19086-2:2018 *[9]* is the other side of enabling interoperability that relies on contracts and on a common understanding of service quality and compatibility. Standardization promotes service composability by making SLAs and metrics themselves reusable. Reusability of metrics also about ensures a common way to measure quality and performance of services and components.

Customers and users expect to find similar measurement definitions across providers and their SLAs. Regulators need consistent and measurable interpretations of system characteristics such as security, privacy or safety. Tools vendors and system administrators will be expected to implement similar monitoring indicators and technologies across IIoT solutions.

A library of metric definitions or templates that providers and operators can share will help this goal. Figure 7 illustrates the use of metrics from a library, serving different purposes over the life of an IIoT solution.

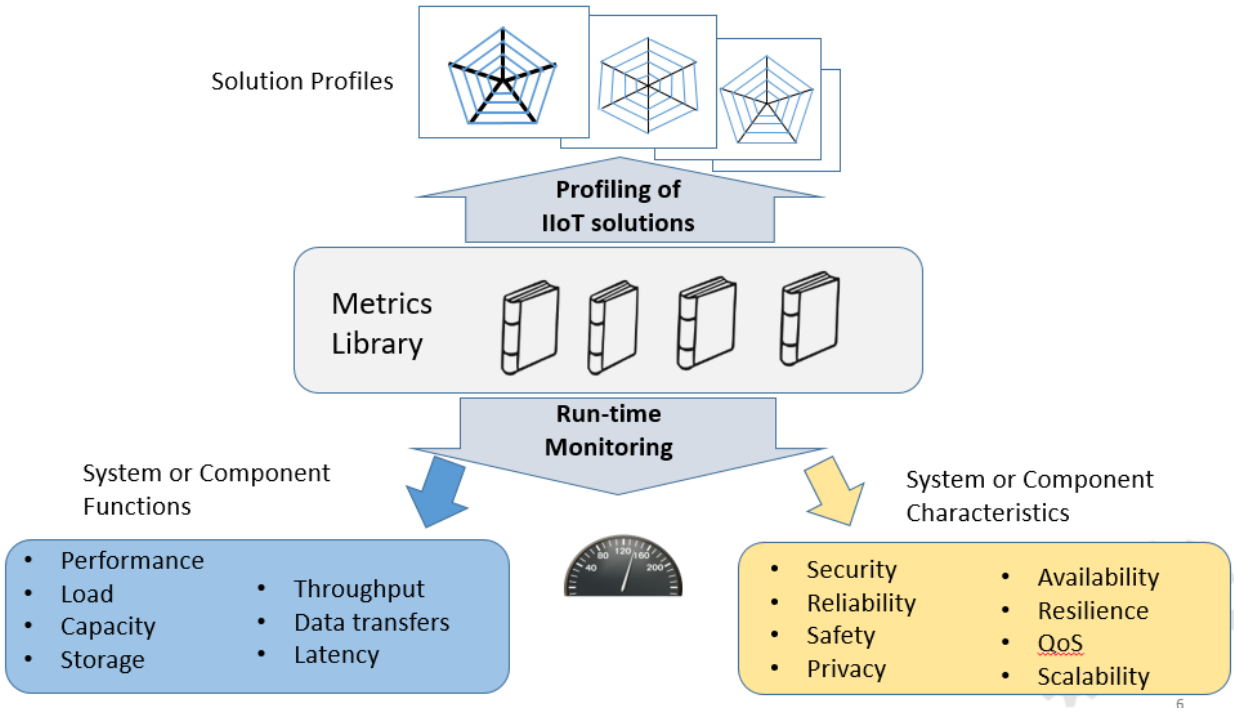

Figure 7: Metrics Library and its usages

A library *of* metric templates and definitions will contain different types of artifacts:

Metric templates are predefined frameworks to create metric definitions. They will follow SLA standards like ISO/SC38 19086 part 1 *[8]* and in particular its metrics model (part 2) *[9]*. They may also follow such templates as IIC-defined Project Explorer for IIoT solution profile assessment (profile metrics).

Metric definitions are precise definitions of metrics, for solutions developers to download for a shared understanding of how services are to be measured.

Metric elements are precise definitions of metrics parts such as rules or calculation logic, as in many cases only these parts are shared by users who need some degree of customization (see metric foundations in a previous section).

Metric implementations are in the form of executable code or monitoring components. They are made available to solution designers. These may be developed and made accessible as open-source.

Metric libraries help normalize the definition of metrics and promote a common understanding of system properties, quality and performance between various stakeholders in an IIoT solution: end-users, vendors, solution developers and system integrators.

THE VALUE OF STANDARDIZING METRICS AND INDICATORS

Ideally a common set of metrics should be used across DX solutions in a given industry sector, based on some canonical ways to measure the business value areas. While this goal may need more time and experience to be realized, a first step is a common way to *define* metrics (common templates, terminology and models, as illustrated in Tables 4, 5, 6 and 7 of the section "The Metric Model Standard of ISO-IEC JTC1/SC38" in the Appendix). Customers and users will expect to find well-established definitions across SLAs. Experts and regulators expect stable and well-understood definitions of system characteristics. Tool vendors and system administrators want to implement similar monitoring indicators and technologies across IIoT solutions.

In cloud computing, consortia and agencies like CSMIC, NIST and ISO/IEC-JTC1/SC38 have been working on standardizing metrics independently of their domain of application. They are identifying a catalog of common metrics (CSMIC) and promoting the metric description model itself (NIST and SC38). As support for this promotion, guidelines for a practical usage of the ISO/IEC metric model standard have been collected in the technical report ISO/IEC TR23951 *[11]*.

The benefits of a using a common metric model are:

To clarify the metric descriptions in SLAs and other agreements and make comparisons easier. It is unlikely that service providers will share exactly the same metrics, even for common measures such as service uptime percentage. Typically, an SLA will define a combination of metrics. Often, the metric definition, say for measuring service availability, is scattered over an SLA narrative and mixed with related information that is not part of the metric definition per se (like performance objectives or targets, remediation and penalties). Extracting the actual definition of what is measured and how, then representing it in a template with explicit structure and terminology has proven to be of great value to understand and compare the metrics used across providers.

To support the creation of metrics. Operation managers, engineers, SLA writers and auditors need some framework to describe and design new metrics. A metric model or structure helps define a sharable representation understandable by all. It also helps detect missing components.

To develop common metric foundations. It is desirable to share the same metric conventions and elements, if not the same metric. These are expressed as a partially developed metric definition, called a metric *foundation.* An example of foundation is a metric definition that abstracts the details and constrains only the general logic of the metric calculation. It allows service providers to define their own parameters and rules, but within some limits. Or, a metric foundation includes predefined metric elements that serve to harmonize concepts and terminology across users. For example, in case there is agreement for sharing across providers common rules defining notions such as "service downtime" or "service misuse", these become shared metric elements. A common metric foundation promotes consistency across metrics definitions.

CONCLUSIONS AND OUTLOOK

This paper described various perspectives and uses for metrics in digital transformation solutions, while reflecting on existing work done with metrics in related areas, from large scale service measurement to standardization.

Metrics and various KPIs have been used for a long time. Operational performance, product or service quality, business effectiveness and properties such as security, safety or reliability have all been commonly measured and subject to metrics and targets. However, these metrics have often been designed for a narrow purpose and have served a rather accessory monitoring role. Digital transformation solutions evolve in complexity and interdependency and the IIoT physical and operational context they derive their insights and value from is always changing. It is expected that metrics and related monitoring functions will play a more active and dynamic role for composing, configuring, controlling and managing these systems as they evolve.

APPENDIX: A SURVEY OF EXISTING WORKS ON METRICS IN DX RELATED AREAS

QUALITY METRICS FOR NETWORK CARRIERS AND MOBILE DEVICES

A STANDARD-BASED, LARGE-SCALE QUALITY MANAGEMENT SYSTEM

This example illustrates an elaborate use of metrics within a large scale QMS deployed to assess various aspects of mobile communication services.

The Information and Communications Technology (ICT) industry realized as early as the 1990's that a sector-specific quality management standard was needed to improve the quality of the telecom network equipment and services. This resulted in TL 9000 standards (see *[5]* for an informal and general introduction) where 'TL' stood for Telecommunications. Since then it has evolved as an ICT sector-specific quality management system (QMS) standard, making it possible to track and compare products and related services across providers.

The TL 9000 standard comprises two parts: a Requirements Handbook and a Measurements Handbook. The Requirements Handbook is founded on ISO 9001 quality requirements *[6]* and more than 80 ICT-specific requirements. It establishes a common set of quality management system requirements for suppliers of ICT products: hardware, software and services. The Measurements Handbook is based on a set of performance quality metrics. Figure 8 shows the scope covered by the two handbooks of TL 9000.

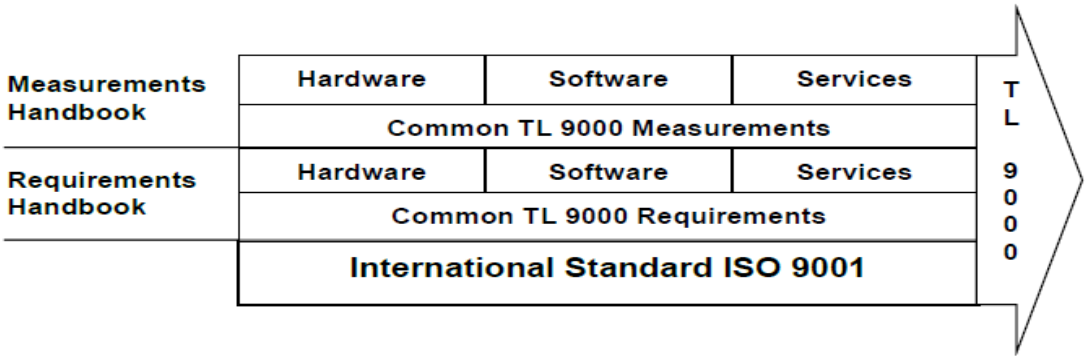

Figure 8: The structure of the TL 9000 standard[1]

Any company claiming compliance to TL 9000 must comply with both requirements and measurements parts. Measurements compliance requires monthly field-performance data to be submitted to a third-party central repository managed by the University of Texas at Dallas (UTD).

Organizations registered to TL 9000 collect, validate and submit data per the defined measurements definition to UTD using a double-blind system that maintains anonymity of data. Since then companies all over the world have submitted monthly data to UTD.

METRICS AND THEIR OUTCOMES

Table 2 shows some examples of metrics the results of which are reported monthly.

Performance Area	Metric
General Category	Number of problem reports Fix response time On-time delivery
Outage measurements	Network element impact outage Support service caused outage Mean time to restore service
Software measurements	Software fix quality Software problem reports
Hardware measurements	Equipment return rate

Table 2: Examples of monthly reported TL 9000 metrics

Detailed calculation formulas and counting rules for each measurement are provided in the TL 9000 Measurements Handbook.

UTD calculates the performance data reports for each measurement by product category using the appropriate data elements for each compared data measurement. Data reports may be published if there are valid data submissions from three or more companies.

[1] Source: QuEST Forum

UTD publishes monthly charts showing best-in-class, worst-in-class and industry-average values without company names. These charts are accessible to member companies and TL 9000 registered companies. Best-in-class is the best performance from a single registration for a product category for a particular measurement. For some measurements, the optimum performance is zero while for others it is 100%. Worst-in-class is the worst performance from a single certified registration. Industry average is the composite average of data from all eligible submissions over a defined period.

Figure 9 shows an example of such a chart.

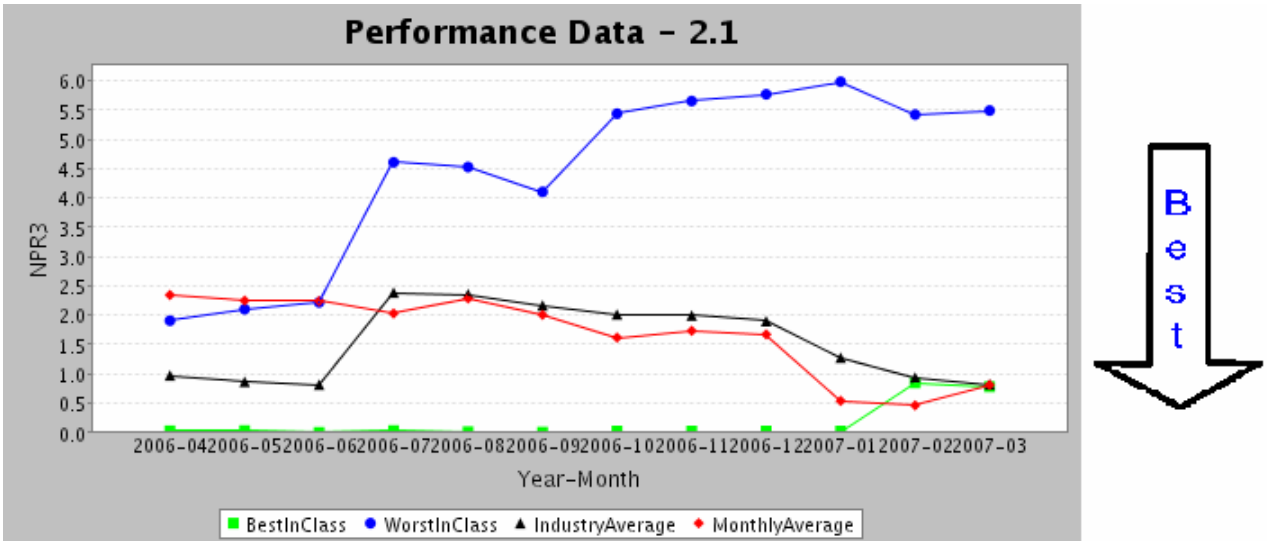

Figure 9: Chart of monthly performance reports[1]

UTD creates benchmarking charts with best-in-class, worst-in-class and industry-average values and makes them accessible to QuEST Forum *[7]* members and TL 9000 registered companies.

Best-in-class, worst-in-class and industry-average calculations for a particular month are based on data submitted from a smoothing window of either 6 or 12 months depending on the measurement. Monthly average is derived from data submitted for a single month.

With this measurement a company can compare its performance with the rest of the industry, prioritize improvements and determine optimum use of resources for performance improvement in areas that matter most to their business. The QMS has been designed to scale, there are over 160 product categories for which data is submitted by over 650 companies globally and from 1500 locations worldwide.

[1] Source: QuEST Forum

MANAGING SERVICE PROVIDERS

In addition to assessing monthly performance, communications service providers use TL 9000 measurements data to manage their supply chains and to benchmark performance against suppliers of similar products. Figure 10 shows an example of a chart prepared by a service provider to monitor performance of their suppliers by tracking key TL 9000 metrics. Green cells show those metrics that meet or exceed target. Yellow cells show values that are showing a trend, which is likely to move in the wrong direction. Red values indicate that the supplier is underperforming in those areas and needs improvements.

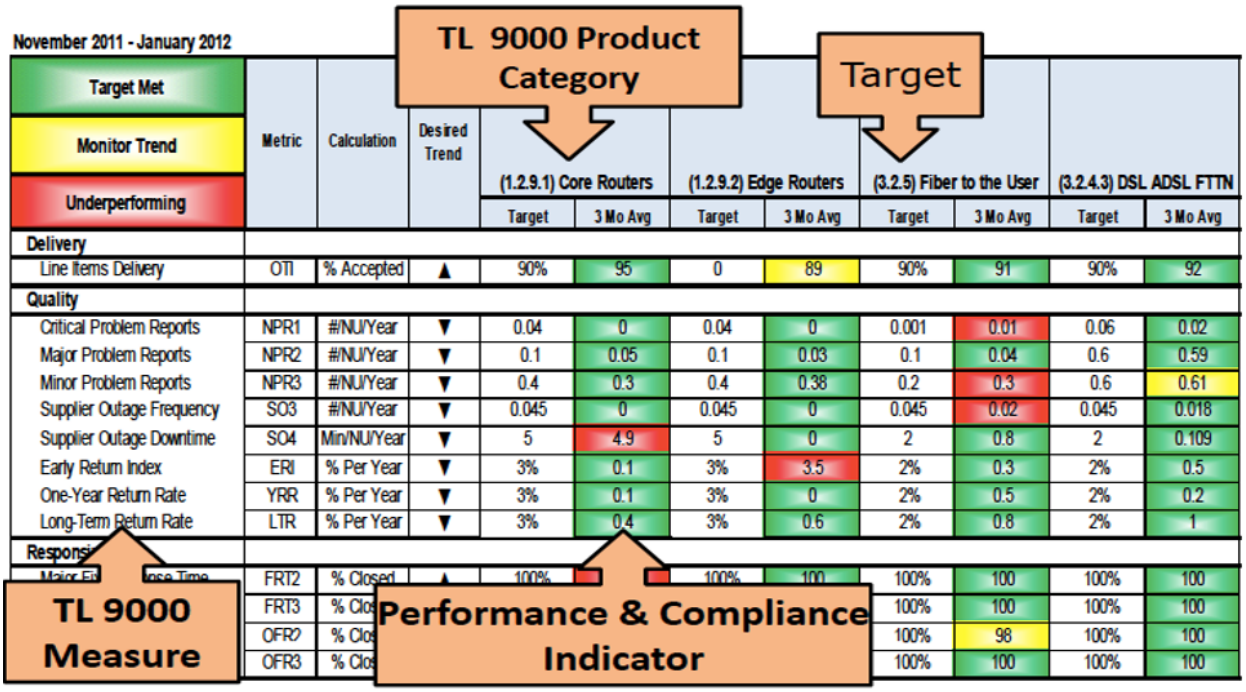

November 2011 - January 2012

Target Met | Monitor Trend | Underperforming

	Metric	Calculation	Desired Trend	(1.2.9.1) Core Routers		(1.2.9.2) Edge Routers		(3.2.5) Fiber to the User		(3.2.4.3) DSL ADSL FTTN	
				Target	3 Mo Avg	Target	3 Mo Avg	Target	3 Mo Avg	Target	3 Mo Avg
Delivery											
Line Items Delivery	OTI	% Accepted	▲	90%	95	0	89	90%	91	90%	92
Quality											
Critical Problem Reports	NPR1	#/NU/Year	▼	0.04	0	0.04	0	0.001	0.01	0.06	0.02
Major Problem Reports	NPR2	#/NU/Year	▼	0.1	0.05	0.1	0.03	0.1	0.04	0.6	0.59
Minor Problem Reports	NPR3	#/NU/Year	▼	0.4	0.3	0.4	0.38	0.2	0.3	0.6	0.61
Supplier Outage Frequency	SO3	#/NU/Year	▼	0.045	0	0.045	0	0.045	0.02	0.045	0.018
Supplier Outage Downtime	SO4	Min/NU/Year	▼	5	4.9	5	0	2	0.8	2	0.109
Early Return Index	ERI	% Per Year	▼	3%	0.1	3%	3.5	2%	0.3	2%	0.5
One-Year Return Rate	YRR	% Per Year	▼	3%	0.1	3%	0	2%	0.5	2%	0.2
Long-Term Return Rate	LTR	% Per Year	▼	3%	0.4	3%	0.6	2%	0.8	2%	1
Responsi											
Major Fi… …ose Time	FRT2	% Closed	▲	100%		100%	100	100%	100	100%	100
	FRT3	% Clo						100%	100	100%	100
	OFR2	% Clo						100%	98	100%	100
	OFR3	% Clo						100%	100	100%	100

Figure 10: TL 9000 supplier executive dashboard[1]

DX RELEVANCE

IIoT systems that enable digital transformation have similarities with the TL 9000 QMS mobile telecommunication ecosystem of equipment and services, in particular those IIoT systems that:

- involve a mix of hardware equipment, devices, networks, servers and software assets,
- involve fleets of devices and physical assets on the edge, possibly largely distributed and managed separately by multiple providers,

[1] Source: QuEST Forum

- require monitoring for assessing overall performance as well as good operational condition of devices and sensors in addition to physical assets (using performance metrics),
- rely on a supply-chain of services and components that are often contracted out to third parties, of which reliability and service quality must be regularly assessed and
- are subject to policies and regulations which involve in turn some auditing and qualitative assessment of the preparedness of a system and its procedures (using readiness metrics).

Such properties seem relevant to industry sectors that involve large scale distribution, fleets of devices and a chain of providers such as energy and utilities and transportation or distributed home care systems. Such IIoT systems can learn from the TL 9000 QMS.

THE SERVICE MEASUREMENT INDEX (SMI) FROM CSMIC

A METRIC FRAMEWORK FOR EVALUATING CLOUD SERVICES

The Cloud Services Measurement Initiative Consortium (CSMIC) *[4]* has developed an evaluation framework for cloud-based services, the *Service Measurement Index* (SMI). It is a set of business-relevant key performance indicators (KPI's) that provide a standardized method for measuring and comparing a business service regardless of whether that service is internally provided or sourced from an outside company. Most of its metrics are *readiness* metrics.

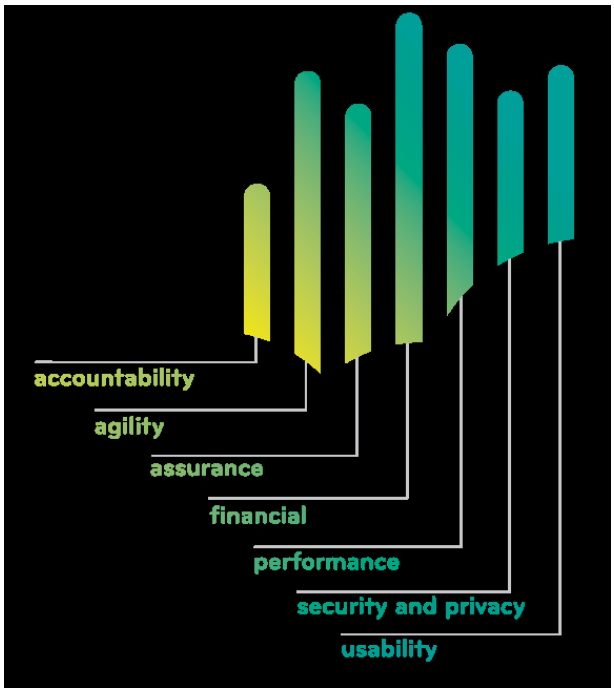

Figure 11: Major business service properties to be assessed[1]

[1] Source: CSMIC

CSMIC has identified seven major service properties shown in Figure 11:

Accountability contains attributes used to measure the properties related to the cloud-service-provider organization. They may be independent of the service being provided. Some of the attributes are auditability, compliance, contracting experience and provider SLA verifiability.

Agility is the impact of a service upon a client's ability to change direction, strategy or tactics quickly with minimal disruption. Some of the attributes are service scalability, extensibility, adaptability, flexibility and portability of a service across providers.

Assurance indicates how likely it is that the service will be available as specified. These involve key system characteristics, such as availability, serviceability and maintainability.

Financial is the amount of money spent on the service by the client. Some of the attributes are cost, responsiveness, billing process integration, predictability and financial agility.

Performance covers the features and functions of the provided services. Some of the attributes are interoperability, accuracy, response time and suitability.

Security and privacy indicate the effectiveness of a cloud-service provider's controls on access to services, service data and the physical facilities from which services are provided. Some of the attributes are access control and privilege management, sensitivity to data geographic and political aspects, data integrity and loss and data retention and disposition.

Usability is ease with which a service can be used. Attributes include accessibility, learnability, operability and customization.

CSMIC has defined a battery of metrics in these categories. These metrics are either:

Quantitative: based on actual measures of events or properties, with outcomes having the semantics of a quantity (count, ratio, measure of speed, volume, time) to which some numerical targets and thresholds can be applied or

Qualitative (based on checklists and scores): a metric that has either nominal or ordinal values. When ordinal, the metric usually expresses a score (e.g. on a scale from 1 to 10). When nominal, it expresses a quality (e.g. good, average and bad).

These metrics address the whole spectrum of the aspects of a service from administrative and business, to the usage perspective, both for performance and experience, as well as typical service-level properties, such as response time and availability. Many of these apply in an IIoT context.

AN EXAMPLE OF READINESS METRIC

Consider the following SMI measure (Table 2) for the ability to assign a service to a different provider without being hindered by questions of legal ownership of intellectual property or data or by the inability to transition the service (legal portability). This is a readiness metric, its value should be established prior to the deployment of an IIoT system or prior to selecting a provider for any of the parts of such a system. It is also an example of qualitative metric, as are many readiness metrics. This metric was designed for a cloud service, but clearly applies to form of contracted data service.

The definition of the metric is covered in four parts, each involving a set of attributes:

Measure identification: these attributes are the first-line attributes to be shown in a catalog of metrics, to help users evaluate relevance to their system requirements. They also describe the overall rationale of the metric.

Purpose of the measure: these attributes define more precisely the scope of use, what is really being measured and for what benefit.

Measurement definition: these attributes define precisely the way a metric output is calculated. It corresponds to the metric *expression*, along with its *rules*, in the ISO/SC38 model.

Other information: these attributes provide useful complements and execution guidance as well as examples.

Table 3 shows an excerpt of the definition of this metric as provided by CSMIC (as more completely defined in *[4]*):

Measure identification

Measure Name	Legal Portability
SMI Attribute	Portability
Type	Qualitative
Rationale/Context	For the purposes of the SMI, legal portability is the ability of the customer to move the provision of the service to a different provider without being hindered by questions of legal ownership, of intellectual property or data or by the inability to transition the service without adequate assistance from the provider. Note that legal expertise is required for the proper determination of a score for legal portability. The SMI measures two aspects of a portable service: • Legal portability • Technical portability Since legal expertise must be relied upon for proper scoring of the Legal Portability measure, the SMI provides only guidelines on what to look for and a methodology for consistent scoring.
Audience/consumer of measure	Cloud services customer/evaluator or consultant acting on their behalf

Purpose of the Measure

Description of Use	This is a high-level measure of the portability of the service, where portability is the ability to move the service to another provider with minimal effort and without risk of loss or damage to the customer. The legal portability measure is used in conjunction with the technical portability measure to provide a score for the portability attribute of the SMI.
Assumptions	To optimize the value of the SMI, each measure in the Portability attribute and each attribute in the agility category will require an appropriate weighting, set by the customer/evaluator.
Business Goal	Selection of the cloud service that will most closely meet the needs of the organization. Generally speaking, cloud services that are provided with a higher degree of portability (easier to accomplish, with lower risk of data or intellectual property loss) will be preferable to services provided with less portability. Even the best relationships between customers and service providers can deteriorate, customer requirements may change over time and the service provider's services or business position may also change, so while the need to change service providers may not be obvious at the outset, it should always be planned for.
Technical Goal	Provision of a reasonable expectation that the service can be continued with a minimal or at least a clearly understood, level of effort and risk should the services need to be supplied by a different provider.

Measurement Definition

Formula	In order to determine how *good* or *bad* the portability associated with each cloud service is, judgment must be applied to determine how well the terms of the agreement meet the specific needs of the customer. A suggested minimum list of questions to answer when reviewing each provider's proposed agreement is as follows: • What is the duration of the agreement and does it allow the customer to reduce billable volumes to zero without penalty? • Is there a reasonable expectation that the customer may wish to transition services to a different provider within that timeframe? • Does the agreement make it clear that any intellectual property owned by the customer prior to the agreement remains the sole property of the customer? • Does the agreement make it clear, if applicable, that any customer products or services that make use of or rely upon the provider's products and services remain the sole property of the customer and that such derivative work is legally permitted without compensation to the provider beyond the service fees agreed upon? • Does the agreement make it clear what data, documentation or other relevant information associated with the service is the property of the customer and must be returned or destroyed, as appropriate, upon termination of the service? • How comprehensive is the transition assistance described in the agreement? What risk is there that data or intellectual property could be lost due to inability to transfer them off of the provider's infrastructure? • Are there well-defined commitments around how much time any necessary transitions should take? Are transition roles defined clearly enough to provide confidence that these commitments can actually be met? Once the agreements have been reviewed for each provider, scores can be assigned according to the guidelines given below. The questions in the list above may be more or less relevant depending on the exact nature of the service and the applications being run.
Unit of Measure	Point score from 0 to 10. Decimal values are not allowed within this range. Point scores are defined as follows (all values, including 1 through 4 and 6 through 9, can be used to differentiate between alternative services from different providers, in proportion to the differences in the degree of legal portability): **Score = 1:** The agreement does not address issues of intellectual property ownership, data ownership or the provider's role in transitioning services should the services be terminated or it addresses them only in ways which

	represent a significant risk to the customer's ownership rights or ability to transition the service to another provider. **Score = 5:** Legal portability is considered to be adequate for the customer's needs. No material risk of loss of intellectual property rights or data ownership exists. Transition assistance in the case of termination may not be available or is not spelled out in the agreement, but the need for assistance is expected to be minimal. **Score = 10:** All legal portability issues have been clearly addressed in the agreement. No risk of loss of intellectual property rights or data ownership exists. The service provider commits to providing transition assistance in the agreement, including clear definition of the procedure, roles and timeline. The contract duration is monthly and/or requires no minimum billable volumes to avoid financial penalties.
Frequency	At inception of the service and as needed thereafter (rebidding, benchmarking, migration from one CSP to another, etc.).
Exclusions	None.

Other Information

Decision Criteria	Decisions are based on weighted scoring analysis, where each measure of each attribute in the SMI is assigned a weight by the customer of the service and scores are multiplied by weights to calculate weighted scores for each measure, each attribute and finally the overall SMI.
Data Collection	Reviews of proposed agreements are performed by qualified legal counsel, typically working in conjunction with customer personnel or consultants charged with overall responsibility for the SMI analysis. Scoring is based on analysis of each provider's proposed contract documentation.
Additional Comments	Given the varied and evolving nature of cloud services, it is not practical to define strict rules for how portability must be measured for every conceivable scenario
Example	(Oversimplified for illustration). Consider various providers of a cloud-based data analytics service: **Provider A** is an in-house solution. Provider A receives a score of 10, since an in-house solution is generally free of legal concerns that would affect portability. **Provider B** is an external vendor. This provider's agreement does not mention intellectual property rights or data ownership and no transition assistance is described. However, the contract does not require any minimum billable volumes to be maintained. Provider B receives a score of 1.

	Provider C is an external vendor. This provider's agreement does clearly spell out that intellectual property rights for the software and any related trade secrets belong solely to the customer, along with any data stored on the provider's infrastructure. However, no transition assistance is described. Provider C receives a score of 4.
Relationship to other measures	Contract duration is also a factor in the financial elasticity measure. Technical portability is the other measure in the portability attribute and must be measured along with legal portability in order to determine the portability score.

Table 3: A metric for service legal portability according to CSMIC

The readiness metric for legal portability in Table 3 follows the general structure of readiness metrics in the CSMIC SMI. The goal is to follow a standard structure that invites the metric designer to provide all necessary information to define and apply the metric in a consistent way across providers.

In summary, the CSMIC SMI is:

- a framework for organizing and classifying service measures,
- a standard way of describing and documenting service measures and
- a mechanism for making decisions regarding the selection of cloud service providers.

This framework applies beyond the context of cloud services, to other forms of qualitative assessment of service from third-party providers, such as involved in IIoT solutions.

THE METRIC MODEL STANDARD OF ISO-IEC JTC1/SC38

STANDARDIZING A METRIC DEFINITION: STRUCTURE AND RATIONALE

The international standards organization JTC1 has been developing a metric model ISO/IEC 19086-2:2018 *[9]* under ISO/IEC-JTC1/SC38, that can serve as a foundation and standard representation for various service level objectives metrics *[8]*. NIST in the US has been using this model for defining formal metrics to be used by services providers. This model is particularly well-suited for *performance metrics*.

Figure 12 gives an overview of the metric model from SC38 standard.

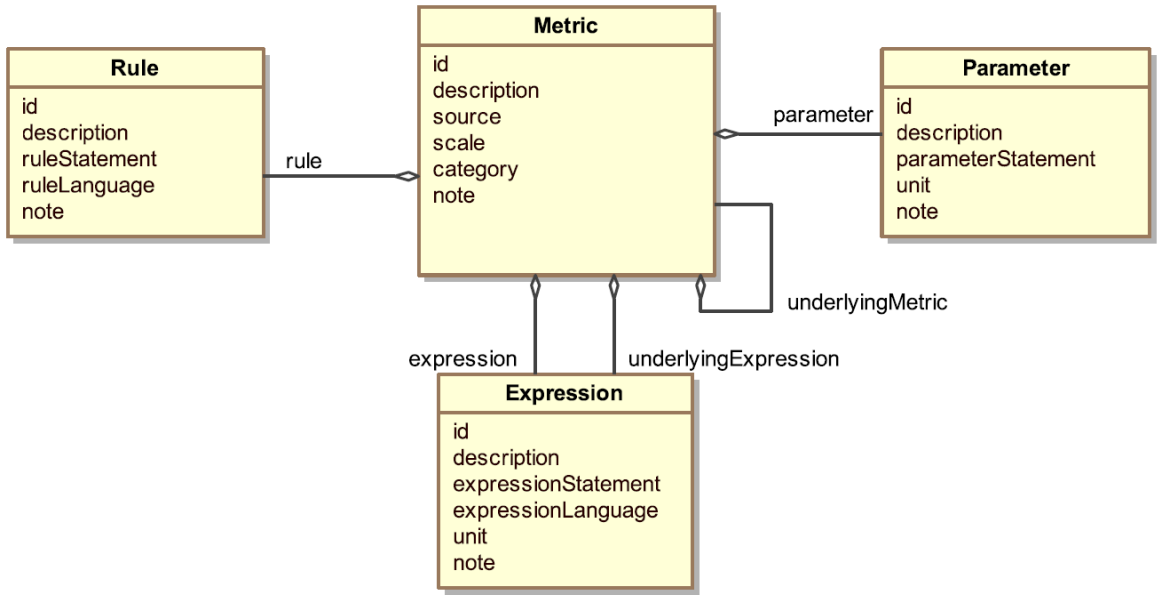

Figure 12: The metric model according to the ISO/IEC 19086-2:2018 standard

Consider a *service availability* metric, a likely metric for parts of an IIoT system such as a service for cloud-hosted data archiving or an event logging service. Operational availability is often measured as uptime percentage over an expected period. A target for such an indicator is often provided in cloud SLAs, for example 99.5%. This target in itself does not belong to the metric definition but to the context of use for the metric (as specified in the SLA). In contrast, a precise definition of *availability* is part of the metric definition. Even for similar services or components, the notion of service availability differs from one provider to the other often in subtle ways, affecting significantly the metric outcome.

Measuring service availability typically depends on:

The calculation logic: a percentage of service uptime over a period of time. This is captured as *expressions* in the metric definition.

A *precise definition of what is a downtime*: includes exception conditions that disqualify downtime periods based on various criteria, such as the nature of the downtime, its context (e.g. did it occur during a scheduled maintenance period?) or its cause (is it the consequence of an uncontrollable natural event?). These definitions and exceptions are captured as *rules* in the metric definition.

A standardization of a metric representation such as provided by the metric model standard is helping to clarify these metric elements and to compare the metrics used by different providers. For example, a service shutdown will not be counted as such by a provider if the shutdown time is less than three minutes, but another provider might exclude shutdowns lasting up to fifteen minutes. This will have a significant effect over the availability outcome. Even if both providers

use the same availability calculation logic (uptime percentage) and same metric measurement methods, it will be meaningless to compare their availability rates due to their different definition of what qualifies as downtime.

The metric model identifies three major components of a metric:

The *metric expressions* define unambiguously the calculations that produce the metric result.

The *metric rules* capture some aspects of the metric computation that are not easily translated into an expression. It is generally about capturing details on *how* or *when* to perform measures. These rules include exclusion conditions, events or occurrences that should not be counted, such as downtimes that are not considered valid.

The *metric parameters* represent values that may vary from one execution to the next. This apparently simple feature of a metric is key to the composability and reusability of metrics. A common parameter is the period of measurement before producing an output such as a month for a service availability metric as would be required by the SLA of a service with monthly billing. Another could be the frequency of measures performed to get availability raw data.

The ISO/SC38 metric model allows for expressing metrics composition. Controlling the properties (such as performance and reliability) of a system or of a sub-system depends on measures done at a lower level, on components or contracted sub-services. For example, the reliability of the analytics performed as a service by a system will depend on the reliability of edge components in providing a data stream with acceptable loss of data, as well as on the reliability of the cloud service that aggregates and stores data from various sources. In such systems, metrics are composed of other metrics. A unified metric model is crucial for understanding and enabling these compositions.

AN EXAMPLE OF SERVICE AVAILABILITY METRIC

This model has been given a tabular representation for practical use. Tables 4, 5, 6 and 7 show a simplified service availability metric definition based on this model, inspired from the SLA of a real-world service provider. Table 1 shows the main element for the “Simple Cloud Service Availability” metric:

Metric (id: SCSA_1, name: Simple Cloud Service Availability)	
attribute	value
description	Evaluates service availability based on the percentage of the time during the measurement period when the cloud service is available.
source	SLA from ShinyCloud service provider
scale	ratio

associated element	reference
parameter	**id**: OPeriod (a *measurement* parameter)
parameter	**id**: MinOutage, value= 5 minute (a *configuration* parameter)
rule	**id:** InvalidMinimum**,**
rule	**id:** DegradedService,
rule	**id:** ScheduledMaintenance**,**
expression	**id**: ServiceUpTimePercentage
underlyingExpression	**id**: ObservedPeriodDuration
underlyingExpression	**id**: TotalDowntimeDuration
underlyingExpression	**id**: DownTimeSequence

Table 4: The main element of the service availability metric

The set of metric rules is captured in table 5:

Set of Rules		
ruleLanguage: plain English text		
id	ruleStatement	note
InvalidMinimum	A downtime incident the duration of which is less than MinOutage is not counted as a valid downtime period.	MinOutage is a parameter
DegradedService	A period when the service is unreachable for at least one minute, or when a response to a request takes longer than one minute, is a downtime.	
ScheduledMaintenance	A downtime or part of a downtime that occurs during a scheduled maintenance period is not counted.	

Table 5: The rules of the service availability metric

The metric parameters are captured in Table 6:

Set of Parameters				
ID	Description	ParameterStatement	Unit	Note
OPeriod	An observation period defined as a time interval.	{start, end} *start* is the beginning of the observation period, *end* is the end of the period.	N/A	

MinOutage	The minimum duration for an outage to be considered as a downtime.	Minimum outage duration	minute	

Table 6: The parameters of the service availability metric

The metric expressions are captured in Table 7:

Set of Expressions			
expressionLanguage: algorithmic, including specific functions			
id	expressionStatement	unit	note
ServiceUpTime Percentage	= ((ObservedPeriodDuration – TotalDowntimeDuration) / ObservedPeriodDuration) * 100	percentage	Main expression.
ObservedPeriod Duration	= duration (OPeriod.start, OPeriod.end)	minute	OPeriod is a parameter.
TotalDowntime Duration	= duration(DownTimeSequence)	minute	Calculates the total duration of all valid downtimes.
DownTime Sequence	= apply_rules (*DegradedService*, *ScheduledMaintenance, InvalidMinimum*,)) to (all observed incidents) over (OPeriod)	N/A	Returns a list of valid downtime intervals, each one characterized with a start and end time.

Table 7: The expressions of the service availability metric

The rules in metric definitions are essential. They define precisely how measurements are made, when and under which conditions. They also define exceptions. In this example (see Table 4) it is clear how significantly these rules affect the outcome of a metric. Different rules may produce different availability outputs and a customer may be better off with a provider committing to 95% availability based on a metric with a tight rule set with few exception rules, than with a provider committing to 99% availability, but with weaker downtime definition rules allowing for many exceptions.

BIBLIOGRAPHY

[1] Kaizen – continuous improvement methodology, *https://en.wikipedia.org/wiki/Kaizen* and *https://www.kaizen.com/what-is-kaizen.html*

[2] Intelligent Transport Solutions for Smart Cities and Regions, IIC Journal of Innovation, *https://www.iiconsortium.org/pdf/June_2017_JoI_Intelligent_Transport_Solutions_for_Smart_Cities_and_Regions.pdf*

[3] The Industrial Internet Security Framework, (IISF), IIC, 2017, *https://www.iiconsortium.org/IISF.htm*

[4] Service Measurement Index Framework (v2.1), CISMIC, Carnegie Mellon University Silicon Valley, 2014, *http://csmic.org/downloads/SMI_Overview_TwoPointOne.pdf*

[5] TL 9000 standard*, https://tl9000.org/*

[6] ISO 9000 standards*, https://www.iso.org/iso-9001-quality-management.html*

[7] QuEST Forum , Quality Excellence for Suppliers of Telecommunications (QuEST) Forum *www.questforum.org*

[8] ISO/IEC 19086-1:2016*, Information Technology – Cloud computing – Service level agreement (SLA) framework – Part 1: Overview and concepts, https://www.iso.org/standard/67545.html*

[9] ISO/IEC 19086-2:2018, *Information Technology – Cloud Computing – Service Level Agreement (SLA) Framework — Part 2: Metric Model, https://www.iso.org/standard/67546.html*

[10] Managing and Assessing Trustworthiness for IIoT in Practice, white paper, IIC, 2019 *https://www.iiconsortium.org/pdf/Managing_and_Assessing_Trustworthiness_for_IIoT_in_Practice_Whitepaper_2019_07_29.pdf*

[11] ISO/IEC TR23951, *Information Technology – Cloud Computing – Guidance for using the Cloud SLA Metric Model, (to be published in 2019)*

[12] IIC Business Strategy and Solution Lifecycle – managing the IIoT Value Chain Transformation, Slama D., Durand J., Morrish J, Journal of Innovation, IIC, 2015, *https://www.iiconsortium.org/news/joi-12-15.htm*

[13] The Industrial Internet Reference Architecture, IIC, 2019, *https://www.iiconsortium.org/IIRA.htm*

[14] ISO/IEC 30141:2018 – Internet of Things (IoT) – Reference architecture, 2018, *https://www.iso.org/standard/65695.html?browse=tc*

[15] A Horizontal Taxonomy for the Industrial Internet, Journal of Innovation, Stan Schneider, IIC, 2015, *https://www.iiconsortium.org/news/joi-12-15.htm*

[16] CERT® Resilience Management Model, Version 1.2, Software Engineering Institute, *https://resources.sei.cmu.edu/asset_files/Handbook/2016_002_001_514462.pdf* , and *https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=30375*

[17] Trustworthiness, IIC Journal of Innovation, IIC, 2018, *https://www.iiconsortium.org/news/journal-of-innovation-2018-sept.htm*

[18] Industrial Internet Consortium: IoT Security Maturity Model: Description and Intended Use, 2020-05-05, retrieved 2020-05-05 *https://www.iiconsortium.org/pdf/SMM_Description_and_Intended_Use_V1.2.pdf*

[19] IoT Security Maturity Model: Practitioner's Guide, Version 1.2, 2020-05-05, retrieved 202-05-05. *https://www.iiconsortium.org/pdf/IoT_SMM_Practitioner_Guide_2020-05-05.pdf*

[20] Office of Electricity Delivery & Energy Reliability: Cybersecurity Capability Maturity Model (C2M2), retrieved 2016-09-26

http://energy.gov/sites/prod/files/2014/03/f13/C2M2-v1-1_cor.pdf
from

http://energy.gov/oe/services/cybersecurity/cybersecurity-capability- maturity-model-c2m2-program/cybersecurity

AUTHORS AND LEGAL NOTICE

This document is a work product of the Industrial Internet Consortium Digital Transformation Working Group, co-chaired by Jim Morrish (Transforma Insights), Dirk Slama (Bosch) and Bassam Zarkout (IGnPower).

AUTHORS
The following persons contributed substantial written content to this document and acted as editors:

Jacques Durand (Fujitsu), Frederick Hirsch (Fujitsu).

CONTRIBUTORS

The following persons contributed content, valuable ideas and feedback that significantly improved the content and quality of this document:

Ashok Dandekar (QuestForum / Fujitsu), Eric Simmon (NIST), Frederic DeVaulx (NIST), Takao Mizutani (Fujitsu), Jim Morrish (Transforma Insights), Bassam Zarkout (IGnPower), Dirk Slama (Bosch), Stan Schneider (RTI), Anish Karmarkar (Oracle).

TECHNICAL EDITOR

Stephen Mellor (IIC staff) oversaw the process of organizing the contributions of the above Editors, Authors and Contributors into an integrated document.

IIC ISSUE REPORTING

All IIC documents are subject to continuous review and improvement. As part of this process, we encourage readers to report any ambiguities, inconsistencies or inaccuracies they may find in this Document or other IIC materials by sending an email to *admin@iiconsortium.org*.