



The Industrial Internet of Things Volume G5: Connectivity Framework

IIC:PUB:G5:V1.0:PB:20170228

Copyright © 2017, Industrial Internet Consortium

ACKNOWLEDGEMENTS

This document is a work product of the Industrial Internet Consortium Connectivity Task Group, co-chaired by Dr. Rajive Joshi (RTI) and Paul Didier (Cisco).

EDITORS

Rajive Joshi (Lead, RTI), Stephen Mellor (IIC), Paul Didier (Cisco)

AUTHORS

The following persons have written substantial portion of material content in this document:

Rajive Joshi (Lead, RTI), Paul Didier (Cisco), Jaime Jimenez (Ericsson), Timothy Carey (Nokia)

CONTRIBUTORS

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

Bob Gessel (Ericsson), Tony Hodgson (Synapse Wireless), Matthew Gilmore (Itron), Edward Eckert (Itron), Aron Semle (Kepware), Jeff Lund (Belden), Cliff Faurer (Enterprise Web), Christoph Gericke (Harting), Stefan Schönegger (BR Automation), Tom Rutt (Fujitsu), Jiyhe Lee (Samsung), Farooq Bari (AT&T), Gerardo Pardo-Castellote (RTI), Stan Schneider (RTI), Reinier Torenbeek (RTI), Brett Murphy (RTI), Norman Finn (Cisco), Kevin White (Distrix), Mark Crawford (SAP), Shi-Wan Lin (Thingswise), Aravind Parandhaman (NEC), Jeff Harding (ABB), Eric Harper (ABB), Brad Miller (GE), Marcellus Buchheit (Wibu-Systems)

We are also grateful to everyone who made comments on the draft version available to the Connectivity Task Group, and will thank in advance anyone who provides further constructive comments on the current version.

We acknowledge the work by the members of the Architecture Task Group led by Shi-Wan Lin for developing the Industrial Internet Reference Architecture and the Vocabulary Team lead by Anish Karmarkar in the Technology Working Group for maintaining the Vocabulary; both those are a companion document to this one.

Finally, we are grateful for the ongoing support of IIC Staff, whose diligence with supporting the tools and the mechanics of the document writing have made this process easier.

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.

CONTENTS

1	Introduction	8
1.1	Purpose.....	9
1.2	Scope.....	9
1.3	Summary.....	10
1.4	Structure.....	11
1.5	Audience.....	12
1.6	Use.....	12
1.7	Relationship with Other IIC Documents.....	12
2	Connectivity Framework	14
2.1	IloT Connectivity Stack Model.....	14
2.2	Architectural Role.....	16
2.3	Key Architectural Qualities.....	17
2.3.1	Performance.....	17
2.3.2	Scalability.....	18
2.3.3	Reliability.....	18
2.3.4	Resilience.....	18
2.3.5	Security.....	18
2.3.6	Longevity.....	19
2.3.7	Integration and interoperability.....	20
2.3.8	Operation.....	20
2.3.9	Safety.....	20
3	Connectivity Reference Architecture	21
3.1	IloT Connectivity Challenge.....	21
3.2	Connectivity Core Standards.....	22
3.3	Core Gateways.....	23
3.4	Core Standards Criteria.....	25
4	Connectivity Framework Layer	27
4.1	Core Functions.....	27
4.1.1	Data Resource Model.....	28
4.1.2	ID and Addressing.....	28
4.1.3	Data Type System.....	28
4.1.4	Data Resource Lifecycle (CRUD).....	29
4.1.5	State Management.....	29
4.1.6	Publish-Subscribe.....	29
4.1.7	Request-Reply.....	30
4.1.8	Discovery.....	30
4.1.9	Exception Handling.....	30
4.1.10	Data Quality of Service (QoS).....	30
4.1.11	Data Security.....	31
4.1.12	API.....	32
4.1.13	Governance.....	32
4.2	Typical Considerations.....	33
4.2.1	System Architecture Considerations.....	33

4.2.1.1	Peer-to-Peer vs. Broker	33
4.2.1.2	Data-Centric vs. Device/App-Centric	33
4.2.1.3	Explicit vs. Implicit Governance	34
4.2.2	Data Considerations	34
4.2.2.1	Content-Based Selection	34
4.2.2.2	Time-Based Selection	34
4.2.3	Performance Considerations	35
4.2.3.1	Real-Time	35
4.2.3.2	Latency and Jitter vs. Throughput.....	35
4.2.4	Scalability Considerations.....	35
4.2.4.1	Data Objects.....	35
4.2.4.2	Apps.....	36
4.2.5	Availability Considerations	36
4.2.5.1	Redundancy.....	36
4.2.5.2	Recovery.....	36
4.2.6	Deployment Considerations	37
4.2.6.1	Platform Constraints	37
4.2.6.2	Incremental Upgrades.....	37
5	Connectivity Transport Layer.....	38
5.1	Core Functions.....	38
5.1.1	Messaging Protocol	38
5.1.2	Communication Modes	39
5.1.3	Endpoint Addressing.....	39
5.1.4	Connectedness	39
5.1.5	Prioritization	40
5.1.6	Timing & Synchronization.....	40
5.1.7	Message Security.....	40
5.2	Typical Considerations.....	41
5.2.1	Network Layer Considerations	41
5.2.1.1	Topology.....	41
5.2.1.2	Span.....	41
5.2.1.3	Segmentation	41
6	How to Assess a Connectivity Technology?	42
6.1	General Info	42
6.2	Business Viewpoint	43
6.2.1	Purpose.....	43
6.2.2	Pedigree.....	43
6.2.3	Variants.....	43
6.2.4	Maturity.....	43
6.2.5	Stability	43
6.2.6	Standards Body.....	43
6.2.7	Openness	43
6.3	Usage Viewpoint	43
6.3.1	Architecture	43
6.3.2	Technology Options.....	43
6.3.3	Applications	43

6.3.4	Typical Usage	43
6.3.5	Operations	43
6.3.6	Security	43
6.3.7	Safety	43
6.3.8	Gateways	43
6.4	Functional Viewpoint.....	44
6.4.1	Core Framework Layer Functions.....	44
6.4.2	Core Transport Layer Functions	44
6.5	Implementation Viewpoint.....	45
6.5.1	System Architecture Considerations	45
6.5.2	Data Considerations	45
6.5.3	Performance Considerations	45
6.5.4	Scalability Considerations.....	45
6.5.5	Availability Considerations	45
6.5.6	Deployment Considerations	45
6.5.7	Network Layer Considerations	45
7	Connectivity Standards.....	46
7.1	IIoT Connectivity Framework Standards	47
7.1.1	Data Distribution Service (DDS).....	47
7.1.2	Web Services Using Hypertext Transfer Protocol (HTTP).....	49
7.1.3	Open Platform Communications Unified Architecture (OPC-UA).....	49
7.1.4	oneM2M	51
7.2	IIoT Connectivity Transport Standards.....	52
7.2.1	TCP and UDP over IP	52
7.2.2	Constrained Application Protocol (CoAP).....	52
7.2.3	MQTT (formerly MQ Telemetry Transport).....	53
7.3	Fieldbus Technologies.....	53
8	Core Connectivity Standards.....	55
9	Other Connectivity Technologies	58
Annex A	Assessment Template: DDS	59
Annex B	Assessment Template: OPC-UA	72
Annex C	Assessment Template: oneM2M	82
Annex D	Assessment Template: HTTP.....	94
Annex E	Assessment Template: CoAP	101
Annex F	Assessment Template: MQTT	112
Annex G	Revision History	120
Annex H	Acronyms.....	121
Annex I	Glossary.....	122
Annex J	References	123
	Use of Information—Terms, Conditions and Notices	127

FIGURES

Figure 1-1:	Connectivity is a crosscutting function in the Industrial Internet Reference Architecture. It provides the ability to exchange data between participants within and across functional domains (control, operations, information, applications, business).....	8
Figure 1-2:	Scope of the Connectivity as a crosscutting function within the IIoT Reference Architecture. Connectivity provides the data sharing mechanisms for the higher-level functions, including “Distributed Data Interoperability and Management”. The “neck” of the hourglass represents the “Internet” network layer, common across industries. This document focuses on the connectivity layers above the neck.....	10
Figure 1-3:	IIC Technical Publication Organization.	13
Figure 2-1:	Industrial Internet Connectivity Stack Model. Each layer builds on the capabilities provided by the layer below. The ‘Connectivity Framework’ layer provides data sharing mechanisms among participants. The ‘Distributed Data Interoperability and Management’ layer relies on the mechanisms provided by the ‘Connectivity Framework’ layer to provide meaningful information sharing.	14
Figure 2-2:	The focus of this document is on connectivity layers above the network layer, namely the connectivity transport and the connectivity framework layers.	15
Figure 2-3:	Connectivity protection building blocks described in the Industrial Internet Security Framework.	19
Figure 3-1:	The fundamental N^2 (N-squared) IIoT connectivity challenge. Each new connectivity technology requires building a bridge to all the existing connectivity technologies, in order to facilitate information exchange between endpoints in different connectivity technologies. This approach does not scale beyond a few (small N) technologies, and results in information silos.	21
Figure 3-2:	Connectivity Gateway Concept. A connectivity core standard technology (baseline) is one that can satisfy all of the connectivity requirements for a functional domain. Gateways provide two functions (1) integrate other connectivity technologies used within a functional domain, (2) interface with connectivity core standards in other functional domains.....	22
Figure 3-3:	A standardized gateway between core connectivity standards can allow domain-specific endpoints connected to one core standard to communicate with domain-specific endpoints integrated over another core standard.	24
Figure 3-4:	Each core connectivity standard requires a standardized gateway to all other core standards. Each additional core standard creates increasing complexity and interoperability challenges. By restricting the design to a few core connectivity standards, we cover the needs of IIoT systems across the functional domains, and attain the goal of horizontal interoperability across industries.	24
Figure 4-1:	Connectivity framework layer functions.	27
Figure 5-1:	Connectivity transport layer functions.	38

Figure 7-1: IIoT connectivity standards. Dotted boxes show the connectivity standards (e.g. oneM2M, OPC-UA) that have originated in respective industry verticals (e.g. telecommunications, manufacturing) to provide enabling features for those industries, and also offer an application focus that is applicable to multiple industries. Others (e.g. DDS, Web Services) have originated in an industry agnostic manner for general-purpose use, and are applied in multiple industries for many different kinds of application areas. Transports that are specific to a framework layer are shown without any spacing between the framework and the transport layer boxes.....	46
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----

TABLES

Table 2-1: Role and scope of the Connectivity functional layers.	17
Table 8-1: IIoT Connectivity Core Standards Criteria applied to key connectivity framework standards.	55
Table 8-2: Non-overlapping system aspect examples addressed by the potential IIoT connectivity core standards.	57

1 INTRODUCTION

Ubiquitous connectivity is one of the foundational technologies enabling data sharing amongst participating components of an Industrial Internet of Things (IIoT) system.

Connectivity provides the ability to exchange data amongst participants within a functional domain, across functional domains within a system and across systems. The data exchanged may include sensor updates, events, alarms, status changes, commands, and configuration updates. Connectivity is a crosscutting function across the functional domains defined by the Industrial Internet Reference Architecture¹, as shown in Figure 1-1.

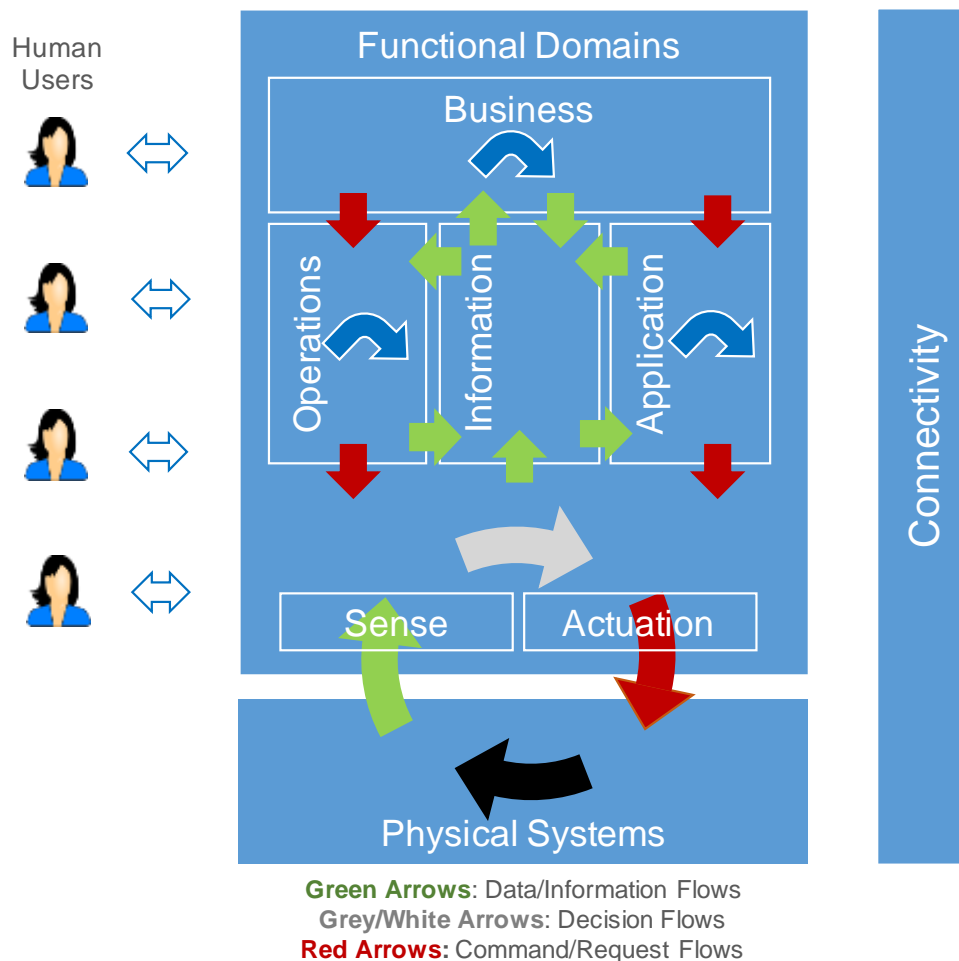


Figure 1-1: Connectivity is a crosscutting function in the Industrial Internet Reference Architecture. It provides the ability to exchange data between participants within and across functional domains (control, operations, information, applications, business).

¹ See [IIC-IIIRA2015]

1.1 PURPOSE

The IIoT landscape is replete with proprietary connectivity technologies and specialized connectivity standards optimized for a narrow set of domain-specific use cases in vertically integrated systems. These domain-specific connectivity technologies, though optimal in their respective domains, can be a hindrance to the sharing of data, designs, architectures, and communications essential to creating new value streams and unlocking the potential of a global IIoT marketplace. The overarching goal of IIoT connectivity is to unlock data in these isolated systems (“silos”) and enable data sharing and interoperability between previously closed components and subsystems (brownfield) and new applications (greenfield), within and across industries.

This document maps the rich landscape of IIoT connectivity. It clarifies the IIoT connectivity stack, defines an open connectivity reference architecture, and helps practitioners navigate their way to categorize, evaluate, and determine the suitability of a connectivity technology for the system at hand. Specifically, it addresses the following questions:

- What connectivity layers to expect for IIoT?
- What core functions to expect from each layer?
- What are the typical considerations and trade-offs at each layer?
- How to open up communication to participants using a domain-specific connectivity technology?
- What is expected from core connectivity standards?
- How to categorize a given connectivity technology?
- How to evaluate a given connectivity technology?
- How to determine suitability of a connectivity technology against system requirements?
- How to determine the most appropriate core connectivity standard?
- What are the most suitable core connectivity standards for a particular (sub)system?

1.2 SCOPE

The interoperability layers of an IIoT system are hourglass shaped. The top is a wide spectrum of data models and functions specific to a particular industry and the “neck” is the “internet” network layer common across industries, as shown in Figure 1-2. Connectivity provides the basic data-sharing mechanisms to support the higher-level functions such as *Distributed Data Interoperability and Management* as a crosscutting function (semantic interoperability, see [Industrial Internet Reference Architecture¹](#)).

The neck of the hourglass in Figure 1-2 is the starting point—it is the “internet” in “IIoT”. But the connectivity layers above the neck are not well understood, so we focus on the connectivity functions and considerations above the “internet” networking layer for building IIoT systems.

¹ See [IIC-IIIRA2015]

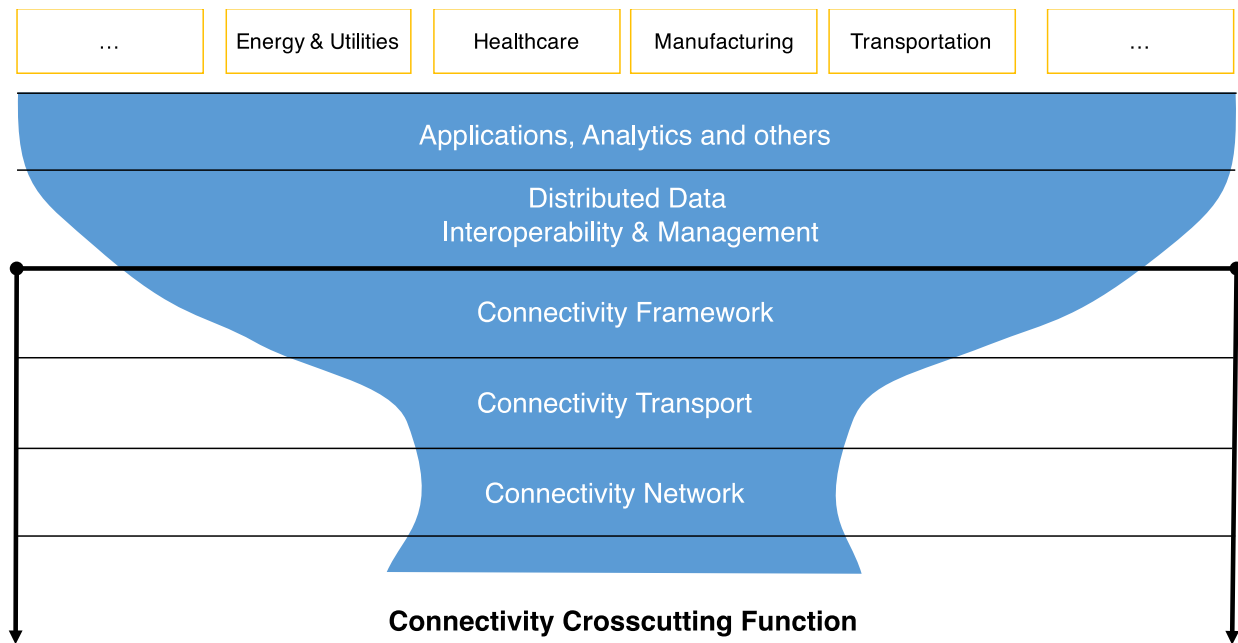


Figure 1-2: Scope of the Connectivity as a crosscutting function within the IIoT Reference Architecture. Connectivity provides the data sharing mechanisms for the higher-level functions, including “Distributed Data Interoperability and Management”. The “neck” of the hourglass represents the “Internet” network layer, common across industries. This document focuses on the connectivity layers above the neck.

1.3 SUMMARY

The connectivity challenges in IIoT include meeting diverse requirements, working over many transports, and connecting a dizzying array of “things” from small devices to huge, intelligent networks of complex subsystems. And challenges are both business and technical: we shall consider the business, function, usage and implementation viewpoints.

The connectivity reference architecture strives for broad applicability across the IIoT and the power to handle challenging, unique applications. It introduces the notions of a *connectivity gateway*¹ and *core connectivity standards*. There are two types of gateways: Core Gateways that connect core standards, and noncore gateways that connect a domain-specific connectivity technology to a core connectivity standard. Rather than building many bridges between many standards, each core connectivity standard need only connect to the other core connectivity standards through *core gateways*. The many domain-specific connectivity technologies need then interface to only one of the core connectivity standards. This strikes a balance between allowing any connectivity technology so requiring many complex bridges and allowing only one core standard that cannot span the IIoT.

¹ Gateway is a base term defined in the Industrial Internet Vocabulary [IIC-IIV2015]: a forwarding component, enabling various networks to be connected. It may be a software component.

The connectivity reference architecture proposes an IIoT connectivity stack model using the OSI and the Internet models as reference. It defines the core functions and key considerations at each layer in the IIoT connectivity stack. It defines an assessment template worksheet to understand and assess any connectivity technology objectively, and then determine the core connectivity standard closest to the technology under assessment.

Assessment templates for dealing with major IIoT system design challenges introduce core connectivity standards. These include the Data Distribution Service (DDS) for systems facing a software integration challenge, OPC-Unified Architecture (OPC-UA) for systems facing device interchangeability issues, HTTP/REST¹ for web and mobile user interfaces, and oneM2M for information and communications technology integration with wide area wireless telecommunication provider network services. The architecture integrates other connectivity technologies by interfacing to a core connectivity standard. This satisfies the range of application challenges with minimum complexity. We also provide assessment templates for common domain-specific connectivity technologies typically used at the network edge.

1.4 STRUCTURE

Chapter 2 defines the IIoT connectivity stack model and introduces the *Connectivity Framework* (framework) and the *Connectivity Transport* (transport) layers. It clarifies the role of connectivity in enabling syntactic interoperability, i.e. exchanging structured data, in system architecture, and introduces the key system characteristics directly affected by connectivity.

Chapter 3 defines the requirements for core connectivity standards and proposes connectivity gateways to bridge a domain-specific connectivity technology to a core connectivity standard and open up hitherto inaccessible endpoints. This approach is tenable with a few core connectivity standards with core gateways for interoperability amongst them, and many domain-specific technologies that can use a gateway to any of those core standards.

Chapter 4 dives into the *connectivity framework* layer. It defines the core functions and the typical considerations and trade-offs to apply when considering a connectivity framework technology.

Chapter 5 dives into the *connectivity transport* layer. It defines the core functions and the typical considerations and trade-offs to apply when considering a connectivity transport technology.

Chapter 6 defines a template for assessing any connectivity technology from a business, usage, functional, and implementation viewpoint. It introduces a worksheet that can be used as a tool to understand, categorize and evaluate any connectivity technology.

Chapter 7 uses the assessment template worksheets to describe the prominent connectivity standards for IIoT. It also describes some of the connectivity standards prominent in specific verticals.

¹ REST stands for Representational State Transfer and is an architectural style for networked applications [WKPD-REST]. It is almost always implemented with the HTTP or CoAP protocols.

Chapter 8 highlights the standards that meet the requirements of core connectivity, and are suitable for serving as core connectivity standards.

Chapter 9 provides guidelines on how to open up domain-specific connectivity technologies, via a core connectivity standard. It recommends completing the worksheets to identify the core connectivity standard closest to the domain-specific connectivity technology. It also makes some suggestions for a core connectivity standard based on the primary functional domain (see Figure 1-1) of applicability for the connectivity technology.

1.5 AUDIENCE

The intended audience of this document is system architects, solution architects, technology evaluators, technology decision makers, business strategists and business investment decision makers.

1.6 USE

The document is intended to be used a guide map, that can be read in its entirety in sequential order. Chapter 7 can be skimmed on the first pass, and serves as a ready reference for a deeper dive into a specific technology.

The worksheet in chapter 6 is meant as a tool for practitioners. It can be used to map out connectivity technologies of interest, in the rich landscape of connectivity.

We do not take a prescriptive approach. Instead, we provide, in the hands of system architects, a tool (see chapter 6, assessment template worksheets) that, given a system's requirements, will help determine the most suitable connectivity technologies and core connectivity standards.

1.7 RELATIONSHIP WITH OTHER IIC DOCUMENTS

The 'Industrial Internet of Things, Volume G5: Connectivity Framework' (IICF) is one of many framework documents of the IIC technical publications (see Figure 1-3) that extend from the Industrial Internet Reference Architecture (IIRA)¹.

¹ See [IIC-IIRA2015]

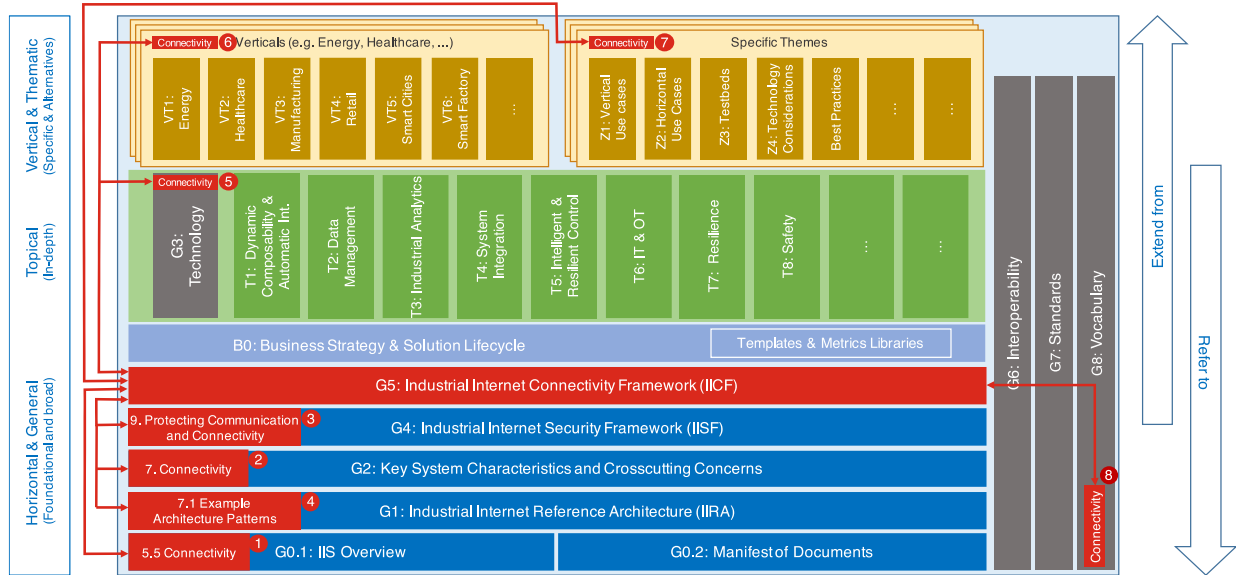


Figure 1-3: IIC Technical Publication Organization.

As shown by ❶ the IICF is part of a series of documents covering Connectivity concerns. The IICF addresses connectivity as a key cross-cutting concern of IIoT as described by ❷. Security considerations for connectivity are described in ❸. Architectural patterns using connectivity and security are described in ❹. This figure also shows how other documents extend from the IICF in covering connectivity related issues in the respective areas.

Connectivity usage will be covered in each of the technology specific documents (“T” series) as illustrated by ❺. Specific connectivity usage in IIoT vertical target segments will be covered in a collection of documents (“V” series), capturing connectivity-relevant topics as part of the use cases, testbeds, solutions and best practices for each of the addressed vertical markets as illustrated by ❻. Connectivity’s implications for system characteristics across vertical markets will be covered as part of system-thematic specific documents as shown in ❼. Finally, connectivity related terms used in this document and their respective definitions are provided in a common Industrial Internet Vocabulary¹ document shown in ❸.

¹ See [IIC-IIV2015]

2 CONNECTIVITY FRAMEWORK

2.1 IIoT CONNECTIVITY STACK MODEL

The seven-layer *Open Systems Interconnect (OSI) Model* (Wikipedia)¹ and the four-layer *Internet Model* (Wikipedia)² do not capture all industrial internet connectivity requirements. IIoT systems require a new connectivity functional layer model to address distributed industrial sensors, controllers, devices, gateways and other aspects. This section proposes an IIoT connectivity stack model using the OSI model and the Internet model as reference.

Figure 2-1 shows the IIoT connectivity stack model, and the scope of the connectivity as a crosscutting function within the *Industrial Internet Reference Architecture (IIRA)*³. The connectivity function provides the data-sharing mechanisms amongst participants within a functional domain and across functional domains within an IIoT system.

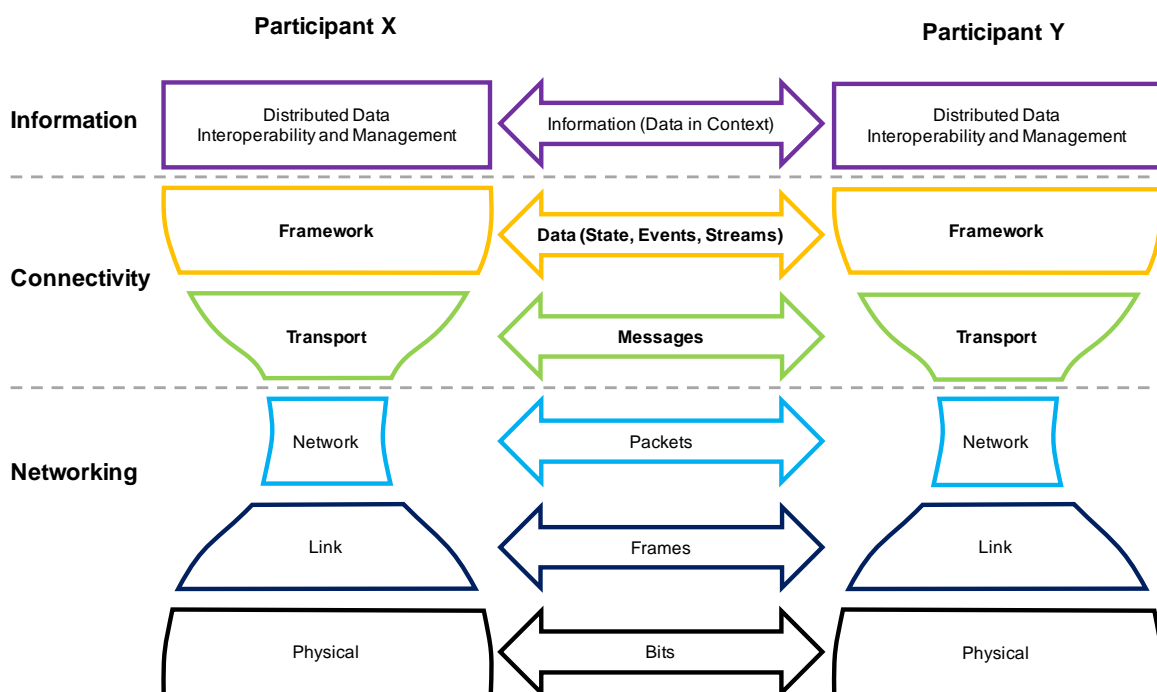


Figure 2-1: Industrial Internet Connectivity Stack Model. Each layer builds on the capabilities provided by the layer below. The 'Connectivity Framework' layer provides data sharing mechanisms among participants. The 'Distributed Data Interoperability and Management' layer relies on the mechanisms provided by the 'Connectivity Framework' layer to provide meaningful information sharing.

¹ See [ISO-7498-1], for overview [WKPD-OSI]

² See [IETF-RFC1122], for overview [WKPD-IPS]

³ See [IIC-IIRA2015]

The lowest layer is the *physical layer*, which refers to the exchange of physical signals (electric, optical, or other) on the physical media (wired or wireless) connecting the participants. Above it is the *link layer*, which refers to the exchange of *frames* using signaling protocols on the shared physical link between adjacent participants. Above it is the *network layer*, which refers to the exchange of *packets* (bounded length), possibly routing them over multiple links to communicate between non-adjacent (remote) participants. Above it is the *transport layer*, which refers to the exchange of *messages* (variable length) between participant applications. Above it is the *framework layer*, which refers to the exchange of structured data (state, events, streams) with configurable quality-of-service between participant applications. Above it, but outside the scope of connectivity, is the *distributed data interoperability and management* layer crosscutting function that relies on the data sharing mechanism provided by the connectivity framework layer.

The Internet Protocol (IP) (Wikipedia)¹ is the prevailing network layer connectivity standard that has given birth to the internet and now IIoT. The IP network layer has enabled independent innovation, both below and above the network layer. The bottom three layers, common to the OSI model and the Internet model, have been in use longer and are generally understood; although evolution of IP and non-IP connections and the multitude of wireless-access technologies coming to market create new choices for the IIoT community.

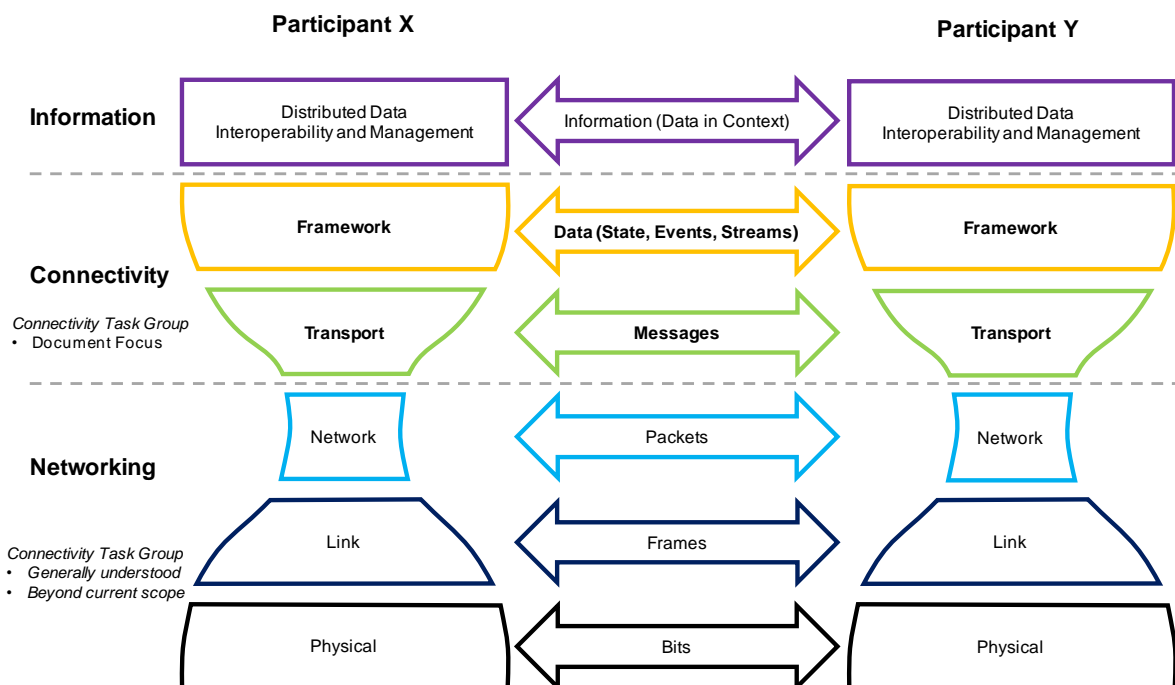


Figure 2-2: The focus of this document is on connectivity layers above the network layer, namely the connectivity transport and the connectivity framework layers.

¹ See [IETF-RFC1122], for overview [WKPD-IPS]

The layers above the network layer have evolved rapidly in the last decade and are not as widely recognized or understood. Therefore, the focus of this document is on the layers above the network layer, namely the transport and framework layers, as shown in Figure 2-2.

2.2 ARCHITECTURAL ROLE

The connectivity function in the IIRA supports exchange of data among endpoints in a system of interest. The information, for example, can be sensor updates, telemetry data, control commands, alarms, events, logs, status changes or configuration updates. Fundamentally, connectivity's role is to provide interoperable communications among endpoints to facilitate component integration.

Interoperability in communication can be achieved at various levels of abstraction, from custom integration to plug-and-play interfaces based on open standards. One common classification of interoperability is as follows (see [Tolk, Wikipedia](#)¹):

Technical interoperability is the ability to exchange information as bits and bytes (e.g. pencil scribbles), assuming that the information exchange infrastructure (e.g. pencil and paper) is established and the underlying networks and protocols are unambiguously defined.

Syntactic interoperability is the ability to exchange information in a common data structure (e.g. using words from a language), assuming that a common protocol to structure the data is used (e.g. the language's alphabet and rules of grammar) and the structure of the information exchange is unambiguously defined (e.g. whitespace, punctuation). Syntactic interoperability requires that technical interoperability be established.

Semantic interoperability is the ability to interpret the meaning of the exchanged data unambiguously as information in the appropriate context. However, the goal of the connectivity function is limited to provide syntactic interoperability between participating endpoints.

For IIoT systems, connectivity comprises two functional layers:

The *connectivity transport* layer provides the means of carrying data between endpoints. It provides technical interoperability between endpoints participating in a data exchange. This function maps to layer 4 (transport) of the OSI model or the transport layer of the Internet model (see Table 2-1).

The *connectivity framework* layer facilitates how data is unambiguously structured and parsed by the endpoints. It provides the mechanisms to realize syntactic interoperability between endpoints. In this context, "common data structure" refers to the structure or schema of the data being exchanged. Familiar examples include data structures in programming languages and schemas for databases. The connectivity framework function spans layers 5 (session) through 7 (application) of the OSI model or the application layer of the Internet Model (see Table 2-1).

¹ See [Tolk-2007], for overview [WKPD-CI]

The *data services framework* in the distributed data interoperability and management function builds on the syntactic interoperability foundation provided by the connectivity framework to provide the foundation for *semantic interoperability* required by the dynamic composition and coordination function of the Industrial Internet Reference Architecture (IIRA)¹.

The role and scope of the IIoT connectivity functional layers are summarized in Table 2-1.

IIoT Connectivity Stack Model	Correspondence to OSI Model (ISO/IEC 7498)	Correspondence to Internet Model (RFC 1122)	Correspondence to Levels of Conceptual Interoperability
Framework Layer	7. Application	Application Layer	Syntactic Interoperability: <i>Structured data types</i> shared between endpoints. Introduces a common structure to share data; i.e., a <i>common data structure</i> is shared. On this level, a common protocol is used to exchange data; the structure of the data exchanged is unambiguously defined.
	6. Presentation		
	5. Session		
Transport Layer	4. Transport	Transport Layer	Technical Interoperability: <i>Bits and Bytes</i> shared between endpoints, using an unambiguously defined communication protocol.
Network	3. Network	Internet Layer	Packets shared between endpoints that may not be on the same physical link. Packets are routed between physical links by a “network router”.
Link	2. Data Link	Link Layer	Digital Frames shared between endpoints on a shared substratum (link).
Physical	1. Physical		Analog signal modulation between endpoints on a shared substratum.

Table 2-1: Role and scope of the Connectivity functional layers.

2.3 KEY ARCHITECTURAL QUALITIES

The connectivity function supports the key architectural qualities of an IIoT system, and they can be used to assess the alternative connectivity choices for concrete architectures. The qualities are described below.

2.3.1 PERFORMANCE

In IIoT systems, high performance connectivity is expected. The spectrum of performance ranges from tight sub-millisecond control loops to supervisory control to analysis at very low frequencies such as daily, weekly or even monthly. The performance characteristic is measured along the following axes.

Latency and jitter. Latency is the time it takes for data to go from source to destination (“time of flight”). Jitter is the variation in latency. The data usually has a limited useful lifetime, so low latency is essential. Low jitter is also needed to ensure the application has integrity and system

¹ See [IIC-IIRA2015]

maintains predictable performance. The connectivity function addresses latency and jitter in the data exchanged between endpoints, possibly in exchange for throughput.

Throughput. Throughput is the load on the network as defined by the volume of data flow per unit of time. Bandwidth is the network capacity of a connectivity technology. In some designs, a large volume of data may be exchanged in a short time on an ongoing basis among endpoints; high throughput would be needed.

In practice, the operational settings that optimize for high throughput are not the same as those that optimize for low latency. Therefore, the connectivity function should support achieving the right balance as per the requirements of the data flow.

In industrial internet applications, particularly at the edge, low latency and jitter are generally more important to performance than throughput and bandwidth. Automation and control of real-world processes require short reaction times or tight coordination to maintain effective control. Industrial devices in the control domain do not produce large amounts of data in short periods and therefore do not require high bandwidth connectivity. Rather, the data needs to be communicated quickly and consistently (with low latency and jitter).

2.3.2 SCALABILITY

Physical things communicate using connectivity endpoints. Therefore, the connectivity function should support *horizontal scaling*, by which we mean the ability to accommodate an increasing number of connectivity endpoints, reaching Internet scale.

2.3.3 RELIABILITY

The needs of the application data, like strict order of data delivery and data loss rates, determine the required level of reliability for connectivity.

2.3.4 RESILIENCE

Because many IIoT systems will operate continually in a real-world environment, the connectivity function should be available (in the logical view), even when there is a temporary physical disconnection. When a broken connection is restored, data exchange should be automatically restored so that the latest updates are available to the consumers along with any relevant missed updates.

The connectivity function should support graceful failure or disconnection of endpoints, by, for example, confining the loss of data exchange only to disconnected endpoints.

2.3.5 SECURITY

Security considerations for IIoT systems are described in detail in the [Industrial Internet Security Framework \(IISF\)](#)¹. The chapter “Communications and Connectivity Protection” describes the following functional building blocks: physical security of connections, communicating endpoints

¹ See [IIC-IISF2016]

protection, information flow protection, network configuration and management, network monitoring & analysis, and cryptographic protection, as shown in Figure 2-3.

Communications & Connectivity Protection

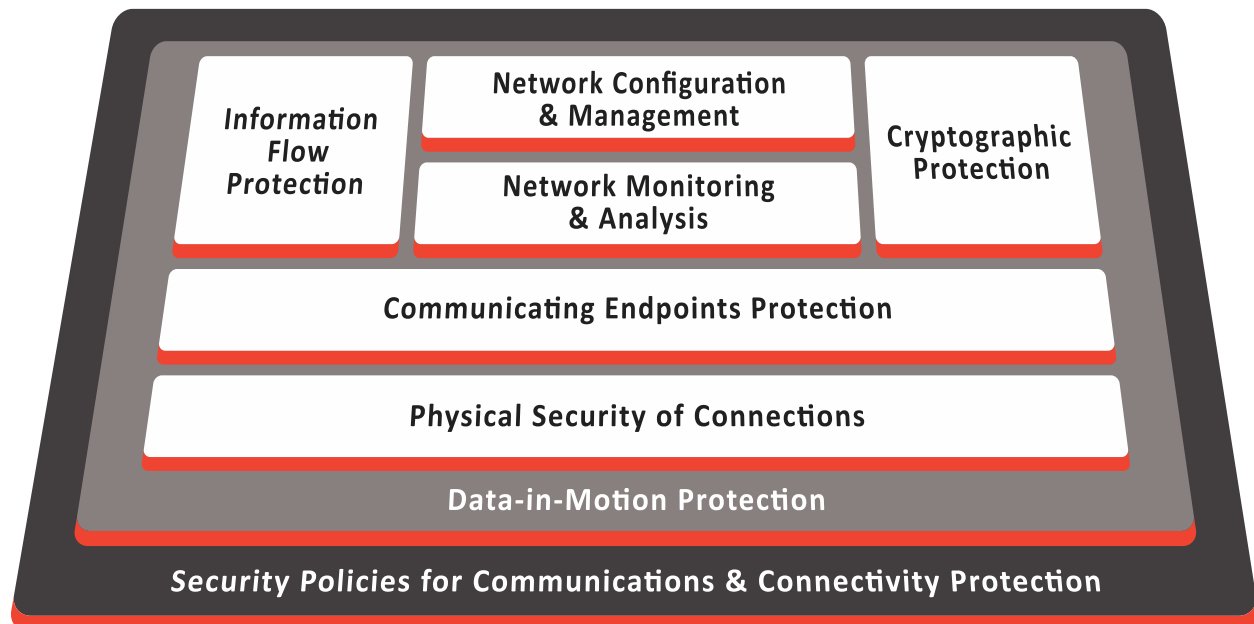


Figure 2-3: Connectivity protection building blocks described in the Industrial Internet Security Framework¹.

The security policies govern connectivity-endpoint data-exchange as part of a broader protection strategy. For example, they specify how to filter and route traffic, how to protect exchanged data and metadata (authenticate or encrypt-then-authenticate) and what access control rules should be used.

Cryptographic protection of connectivity endpoints relies on:

- explicit endpoint data exchange policies,
- strong mutual authentication between endpoints,
- authorization mechanisms that enforce access control rules derived from the policy, and
- mechanisms for ensuring confidentiality, integrity, and freshness of the exchanged data.

Adequate cryptographic protection should be considered for each of the layers shown in Table 2-1.

2.3.6 LONGEVITY

Connectivity components, especially those in the network layer and below, are built into the hardware and hence are not easily replaceable. Where possible and feasible, the connectivity

¹ See [IIC-IISF2016]

software components should support incremental evolution including upgrades, addition and removal of components. The connectivity function should also be able to support incremental evolution of the data exchange solutions during the lifecycle of a system.

2.3.7 INTEGRATION AND INTEROPERABILITY

IIoT systems comprise components that are often systems in their own right. The connectivity function should support the integration and the interoperability of system components, isolation and encapsulation of data exchanges internal to a system component, and hierarchical organization of data exchanges. In dynamic systems, the connectivity function should also support discovery of system components and the discovery of relevant data exchanges for system composition.

2.3.8 OPERATION

IIoT systems generally operate non-stop in a real-world environment. To support a system's operational needs, it should be possible to monitor, manage and dynamically replace connectivity elements. Monitoring includes health, performance and service-level characteristics of the connectivity function; management includes configuring and administering the capabilities; dynamic replacement requires replacement of hardware and or software while a system is operating.

2.3.9 SAFETY

A high degree of assurance is required in life- and mission-critical systems¹ to avoid unintended consequences during system operation. The connectivity function should be able to support safety evaluations and provide evidence required to make informed safety assessments.

¹ For example, autonomous vehicles or medical systems

3 CONNECTIVITY REFERENCE ARCHITECTURE

3.1 IIoT CONNECTIVITY CHALLENGE

The goal of the industrial internet is to enable seamless information sharing across domains and industries. Historically, there have been a plethora of domain specific connectivity technologies, tightly integrated and optimized to solve domain specific connectivity needs. IIoT systems typically include integration of brownfield technologies to preserve the legacy investments, and greenfield technologies to spur innovation.

Figure 3-1 shows the challenge of building applications that require information exchange across different connectivity technologies. To facilitate information exchange, one has to build bridges to each of the other connectivity technologies. Given N connectivity technologies, this requires building $N*(N-1)/2 = O(N^2)$ bridges. That quickly becomes impractical for large N (> 3 or 4). The result is information silos, making it impossible to realize the vision of the Industrial Internet to create new value stream from heretofore locked up information flows.




Figure 3-1: The fundamental N^2 (N-squared) IIoT connectivity challenge. Each new connectivity technology requires building a bridge to all the existing connectivity technologies, in order to facilitate information exchange between endpoints in different connectivity technologies. This approach does not scale beyond a few (small N) technologies, and results in information silos.

In this document, we use the term “*domain-specific*” connectivity technology to refer to a connectivity technology that is especially suited to a particular application area. Domain-specific connectivity technologies include emerging technologies, optimized for certain use cases.

We accept that an IIoT system may require multiple connectivity technologies. Mandating a single connectivity standard across all domains and across all industries is neither realistic nor

feasible. We need connectivity architectures that can address the diversity of IIoT systems, while tackling the N^2 challenge and enabling the vision of the industrial internet.

The rest of this section describes a connectivity reference architecture that achieves near linear scalability, $O(N)$, with respect to the number of connectivity technologies. It accomplishes this by defining a small set of connectivity core standards. *Standardized core gateways* bridge the connectivity core standards. Domain-specific connectivity technologies need a gateway to just one of the connectivity core standards, to participate in an information exchange with the rest of the IIoT ecosystem.

3.2 CONNECTIVITY CORE STANDARDS

New connectivity technologies will need to be integrated with legacy technologies during a system's lifetime. A connectivity architecture shall allow a plethora of connectivity technologies to interoperate within an industry, and across industries to support the vision of an IIoT that spans industries.

A *connectivity gateway* bridges one or more connectivity technologies, as shown in Figure 3-2.




Figure 3-2: Connectivity Gateway Concept. A connectivity core standard technology (baseline) is one that can satisfy all of the connectivity requirements for a functional domain. Gateways provide two functions (1) integrate other connectivity technologies used within a functional domain, (2) interface with connectivity core standards in other functional domains.

To keep the connectivity architecture manageable, a connectivity technology standard is chosen as the baseline within a functional domain, and referred to as the “connectivity core standard” (see Figure 3-2). Gateways are used to bridge other connectivity technologies within the domain and to the connectivity core standards used in other functional domains. Connectivity between functional domains, often implemented in a tiered manner, can be intermittent. Connectivity gateways can help mitigate this intermittent connectivity. Applications are simpler and easier to maintain if logic is not needed to react to failed data exchanges.

As shown in Figure 3-2, some endpoints can connect directly to a core standard. Other endpoints and subsystems connect through gateways. A core standard then connects them all together, allowing multiple connectivity technologies to be integrated without having to bridge between all possible pairs, so avoiding the dreaded N-squared bridging problem (see Figure 3-1). Each domain-specific connectivity technology needs only a gateway to just one connectivity core standard.

Connectivity gateways enable incorporation of new connectivity technologies. They provide a stable foundation anchored in the “best-of-breed” technologies available today, yet can pivot in the future to a new baseline core standard that better satisfies the requirements.

There are several kinds of connectivity gateways:

- *Framework gateways* expand the logical span of communications across connectivity framework technologies. They preserve the syntactic structure of data, but may change the technical representation.
- *Transport gateways* expand the logical span of communications across transport technologies. They do not make any logical changes to the byte sequence (payload) and are transparent to it.
- *Physical/link/network gateways* convert the communications between different physical, link, and networking technologies.

In practice, connectivity gateways may span multiple layers of the connectivity stack (see Figure 2-1).

3.3 CORE GATEWAYS

Using a gateway to a core connectivity standard, a domain-specific endpoint can communicate with endpoints on other domain-specific technologies also connected via gateways to the core connectivity standard (see Figure 3-2). Core connectivity endpoints can directly communicate with each other, and via gateways with domain-specific connectivity endpoints.

Different functional domains may have different choices of core connectivity standards, due to different priorities on technical requirements, tradeoffs and ecosystems. To enable communication between different connectivity core standards, standardized gateways are needed. A standardized gateway between core connectivity standards is referred to as a *core gateway*. It allows domain-specific endpoints connected to one core standard to communicate with domain-specific endpoints integrated over another core standard, as shown in Figure 3-3. Also, it allows endpoints on the two core connectivity standards to interoperate.




Figure 3-3: A standardized gateway between core connectivity standards can allow domain-specific endpoints connected to one core standard to communicate with domain-specific endpoints integrated over another core standard.

To realize the goals of communication across functional domain and horizontal interoperability across industries, a standardized Core Gateway shall be defined between each of the core connectivity standards, as shown in Figure 3-4.




Figure 3-4: Each core connectivity standard requires a standardized gateway to all other core standards. Each additional core standard creates increasing complexity and interoperability challenges. By restricting the design to a few core connectivity standards, we cover the needs of IIoT systems across the functional domains, and attain the goal of horizontal interoperability across industries.

Let K be the number of core standards. Then the number of core gateways that will be standardized is $K * (K - 1) / 2$, as shown in Figure 3-4. If N is the total number of connectivity technologies (see Figure 3-1), then only additional $(N-K)$ gateways are needed with the introduction of core standards. The total number of gateways required becomes $K*(K-1)/2 + (N-$

K) vs. the original $N*(N-1)/2$ shown in Figure 3-1. Assuming $K \ll N$, the number of gateways goes from $O(N^2)$ to $O(N)$, which is much more tractable.

Each additional core standard creates increasing complexity and interoperability challenges with the square of the number of core standards. A few (small K) core connectivity standards should suffice to cover the needs of IIoT systems across the functional domains and industries to attain the goal of horizontal interoperability.

3.4 CORE STANDARDS CRITERIA

A connectivity core standard should align with the priorities on the requirements, engineering tradeoffs and ecosystems in its functional domain. It should not get in the way of providing seamless interoperability between domain-specific endpoints connected to it via gateways. This means meeting not only the functional requirements, but also the non-functional requirements of reliability, performance, scalability, availability, security and safety. Below, we define the criteria for qualifying as a connectivity core standard.

A connectivity core standard shall:

- provide syntactic interoperability,
- be an open standard with strong independent, international governance, and with support for certifying or validating or testing interoperability of implementations,
- be horizontal and neutral in its applicability across industries,
- be stable and deployed across multiple vertical industries,
- have standards-defined Core Gateways to all other connectivity core standards.

A connectivity core standard shall provide syntactic interoperability (see section 2.2). It is not simply sending opaque blobs. Applications not only get the data, but they can also discover the data types to unambiguously parse and manipulate it as structured data. So, an application will, for instance, know that the data it received is a structure with three floating-point number fields and a string field. The connectivity stack (see section 2.1) does not provide semantics—the interpretation of the fields, such as the units, ranges, and context is important for IIoT systems, but outside the scope of connectivity, and covered by the Distributed Data Interoperability and Management layer in the Industrial Internet Reference Architecture.

A core connectivity standard shall be an open standard managed by a recognized standards development organization (SDO). The SDO should provide independent, international governance. There should be support for validating or certifying or testing interoperability of implementations adhering to a specification from the SDO.

A core connectivity standard shall be stable and deployed in systems across multiple industries. It should not qualify until it has been fielded and has operational proof points in fielded systems. Connectivity standards that are not proven deployed across multiple industries or in fielded systems can be considered a common connectivity standard in one or more specific industries. We should strike a balance between leading the industry and lowering risk. We set that balance at the point of deployed applications across industries.

A core connectivity standard shall have commitments from SDOs to build standards-based core gateways to the other core connectivity standards. This ensures syntactic interoperability between the core connectivity standards.

A core connectivity standard should support all the core functions of a connectivity framework. It should be fast, flexible, and impose minimal overhead. It should be a proven, well-established technology, and be open and extensible to future needs of the most demanding IIoT systems.

Specifically, it should meet the following technical criteria:

- the connectivity framework functional requirements described in section 4.1, within each functional domain and across functional domains,
- the non-functional requirements of performance, scalability, reliability, resilience, within and across functional domains,
- security and safety requirements within and across functional domains,

and the following business criteria:

- not require any single component from any single vendor (consistent with the internet model) and
- have readily-available, professionally-supported Software Development Kits (SDKs) from multiple vendors, ideally including both commercial and open source.

The technical and business criteria ensure that an endpoint can use a gateway to any core connectivity standard to communicate with other endpoints connected via a gateway to another core connectivity standard.

The design of specifying only a few core standards with core gateways amongst them mitigates the “N-squared” problem (see section 3.1, Figure 3-1). The core connectivity standards bear the burden of mapping to all other core connectivity standards. Core connectivity standards allow all other domain-specific connectivity technologies (standard or non-standard) prevalent within a domain to continue to be used, while providing a pathway for an open architecture to communicate with the larger IIoT ecosystem. Domain-specific connectivity technologies will need a gateway to one of the core connectivity standards. Those gateways can be products, hardware or software, standard or not. There is a practical need to limit the number of core standards to just a few, and judiciously allow for new ones to be added, if there is clearly no significant overlap with existing core standards. Otherwise we would be back again to an N^2 problem (see section 3.1, Figure 3-1) amongst the connectivity core standards.

4 CONNECTIVITY FRAMEWORK LAYER

The *connectivity framework* layer provides a logical data exchange service to the endpoints participating in an information exchange. It can observe and “understand” the data exchanges, and use that knowledge to optimize data delivery. It is a logical functional layer on top of the connectivity transport layer (see Figure 2-1) and should be agnostic to the technologies used to implement connectivity transports.

The key role of the connectivity framework layer is to provide syntactic interoperability among the endpoints. Data that is exchanged is structured in a common, unambiguous data format, independent of endpoint implementation, and decoupled from the hardware and programming platform. Depending on the application logic behind endpoint, one or more data exchange patterns may be required. There are two predominant data exchange pattern styles: publish-subscribe (see section 4.1.6) and request-reply (see section 4.1.7).

A key benefit of the connectivity framework is to abstract and hide the implementation of the various functions so that the applications that use the connectivity framework won’t need to know the implementation, just use its capabilities. It reduces the cost of development and increases productivity and quality.

4.1 CORE FUNCTIONS

The key connectivity framework functions include a data resource model, publish-subscribe and request-reply data exchange patterns, data quality of service, data security, and a programming API. These are summarized in Figure 4-1 and described below.

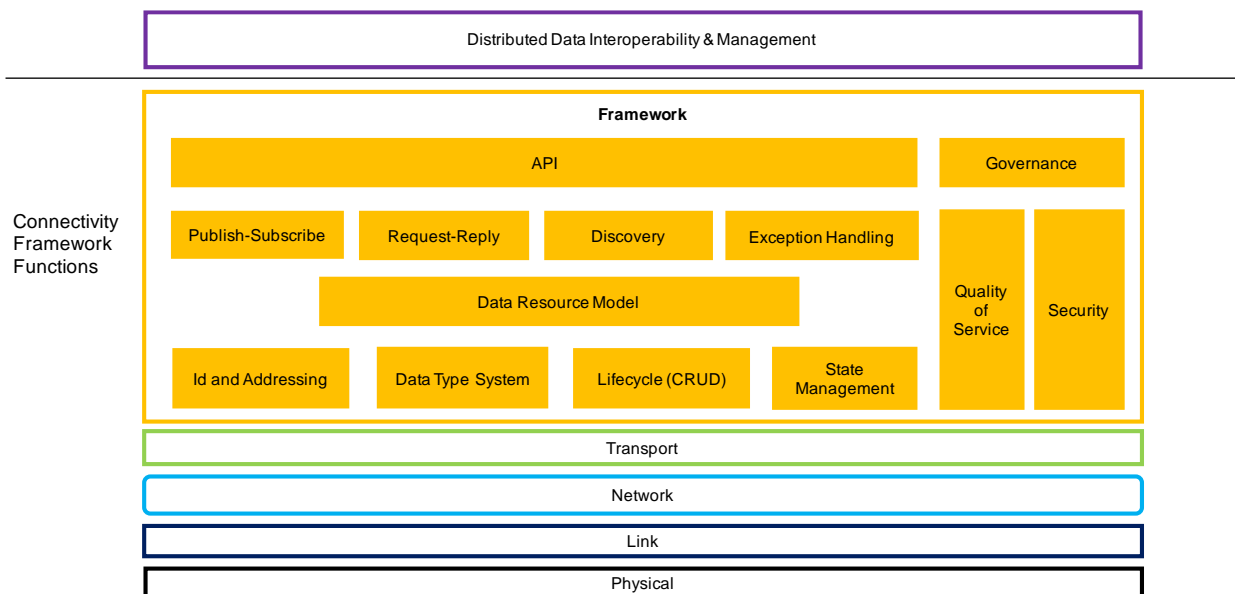


Figure 4-1: Connectivity framework layer functions.

4.1.1 DATA RESOURCE MODEL

Connectivity frameworks provide a way of representing data objects¹ that can change state over time. A data-object is a structured collection of fields, as in a programming language. It may be hierarchical and may be statically or dynamically typed.

A connectivity framework distributes changes to data-object amongst the participants.

Data models for different application areas or industries are usually mapped into the abstract data-objects provided by a connectivity framework.

4.1.2 ID AND ADDRESSING

A connectivity framework provides the means to identify and address each data object. The *id* is used to address a data object and read and write fields in the data-object representation. It could be:

- an explicit id field in the data-object representation, or
- an implicit id based on specially marked fields in the data object representation or
- a uniform resource identifier (URI) within the namespace of a device or application or network endpoint.

4.1.3 DATA TYPE SYSTEM

To ensure syntactic interoperability, a connectivity stack (see section 2.1) will provide a way to describe the data syntax. A *data type* is a syntactic constraint placed upon the interpretation of data. It is not possible to connect systems without sharing or mapping data types, either implicitly (e.g. in code) or explicitly.

A connectivity framework provides a *data type system* for representing data objects as structures in a programming environment and for formatting data to be communicated on the wire. The data type system may be *object-oriented*, like the data types found in statically-typed programming languages (e.g. C, C++, C# or Java), or *object-based* like the dynamic data types found in dynamically-typed programming languages (e.g. JavaScript, Python, Lua).

The data type system should provide a means of managing the evolution of data types. This includes versioning and assignability rules across versions, so that applications using newer versions of a data type can communicate with applications using older versions of a data type to the maximum extent possible.

The data type system also defines the serialized data format in communication (in motion) and in storage (at rest), and operations to serialize from a programming language representation into the serialized format, and to de-serialize back into the programming language representation.

Requiring explicit data syntax definitions using a data type system enables functional standardized gateways (see section 3.2). Also, generic tools can collect and traverse syntactically

¹ a.k.a. data resource or data item or data point or tag

meaningful data. Explicit data typing allows intelligent or standard technologies that can map data representations.

4.1.4 DATA RESOURCE LIFECYCLE (CRUD)

A connectivity framework should provide a means to manage the lifecycle of a data object. These include the following four critical operations, abbreviated as “CRUD”:

- create (C): a means of creating or introducing a new data object,
- read (R): a means of observing the state of a data object,
- update (U): a means of updating the state of a data object and
- delete (D): a means of deleting a data object.

4.1.5 STATE MANAGEMENT

IIoT components need access to high frequency, high-volume data beyond when it was initially produced. For example, a component may require the last n updates to plan the next action or make a prediction.

A connectivity frameworks can manage the historical state of the data objects. They can cache the last n updates to a data object so that applications can simply examine the historical state.

Connectivity frameworks can get the current state, even though the state may have changed long before a participant joined the system or reconnected. It may maintain a virtual view of the data objects associated with an endpoint and synchronize it when a connection is re-established.

4.1.6 PUBLISH-SUBSCRIBE

A connectivity framework should support the publish-subscribe data-exchange pattern in which a component publishes data on a well-known topic without regard to subscribers, and a component subscribes to data from the well-known topic without regards to publishers. This decouples publishers from subscribers, using only the channel for the topic at hand, so that components are loosely coupled and can be replaced independently of one another. An endpoint may operate in both a publisher and subscriber role. This data exchange pattern is also called the *push* data-exchange pattern.

The publish-subscribe data exchange pattern is useful for one-to-many and many-to-many data distribution scenarios, including streaming, alarms and events, command and control, and configuration ([Industrial Internet Reference Architecture, IIRA](#))¹.

The decoupling in space (location) and time (asynchronous delivery) provided by publish-subscribe data exchange pattern achieves the reliability, performance and scale demanded by IIoT systems. It also decreases the likelihood of fault propagation and simplifies incremental updating and evolution.

¹ See [IIC-IIRA2015]

4.1.7 REQUEST-REPLY

A connectivity framework for IIoT should support the request-reply data exchange pattern. This data exchange pattern uses *requestors* that can initiate a service request to be fulfilled by endpoints in the *replier* role. An endpoint may operate in both a requestor and a replier role. This pattern is also called a *pull* or *request-response* data exchange pattern.

The request-reply data exchange pattern is useful when working with a sparse subset of large data—for example to query specific data objects or invoke specific services.

The request-reply data exchange pattern permits synchronous or asynchronous exchange of data between endpoints. In synchronous request-reply, a requestor waits for the replies before issuing the next request. In asynchronous request-reply, a requestor can have multiple outstanding requests and replies are processed as they are received.

4.1.8 DISCOVERY

To support more intelligent decisions, the discovery, authentication and access to services (including data exchanges) should be automated.

Connectivity frameworks provide mechanisms to discover the:

- publish-subscribe topics and the associated quality of service
- request-reply services and their associated quality of service,
- data types associated with the topics and services, and
- endpoints participating in a data exchange.

4.1.9 EXCEPTION HANDLING

A connectivity framework should also provide for exception handling, for example when there are disruptions in connectivity. This could happen because of:

- disconnected or intermittent links (at the lower layers),
- switching network interfaces (e.g. between wired and wireless links),
- changes in network configuration (e.g. cable replaced, network ports moved),
- data quality of service needs not met,
- remote endpoint or component failure, or
- non-responsive participants.

A connectivity framework should shield the data flows from the impact of such exceptions, and should provide a means of informing the applications when an exception cannot be automatically managed by the connectivity framework.

4.1.10 DATA QUALITY OF SERVICE (QoS)

IIoT data exchanges can have varying requirements for how the data is delivered. Those aspects are referred to as the data quality of service (QoS).

A connectivity framework should support these data exchange QoS categories.

Delivery refers to the delivery aspects of the data including:

- *Best-efforts delivery*: An update is sent once, regardless of whether the receivers get it. Also called a *fire-and-forget* scheme, this is a form of “at most once” delivery. It is suitable when high-frequency periodic updates need to be distributed in a system and out-of-order or missing updates can be tolerated.
- *Reliable delivery*: An update is sent and also cached by the sender for later redelivery, in case receiver(s) don’t get it in a timely fashion. The amount of caching and timing can be configured based on the application and data flow requirements. Acknowledgements from a receiving endpoint can be automatic at the connectivity framework level, or may require explicit response from the application. This is a form of “at least once delivery”. It is suitable for low frequency status updates, events and notifications and also for commands when updates from a source are expected in-order.

In addition:

- *Timeliness* is the ability of the connectivity framework to establish end-to-end timing constraints, adaptively reconfigure to either guarantee specified timing or minimize timing violations, and to notify the application if a timing constraint has been violated.
- *Ordering* is the ability of the connectivity framework to present the data in the order it was produced, or received, and collate updates from different sources in the system.
- *Durability* is the ability of the connectivity framework to make data available to late joiners, and extend the lifecycle of the data beyond that of the source when so desired, and survive failures in the infrastructure.
- *Lifespan* is the ability of the connectivity framework to expire stale data.
- *Fault tolerance* is the ability of the connectivity framework to ensure that redundant connectivity endpoints are properly managed, and appropriate failover mechanisms are in place when an endpoint or a connection is lost.

The underlying transport layer will ultimately bound a connectivity framework’s performance and scalability limits. The connectivity framework should introduce minimal overhead in providing the data exchange QoS and should have minimal impact on the overall performance and scalability.

4.1.11 DATA SECURITY

A connectivity framework should provide the ability to ensure confidentiality, integrity, authenticity and non-repudiation of the data exchange, when so desired.

The connectivity framework security mechanisms should provide a means to:

- upon discovery, authenticate endpoints before allowing them to participate in a data exchange,
- authorize permissions (read, write) granted to the endpoints participating in a data exchange, to ensure that endpoints cannot write or read data that they have not been given access to,

- ensure data integrity and trustworthiness of the data delivery, so that received data is not tampered with while stored or in transit and
- selectively encrypt sensitive data flows.

This last point is important, since certain high volume data flows may not be sensitive enough to warrant the extra overhead of encrypting and decrypting the data. The decision to encrypt should be based on a risk-impact assessment.

The connectivity framework access-control-model should be sufficiently fine-grained to limit the permissions of each endpoint narrowly to the operations and services needed for performing their intended functions. This enables the application of the *principle of least privilege* that is essential to limit the consequence of security breaches and insider attacks.

The connectivity framework security mechanisms should provide secure logging and auditing capabilities to detect security attacks and assess their consequences.

For more details, please refer to the [Industrial Internet Security Framework \(IISF\)](#)¹.

4.1.12 API

IIoT systems involve multiple software components, developed by multiple parties over time, with a variety of programming languages. Therefore, IIoT software development requires an *Application-Programming Interface* (API) to support the design and implementation of application-specific data exchanges.

Some connectivity frameworks provide standardized APIs in various programming languages (e.g. C, C++, C#, Java, Python, Lua, Javascript, and so on), to ease the portability of application code from one implementation to another and to decouple the application from the framework implementation. Others define a protocol interface, and let the implementers define the programming API. This makes it harder to switch implementations, but allows the APIs to be customized to taste.

4.1.13 GOVERNANCE

A connectivity framework should provide a means to configure, administer, and monitor its operation. These include all aspects of the connectivity framework functions, including data types, data quality of service, data security policies, resource management, and timing.

Some connectivity framework standards define the mechanisms for configuration, and administration. Others do not standardize on the mechanisms and leave it up to the implementations. Mechanisms may be file based or API based or both.

Monitoring is useful for diagnostics and troubleshooting of an operational IIoT system. It should be configurable to the desired level of detail. Connectivity framework standards may define the mechanisms for monitoring or may leave them up to the implementations.

¹ See [IIC-IISF2016]

4.2 TYPICAL CONSIDERATIONS

Typical considerations for choosing a connectivity framework can be grouped into system, data, performance, scalability, availability, deployment and operational considerations. The tradeoffs in each should be carefully evaluated.

4.2.1 SYSTEM ARCHITECTURE CONSIDERATIONS

4.2.1.1 PEER-TO-PEER VS. BROKER

Peer-to-peer is a symmetric data exchange pattern between endpoints without any intermediary or broker. A peer-to-peer architecture provides the lowest latency and jitter data exchange between endpoints. It can also avoid startup dependencies, as peers can come up in any order. There is a no single point of failure or vulnerability. However, a distributed peer-to-peer based system requires more careful planning---for example, one may need relays to avoid undue load on extremely resource constrained peers.

On the other hand, a *broker-based architecture* requires running a centralized process on a host in the system. Data exchanges flow through the broker. It needs to be started and run before the endpoints can communicate. A broker can become a choke point and a single point of failure, if not mitigated by redundancy and load balancing. Latency and determinism can suffer, especially as data volume goes up. It can increase vulnerability from a security perspective, but provisioning and management can be centralized.

4.2.1.2 DATA-CENTRIC VS. DEVICE/APP-CENTRIC

In a *data-centric* architecture, the endpoint application code does not need to be aware of who produces or consumes the data. Data is regarded as a first-class citizen that can be exchanged, stored, transformed and manipulated in its own right, independently of the endpoints that produce or consume it. There is an analogy with databases, which provide a data-centric abstraction for *data at rest*. Data-centric connectivity frameworks provide a data-centric abstraction for *data in motion*. Integrating new applications requires them to have knowledge of the data-centric abstraction.

In a *device-centric* architecture, the endpoint application code is aware of the devices or application endpoints responsible for producing or consuming data. Devices or application endpoints are regarded as the first-class citizen, and are modeled as such; data cannot exist without the context of a device or application. Integrating new types of devices or applications requires integrating new interfaces.

Data-centric connectivity frameworks provide *location, device and application independence*. They allow components to be decoupled from one another, developed and integrated independently. Device-centric connectivity frameworks require application components to understand the device context to use the data meaningfully. Each kind of device is be integrated separately, and the applications are aware of their behavioral interfaces.

Data-centric connectivity frameworks fit well with multi-party development, simplify the integration effort and reduce the overall effort and time to debug and integrate components into systems. Because the interfaces to the data are explicit in the design, a data-centric approach results in an *open architecture*. Components become independent of the data, thereby simplifying the system development and evolution and increasing reusability. Data-centric systems also scale well to large systems. However, it can be challenging for IIoT systems integrators to require their vendors use a common data-centric abstraction.

For IIoT systems, open architecture is highly desirable, since it allows multiple data sources to be combined together to generate new value, insights and applications.

Many system designers choose between device-centric or data-centric core connectivity standards based on the relative need of integrating similar brands of devices vs. reducing complexity and easing development of the software. Regardless of the chosen method, connectivity frameworks should allow for both approaches.

4.2.1.3 EXPLICIT VS. IMPLICIT GOVERNANCE

Connectivity framework governance (see section 4.1.13) may be explicit or implicit or a mix. When governance is *explicit*, configuration elements can be controlled independently of the applications; they can be shared and managed through a common repository. When governance is *implicit*, configuration elements are embedded within the application code across the various system components.

Explicit governance allows data architecture evolution and upgrades in a controlled fashion, independently of the application code. This is beneficial for large teams working on safety-critical systems where the development needs to be carefully managed, while enabling multiple sub-teams to work independently.

Implicit governance works best for organic evolution and requires that the data architecture be discoverable via dynamic APIs.

4.2.2 DATA CONSIDERATIONS

4.2.2.1 CONTENT-BASED SELECTION

IIoT systems involve movement of large volumes of data. Components are only interested in a specific subset of the data at a given time, although that interest set may change. Given the data interest set across the components, connectivity frameworks can optimize the data distribution, resulting in lower overall system resource footprint, and so lower system cost.

For IIoT systems, the ability to specify a content-based data subset of interest and automatically optimize the data flows is highly desirable.

4.2.2.2 TIME-BASED SELECTION

IIoT systems typically involve distribution of high frequency data. A component may produce data faster than some consuming components desire or are able to handle. In this situation, time-

based filtering in the connectivity infrastructure is required. For example, a sensor may generate data at a rate of 1000Hz rate, but a user display may not require data at a rate faster than the display refresh rate of, say, 30Hz.

A consumer's desired data rate may change over time or for different data items. By knowing the desired data-rate needs across components, connectivity frameworks can optimize the use of system resources for data distribution. This can result in lower overall system resource footprint and lower system cost.

For IIoT systems, the ability to specify a time-based data subset of interest and automatically optimize the data flows is highly desirable.

4.2.3 PERFORMANCE CONSIDERATIONS

4.2.3.1 REAL-TIME

“Real time” is more about deterministic response than it is about fast response. Many systems require low average latency, but real-time systems succeed only if they always respond “on time”. This is the maximum latency, and can be expressed as the average delay plus the variation or jitter. Even a fast server with low average latency can experience large jitter under load. For real-time operation, the latency needs to be predictable (i.e. the jitter should be consistently small).

4.2.3.2 LATENCY AND JITTER VS. THROUGHPUT

Throughput refers to the volume of data distributed per unit time. The throughput demands can vary widely—for example, under load or stress in an emergency situation, there may be a lot more communication compared to normal or steady state operation. Latency and jitter can suffer when throughput demands increase on the connectivity infrastructure without increased capacity. A connectivity framework should be able to meet the latency and jitter requirements for real-time performance as with increasing throughput demand, a core consideration of the quality-of-service function.

For IIoT systems, the latency and jitter vs. throughput tradeoffs should be carefully evaluated, and the limiting factors for throughput and latency should be understood.

4.2.4 SCALABILITY CONSIDERATIONS

4.2.4.1 DATA OBJECTS

When the number of data objects increases, it is no longer practical to send every update to every possible consumer. Connectivity frameworks should support data-object scaling by offering run-time introspection so consumers can choose data objects of interest, and configure producer update distribution to a sparser set currently of interest. A producer can also batch multiple data-object updates destined for the same consumer to make the data distribution efficient and scalable.

For IIoT systems, a connectivity framework should effectively handle an increasing number of data objects as more memory resources are added, and should support data objects of varying sizes.

4.2.4.2 APPS

IIoT systems comprise independently developed applications, each with evolving interfaces and data formats that should continue to interoperate with older versions of that interface. Forcing all apps to coordinate their update simultaneously to a new interface and data format version is unrealistic; there needs to be version negotiation between components. Interacting teams need tools, processes and eventually architectural support to solve the system-integration problem.

Data-centric connectivity frameworks allow applications to control data-oriented interfaces directly. Applications interact with shared data objects described by explicitly defined data types. Differences in data-oriented interfaces between apps can be automatically detected and adapted to match a participant's expectations, so as to decouple application interface dependencies and allow large projects to evolve interfaces and make parallel progress on multiple fronts.

For IIoT systems, a connectivity framework should support component interface evolution so that new capabilities can be added over time, without impacting the already existing components.

4.2.5 AVAILABILITY CONSIDERATIONS

4.2.5.1 REDUNDANCY

Failure of IIoT systems during operation can have fatal consequences. The system-relevant time period of continuous operation is dependent on the context of the system. For a power plant, the relevant time period could span years. For a medical imaging machine, the relevant time period could be only a few seconds.

But hosts and networks do fail, so redundant infrastructure (duplicate or triplicate or more) and failover mechanisms should be put in place. They rely on the connectivity infrastructure to communicate fault conditions and effect the appropriate state changes. To provide continuous system availability, a connectivity framework should support redundant endpoints and networks, and remove duplicate data transparently when the same update is received over multiple paths.

For IIoT systems, a connectivity framework should support continuous availability over a system-relevant time period.

4.2.5.2 RECOVERY

An IIoT connectivity framework should be continuously available. It should not have single points of failure, and it should provide mechanisms for timely detection of system component failures. There should be mechanisms for component state durability and for state recovery and failover.

For IIoT systems, a connectivity framework should provide mechanisms for data durability and state recovery from fault conditions.

4.2.6 DEPLOYMENT CONSIDERATIONS

4.2.6.1 PLATFORM CONSTRAINTS

IIoT system components run on a variety of platforms, from small resource-constrained devices to enterprise-class machines. Generally, the development environment is different from the deployment environment. The memory footprint, CPU, programming language and environments can vary greatly across these hosts. The connectivity platform should be supported on the development and deployment compute platforms used in the system.

For IIoT systems, a connectivity framework should support the operating system, the CPU and the resource constraints on the platform(s) being used.

4.2.6.2 INCREMENTAL UPGRADES

For IIoT systems that have long lifespans, components are upgraded incrementally. Newer or updated components may use newer versions of connectivity framework software. Connectivity framework that support backwards and forward version compatibility can facilitate the upgrade process.

For IIoT systems, a connectivity framework should provide backwards compatibility of communication protocols and data structures, so that components can be upgraded incrementally.

5 CONNECTIVITY TRANSPORT LAYER

The connectivity transport layer provides a logical transport network connecting the endpoints. The connectivity transport is akin to a pipe, opaque to the data flow amongst endpoints.

The key role of the *connectivity transport* layer is to provide *technical interoperability* among the endpoints.

5.1 CORE FUNCTIONS

The key connectivity transport functions include endpoint addressing, modes of communication, network topology, connectedness, prioritization, timing and synchronization, and message security. These are summarized in Figure 5-1 and described below.

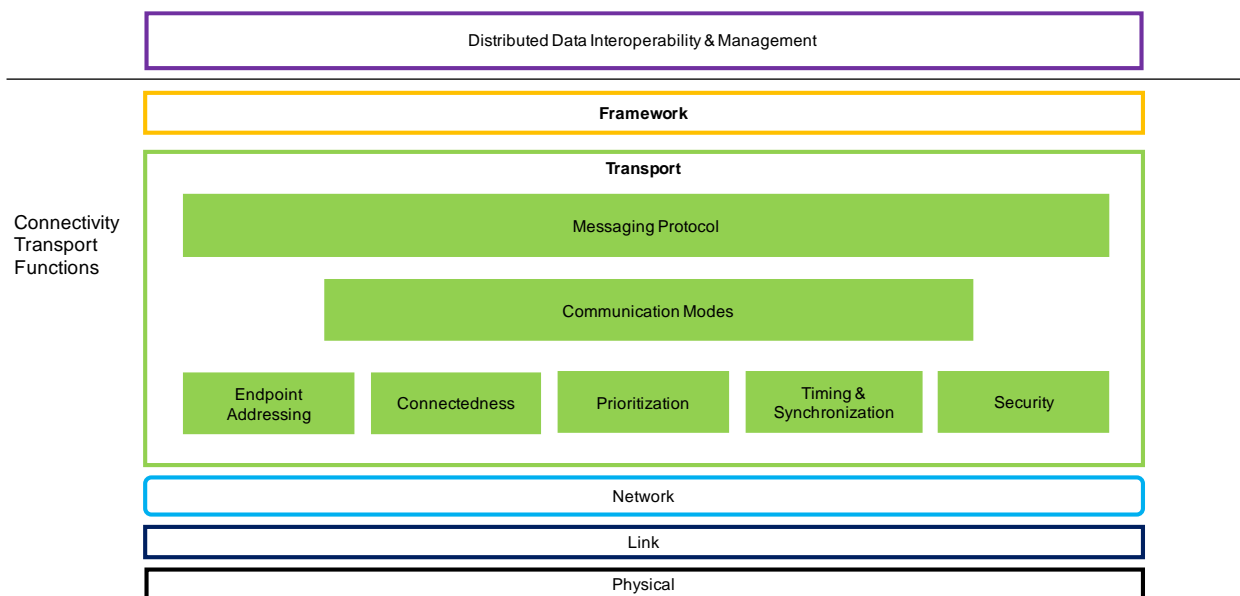


Figure 5-1: Connectivity transport layer functions.

5.1.1 MESSAGING PROTOCOL

The messaging protocol is the wire protocol that describes the format and behavior of the messages exchanged between the endpoints. A messaging protocol should support the implementation of connectivity framework layer functions (see section 4.1). It may be directly exposed for use by applications, possibly using ad hoc (and unnamed) application-specific connectivity frameworks.

The messaging protocol may include discovery, authentication, session establishment, message retry and acknowledgment, fragmentation and re-assembly of large messages, data encoding and serialization, message reorder and de-conflicting across connectivity transports.

Messaging protocols may be configured and optimized for different network layer configurations. Network layer parameters such as bandwidth, round-trip time and maximum message size should inform the selection of the messaging protocol quality of service.

5.1.2 COMMUNICATION MODES

A connectivity transport may support the following communication modes:

- unicast—suitable for one-to-one communication between two endpoints,
- multicast—suitable for one-to-many communication between endpoints and
- broadcast—suitable for one-to-all communication between endpoints, where “all” refers to all the endpoints present on the communication transport network at the time of transmission.

5.1.3 ENDPOINT ADDRESSING

Any of the nodes (for example, a device or an application host) in IIoT systems can house one or more components, each with one or more connectivity endpoints. An address identifies a node for network-level communication purposes. This address could be locally unique and possibly globally unique. A node and hence the endpoints residing on it may be reachable over multiple addresses.

The addressing scheme and associated infrastructure should support endpoints on the Internet scale.

5.1.4 CONNECTEDNESS

Network layer protocols (see Figure 2-1) offer either connection-oriented or connectionless services for delivering packets across the network. Connectionless services are more common at the network layer. In many protocol suites, the network-layer protocol is connectionless, and the transport layer provides connection-oriented services. For example, in TCP/IP, the Internet Protocol (IP) and the User Datagram Protocol (UDP) layered on top of it are connectionless, while the Transmission Control Protocol (TCP) is connection-oriented.

A connectionless transport is best for low latency and jitter applications or when a high degree of scalability is required in a local area network. The connectionless UDP transport has proven itself for use real-time applications.

A connection-oriented transport is best suited for high throughput applications in a network with complex topology and high variation of traffic loads, since it provides a “virtual circuit” that reduces the variation in routing path. The connection-oriented TCP transport is battle tested for transiting through *firewalls* and *network address translation* (NAT) routers, and connecting across wide area networks. New applications may call for connection-oriented connectivity transports that do not suffer the drawbacks that we find in TCP today, such as unbounded retransmission delays.

When using a connectionless transport, the connectivity framework design needs to handle failures in the transport caused by loss or out-of-order packets. Consequently, designing a

connectivity framework based on the connection-oriented transport may preclude it from providing a connectionless data exchange.

5.1.5 PRIORITIZATION

IIoT systems need to ensure that critical data is delivered ahead of non-critical data.

The connectivity transport function can provide the ability to prioritize some messages over others in the data exchange between endpoints.

5.1.6 TIMING & SYNCHRONIZATION

IIoT systems need a way to synchronize local endpoint clocks over a connectivity transport network. Many methods are in use today, including NTP- or PTP-based time synchronization and GPS clocks, and new approaches are in development.

The connectivity transport function may provide the ability to synchronize time across the network.

5.1.7 MESSAGE SECURITY

The security mechanisms provided by the connectivity transport layer should implement and enforce the connectivity-framework-layer data security function (see section 4.1.11).

Transport layer security involves both the messaging protocol and the network layer security. Both should provide mechanisms for endpoint authentication, message encryption and message authentication. Security implemented by each function may provide controls with different granularity and be separately administered.

At the network level, network endpoint security mechanisms can grant access based on policy and enforce security by means of encrypted virtual local area networks (VLANs) and firewalls.

At the messaging protocol level, message oriented security mechanisms based on policy can enforce permissions by fine-grained cryptographic means. For example, different data flows may be configured to use different cryptographic keys such that permissions granted to an application to access one flow does not allow it to observe a different flow.

There may be multiple transport and network hops between endpoints. End-to-end security is desired, and security should not be compromised when crossing gateways, proxies and bridges between the endpoints.

For more details, please refer to the [Industrial Internet Security Framework \(IISF\)](#)¹.

¹ See [IIC-IISF2016]

5.2 TYPICAL CONSIDERATIONS

5.2.1 NETWORK LAYER CONSIDERATIONS

5.2.1.1 TOPOLOGY

A transport may require or preclude a specific network topology. Network topologies in IIoT systems include:

- point-to-point,
- hub-and-spoke,
- meshed,
- hierarchical and
- a combination of the above.

Connectivity gateways can be used to bridge multiple networks and topologies, and to form more complex topologies, as needed.

For IIoT systems, a transport should not restrict the network topology.

5.2.1.2 SPAN

A transport communication path may span across different physical geographies. A logical transport layer can span the local area networks (LAN), large geographic distances (wide area networks i.e. WAN), or somewhere in between (metropolitan area networks i.e. MAN).

For IIoT systems, it is desirable for a transport to support a variety of network spans, including, LAN, MAN, WAN, and possibly space networks.

5.2.1.3 SEGMENTATION

IIoT systems need a way to separate data from different functional domains over the same network.

The transport may provide the ability to segment a network, to isolate different functional domains and one set of data exchanges from another.

For IIoT systems, it is highly desirable that the connectivity transport be able to support multiple independent and isolated communication paths between the same network endpoints.

6 HOW TO ASSESS A CONNECTIVITY TECHNOLOGY?

We apply the Industrial Internet Reference Architecture (IIRA)¹ viewpoints to create an *assessment template* for use in evaluating any connectivity technology. The connectivity functions (see sections 4.1 and 5.1) constitute the functional viewpoint; the typical considerations (see sections 4.2 and 5.2) constitute the implementation viewpoint. The business and usage viewpoints are described below.

The assessment template is intended to be a tool for understanding any connectivity technology in the context of the IIoT needs. The worksheet is helpful for:

- understanding how a connectivity technology supports specific IIoT functional needs,
- evaluating a connectivity technology's trades-offs for typical IIoT considerations and
- determining a connectivity technology's suitability for a particular use case (once the specific requirements are understood).

The worksheet helps categorize objectively a connectivity technology across the layers of the IIoT connectivity stack model (see Figure 2-1) based on the functions it supports: is it a connectivity framework (Figure 4-1) or a connectivity transport (Figure 5-1)? Some technologies span multiple layers of the connectivity stack (Figure 2-1).

Connectivity technologies can be compared objectively, and the most applicable connectivity technology can be easily identified.

The worksheet is described below.

6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i>
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i>
Description	<i>Short synopsis of the technology.</i>
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i>
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i>
References	<i>Website and other useful links to the technology.</i>

¹ See [IIC-IIRA2015]

6.2 Business Viewpoint	
6.2.1 Purpose (Section 6.2.1)	<i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i>
6.2.2 Pedigree (Section 6.2.2)	<i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i>
6.2.3 Variants (Section 6.2.3)	<i>Describe the options and variants from the original generic description of the technology.</i>
6.2.4 Maturity (Section 6.2.4)	<i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i>
6.2.5 Stability (Section 6.2.5)	<i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i>
6.2.6 Standards Body (Section 6.2.6)	<i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i>
6.2.7 Openness (Section 6.2.7)	<i>Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?</i>

6.3 Usage Viewpoint	
6.3.1 Architecture (Section 6.3.1)	<i>Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.</i>
6.3.2 Technology Options (Section 6.3.2)	<i>List the choices to be made for using the connectivity technology in a system.</i>
6.3.3 Applications (Section 6.3.3)	<i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i>
6.3.4 Typical Usage (Section 2.2)	<i>What function or where in the system this technology is typically used?</i>
6.3.5 Operations (Section 2.3.8)	<i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i>
6.3.6 Security (Section 2.3.5)	<i>What are the system security implications of this connectivity technology?</i>
6.3.7 Safety (Section 2.3.9)	<i>For systems that need it, are certifiable implementations available?</i>
6.3.8 Gateways (Section 3.3)	<i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i>

6.4 Functional Viewpoint	
6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<i>Does it provide a data resource model? Summarize the salient aspects.</i>
ID & Addressing (Section 4.1.2)	<i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i>
Data Type System (Section 4.1.3)	<i>Does it provide a data type system? Summarize the salient aspects.</i>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<i>Does it provide a means of managing a data object's lifecycle? Summarize the salient aspects.</i>
State Management (Section 4.1.5)	<i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i>
Publish-Subscribe (Section 4.1.6)	<i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i>
Request-Reply (Section 4.1.7)	<i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i>
Discovery (Section 4.1.8)	<i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i>
Exception Handling (Section 4.1.9)	<i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i>
Data Quality of Service (QoS) (Section 4.1.10)	<i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i>
Data Security (Section 4.1.11)	<i>Does it provide a data object security model? Summarize the salient aspects.</i>
API (Section 4.1.12)	<i>Is there a standard API? Which programming languages is it available for?</i>
Governance (Section 4.1.13)	<i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i>
6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i>
Communication Modes (Section 5.1.2)	<i>Which communication modes does it support?</i>
Endpoint Addressing (Section 5.1.3)	<i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i>
Connectedness (Section 5.1.4)	<i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i>
Prioritization (Section 5.1.5)	<i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i>
Timing & Synchronization (Section 5.1.6)	<i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i>
Message Security (Section 5.1.7)	<i>Does it provide mechanisms for message security? Summarize the salient aspects.</i>

6.5 Implementation Viewpoint	
6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<i>Does the connectivity framework require running a special process or broker?</i>
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i>
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<i>Is the governance explicit and shareable?</i>
6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<i>Can a content-filter specify the data subset of interest?</i>
Time-Based Selection (Section 4.2.2.2)	<i>Can sub-sampling specify the data subset of interest?</i>
6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i>
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<i>How does the latency and jitter change with throughput? What limits the throughput?</i>
6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i>
Apps (Section 4.2.4.2)	<i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i>
6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i>
Recovery (Section 4.2.5.2)	<i>Can the connectivity framework support recovery when fault conditions occur?</i>
6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i>
Incremental Upgrades (Section 4.2.6.2)	<i>Does the connectivity framework facilitate incremental upgrades?</i>
6.5.7 Network Layer Considerations	
Topology (Section 5.2.1.1)	<i>What network topologies are allowed?</i>
Span (Section 5.2.1.2)	<i>What is the span of the transport: LAN vs. WAN?</i>
Segmentation (Section 5.2.1.3)	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i>

7 CONNECTIVITY STANDARDS

This chapter lists the prominent connectivity standards for the framework and transport layers, using the assessment template defined in chapter 6, so that the standards can be understood in the context of IIoT needs. As additional standards are identified, they should be added to this list. The assessment template allows us to capture and describe the technologies in a uniform and objective manner.

Figure 7-1 summarizes the prominent IIoT connectivity frameworks and transports. It shows connectivity frameworks that have originated in an industry agnostic manner for general-purpose use. The dotted boxes show the connectivity standards that have originated in certain industry verticals with a specific application focus that has applicability across multiple industries. Some connectivity frameworks define their own transport protocols (e.g. DDS, OPC-UA), and in the diagrams those are shown without any gap between the framework and the transport layer boxes. Others (e.g. Web Services, oneM2M) rely on general-purpose transport protocols. The network (IP) and lower layers (wired and wireless) are also shown for completeness, but are outside the scope of this document.

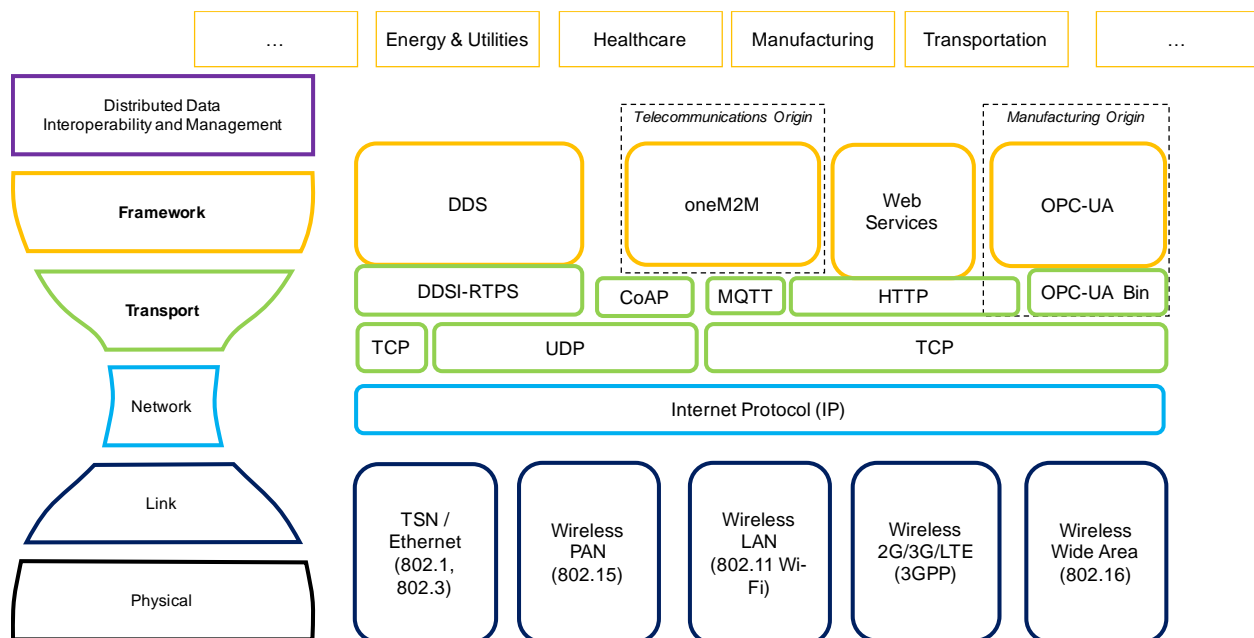


Figure 7-1: IIoT connectivity standards. Dotted boxes show the connectivity standards (e.g. oneM2M, OPC-UA) that have originated in respective industry verticals (e.g. telecommunications, manufacturing) to provide enabling features for those industries, and also offer an application focus that is applicable to multiple industries. Others (e.g. DDS, Web Services) have originated in an industry agnostic manner for general-purpose use, and are applied in multiple industries for many different kinds of application areas. Transports that are specific to a framework layer are shown without any spacing between the framework and the transport layer boxes.

The distinction between transport and framework layers is important. To be considered a connectivity framework at a minimum, a connectivity transport would have to be paired with a data type system. For instance, a connectivity transport such as MQTT could be paired with data type system technology such as [protocol buffers](#)¹, and could be used to create a custom connectivity framework. However, there is currently no standard that describes such a pairing.

The connectivity framework and transport standards shown in Figure 7-1 are discussed below. Detailed assessments are provided, starting from Annex A.

7.1 IIOT CONNECTIVITY FRAMEWORK STANDARDS

7.1.1 DATA DISTRIBUTION SERVICE (DDS)

Data Distribution Service (DDS) is an open connectivity-framework standard specifically targeted at IIoT applications. The Object Management Group (OMG) maintains the DDS family of specifications in the [DDS Portal](#), including [Remote Procedure Call over DDS](#)².

DDS is generally used in the control, application, information, operations domains, and sometimes in the business domain (see Figure 1-1). DDS's main purpose is to connect components (devices or gateways or applications) to other components to enable real-time systems and system-of-systems. Components interact with a shared data space, and never directly with each other. Therefore, it is referred to as a *data-centric* middleware standard. It has roots in high-performance defense, industrial, and embedded applications.

DDS implements direct component-data-component communication via a relational data model. DDS is also referred to as a *databus* because it is the data-in-motion analog to a database that manages for "data-at-rest". Both a database and a databus implement the "data-centric" abstraction; applications interact with the infrastructure, not directly with each other. The difference is that a database saves *past* data that can later be *searched* by relating properties of the stored data. A *databus* manages *future* data by *filtering* on properties of the incoming data. Data centrality makes a database essential for large storage systems. Data centrality makes a databus a fundamental technology for large IIoT software integration and autonomous operation.

The DDS wire protocol is [The Real-Time Publish-Subscribe Protocol \(RTPS\) DDS Interoperability Wire Protocol Specification \(DDSI-RTPS\)](#)³ connectivity transport standard, as shown in Figure 7-1. DDSI-RTPS is independent of the underlying transport. Many applications use the UDP transport. DDS can also run over TCP, shared memory, backplane connections, hypervisor transports and many others. DDS does not require any special transport properties such as reliability support.

¹ See [GOO-PB]

² See [OMG-DDS] and [OMG-DDSRPC]

³ See [OMG-DDSI-RTPS]

Similar to the way a database controls access to stored data, a databus controls data access and updates by many simultaneous components. At its core, DDS is built around a data-centric publish-subscribe data exchange pattern. However, the standard also defines a request-reply data exchange pattern, and vendors offer queuing. The key abstraction is that applications interact with the databus itself, not directly with other applications participating in that interaction. DDS offers precise data-centric quality-of-service (QoS) control, reliable multicast, configurable delivery, multiple levels of data durability, history, component and transport redundancy, automatic discovery, connectivity management, and transport agnostic fine-grained data-centric security. In addition, one-to-many and many-to-one communications is a key strength. DDS offers powerful ways to filter and select exactly which data goes where, and “where” can be thousands of simultaneous components. To support small, edge devices, there are lightweight versions of DDS that run in constrained environments. The DDS databus ensures ultra-reliable operation and simplifies application code. It does not require servers, greatly easing configuration and operations while eliminating failure and choke points.

A DDS-based system has no hard-coded interactions between components. The DDS databus automatically discovers and connects publishing and subscribing components. No configuration changes are required to add new components (e.g. a smart machine) to a system. Components can be developed or sourced from independent parties. DDS overcomes problems associated with point-to-point system integration, such as lack of scalability, interoperability and the ability to evolve the architecture. It enables plug-and-play simplicity, scalability and exceptionally high real-time performance.

DDS is commonly used for system integration and for building autonomous systems, because of the flexibility, reliability and speed necessary to build complex or real-time applications. DDS is a proven technology for reliable, high performance, large-scale IIoT software systems across many vertical industries. IIoT applications using DDS include wind farms, hospital integration, medical imaging, autonomous planes and cars, rail, asset tracking, automotive testing, smart cities, communications, data center switches, video sharing, consumer electronics, oil & gas drilling, ships, avionics, broadcast television, air traffic control, SCADA, robotics and defense.

DDS Gateways exist for many other connectivity technologies, including DNP3, C37.118, Modbus, HLA, JMS and so on. The [DDS-Web v1.0 specification](#)¹ defines a standardized gateway for Web Services. A standard for a gateway between OPC-UA and DDS is underway at the OMG. The OPC Foundation is developing an [OPC-UA-DDS pubsub profile](#)² with the goal of adding DDS as an additional publish-subscribe communication option to OPC-UA. Work is underway at oneM2M, investigating an interworking gateway between oneM2M and DDS, a DDS protocol binding for oneM2M, and DDS based direct exchange of data between [oneM2M](#) entities³.

¹ See [OMG-DDSWEB]

² See [OPC-DDS]

³ See [ONEM2M-27]

For more details, and to determine the suitability of DDS for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex A.

7.1.2 WEB SERVICES USING HYPERTEXT TRANSFER PROTOCOL (HTTP)

Web services (Wikipedia)¹ using Hypertext Transfer Protocol (HTTP) refers to the application-specific connectivity frameworks, primarily devised for human user interaction interfaces. They rely on a RESTful² style of architecture (Fielding, Wikipedia³) using the HTTP connectivity transport standard to exchange textual data, as shown in Figure 7-1. It requires the TCP transport. The IETF⁴ maintains the HTTP open standard specification; the W3C⁵ maintains the web (HTML5) specifications.

Web services using HTTP are generally used in the application domain (see Figure 1-1). Data is represented in textual form (either as JSON or XML), and embedded in a hypermedia (HTML) context. A Uniform Resource Identifier (URI) represents a data object on a server. A client (app) sends a request to a web server, specifying a data object URI, an operation and a payload. The server replies with failure or success and a response payload. The communication is text-based and designed for human speeds. It is not efficient for device-to-device communications and not suitable for real-time communications.

For more details, and to determine the suitability of web services using HTTP for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex D.

7.1.3 OPEN PLATFORM COMMUNICATIONS UNIFIED ARCHITECTURE (OPC-UA)

OPC Unified Architecture (OPC-UA⁶) is a connectivity framework standard used in the manufacturing industry. OPC-UA is designed to support multiple transports. Currently transport mappings are defined for TCP with a OPC-UA Binary encoding connectivity transport standard or the HTTP connectivity transport, as shown in Figure 7-1. Work on a Web Socket transport mapping has started. All current transport mappings require the TCP transport. The OPC Foundation maintains the OPC-UA family of specifications.

At its core, OPC-UA targets device interoperability. Before OPC-UA (or its predecessor OPC), applications simply accessed devices directly through proprietary APIs provided by their vendors. Unfortunately, this meant that applications became dependent on the particular device they controlled. Worse, higher level applications such as Human-Machine Interfaces (HMI) had no easy way to find, connect to, or control the various devices in factories. OPC-UA is generally used

¹ See [W3C-WSA], for overview [WKPD-WS]

² RESTful means to adhere to a REST communications architecture style.

³ See [Fielding-2000], for overview [WKPD-REST]

⁴ See [IETF]

⁵ See [W3C]

⁶ See [OPC-UA]

in the operations domain and is also being applied to information and application domains (see Figure 1-1).

OPC-UA is an evolution of the classic OPC (Object Linking and Embedding for Process Control) standards. It unifies the various original OPC specifications and is an evolution from an API to a network protocol. Adapters are available to bridge between OPC-UA and classic OPC. OPC is operational in thousands of factories globally. Traditionally, OPC was used to configure and query plant-floor servers (usually Programmable Logic Controllers (PLCs)). Actual device-device communication was then effected via a hardware-based fieldbus¹ such as Modbus or Profinet.

OPC-UA retains some of that flavor; it connects and configures plant-floor servers. The UA version adds better syntactical data typing (see section 4.1.3) and semantic information modeling capabilities. There are many companion specifications that define information models for various device types. For example, Field Device Integration (FDI) defines a model that represents all fieldbus device types. A remote client such as a graphical interface can browse the device data controlled by a server on the floor. By allowing this introspection across many servers, clients can build a directory with cross-references of all the devices on the floor. Additionally, OPC-UA also addresses the specific needs of device-device communication and therefore does not anymore rely on additional fieldbus solutions. Its scalability allows for implementation on devices with very restricted hardware resources, such as sensor and actuator devices.

OPC-UA divides system software into clients and servers. The servers usually reside on a device; they provide a way to access the device through a standard “device model”. There are device models for dozens of types of devices from sensor to feedback controllers. Each manufacturer is responsible for providing the server that maps the generic device model to its particular device. The servers expose an object-oriented, remotely-callable API that implements the device model.

Generic device models are central to the OPC-UA architecture. For example, the object model for a motor starter includes methods for setting parameters, reading data and operating the starter. Thus, applications can control a starter directly without being dependent on the manufacturer’s particular implementation.

OPC-UA is developing a “pub-sub” capability. This will provide direct device-to-device connection. There will be several “profiles” using different underlying protocols. The UDP profile supports multicast for efficiency. It targets simple implementation and does not attempt advanced functions like fragmentation control (sending large data types in pieces), reliability or quality-of-service control. Another profile is designed for connection to cloud-based data analytics. There is work on a DDS profile that will provide more sophisticated pub-sub functionality.

OPC-UA targets all kinds of manufacturing, including automotive, oil and gas, pharmaceuticals, food & beverage, medical machines, machine tools. It connects applications at the factory-floor

¹ A “fieldbus” is the name of a family of industrial *computer network* protocols used for real-time distributed control.

level as well as between the factory-floor and the enterprise IT cloud. There are explorations of use outside of manufacturing, although no deployed implementations.

For more details, and to determine the suitability of OPC-UA for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex B.

7.1.4 ONEM2M

oneM2M is a relatively new standard (2015) managed by a partnership of regional international standards industry organizations in the telecommunications industry.

oneM2M provides a common service layer that sits between applications and connectivity transport. It offers functions that IoT applications across different industry segments commonly need. Those functions are exposed to applications via RESTful APIs.

oneM2M standards comprise a horizontal platform architecture that fits within a three-layer model comprising of applications, middleware services and networks. oneM2M's connectivity standards permit applications that are hosted on connected machines and devices, enterprise systems and mobile devices to communicate with each other in an efficient, secure manner. The oneM2M horizontal platform is scalable as the common service elements are able to be deployed on hosts, at the proximal network edge or within the enterprise cloud.

Connectivity services provide capabilities that allow for efficient communication between application endpoints. It provides interworking mechanisms that adjust the underlying network (e.g., mobile, wireless) quality of service to meet the needs of current application data exchange. Currently, the oneM2M service layer can use HTTP, CoAP, MQTT and WebSockets for connectivity transports. DDS is being explored as another option for providing real-time connectivity between the oneM2M service layer entities. Also, an interworking gateway between oneM2M and DDS and OPC-UA is under investigation.

Typical usage of oneM2M includes registration and subscription of devices and applications, service charging and accounting, management of application and devices, and monitoring.

oneM2M has commercial deployments in home-automation applications. It is suitable for large-scale consumer IoT applications. oneM2M is also actively engaged in deployments in other IIoT application domains, including telematics and intelligent transportation, home automation, utilities, healthcare, smart cities, and industrial automation. In all these domains, oneM2M provides semantic enablers (in the scope of the Distributed Data Interoperability and Management layer). As a key part of the telecommunication industry's existing initiatives and its new "5G" initiative, oneM2M provides enablers that focus on the connection between device and the cellular network.

To date, the real-time data distribution performance, latency, jitter benchmarks of oneM2M deployments have not been documented and made publicly available.

For more details, and to determine the suitability of oneM2M for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex C.

7.2 IIoT CONNECTIVITY TRANSPORT STANDARDS

7.2.1 TCP AND UDP OVER IP

In the context of IIoT, the network layer is the internet protocol (IP). The IP suite also provides the UDP¹ and the TCP² transport on top of the IP layer. These IP transports provide the foundation for the other connectivity transports and frameworks.

UDP, Universal Datagram Protocol is a connectionless transport (see section 5.1.4) that provides best-effort delivery quality of service. A message is not resent if it is lost in the transmission. Messages may be received out-of-order. Messages are sent as quickly as possible, and so it is suitable for low latency real-time communications. A message shall be less than 64KB long. A connectivity transport or framework on top of UDP should therefore deal with fragmentation by caching and assembling portions of larger messages.

TCP, Transmission Control Protocol is a connection-oriented (see section 5.1.4) transport that provides reliable and ordered delivery quality of service. A message is resent if it is lost in transmission. Messages are delivered in order. This can lead to head-of-line blocking—high priority, time-critical messages may be blocked behind low priority, non-critical messages. Retries hold up all messages in the channel. Thus, message latencies can vary greatly, leading to large jitter, especially when messages are lost in transmission. The connection sequence can be expensive in time and resources. There is no inherent limit on the message size.

The choice between UDP and TCP at the connectivity transport level has significant implications for the connectivity framework and its suitability (see Figure 1-1). As shown in Figure 7-1, some connectivity frameworks require TCP and inherit its characteristics; some require UDP and inherit those characteristics; and some can use either to support the varying application requirements.

7.2.2 CONSTRAINED APPLICATION PROTOCOL (CoAP)

Constrained Application Protocol (CoAP) is a connectivity transport standard inspired by HTTP, but designed to be more lightweight and efficient (see Figure 7-1). It was established using the UDP transport. The IETF maintains the CoAP open standard specification. Since its definition alternative transports using TCP with TLS³, SMS, and Web Sockets have been developed.

CoAP is generally used in the operations domain (see Figure 1-1). Like HTTP, a client sends a request to a server, specifying a data object, an operation, and a payload. The server replies with failure or success and a response payload. In addition, a client can also register to be notified of any changes in data object. Unlike HTTP, it is suitable for device-to-device queries. However, retries and reordering are implemented in the application stack. CoAP is designed to interoperate with HTTP and the RESTful web services through simple proxies.

¹ See [IETF-RFC768]

² See [IETF-RFC793]

³ See [IETF-RFC4279]

For more details and to determine the suitability of CoAP for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex E.

7.2.3 MQTT (FORMERLY MQ TELEMETRY TRANSPORT)

MQTT is an open connectivity transport standard, maintained by OASIS. It requires the TCP transport.

MQTT is generally used in the information domain (Figure 1-1). It targets device data collection. As the name indicates, the main purpose is telemetry or remote monitoring. Its goal is to collect data from many devices and transport that data to the IT infrastructure. It targets large networks of small devices that need to be monitored or controlled from the cloud.

MQTT implements a hub-and-spoke architecture. Typically, all the devices connect to a data concentrator server. The protocol generally works on top of TCP, which provides a simple, reliable stream. Since the IT infrastructure uses the data, the entire system is designed to transport data easily into enterprise technologies. MQTT has also been adapted for UDP in a separate protocol called MQTT-SN.

MQTT is suited for many-to-one data collection. It is not commonly used for device-to-device transfer or for one-to-many data distribution. MQTT is a simple protocol with few control options. Most applications don't need to be particularly fast; latency specifications are often measured in seconds.

MQTT targets applications such as monitoring an oil pipeline for leaks or vandalism, that require message feeds from thousands of sensors to be concentrated into a single location for analysis. When the system finds a problem, it can take action to correct that problem. Other applications for MQTT include power usage monitoring, lighting control and even intelligent gardening and agriculture. They share a need for collecting data from many sources and making them available to the IT infrastructure.

For more details and to determine the suitability of MQTT for a specific set of system requirements, please refer to the assessment template (see chapter 6) in Annex F.

7.3 FIELDBUS TECHNOLOGIES

Fieldbus ecosystems are well developed and extensively deployed in many industries. Most originated with special-purpose hardware and protocols. Well-known fieldbuses include Profibus (Profinet), EtherNet/IP, Modbus & Modbus/TCP, HART & HART wireless, and the Foundation Fieldbus family. Each has developed extensive ecosystems of vendors and customers.

The industrial internet will bring benefits of common connectivity standards based on the Internet Protocol (IP). This is a significant transition; today's industrial ecosystems use a wide variety of communication and connectivity standards.

Interoperability between fieldbus variants is, in general, poor. Many of these have been adopting IP-based networking models and Ethernet transports. This is improving technical interoperability. Syntactic or higher levels of interoperability are only available with special point solutions.

Fieldbuses implement parts of the connectivity transport and framework functions. None satisfies all of the connectivity core standards criteria (see section 3.4). However, they are well deployed with extensive experience. Most of these protocols support the following types of communication:

- *Management*: Request-Reply pattern, with explicit schemas, for example a RESTful architectural style, used for resource lifecycle communication for management and status of the device/thing
- *Operational*: Peer-to-peer publish-subscribe pattern, with implicit schemas, often time-series, used for data streaming of key operational data used for sense-control-actuate control loop processing that maintains the system's operational integrity.

This suggests that the core connectivity standards should support both forms of data exchange patterns to support the range of functions in the overlying applications.

Most automation and control applications in operations and being deployed for the foreseeable future will rely upon a Fieldbus-based ecosystem protocol. Most "Industrial Internet" applications will have to access data from these protocols to acquire data from and actuate decisions or insights into these protocols. Integration is important.

8 CORE CONNECTIVITY STANDARDS

The assessment of the connectivity standards listed in chapter 7 confirms that not all of the connectivity standards need to support the applications across IIoT to the same degree. Some are more suited to one functional domain (see Figure 1-1), while others are applicable across multiple functional domains. Some are vertically focused, specific to certain industries, while others are horizontally focused, and used in multiple industries.

The result of applying the connectivity core standards criteria (see section 3.4) to the IIoT connectivity framework standards identified in chapter 7 is summarized in Table 8-1.

	Core Standard Criterion	DDS	Web Services	OPC-UA	oneM2M
1	Provide syntactic interoperability #	✓	Need XML or JSON	✓	✓
2	Open standard with strong independent, international governance #	✓	✓	✓	✓
3	Horizontal and neutral in its applicability across industries#	✓	✓	✓	✓
4	Stable and deployed across multiple vertical industries#	Software Integration & Autonomy	✓	Manufacturing	Home Automation
5	Have standards-defined Core Gateways to all other core connectivity standards#	Web Services, OPC-UA*, oneM2M*	DDS, OPC-UA, oneM2M	Web Services, DDS*, oneM2M*	Web Services, OPC-UA*, DDS*
6	Meet the connectivity framework functional requirements	✓	✗	Pub-Sub in development	✓
7	Meet non-functional requirements of performance, scalability, reliability, resilience	✓	✗	Real-time in development	Reports not yet documented or public
8	Meet security and safety requirements	✓	✓	✓	✓
9	Not require any single component from any single vendor	✓	✓	✓	✓
10	Have readily-available SDKs both commercial and open source	✓	✓	✓	✓

green = Gating Criteria

* = work in progress, ✓ = supported, ✗ = not supported

Table 8-1: IIoT Connectivity Core Standards Criteria applied to key connectivity framework standards.

DDS is managed by the OMG, has an established standard gateway mapping to web services (and many others), was established in 2004 and is widely deployed in many types of systems in multiple industries. Standardized gateway mappings to OPC-UA and oneM2M are in development. DDS is suitable for multiple functional domains (see Figure 1-1). It is widely used in the control domain, and increasingly being used in the information, operation and business domains. Security is a recent addition to the specification, now nearing deployment. For more details on the applicability of DDS see the detailed assessment template mapping in Annex A.

Web Services using HTTP are important in any IIoT system, especially from the IT perspective. HTTP is managed by the IETF, has standard mappings to DDS, OPC-UA and oneM2M, and is well established and widely deployed across multiple industries. Web services are suitable for the application and business domains. Web services are typically used for administrative

communication on many IIoT devices, but those devices generally use other mechanisms for the operational communication. For more details on the applicability of HTTP see the detailed assessment template mapping in Annex D.

OPC-UA is managed by the OPC Foundation, and has a standard mapping to web services. A standard gateway mapping of OPC-UA and DDS is under development at the OMG, a standard mapping of OPC-UA on top of DDS is under way at the OPC Foundation, and a oneM2M mapping is under development at oneM2M. OPC-UA is being deployed across manufacturing. This last requirement is not quite met. The precursor technology (“Classic OPC”) is certainly widespread. Because of that we consider the risk of OPC-UA deployment to be low enough in manufacturing to qualify. Other related industries are beginning to deploy OPC-UA as well, so its applicability will likely expand in the future. OPC-UA is suitable for the operation and application domains. For more details on the applicability of OPC-UA see the detailed assessment template mapping in Annex B.

oneM2M is managed by a partnership of regional international standards organizations to include: ARIB, ATIS, CCSA, ETSI, TIA, TSDSI, TTA, TTC, uses a RESTful style architecture with transport bindings for HTTP and Web Sockets and is scheduled to have established standard mappings to DDS and OPC-UA in the next release, and has been recently deployed across multiple industries. As the standard is quite new (released 2015), deployments are also new and less than what one might require for it to be considered stable and proven across multiple industries. For more details on the applicability of oneM2M see the detailed assessment template mapping in Annex C.

Considering the current state-of-the-art, the criteria for IIoT connectivity core standards, the practicality of maintaining a minimal set of core standards that support many connectivity technologies, and the need to provide effective guidance; the following emerge as potential IIoT connectivity core standards: DDS, OPC-UA, Web-Services, and oneM2M. These standards span aspects of IIoT systems, as shown in Table 8-2. The set of potential connectivity core standards may expand in the future as other applicable standards mature to meet the criteria and address IIoT system aspects. Table 8-2 shows examples of some applications where these standards are often applied.

System Aspect	Example User	Approach	Targeting Standard
Software Integration and Autonomy	You are a software architect. You are building a system or product line, and you control the architecture. You critically need to integrate components written by different programmers or even entire teams.	A data-centric approach will define the interfaces, capture the dataflow, enable module evolution, and enforce interoperation between teams. This approach also eases redundancy, fast complex data flow, and selective data filtering.	DDS
Device Interchangeability	You are a device manufacturer, with the goal of making devices that will sell into many applications. The device offers services, such as configure, start, stop, etc. You have no idea how the device will eventually be used. Your users are likely not software experts; they just want to add or integrate the device into a work cell.	A device-centric approach will allow the device users to write generic software that will interoperate with competitor's devices.	OPC-UA
Web and Mobile User Interfaces	You are building mobile apps or web browser based applications to provide the human machine interface. You need an easy way to support clean human interaction and access to backend services.	A RESTful approach will make it easy to connect to many types of enterprise systems and UI devices.	Web Services
Information & Communications Technology (ICT) Integration	You are building a wide-area wireless system that needs to allow applications and devices to share data and information. The devices use various technology and domain-specific protocols. The applications and devices you integrate rely on leveraging the services provided by the communications provider network.	A common, standard services-layer approach enables applications and device to share data and information without forcing the application to understand multiple protocols implemented on the devices. The applications can thus run in the Platform Tier and seamlessly connect to diverse IoT devices in the field.	oneM2M

Table 8-2: Non-overlapping system aspect examples addressed by the potential IIoT connectivity core standards.

Core gateways enable horizontal data interoperability between components across functional domains, as shown in Figure 3-4. Other connectivity technologies can integrate into the system architecture using a gateway to one of the core connectivity standards. This satisfies the range of IIoT system architecture challenges with minimum complexity.

9 OTHER CONNECTIVITY TECHNOLOGIES

Historically, specialized industrial connectivity technologies (see section 7.3) have evolved to meet the specific needs of a particular application area. The goal of the IIoT connectivity reference architecture (see chapter 3) is to enable endpoints using one connectivity technology to communicate with endpoints using another unrelated connectivity technology, possibly in a different functional domain. Since gateways exist between the core connectivity standards, endpoints from the originally unrelated technologies can now communicate.

A domain-specific connectivity technology needs to provide a gateway to only one of the core standards. However, the choice of the core connectivity standard has a direct impact on the fidelity and the quality of service of the communication, as the core connectivity standards vary widely in their characteristics (see chapter 7). The most suitable core connectivity standard should be selected for the gateway. We recommend filling out the assessment template defined in chapter 6 for the specific technology under consideration, and then picking out the core connectivity standard that is most aligned with the connectivity technology under consideration.

Some guidelines follow, based on the primary functional domain (see Figure 1-1) of applicability for the connectivity technology.

Control domain connectivity technologies will support high reliability, fast real-time performance, scaling to large number of data objects, and rich quality of service.

Operations domain connectivity technologies will support monitoring and management of devices and applications.

Information domain connectivity technologies will support selectively moving large volume and variety of real-time data to feed streaming analytics and real-time decision processes.

Application domain connectivity technologies will support external APIs and User Interfaces (UIs), including web browsers and mobile handhelds.

Business domain connectivity technologies will support traditional IT applications and data centers.

These guidelines are starting points, and do not substitute for filling out the assessment template (see chapter 6) to select the closest core connectivity standard.

With the gateways to the core connectivity standards in place for the connectivity technologies of interest, the IIoT connectivity architecture enables communication between hitherto isolated endpoints. It can open up new value streams and help realize the full potential of IIoT.

Annex A ASSESSMENT TEMPLATE: DDS

This Annex contains the assessment template for Data Distribution Service (DDS).

A.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> DDS (Data Distribution Service)
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> Object Management Group (OMG)
Description	<i>Short synopsis of the technology.</i> DDS is a connectivity framework for Industrial IoT. It enables network interoperability for connected machines, enterprise systems and mobile devices. It provides scalability, performance, and Quality of Service required to support IoT applications. DDS can be deployed in platforms ranging from low-footprint devices to the Cloud and supports efficient bandwidth usage as well as agile orchestration of system components. It provides a global data space for analytics and enables flexible real-time system integration.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> DDS is suitable for both the Industrial Internet of Things (IIoT) and large-scale Consumer IoT applications. DDS specifically targets the IIoT application domains, including transportation, energy, healthcare, industrial automation, simulation & test, smart cities, military and aerospace.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> Connectivity transport options include: <ul style="list-style-type: none"> • UDP/IP (optional: DTLS) • TCP/IP (optional: TLS) in progress
References	<i>Website¹ and other useful links to the technology.</i>


¹ See [OMG-DDS]

A.6.2 Business Viewpoint

A.6.2.1 Purpose (Section 6.2.1)

Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.

The OMG DDS standard helps users reliably and securely harness ever-increasing amounts of device generated data while processing the data in real-time, and acting on events as quickly as they occur. As a result, it enables smarter decisions, new services, additional revenue streams and reduced costs. The OMG DDS middleware standard can also simplify the development, deployment and management of IoT applications, speeding time-to-market.



Seamless data sharing regardless of:

Proximity	Physical network
Platform	Transport protocol
Language	Network topology

According to the OMG DDS portal, based on the use of DDS in thousands of applications, one can predict the need for DDS in new projects. If you answer yes to most of the following questions, DDS is likely to be your go-to solution.

- Are the consequences of short downtime severe? If your system goes offline for 5 minutes (or even 5 milliseconds), is it a serious problem? Since DDS does not require servers that could fail and supports redundancy, it makes "fast" reliability and availability much easier. DDS also eliminates struggles with server configuration, startup order, or failover to backup servers.
- Do you require sub-second response? DDS direct peer-to-peer messaging can deliver in milliseconds or even a few microseconds.
- Do you have more than a few software modules or software teams? The databus abstraction will define the interfaces, capture the dataflow, enable module evolution and enforce interoperation between teams.
- Do you have to supply data to many modules, or have too much data to send it all to one place? DDS selective filtering makes it easy to find and deliver exactly the right data.
- Are you building a new IIoT architecture? DDS is not usually used in "retrofit" applications. Typical systems are software projects take more than a year to write, last more than three years, go through multiple versions.

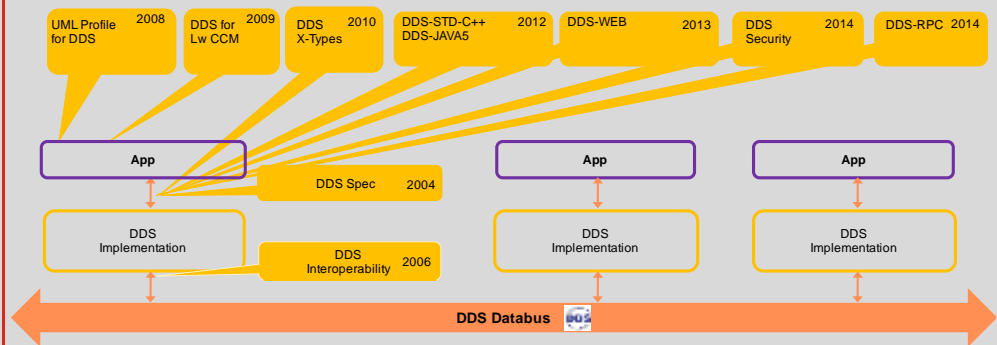
These questions help identify your critical performance, reliability, and integration needs. If you answer yes to any of these questions, you should evaluate DDS as a solution, since it offers many additional benefits.

A.6.2 Business Viewpoint

A.6.2.2 Pedigree
(Section 6.2.2)

Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.

- DDS was designed for real-time, scalable, continuously available, peer-to-peer real-world systems. A few proprietary DDS systems had been available for several years. Starting in 2001, there was a focused industry effort to create an open standard under the auspices of the OMG, resulting in Version 1.0 in 2004. Since then DDS has grown into a family of specifications. The essential specifications include:
- *DDS v1.4* (2015)—defines a data-centric publish-subscribe model for distributed application communication and integration.
- *DDSI-RTPS v2.2* (2014)—defines the Real-time Publish-Subscribe Protocol (RTPS) DDS Interoperability Wire Protocol.
- *DDS-XTypes v1.1* (2014)—defines Extensible and DynamicTopic Data Types for DDS.
- *DDS-RPC v1.0* (2016) – defines a distributed services framework providing language-independent service definition and service/remote procedure invocation using DDS. Supports automatic discovery, synchronous and asynchronous invocations, and QoS.
- *DDS-Security v1.0* (2016)—defines the Security Model and Service Plugin Interface (SPI) architecture for compliant DDS implementations.



DDS enjoys an active and vibrant community continuously working to extend its applicability. The full list of the DDS family of specifications can be found at [website](#)¹.

Multiple independent DDS implementations are available, including both open-source and commercial.

A.6.2.3 Variants
(Section 6.2.3)

Describe the options and variants from the original generic description of the technology.

None. Implementations may differ in their support and coverage of the DDS specifications or compliance profiles.

¹ See [OMG-DDSSTD]

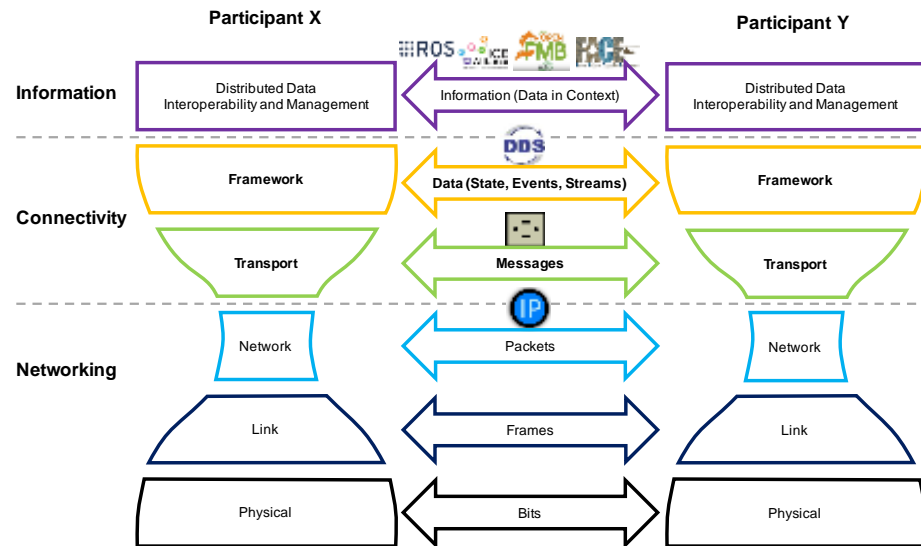
A.6.2 Business Viewpoint	
A.6.2.4 Maturity (Section 6.2.4)	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>The core DDS specifications are mature and have been refined through thousands of applications across multiple industries. DDS vendors collaborate and regularly hold plug-fests to show interoperability between independent DDS implementations.</p> <p>DDS has been applied in multiple verticals to realize higher domain-specific interoperable open architecture specifications. These include:</p> <ul style="list-style-type: none"> • <i>SGIP OpenFMB v1.0</i> (uses CIM extensions over DDS) - NAESB Standard • <i>MDPnP OpenICE</i> Integrated Clinical Environment for Medical Device Interoperability • <i>ROS: Robot Operating System</i> (Open Source) • <i>EUROCAE ED-133</i> flight data exchange between air traffic control centers • <i>Generic Vehicle Architecture (GVA)</i> • <i>Future Airborne Capability Environment (FACE)</i> • <i>Open Mission Systems (OMS)</i> • <i>Open Architecture Radar Interface Standard (OARIS)</i> • <i>Unmanned Aircraft Systems Control Segment (UCS)</i> • <i>Joint Architecture for Unmanned Systems (JAUS) over DDS</i> • <i>Layered Simulation Architecture</i> • <i>Navy Open Architecture</i>
A.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>Yes, DDS is very stable. It has been used for 15 years (counting precursors) across multiple industries. The DDS specifications have also been continuously updated to incorporate the lessons from actual deployments. The core specifications are stable, and easy to use for professionals.</p> <p>The DDS community has continued to innovate actively and expand the breadth and depth of specifications across all aspects of IIoT data connectivity middleware.</p>
A.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>Object Management Group (OMG)</p>
A.6.2.7 Openness (Section 6.2.7)	<p><i>Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?</i></p> <p>Yes, DDS is an open standard. The specifications are openly available to anyone at no cost. Anyone is free to download and implement them. The specifications process is open to participation by both vendors and users.</p> <p>Open source and commercial implementations are available.</p> <p>No, the DDS specifications do not rely on any single component from any single vendor.</p>

A.6.3 Usage Viewpoint

A.6.3.1 Architecture
(Section 6.3.1)

Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.

DDS provides a “middleware” software layer that abstracts an application from the details of the operating system, network transport, and low-level data formats. An application links to a DDS middleware library to participate in an data exchange. The same concepts and APIs are provided in different programming languages allowing applications to exchange data across of operating systems, languages and processor architectures. Low-level details like data wire format, discovery, connections, reliability, timing and QoS management are managed by the middleware layer. It integrates the components of a system together, providing low-latency data connectivity, extreme reliability and a scalable architecture that business and mission-critical Internet of Things (IoT) applications need.



DDS organizes the data exchange between applications as a shared data space. A DDS-DomainParticipant represents an application’s participation in a data space. Within that data space, a collection of data objects with the same structure (data type) and behavior (QoS) is represented by a named DDS-Topic. A DDS-DataWriter is used to publish updates to one or more data objects on a DDS-Topic. A DDS-DataReader is used to subscribe to updates to one or more data objects on a DDS-Topic.

Updates to data objects on a DDS-Topic originate at a DDS-DataWriter, and are sent directly to all the DDS-DataReaders associated with that DDS-Topic. The data paths are direct and peer-to-peer, with no server or broker in the middle. An application can participate in one or more data spaces. A content-filter can filter can be used to specify the subset of updates of interest to a DataReader. Only the relevant updates are delivered, as defined by the content-filter and the QoS configuration specific to the DDS-DataWriter and DDS-DataReader pair.

A.6.3.2 Technology Options
(Section 6.3.2)

List the choices to be made for using the connectivity technology in a system.

Selection of the DDS implementation. Note that the applications participating in a data exchange may use different implementations.

A.6.3 Usage Viewpoint	
A.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>DDS is already proven in mission-critical systems in industries ranging from smart transportation to healthcare to smart energy, and also aerospace & defense. Reasons for using DDS include:</p> <ul style="list-style-type: none"> • <i>Ease of Integration:</i> The data-centric approach used by DDS allows the definition of common and extensible data models for seamless Information Technology (IT) / Operational Technology (OT) interoperability. Its loose and anonymous data-sharing abstraction completely hides connectivity and topology details from applications. • <i>Performance Efficiency and Scalability:</i> DDS implementations can achieve point-to-point latencies that are as low as 30 μsec. and throughput of several million messages per second. It uses a very efficient wire protocol, content- and time-based filtering. When properly architected, DDS-based systems can achieve near-linear scalability. • <i>Advanced Security:</i> The OMG DDS Security Specification defines a comprehensive Security Model and Service Plugin Interface (SPI) architecture for compliant DDS implementations. DDS provides standardized authentication, encryption, access control and logging capabilities to enable secure data connectivity end-to-end in an IoT system. • <i>QoS-Enabled:</i> A rich set of QoS policies allows DDS to control of all aspects of data distribution, such as timeliness, traffic prioritization, reliability and resource usage. • <i>Scalable Discovery:</i> For large-scale dynamic systems, DDS offers automatic discovery that provides plug-and-play functionality to simplify system integration and orchestration. • <i>Applicability:</i> DDS can transparently address peer-to-peer, device-to-device, device-to-cloud and cloud-to-cloud communication. Implementations are available for embedded, mobile, web, enterprise and cloud applications. • <i>Future Proof:</i> The OMG DDS middleware specification enables end-to-end vendor interoperability and eases IoT system development and integration through fully open, future-proof APIs with no vendor lock in.
A.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <ul style="list-style-type: none"> • Data plane (Data Collection; Data Distribution; Data Streaming) • Control plane (Commands & Status; Orchestration) • Monitoring plane (Remote Monitoring & Diagnostics) • Management plane (Events, Analytics & Alarms; User Interfaces)
A.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>Yes, one can monitor and manage a data-space simply by adding an application to participate in data-space. One can also simply replace a participant in the data space by another.</p>

A.6.3 Usage Viewpoint	
A.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>DDS-Security v1.0 specification defines a fine-grained security model at the level of data objects that includes authentication, encryption, access control, data integrity and logging capabilities to enable secure data connectivity end-to-end in an IoT system. This security model applied on top of the network layer, and can therefore support secure multicast, when needed.</p> <p>In addition, DDS the transport layer security mechanisms such as TLS (Transport Layer Security) and DTLS (Datagram Transport Layer Security) can also be used, although they may be unnecessary when the DDS security model is used.</p>
A.6.3.7 Safety (Section 2.3.9)	<p><i>For systems that need it, are certifiable implementations available?</i></p> <p>Yes, certifiable DDS implementations are available, including for example, DO-178C Level A for flight safety critical systems. Certifications for IEC 60601 (class 3 medical devices) and ISO 26262 (automotive road functional safety) are in process. HIPAA-compliant security is available for the medical industry.</p>
A.6.3.8 Gateways (Section 3.3)	<p><i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i></p> <p>Standardized gateways are available to the following connectivity standards:</p> <ul style="list-style-type: none"> • <i>DDS-Web v1.0</i>—access to the data space via RESTful or HTTP technologies and • OPC-UA-DDS—access to the DDS data space via OPC-UA, and access to data objects on the OPC-UA address space. • oneM2M-DDS gateway is under development <p><i>IloT System</i></p> <p>Bespoke gateways to many IloT vertical specific connectivity technologies exist, including DNP3, C37.118, Modbus, HLA, JMS and so on.</p>

A.6.4 Functional Viewpoint	
A.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a flexible resource model. User-defined data types define the structure of data objects. A DDS-Topic is a named collection of data objects that all have the same data type. Topics are user defined, and there can be any number of them.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>Yes, DDS provides a flexible user defined scheme for identifying and addressing data objects. A user can mark certain fields of the data type as “key” fields—those are used to identify and address the data objects with a DDS-Topic. A data object is uniquely addressed by a domain id (which identifies the data space), the topic name within the data space and the key fields within that topic.</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a very flexible data type system, analogous to that of a modern programming languages such as C, C++, Java and .NET. The <i>DDS-XTypes v1.1</i> specification defines extensible and mutable data types that can evolve over a system’s lifecycle. Data types may be defined in a programming language neutral manner, such as Interface Definition Language (IDL) or XML.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object’s lifecycle? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a means to manage the full data object lifecycle, including operations to create, read, update and delete data objects.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a rich set of QoS policies to manage the recent history of data objects. This includes caching user defined history of data objects independently at DDS-DataReaders and DDS-DataWriters, delivering historical data for late joiners, and caching the state in the data space, using the DDS-Persistence Service (an optional part of the DDS specification).</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, the fundamental interaction pattern is publish/subscribe. DDS-DataWriters are used to publish update to data objects; DDS-DataReaders are used to subscribe to data object updates.</p>
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a means to request the state of data object and receive appropriate response(s) via a published design pattern. The <i>DDS-RPC v1.0</i> (2015) specification formalizes the design pattern and defines a distributed services framework providing language-independent service definition and service/remote procedure invocation using DDS. It supports automatic discovery, synchronous and asynchronous invocations and QoS.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a means to discover the DDS-DomainParticipants, DDS-DataWriters, and DDS-DataReaders, and the DDS-Topics in a data space. Discovery is automatic and continuous as applications come and go. An application can access the discovery data simply by subscribing to pre-defined built-in discovery topics.</p>

A.6.4 Functional Viewpoint	
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a rich set of exception handling capabilities. The data delivery contracts are defined via QoS policies, and when any QoS policy is violated, a corresponding flag is raised to signal the exception to the application. Specifically, the LIVELINESS QoS policy is used to monitor the connectivity. Upon loss of connectivity, an exception is signaled by setting the corresponding flag, and also by changing the state of the disconnected data objects.</p>
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>Yes, the DDS specification defines a rich set of 21+ data QoS policies. These include data delivery (best-efforts vs. reliable), timeliness (deadlines), ordering, durability, lifespan, fault tolerance, history, liveliness, ownership, latency, priority and so on.</p>
Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a very fine-grained data object security model, and it is detailed in the <i>DDS-Security v1.0</i> (2016) specification. It defines security policies for authentication, access control (read, write, read-write), confidentiality (encryption), data integrity, data tagging, and logging (when violations occur). The security policies are applied on top of the data object resource model, and can be activated or updated by reconfiguration (i.e. no code changes) at any stage of the development or deployment cycle.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>Yes, the DDS provides a standardized API in multiple programming languages:</p> <ul style="list-style-type: none"> • <i>ISO/IEC C++ 2003 Language PSM for DDS</i>—defines a C++ API only for the DDS specification. • <i>Java 5 Language PSM for DDS</i>—defines a Java API for the DDS specification. • <i>Other language APIs</i> for C, Java, traditional C++, and other languages are derived from the DDS API in IDL using the respective IDL to language mappings.
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>DDS provides standardized APIs to configure the data types, QoS policies, security, resource management, and timing. It also standardizes the APIs to monitor QoS policy violations. Implementations may also support file based mechanisms for configuration and administration. Monitoring of the endpoint internal state is implementation specific, but can be offered and discovered via standardized DDS APIs.</p>

A.6.4 Functional Viewpoint	
A.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>The DDS messaging protocol (DDSI-RTPS v2.2) assumes connectionless (UDP) messaging. The messaging protocol can allow messages larger than the UDP (64kB) limit. The protocol is designed to support widely varying message sizes.</p> <p>DDS implementations can also support TCP. There is an ongoing effort to standardize the mapping under the auspices of the TCP/IP DDSI-RTPS specification.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>DDS supports both unicast (default) and multicast (when available).</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>DDS transport endpoints correspond to a DDS-DataWriter and DDS-DataReader. The endpoints are addressed using a globally unique ID (GUID) for the endpoints. The number of endpoints within a domain is bounded by the number of unique GUIDs.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>No, DDS does not require a connected circuit between the endpoints.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>Yes, DDS provides a means to prioritize messages.</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, DDS does not provide a way to synchronize time between endpoints. In systems using DDS, this is typically accomplished using a separate time synchronization protocol.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>Yes, DDS provides mechanisms for message security. It provides support for authentication of endpoints, message encryption and message integrity.</p>

A.6.5 Implementation Viewpoint	
A.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<p><i>Does the connectivity framework require running a special process or broker?</i></p> <p>This is implementation specific. There are DDS implementations that do not require running a separate process or broker. An application, by linking to a DDS library becomes an active participant in the data exchange. There are no other process dependencies.</p>
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<p><i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i></p> <p>No, the application code (or business logic) does not have to be aware of other endpoints to participate in a data exchange. Applications interact directly with the databus (data objects organized into DDS-Topics) and never directly with each other.</p>
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<p><i>Is the governance explicit and shareable?</i></p> <p>DDS does not require the governance to be implicit, and allows system architects to choose the style of governance. The data types are always explicitly defined, the data flows and the quality of service configuration may be defined implicitly or explicitly, and the data security configuration is always explicitly defined.</p>
A.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<p><i>Can a content-filter specify the data subset of interest?</i></p> <p>Yes, a DDS-ContentFilteredTopic can be used to subscribe to only a subset of data from a DDS-Topic.</p>
Time-Based Selection (Section 4.2.2.2)	<p><i>Can sub-sampling specify the data subset of interest?</i></p> <p>Yes, a TIMEBASEDFILTER QoS policy can be used to subscribe to a subsampled subset of the data.</p>
A.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<p><i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i></p> <p>This is dependent on the underlying hardware transport. Within the transport limits, DDS supports real-time data distribution. It was specifically designed to support the needs of real-time distributed systems and includes several QoS policies real-time data distribution. DDS also can notify applications of delays, allowing the application to adapt to the situation.</p> <p>Several DDS implementations have been documented have very low latency (< 1ms) and very low jitter (μs).</p>
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<p><i>How does the latency and jitter change with throughput? What limits the throughput?</i></p> <p>The variation of latency and jitter with throughput will be implementation specific, based on the design trade-offs made by that implementation. Leading DDS implementations have been documented to have minimal change in jitter as throughput increases. Implementations can achieve throughput as high as 95% of the theoretical network bandwidth.</p>

A.6.5 Implementation Viewpoint	
A.6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<p><i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i></p> <p>Yes, DDS can handle an effectively increasing number of data objects. Every data object is identified by a GUID—the number of unique GUIDs limits the number of data objects in a domain. The port numbers available on a host limits the number of domains.</p> <p>There is no theoretical limit on the data object size. In practice, it will be limited by the amount of memory available on a host.</p>
Apps (Section 4.2.4.2)	<p><i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i></p> <p>Yes, DDS can effectively support interface evolution for an increasing number of distributed application components. Application components are loosely coupled—they interact with the data, not with each other; thus, the interfaces are data-oriented and can evolve independently. The data types in a data-oriented interface can also evolve through extension or mutation—the rules are defined in the <i>DDS-XTypes v1.1</i> specification.</p>
A.6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<p><i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i></p> <p>Yes, DDS can support continuous availability over a defined system relevant time period. DDS accomplishes this by having a continuous ongoing automatic discovery, so that components can be added or removed at any time, and by providing an optional DDS-Persistence Service to cache the data outside of specific application components.</p>
Recovery (Section 4.2.5.2)	<p><i>Can the connectivity framework support recovery when fault conditions occur?</i></p> <p>Yes, DDS can support recovery when fault conditions occur. It accomplishes this by signaling exceptions the application layer, by allowing application to change certain QoS policies, and by providing access to the automatic discovery data.</p>
A.6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<p><i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i></p> <p>Yes, DDS implementations are available for most commonly used operating system and CPUs. DDS implementation can run on devices with limited memory resources (<100kB).</p>
Incremental Upgrades (Section 4.2.6.2)	<p><i>Does the connectivity framework facilitate incremental upgrades?</i></p> <p>Yes, DDS can support incremental upgrades. DDS accomplishes this by means of automatic ongoing discovery when components are added or removed, data-oriented interfaces and support for data type evolution over a system’s lifecycle.</p>

A.6.5 Implementation Viewpoint

A.6.5.7 Network Layer Considerations

Topology <i>(Section 5.2.1.1)</i>	<i>What network topologies are allowed?</i> DDS is agnostic to network topologies. All possible network topologies can be used with DDS.
Span <i>(Section 5.2.1.2)</i>	<i>What is the span of the transport: LAN vs. WAN?</i> A DDS data space is agnostic to the network constraints, and can therefore span both the LAN and the WAN. DDS implementations allow application components to be located anywhere—on the LAN or across the WAN. DDS implementations provide mechanisms to deal with firewalls and other restrictions encountered when going across the WAN.
Segmentation <i>(Section 5.2.1.3)</i>	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Yes, DDS can support multiple independent and isolated communication paths between the same network endpoints. DDS accomplishes this by providing a PARTITION QoS policy, which allows communication between DDS endpoints tagged with the same partition label.

Annex B ASSESSMENT TEMPLATE: OPC-UA

This Annex contains the assessment template for Open Platform Communications Unified Architecture (OPC-UA).

B.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> OPC-UA
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> OPC Foundation
Description	<i>Short synopsis of the technology.</i> OPC-UA is an industrial communication architecture for platform independent, high performance, secure, reliable, and semantic interoperability between sensors, field devices, controllers, and applications at the shop-floor level in real-time as well as between the shop floor and the enterprise IT cloud.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> Automation for manufacturing, process control, energy.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> Current technology mapping options include: <ul style="list-style-type: none"> • TCP/IP for the transport/network layers • HTTP/TCP/IP for transport/network layers • TLS for security
References	<i>Website¹ and other useful links to the technology.</i>

¹ See [OPC-UA]

B.6.2 Business Viewpoint	
B.6.2.1 Purpose (Section 6.2.1)	<p><i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i></p> <p>Defines a comprehensive information modeling mechanism, which is fully extensible by specific vertical markets. Defines a standard set of services, which act on the information model.</p> <p>Expose information about the system, its configuration, topology and data context (the meta data) in the collective "address space" of the individual OPC-UA servers. Allow this address space to be accessed by authorized OPC-UA clients such that they can see what is available and choose what to access.</p>
B.6.2.2 Pedigree (Section 6.2.2)	<p><i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i></p> <p>OPC-UA is the next generation of the OPC protocol, which is widely deployed in industrial automation.</p> <p>OPC was introduced in 1996 based on Microsoft DCOM. This specification is now referred as "Classic OPC".</p> <p>OPC-UA was first introduced in 2006 (version 1.0). It no longer depends on DCOM and uses Web-Services and Binary TCP protocols instead.</p> <p>OPC-UA version 1.03 was released in 2013. It has been endorsed as a key specification for Industry 4.0.</p>
B.6.2.3 Variants (Section 6.2.3)	<p><i>Describe the options and variants from the original generic description of the technology.</i></p> <p>The OPC-UA specification defines many optional profiles and services, notably: Discovery, View, Query, Attribute, Method, Data Monitoring, Data Access and Events & Conditions.</p>
B.6.2.4 Maturity (Section 6.2.4)	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>A website¹ maintains a list of notable project that use OPC-UA.</p> <p>OPC-UA has broad industrial support. Its focus is to allow information to be easily and securely exchanged between diverse platforms from multiple vendors and to allow seamless integration of those platforms without costly, time-consuming software development.</p> <p>There are SDKs from multiple companies that can be used to build OPC-UA compliant systems.</p>
B.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>"Classic OPC" has been widely deployed in the industry. The OPC-UA specification has been stable for many years as has the standard stack implementations and SDKs.</p>
B.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>OPC-UA enjoys broad industry support. The OPC foundation has over 450 members².</p>

B.6.2 Business Viewpoint**B.6.2.7 Openness**
(Section 6.2.7)

Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?

OPC-UA is specified as IEC 62541 standard and therefore allows for individual, royalty-free, implementation according to the standard. Certification and technology contribution No, OPC-UA is not an open standard. Access to specifications and developer resources are available to OPC Foundation members only. Membership requires payment of annual dues at selected membership level. The specifications process is open to participation by both vendors and users.

Open source and commercial implementations are available.

No, the OPC-UA specifications do not rely on any single component from any single vendor.

¹ See [OPC-CS]

² See [OPC-MEM]

B.6.3 Usage Viewpoint	
B.6.3.1 Architecture (Section 6.3.1)	<p><i>Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.</i></p> <p>OPC-UA consists of multiple OPC-UA-Clients connected to a OPC-UA-Server. A OPC-UA-Server holds an address space, which is a collection of data objects organized in a linked graph.</p> <p>Requests originate at a OPC-UA-Client and are sent to an OPC-Server; the OPC-Server processes the request, and sends a reply back to the OPC-UA-Client. Requests are addressed to a specific data object in the server's address space. Structured data is used for the request and reply.</p> <p>A OPC-UA specification for publish and subscribe architectures is currently under development.</p>
B.6.3.2 Technology Options (Section 6.3.2)	<p><i>List the choices to be made for using the connectivity technology in a system.</i></p> <ul style="list-style-type: none"> • Selection of SDK used to implement OPC-UA clients and servers supporting the desired variants (OPC-UA profiles). • Selection of the underlying transport: OPC-UA Binary/TCP or XML/HTTP.
B.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>Industrial automation and process control applications. Client-server interactions between components such as devices or applications. Expose the address space of systems and devices to facilitate configuration, browsing and data access.</p>
B.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <p>OPC-UA is deployed on devices to allow device configuration and data-access.</p> <p>For existing brown field installations, OPC-UA is typically deployed at system boundaries to expose the system address space, support browsing, configuration, monitoring and service invocation. Newer devices and systems are building in OPC-UA.</p>
B.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>OPC-UA discovery services are defined to allow dynamic discovery of components.</p>
B.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>Security is provided at the transport level between each client and server. Clients are authenticated via name and password, PKI certificate, of WS-Security Tokens.</p> <p>Each server enforces access control. Servers may support fine-grained access control to individual variable and operations.</p>
B.6.3.7 Safety (Section 2.3.9)	<p><i>For systems that need it, are certifiable implementations available?</i></p> <p>There are currently no safety-certified OPC-UA implementations.</p>
B.6.3.8 Gateways (Section 3.3)	<p><i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i></p> <ul style="list-style-type: none"> • An OPCUA-DDS gateway standard is under development by the OMG. • An OPC-UA DDS profile is under development by the OPC Foundation. • An OPC-UA gateway standard is under development by oneM2M • OPC-UA clients can connect to OPC-UA servers via HTTP. <p>There are commercially-available gateways between OPC-UA and many industrial protocols such as Modbus, Profibus, Foundation fieldbus, etc.</p>

B.6.4 Functional Viewpoint	
B.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>OPC-UA resources are called <i>nodes</i>. They can be individually addressed using a NodeId. Nodes contain data-elements, operations and references to other nodes.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>OPC-UA nodes have a unique identifier within a server, called the NodeId.</p> <p>Addressing a node also requires addressing the OPC-UA server using its network IP address and port, plus the NodeId.</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>OPC-UA defines a full data type system. Data-variables within nodes can be of simple or complex (structured) data types.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object's lifecycle? Summarize the salient aspects.</i></p> <p>There is no pre-defined resource lifecycle in OPC-UA. However, applications can define their own lifecycles and operations to control the resources.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>The variables within each node constitute the state of the Node.</p> <p>Each OPC-UA server manages its own state that is accessible via Query and Browsing services. They cache the last value of every variable they contain. Those can be queried using the Query service.</p> <p>There is also a full set of historical data access services defined.</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>Plans are underway to include publish-subscribe data exchange. Prototypes are already in field-test.</p>
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, this is the core primary communication pattern in OPC-UA.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>OPC-UA servers can implement a discovery service allowing client applications to discover the nodes they contain.</p> <p>An OPC-UA server may provide a global registration and discovery service, allowing discovery of all the OPC-UA servers in a system.</p>
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>Exceptions are supported and communicated via events and alarms.</p>
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>OPC-UA offers limited quality of service options beyond the ability to specify an update frequency for the monitored data. There are many QoS features built into the services of OPC-UA. The services are purpose built to provide appropriate QoS.</p>

B.6.4 Functional Viewpoint

Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>OPC-UA authenticates clients using either a username/password, or a PKI X509 certificate, or a WS-SecurityToken.</p> <p>Servers may support fine-grained access control to individual variable and operations.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>OPC-UA is a reference architecture specification. A typical user is expected to use software tools to integrate devices adhering to the OPC-UA specification. There is no expectation of software development, and therefore no need for standardized APIs.</p>
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>OPC-UA provides standardized means for configuring and administering the data types, information models, and security. Monitoring of OPC-UA servers is implementation specific, but can be offered and discovered via standardized OPC-UA mechanisms.</p>

B.6.4 Functional Viewpoint	
B.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>OPC-UA specifies the use of two alternative protocols:</p> <ul style="list-style-type: none"> • OPC-UA Binary: A binary protocol on top of TCP • An XML-based protocol on top of HTTP (which runs on top of TCP) • Web Sockets (in progress) <p>There are no explicit message-size limits in OPC-UA when using the TCP protocol.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>Current OPC-UA transport mappings rely on unicast over TCP.</p> <p>Future versions of OPC-UA plan to include mappings for UDP and multicast UDP.</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>Endpoints are the OPC-Server and the OPC-Client. OPC-Clients initiate requests to OPC-Servers.</p> <p>The specification relies on the endpoint-addressing scheme provided by the underlying transport mapping. For current transport mappings, an IP address and a port number identify an endpoint.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>The OPC-UA specification does not require a connection-oriented transport. Secure sessions are established above the transport layer.</p> <p>Current technology mappings rely on a connection-oriented transport (TCP).</p> <p>Future version of OPC-UA will also support connectionless UDP transport.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>No, prioritization is not supported in the specification. However, a OPC-UA Server implementation could easily prioritize the processing of requests. For example, a OPC-UA Server may give certain clients (possibly based on credentials) priority over others. Or, for example, a OPC-UA-Server may process subscription advise requests at a higher priority than browse requests.</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, timing and synchronization services are not currently provided. However, work is underway to provide timing and synchronization services through Time Sensitive Networking (TSN) support.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>OPC-UA was architected from the beginning with security as a top priority requirement. Current mappings use TLS (or HTTPS) for message security.</p>

B.6.5 Implementation Viewpoint	
B.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<p><i>Does the connectivity framework require running a special process or broker?</i></p> <p>The OPC-UA specifications are designed to support aggregated servers, which are a type of broker. With an aggregating server, a client connects to the aggregating server and that server acts as a proxy to one or more servers. The information models and address spaces are aggregated.</p>
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<p><i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i></p> <p>Data is accessible as variables, abstracted away from the physical endpoints. Clients can use OPC-UA discovery to locate which server or servers provides a variable of interest. Aggregated OPC-UA Servers are used to provide applications with a single unified address space and abstract away the physical server providing the variable. Thus, OPC-UA is node-centric which maps closely to a device-centric approach.</p>
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<p><i>Is the governance explicit and shareable?</i></p> <p>The OPC-UA service definition and the information model provide governance. Thus, governance is explicit and shareable.</p>
B.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<p><i>Can a content-filter specify the data subset of interest?</i></p> <p>Yes, the data subset of interest can be specified by content. For example, event data is subscribed to using a filter (essentially a stream filter) that compares the content of each event with a set of criteria provided by the client and only sends the subset that matches. Data Subscriptions are also based on filtering.</p>
Time-Based Selection (Section 4.2.2.2)	<p><i>Can sub-sampling specify the data subset of interest?</i></p> <p>Yes, each client can request its own client specific sampling rate.</p>
B.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<p><i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i></p> <p>Current transport mappings rely on TCP, which is known to have non-deterministic latency. Work is underway to add support for Time Sensitive Networking (TSN), and that is expected to provide deterministic latency and jitter for real-time applications.</p>
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<p><i>How does the latency and jitter change with throughput? What limits the throughput?</i></p> <p>Expected to be similar to that of TCP (see section 7.2.1).</p> <p>Use of binary protocols and direct client-server connections expected to result in throughput limited only by the network bandwidth and CPU of client and server computers.</p>

B.6.5 Implementation Viewpoint	
B.6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<p><i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i></p> <p>Yes, OPC-UA can effectively handle an increasing number of data objects. In OPC-UA, the number of data objects on a server would be limited by the memory on the server and the average size of the data objects in that's server. There is no upper limit on the size of a request or reply.</p> <p>Scalability for a server (with respect to the number of clients) is limited by the number of TCP connections it can sustain as well as the number of independent monitor streams it can produce.</p>
Apps (Section 4.2.4.2)	<p><i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i></p> <p>OPC-UA protocol can easily be extended with new services and data types in a backward compatible manner. OPC-UA is a carefully thought out set of services that address the needs of device integration. Unlike approaches (such as SOAP, REST) where each application defines a new service API for specific purposes, resulting in an explosion of services, the advantage of OPC-UA is that once an application supports the service set it needs, it can interface with any device or system.</p>
B.6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<p><i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i></p> <p>OPC-UA defines several redundancy features that allow seamless client and server failovers. For example, subscriptions defined with a server can be transferred to a redundant server without the need of the application to recreate a new subscription.</p>
Recovery (Section 4.2.5.2)	<p><i>Can the connectivity framework support recovery when fault conditions occur?</i></p> <p>The ability to support redundancy and then monitor the current operating state are defined in the standard, and applications can discover these abilities.</p>
B.6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<p><i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i></p> <p>OPC-UA implementations are available for a variety of operating systems, CPUs and resource constraints.</p>
Incremental Upgrades (Section 4.2.6.2)	<p><i>Does the connectivity framework facilitate incremental upgrades?</i></p> <p>Yes, OPC-UA supports incremental upgrades by allowing OPC-UA Servers to be updated independently of the Clients when a device is upgraded. A server can be upgraded while a redundant server provides continuous support to its clients with no downtime.</p>

B.6.5 Implementation Viewpoint**B.6.5.7 Network Layer Considerations**

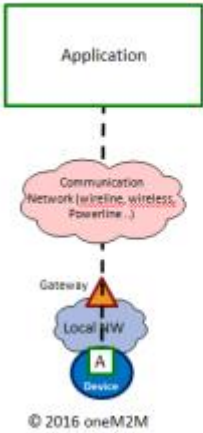
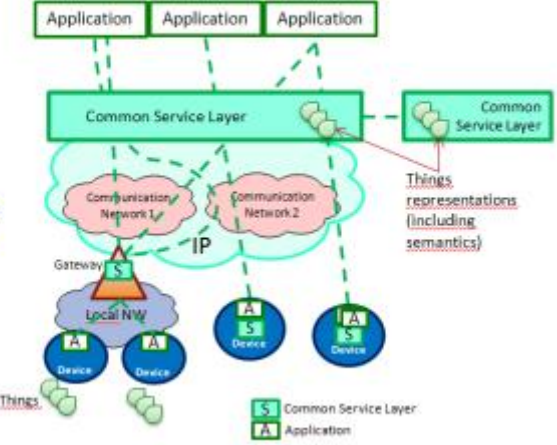
Topology (Section 5.2.1.1)	<i>What network topologies are allowed?</i> Any OPC-UA client can connect directly to any OPC-UA server. OPC-UA is agnostic to network topologies.
Span (Section 5.2.1.2)	<i>What is the span of the transport: LAN vs. WAN?</i> OPC-UA can span across both the LAN and the WAN, as long as the servers are accessible via a TCP/IP network.
Segmentation (Section 5.2.1.3)	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Strictly speaking, there is no segmentation of the information, but servers can be browsed individually and configured to expose different parts of the system. Servers are free to expose multiple endpoints on the same network segment or different segments. Each endpoint can offer different security requirements and expose different address spaces. In this way, there are not limitations.

Annex C ASSESSMENT TEMPLATE: ONEM2M

This Annex contains the assessment template for oneM2M.

C.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> oneM2M
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> oneM2M is a partnership project that includes partners from major regional SDOs and other fora (i.e., ARIB, ATIS, CCSA, ETSI, TIA, TSDSI, TTA, TTC, Broadband Forum, CEN, CENELEC, Global Platform, New Generation M2M Consortium, Open Mobile Alliance).
Description	<i>Short synopsis of the technology.</i> oneM2M provides a common service layer that sits between applications and connectivity transport. It offers functions that IoT applications across different industry segments commonly need. Those functions are exposed to applications via RESTful APIs. oneM2M standards comprise a horizontal platform architecture that fits within a three-layer model comprising of applications, middleware services and networks. oneM2M's connectivity standards permit applications that are hosted on connected machines and devices, enterprise systems and mobile devices to communicate with each other in an efficient, secure manner. The oneM2M horizontal platform is scalable as the Common Service Elements are able to be deployed on hosts, at the proximal network edge or within the enterprise cloud. Connectivity services provide capabilities that allow for efficient communication between application endpoints. It provides native QoS as well as interworking mechanisms that adjust the QoS of the underlying network (e.g., mobile, wireless) to meet the needs of current application data exchange.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> oneM2M service layer is suitable for both the Industrial Internet of Things (IIoT) and large-scale Consumer IoT applications. oneM2M specifically targets the IIoT application domains, including telematics and intelligent transportation, home automation, utilities, healthcare, smart cities, industrial automation.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> oneM2M service layer supports direct bindings to the following network layer/connectivity protocols: <ul style="list-style-type: none"> • CoAP • HTTP • MQTT • WebSockets
References	<i>Website¹ and other useful links to the technology.</i>

¹ See [ONEM2M]

C.6.2 Business Viewpoint	
<p>C.6.2.1 Purpose (Section 6.2.1)</p>	<p><i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i></p> <p>oneM2M standards that constitute the horizontal IoT platform allow applications from various previously siloed domains and for applications within a domain to communicate effectively, reliably and securely.</p> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>Pipe (vertical): 1 Application, 1 NW, 1 (or few) type of Device Point to point communications</p>  <p>© 2016 oneM2M</p> </div> <div style="text-align: center;"> <p>Horizontal (based on common Layer) Applications share common service and network infrastructure Multipoint communications</p>  <p>© 2016 oneM2M</p> </div> </div> <p>The oneM2M interoperable, platform architecture offers significant strategic benefits by consolidating the resources needed to deliver a variety of IoT applications and opening up new service and business opportunities by allowing applications to share resources and data.</p>
<p>C.6.2.2 Pedigree (Section 6.2.2)</p>	<p><i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i></p> <p>oneM2M service layer was developed specifically to address solutions in the M2M/IoT community. The core of the specification was initially developed by ETSI as ETSI M2M. The initial specification for oneM2M was published in January 2015 as release 1.0. Release 2.0 of the specification set is published in August 2016.</p> <p>oneM2M enjoys an active and vibrant community (over 200 member companies) continuously working to extend its applicability. The full list of the oneM2M family of specifications can be found at a website¹.</p> <p>Multiple independent oneM2M implementations are available, including both open-source and commercial.</p>
<p>C.6.2.3 Variants (Section 6.2.3)</p>	<p><i>Describe the options and variants from the original generic description of the technology.</i></p> <p>None. Implementations may differ in their support and coverage of the oneM2M specifications or compliance profiles.</p>
<p>C.6.2.4 Maturity (Section 6.2.4)</p>	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>oneM2M specifications have only been published since January 2015. Since then there have been multiple interoperability events and commercial implementations. Certification of oneM2M-compliant nodes is in the process in certain regions.</p>

C.6.2 Business Viewpoint	
C.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>Since publication in January, 2015 oneM2M has incorporated fixes and clarifications that came from various interoperability and commercial deployments. The rate of fix and clarification requests has dropped off with the majority of the requests being clarification of the standard. As part of its ecosystem development, oneM2M has produced guides to assist application developers use the system. New features are consistently being released with the focus of the work items moving from the core connectivity and service functions to the specifications for new types of gateways needed for interoperability and development of domain-specific resource models.</p>
C.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>oneM2M partner organizations and the ITU publish the specifications for their specific needs. oneM2M collaborates on features with a number external organizations beyond the partnership organizations.</p>
C.6.2.7 Openness (Section 6.2.7)	<p><i>Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?</i></p> <p>Yes, oneM2M is an open standard. The specifications are openly available to anyone at no cost. Anyone is free to download and implement them. The specifications process is open to participation by both member companies of the partner type 1 organizations as well as the partner type 2 organizations themselves.</p> <p>Open source and commercial implementations are available.</p> <p>No, the oneM2M specifications do not rely on any single component from any single vendor.</p>

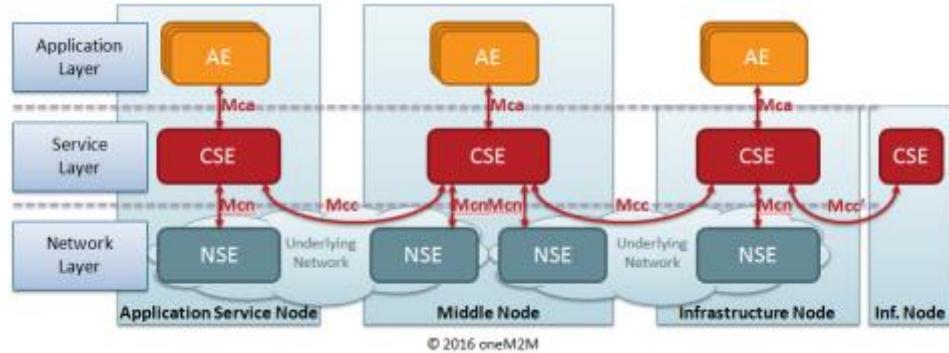
¹ See [ONEM2M-PS]

C.6.3 Usage Viewpoint

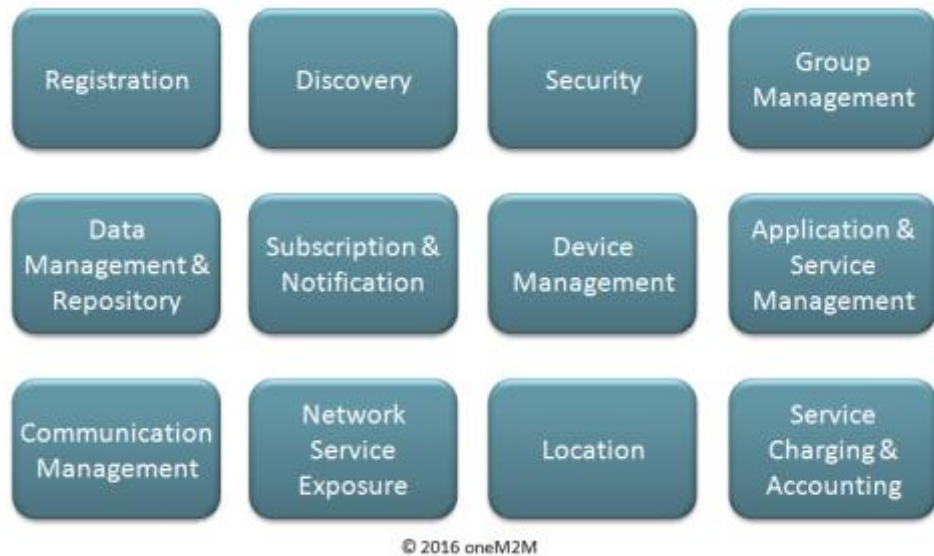
C.6.3.1 Architecture
(Section 6.3.1)

Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.

oneM2M standards comprise a horizontal platform architecture that fits within a three-layer model comprising of applications, middleware services and networks.



Application Entities (AEs) are hosted on nodes (e.g., enterprise server, device). These nodes may be virtualized or physical instances. These AEs communicate with each other by sending requests to a Common Service Entity (CSE) that, in turn routes the request to the target AE while providing services based on the request.



CSEs are hosted on the nodes that may be virtualized or physical nodes. In many instances, the AE and CSE share the same node (e.g., device).

C.6.3.2 Technology Options
(Section 6.3.2)

List the choices to be made for using the connectivity technology in a system.

- Selection of the oneM2M Deployment: Applications interact with CSEs where the network of CSEs makes up the deployment architecture. The flexibility of placing the CSE on the end device, on the edge of a local network; in the cloud or somewhere in between (e.g., Fog) are all deployment options.
- Selection of network layer binding(s) for a CSE.
- Selection of interoperability gateways for a CSE.

C.6.3 Usage Viewpoint	
C.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>oneM2M has commercial deployments home automation applications and is being actively developed for applications in the automotive and smart-city domains. Reasons for using oneM2M include:</p> <ul style="list-style-type: none"> • <i>Ease of Integration:</i> The RESTful architectural approach used by oneM2M allows the definition of common and extensible data models for seamless Information Technology (IT)/Operational Technology (OT) interoperability. • <i>Performance efficiency and scalability:</i> oneM2M implementations allow for deployment configurations that place processing at locations where it can be used in the most efficient manner. This allows for localization of messaging traffic to affected area networks. When properly architected, oneM2M-based systems can achieve near-linear scalability. • <i>Advanced security:</i> oneM2M defines a comprehensive security model for segment and end-to-end authentication, encryption, access control and logging capabilities to enable secure data connectivity end-to-end in an IoT system. • <i>QoS-enabled:</i> Determination of how a CSE treats, in terms of message delivery, is configurable based on a rich set of delivery policies including policies that allow the underlying network to be tuned to the message delivery characteristics of communicating applications. • <i>Scalable discovery:</i> For large-scale dynamic systems, oneM2M is in the process of developing automated onboarding and discovery of applications. • <i>Applicability:</i> oneM2M can transparently address applications that require data to be exchange between applications that transit through a CSE. Implementations are available for embedded, mobile, web, enterprise and cloud applications. • <i>Future proof:</i> The oneM2M specification enables end-to-end vendor interoperability and eases IoT system development and integration through fully open, future-proof APIs with no vendor lock in.
C.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <ul style="list-style-type: none"> • Registration and Service Subscription (Device and application onboarding and discovery): Used in infrastructure CSE. • Discovery of resources: Used in all CSEs. • Service charging and accounting: Used in the infrastructure CSE. • Data plane (data collection; subscription and notifications, data delivery, group management): Used in all CSEs. • Management plane (administration of applications and CSEs, device management): Used in Infrastructure CSEs. • Integration with the underlying network layer services: Used in the Infrastructure CSEs.
C.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>oneM2M service layer provides the capability to monitor and manage applications and CSEs. These components are the building blocks of any oneM2M deployment.</p>
C.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>oneM2M service layer defines a security model to authenticate applications and CSEs. All communication can be securely encrypted as a segment or end-to-end using the underlying network layer security mechanisms (e.g., TLS, DTLS). The security model is applied on top of the network layer.</p>

C.6.3 Usage Viewpoint**C.6.3.7 Safety**
(Section 2.3.9)

For systems that need it, are certifiable implementations available?

While certain regions (e.g., Korea) have started certification processes for applications in that region for functional aspects of the oneM2M service layer and resource interaction, the certification process doesn't provide explicit references to which elements are directly related to the safety of a system.

C.6.3.8 Gateways
(Section 3.3)

List of gateways to core connectivity standards and other relevant connectivity technologies.

oneM2M service layer supports interworking gateways with the following connectivity technologies:

- OSGi (in progress)
- Alljoyn
- OIC (Open Interoperability Consortium)
- LWM2M (Open Mobile Alliance)

DDS is expected in the next release once a determination is made if the support will be for a direct binding or if the support will be through an interworking gateway.

OPC-UA interworking is expected in the next release.

C.6.4 Functional Viewpoint	
C.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer provides a data resource model that contains the user defined structured data-objects.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>Yes, oneM2M service layer data-objects resources can be identified by the system or by the user. The data-object resources can be semantically annotated by the user in order to be discoverable by users using various tools (e.g., data queries, ontological queries).</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer defines a data type system where the resources defined within the oneM2M system are encoded using XML or JSON encoding formats. User data-object resource can either retain their original encoding structures or can be interworked into abstract data-object resources defined within oneM2M.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object's lifecycle? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer is a RESTful resource based system and provides a means to manage the full data-object lifecycle, including operations to create, read, update and delete data-objects.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer data-object resources provides mechanisms to manage the versions and histories of a data-object including capabilities (e.g., number of versions, expiration dates, size constraints) for retaining the versions of the data-objects. For resources that are communicated via the connectivity layer, oneM2M provides a rich set of policies used to determine when to transmit the data-object resources.</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer defines a mechanism where applications can subscribe to the events (e.g., creation, deletion, modification) of data-object resources. As part of the notification procedure, the application can receive the modified data-object resource or can be simply notified of the event.</p>
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer is a RESTful architecture, which is fundamentally a request-reply architecture. oneM2M provides a rich set of capabilities for communicating the request and receiving the reply.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer resources are discoverable using query mechanisms where the elements of the query criteria are defined by the owning applications or their delegates. In addition, oneM2M resources can be semantically annotated for use in ontological queries.</p>
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer provides a rich set of communication exception handling policies that all includes policies of what to do if a communication fails or the recipient is not available.</p>

C.6.4 Functional Viewpoint	
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>Yes, oneM2M service layer provides a set of communication QoS policies to determine the priority of delivering request (e.g., recipient, time-of-day, capacity limits).</p>
Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer provides a data-object security model. It defines security policies for access control (read, write, create, delete, notify), confidentiality (encryption). The owners or their delegated representatives administer security. All communication between end-point can be authenticated and communication encrypted either end-to-end or by communication segment.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>Yes, oneM2M service layer uses RESTful architecture patterns (CRUD), extending the traditional RESTful architectural patterns for subscriptions, notifications and the ability to execute operations.</p>
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>oneM2M service layer provides standardized management APIs for configuring, administering and monitoring CSEs and applications.</p>

C.6.4 Functional Viewpoint	
C.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>The oneM2M service layer runs on top of multiple transport protocols that are based on IP (i.e., CoAP, HTTP, Web Sockets, MQTT). oneM2M messages do not have a defined size limitation but are either limited by the underlying transport protocol's limitation (e.g., HTTP, Content-Length, MQTT 256Meg) or uses the underlying protocols (e.g., CoAP block) mechanisms for message fragmentation and reassembly.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>oneM2M service layer supports both unicast (default) and multicast (when available by the underlying transport) for group-based operations.</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>oneM2M service layer is a RESTful architecture where all resources are addressed using URIs, this includes the endpoints represented by applications and CSEs. The URIs can be defined within the address space of the M2M service provider or they can be globally unique id (GUID). Endpoints follow the addressing scheme defined in the oneM2M specification.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>No, oneM2M service layer is a RESTful client/server architecture and does not require a connected circuit between the application endpoints. oneM2M does require endpoints to register with CSE's to which it connects, but the registration is not reliant on a connected connection between the endpoints.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer provides communication QoS policies to determine the priority of delivering request (e.g., recipient, time-of-day, capacity limits).</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, oneM2M service layer does not provide a way to synchronize time between endpoints. In oneM2M systems this is typically accomplished using a separate time synchronization protocol.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>Yes, oneM2M service layer provides mechanisms for message security. It provides support for authentication of endpoints either end-to-end or endpoint-to-CSE, message encryption (both end-to-end or endpoint-to-CSE), and message integrity.</p>

C.6.5 Implementation Viewpoint	
C.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<i>Does the connectivity framework require running a special process or broker?</i> Yes—oneM2M service layer requires applications to connect to a CSE to communicate with other applications. The target applications are not required to be connected to the same CSE as the source applications. CSEs are connected in a hierarchical tree topology with a root CSE (i.e., IN-CSE) that is in the domain of the M2M service provider.
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i> No, the application does not have to be aware of other endpoints to participate in an information exchange. Applications interact directly with the data-object resources organized by the owning application and never directly with each other.
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<i>Is the governance explicit and shareable?</i> With the exception that oneM2M has defined the domain of data types and resources used by applications and CSE, oneM2M service layer does not require the governance to be implicit, and allows system architects to choose the style of governance including how data-object resources are structured and identified.
C.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<i>Can a content-filter specify the data subset of interest?</i> Yes, oneM2M service layer provides a number of content-filters that can be used for discovery of resources (e.g., time, size, content type, user defined, semantic criteria).
Time-Based Selection (Section 4.2.2.2)	<i>Can sub-sampling specify the data subset of interest?</i> Yes, oneM2M service layer has a number of filters for time-based selection based on creation, modification and expiration times. In oneM2M there is a specialized content based resource, called timeSeries, which provides various criteria for collection (sampling) and reporting of data.
C.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i> To date, the real-time data distribution performance of oneM2M deployments have not been documented and made publicly available. Real-time characterization is not currently available.
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<i>How does the latency and jitter change with throughput? What limits the throughput?</i> To date, the latency and jitter aspects of throughput of oneM2M deployments have not been documented and made publicly available. Latency and jitter characterization is not currently available.

C.6.5 Implementation Viewpoint	
C.6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<p><i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i></p> <p>Yes, oneM2M service layer can handle an increasing number of data-objects; however, the size of the object identifiers is limited by the size of string data type and that defines an upper bound on the maximum number of data objects.</p> <p>There is no theoretical limit on the data-object size. In practice, it will be limited by the amount of memory available on a host or the constraints of the underlying transport.</p>
Apps (Section 4.2.4.2)	<p><i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i></p> <p>Yes, oneM2M service layer can effectively support interface evolution for an increasing number of distributed application components. Application components are loosely coupled--they interact with resources in CSEs not with each other; the interfaces are data-oriented and can evolve independently. Applications are in complete control of how the data-object resources are organized and identified.</p>
C.6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<p><i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i></p> <p>Yes, as a service layer oneM2M does not place constraints on the availability of CSEs or applications. CSEs can use standard techniques for redundancy (load-balancers, clusters, virtualized environments).</p>
Recovery (Section 4.2.5.2)	<p><i>Can the connectivity framework support recovery when fault conditions occur?</i></p> <p>Yes, oneM2M service layer can support recovery when fault conditions occur. It accomplishes this by informing applications of errors and by allowing applications to change the behavior of how the system treats communication with endpoints.</p>
C.6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<p><i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i></p> <p>Yes, oneM2M service layer does not require a specific type of OS, CPU or even the database management system. oneM2M has been architected to work with constrained devices in mind.</p>
Incremental Upgrades (Section 4.2.6.2)	<p><i>Does the connectivity framework facilitate incremental upgrades?</i></p> <p>Yes, the oneM2M service layer does not place constraints on the upgradability of CSEs or applications. CSEs can use standard techniques for upgrading the CSE (load-balancers, clusters, virtualized environments).</p>

C.6.5 Implementation Viewpoint**C.6.5.7 Network Layer Considerations**

Topology (Section 5.2.1.1)	<i>What network topologies are allowed?</i> oneM2M service layer is agnostic to network topologies.
Span (Section 5.2.1.2)	<i>What is the span of the transport: LAN vs. WAN?</i> oneM2M service layer can be used within the LAN or across the WAN. CSEs and application can be located in either the LAN or WAN. The root CSE of the hierarchical tree is usually located within the M2M service providers WAN environment. oneM2M service layer implementations provide mechanisms to deal with firewalls and other restrictions encountered when going across the WAN.
Segmentation (Section 5.2.1.3)	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Yes, since the oneM2M service layer runs on IP based communication protocols. As such it can support multiple independent and isolated communication paths between the same network endpoints using the underlying IP network's mechanisms for path redundancy and isolation.

Annex D ASSESSMENT TEMPLATE: HTTP

This Annex contains the assessment template for Hypertext Transfer Protocol (HTTP).

D.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> Hypertext Transfer Protocol (HTTP)
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> Internet Engineering Task Force (IETF) and World Wide Web Consortium (W3C)
Description	<i>Short synopsis of the technology.</i> HTTP is the connectivity transport designed for the World Wide Web (WWW). Its primary goal is to serve the needs of web browsing. It is also used to expose application server interfaces as web services.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> World Wide Web (WWW). User Interfaces. Collaborative hypermedia applications.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> <ul style="list-style-type: none"> • TCP/IP • TLS for Transport Level Security
References	<i>Website¹ and other useful links to the technology.</i>

¹ See [HTTPWG]

D.6.2 Business Viewpoint	
D.6.2.1 Purpose (Section 6.2.1)	<p><i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i></p> <p>HTTP is the core connectivity transport of the World Wide Web (WWW). It was developed to support browsing the web of interconnected pages of hypertext markup language (HTML) and associated resources required to render a web page. As a result of its widespread availability, it has been used for exposing application sever interfaces, commonly referred to as web service APIs.</p>
D.6.2.2 Pedigree (Section 6.2.2)	<p><i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i></p> <p>HTTP originated in 1990-92 with informal drafts describing it as a protocol for collaborative hypermedia applications. HTTP/1.0 appeared as an IETF informational RFC in 1996. HTTP/1.1 appeared as a draft standard in 1999. It expanded upon the TCP/IP binding, and was finalized in 2014. HTTP/2 appeared in 2015. It further optimized the TCP/IP binding, while preserving the semantics.</p>
D.6.2.3 Variants (Section 6.2.3)	<p><i>Describe the options and variants from the original generic description of the technology.</i></p> <p>None.</p>
D.6.2.4 Maturity (Section 6.2.4)	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>HTTP is a mature technology. It forms the basis of the World Wide Web.</p>
D.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>Yes, HTTP is stable. HTTP/1.x is widely deployed. Toolkits for writing HTTP clients are available in nearly all the popular programming languages. There is a large selection of HTTP server implementations to choose from. Open-source and proprietary implementations are available.</p>
D.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>Internet Engineering Task Force (IETF)</p>
D.6.2.7 Openness (Section 6.2.7)	<p><i>Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?</i></p> <p>Yes, HTTP is an open standard managed by the IETF. Participation in the standards process is open to all. There are no annual dues and the IETF standards are available free of charge. The specifications process is open to participation by both vendors and users.</p> <p>Open source and commercial implementations are available.</p> <p>No, the specifications do not rely on any single component from any single vendor.</p>

D.6.3 Usage Viewpoint	
D.6.3.1 Architecture (Section 6.3.1)	<p><i>Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.</i></p> <p>HTTP defines a request-reply application protocol to exchange application state represented as hypertext with embedded resource identifiers. A client can request some action on a server resource, and be informed of the outcome of that request.</p> <p>A client request consists of an action method and a resource (path) on which the method is to be applied. The server replies to the request with a status code, which informs the client of the outcome of the method or the reason the method was not performed.</p> <p>Each request or reply message can have associated resource representation metadata header fields, which are name-value pairs, and provide additional information about the operation. Some headers are pre-defined, and applications can add their own headers.</p> <p>Each request or reply message can also have an optional body to hold a resource representation, which is hypertext with embedded links to resources.</p>
D.6.3.2 Technology Options (Section 6.3.2)	<p><i>List the choices to be made for using the connectivity technology in a system.</i></p> <ul style="list-style-type: none"> • Selection of resource representation format. • Multiple implementations choices for client and server libraries are available, including open source and proprietary, in a variety of programming languages. Implementations vary in their quality, performance, scalability, availability and security characteristics.
D.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>HTTP is most commonly known for navigating web pages and building application server interfaces. HTTP based applications are typically driven by a human end user. In the context of IoT, HTTP is also used as a connectivity transport for Information Technology (IT) applications.</p>
D.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <p>HTTP is typically used for serving web pages, and for exposing application server interfaces, and as a connectivity transport layer for some connectivity frameworks.</p>
D.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>Yes, another that can serve the same resources can replace a server. A server can support multiple clients.</p>
D.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>HTTP uses transport layer security (TLS) to provide end-to-end authentication, encryption and integrity. HTTP over TLS is referred to as HTTPS.</p>
D.6.3.7 Safety (Section 2.3.9)	<p><i>For systems that need it, are certifiable implementations available?</i></p> <p>There are no known safety certifiable implementations of HTTP.</p>
D.6.3.8 Gateways (Section 3.3)	<p><i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i></p> <p>Gateways to HTTP are defined by other core connectivity standards:</p> <ul style="list-style-type: none"> • OMG's Web-Enabled DDS defines a gateway for DDS. It allows HTTP clients to participate in a DDS data space. • OPC-UA supports a web service protocol using HTTP. • oneM2M uses HTTP as a connectivity transport option

D.6.4 Functional Viewpoint	
D.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>Yes, HTTP provides a data resource model. A data object is represented via an HTTP resource, formatted as a uniform resource identifier (URI) path string that is meaningful in context of the server. HTTP defines a core set of methods, GET, POST, PUT, DELETE that can be applied to the resources on a server. Resource representation returned by a server is hypertext that provides the context and links to other resources. To drive the application state, application architects will define the hypertext and the resource organization.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>Yes, the HTTP resource URI path string provides a way of identifying and addressing a data object within a server. The server itself is addressed as a network IP address and port number. The result of combing the URI with the network endpoint is a Uniform Resource Locator (URL), expressed as <i>http://</i> or <i>https://</i> (when TLS is used for security).</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>No, HTTP does not provide a data type system.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object's lifecycle? Summarize the salient aspects.</i></p> <p>Yes, it provides HTTP a means for managing data object lifecycles. A client can use the POST or the PUT method to create a data object; the GET method to retrieve the data object's representation; the PUT method to update a data object's representation; and the DELETE method to delete its representation. The server controls which methods are applicable on a data object, via the response for each operation.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>HTTP provides cache-mechanisms for proxies and clients to maintain responses of previous requests. These responses may contain the representations of the resources. The HTTP specification defines mechanisms for determining the freshness of the caches and provides rules for access control and applicability of a cached response.</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>No, HTTP does not provide a means to publish and subscribe to the state of data objects.</p>
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, HTTPS provides a means to request the state of data objects. This is the fundamental means of communicating using HTTP.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>Yes, data objects can be discovered via the embedded resource links in the hypertext.</p>
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>HTTP does not assume that connection will be continuously available. It supports request timeout error code, when a server does not receive a complete request within the time it was prepared to wait. When service on a resource is unavailable, a server can inform the client to retry after a certain amount of time.</p>

D.6.4 Functional Viewpoint	
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>HTTP does not provide data quality of service as described in section 4.1.10.</p>
Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>HTTP does not provide a data object security model.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>HTTP does not provide a standardized programming API. However, libraries are available in most popular programming languages and provide user friendly APIs.</p>
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>HTTP does not define a standardized way to configure, administer, and manage a server. Configuration, administration, and monitoring of HTTP servers is implementation specific. It is common practice to use configuration files for administration and log files for monitoring.</p>
D.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>HTTP relies on TCP. It requires reliable, ordered delivery of requests and responses. It can support partial or chunked delivery of requests and responses. There are no inherent message size limitations.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>Unicast.</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>A transport endpoint is a server IP address and a port number. There is no inherent limitation on the number of endpoints.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>No, HTTP does not require a connected circuit between a client and server. TCP connections may be torn down after a request-response, and reestablished for the next one.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>No, HTTP does not provide a means to prioritize messages.</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, HTTP does not provide the ability to synchronize time.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>Yes, HTTP can use Transport Layer Security (TLS) over TCP to provide message security.</p>

D.6.5 Implementation Viewpoint	
D.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<i>Does the connectivity framework require running a special process or broker?</i> No, HTTP does not require running a special process or broker to communicate between the client and the server.
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i> No, the client application code does not have to be aware of the server implementation details in order to participate in a data exchange. The server responses indicate the available resources and the methods allowed on them.
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<i>Is the governance explicit and shareable?</i> The governance is implicit, embedded in the request and response headers and data exchanged between a client and a server.
D.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<i>Can a content-filter specify the data subset of interest?</i> No, HTTP does not provide a content filtering mechanism to specify a data subset of interest. However, it does support the concept of “content negotiation” between a client and a server. It is left to the server to define the results of the content negotiation.
Time-Based Selection (Section 4.2.2.2)	<i>Can sub-sampling specify the data subset of interest?</i> No, HTTP does not provide a sub-sampling mechanism to specify a data subset of interest.
D.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i> No, HTTP is not designed to support real-time data distribution. The latency is not deterministic. The use of TCP can result in unbounded latency and jitter.
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<i>How does the latency and jitter change with throughput? What limits the throughput?</i> Latency and jitter can suffer as throughput increases. The throughput is limited by the message size, network bandwidth and available memory.
D.6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i> Yes, HTTP can effectively handle an increasing number of data objects. There is no inherent limitation on the representation size of a data object. A data object (i.e. resource) representation may be finite or may be unbounded.
Apps (Section 4.2.4.2)	<i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i> Yes, HTTP can effectively support interface evolution for an increasing number of distributed application clients, since the hypertext is used to decouple the clients from the server state. The hypertext response from a server defines its interface, and controls the resources and methods available to its clients.

D.6.5 Implementation Viewpoint	
D.6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i> Yes, HTTP can support continuous availability over a system relevant time period, as evidenced by the WWW.
Recovery (Section 4.2.5.2)	<i>Can the connectivity framework support recovery when fault conditions occur?</i> Yes, HTTP can support recovery when fault conditions occur.
D.6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i> HTTP is generally available on a wide variety of operating systems on a variety of CPUs, including embedded devices.
Incremental Upgrades (Section 4.2.6.2)	<i>Does the connectivity framework facilitate incremental upgrades?</i> Yes, HTTP facilitates incremental upgrades. Since the hypertext from a server controls the client interface, it can be upgraded any time. This is evidenced by the success of the WWW.
D.6.5.7 Network Layer Considerations	
Topology (Section 5.2.1.1)	<i>What network topologies are allowed?</i> HTTP is agnostic to network topologies, as it runs above the network layer.
Span (Section 5.2.1.2)	<i>What is the span of the transport: LAN vs. WAN?</i> HTTP can be used over LAN and WAN. It is typically used over WAN, as is evidenced by the WWW. IT infrastructure and firewall are friendly to HTTP.
Segmentation (Section 5.2.1.3)	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Yes, HTTP can support multiple independent isolated communication paths between the same network endpoints.

Annex E ASSESSMENT TEMPLATE: COAP

This Annex contains the assessment template for Constrained Application Protocol (CoAP).

E.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> Constrained Application Protocol (CoAP)
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> Internet Engineering Task Force (IETF)
Description	<i>Short synopsis of the technology.</i> The Constrained Application Protocol (CoAP) is a specialized web transfer protocol for use with constrained nodes and constrained networks in IoT. The protocol is designed for machine-to-machine (M2M) applications such as smart energy and building automation.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> IoT scenarios where devices are very constrained (in memory or CPU or both). Scenarios that require interoperability between web technologies and the general Internet with the IoT device domain.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> <ul style="list-style-type: none"> • UDP/IP (DTLS) • TCP/IP and Web Sockets (TLS) in progress
References	<i>Website¹ and other useful links to the technology.</i>

¹ See [CoAP]

E.6.2 Business Viewpoint	
E.6.2.1 Purpose (Section 6.2.1)	<p><i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i></p> <p>CoAP is specifically designed to communicate with resource-constrained devices or with devices across constrained (lossy or low throughput) IP networks. The design goals of CoAP are to provide a generic web protocol that keeps message overhead small, thus limiting the need for fragmentation. It offers features such as built-in discovery, multicast support and asynchronous message exchanges.</p>
E.6.2.2 Pedigree (Section 6.2.2)	<p><i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i></p> <p>CoAP was developed by the IETF as an internet standard.</p> <p>Work started on 2009 and culminated on RFC 7252 in 2014. There are several other supporting drafts and standards that relate to it. The IETF CoRE working group that maintains and enhances features related to CoAP enjoys an active and vibrant community with member continuously working to extend its applicability.</p> <p>Multiple independent CoAP implementations are available, including both open-source and commercial.</p>
E.6.2.3 Variants (Section 6.2.3)	<p><i>Describe the options and variants from the original generic description of the technology.</i></p> <p>There are no variants as such, but CoAP supports multiple transports UDP/TCP/SMS.</p>
E.6.2.4 Maturity (Section 6.2.4)	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>CoAP specifications have only been published since June 2014. Multiple interoperability events and commercial implementations have been deployed prior to the release of the publication as per the IETF process.</p> <p>Some of the many implementations are available at a website¹.</p>
E.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>The RFC is relatively recent (2014) but has been solid and stable. Newer documents that specify new functionality, like operation over TCP or HTTP mapping for the browser are more recent and still evolving.</p>
E.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>IETF² is where the CoAP standard is developed and maintained.</p>

¹ See [CoAP-Impl]

² See [IETF]

E.6.2 Business Viewpoint**E.6.2.7 Openness**
(Section 6.2.7)

Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?

CoAP is an open standard. The specifications are openly available to anyone at no cost. Anyone is free to download and implement them. The specifications process is open to participation by individuals.

Open source and commercial implementations are available.

CoAP specifications do not rely on any single component from any single vendor.

E.6.3 Usage Viewpoint	
E.6.3.1 Architecture (Section 6.3.1)	<p><i>Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.</i></p> <p>CoAP aims to provide more than plain connectivity or message passing functionality. Like HTTP it brings the RESTful architectural style of the World Wide Web (WWW) to the constrained space. Servers make resources available under a uniform resource identifier (URI), and clients access these resources using methods such as GET, PUT, POST, and DELETE.</p> <p>A device (endpoint) will run a CoAP Server and often a Client too. Clients elsewhere (i. e. other devices, browsers, applications) can request resources on the device as well as discover new devices and functionality.</p> <p>From a developer point of view, CoAP feels very much like HTTP. Obtaining a value from a sensor is not much different from obtaining a value from a Web API. For more details, please refer to page 10 of RFC 7252¹.</p>
E.6.3.2 Technology Options (Section 6.3.2)	<p><i>List the choices to be made for using the connectivity technology in a system.</i></p> <p>CoAP is a client/server model where the options include:</p> <ul style="list-style-type: none"> • Selection of resource representation format. • Selection of transport layer binding(s): UDP/IP or SMS or TCP/IP (in progress) and Web Sockets (in progress). • Selection of client and server implementation libraries. • Optional: Selection of HTTP proxy (CoAP-HTTP gateway). • Optional: Selection of resource directory server for resource discovery in constrained environments.
E.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>CoAP is a generic REST protocol upon which other technologies have been built. For device management, for example, the Open Mobile Alliance has created LWM2M, which supports management and operations of devices.</p>
E.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <p>The protocol is very versatile. It is suited for data collection, managed and unmanaged systems, systems that require scalability and systems that require security.</p>
E.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>CoRE specifications typically focus on protocol interactions and do not generally specify how elements of the connectivity functions are managed, monitored or replaced.</p>
E.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>CoAP defines a security model to authenticate and encrypt the interaction between CoAP clients and servers based on the underlying network datagram transport layer (DTLS/TLS) security mechanisms.</p> <p>CoAP specifications provide different types of end-to-end security and analysis of several possible attack vectors; please refer to page 80 of RFC 7252².</p> <p>A robust authentication and fine-grained access control security model is currently being defined by the IETF ACE working group for CoAP.</p>

¹ See [IETF-RFC7252]

² See [IETF-RFC7252]

E.6.3 Usage Viewpoint

E.6.3.7 Safety (Section 2.3.9)	<i>For systems that need it, are certifiable implementations available?</i> There are no known safety certifiable implementations of CoAP.
E.6.3.8 Gateways (Section 3.3)	<i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i> CoAP supports interworking gateways with the following connectivity protocols: <ul style="list-style-type: none">• HTTP to CoAP forward and reverse proxy (in progress).

E.6.4 Functional Viewpoint	
E.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>CoAP provides a data resource model, following the approach taken by HTTP. A data object is represented via a resource, formatted as a uniform resource identifier (URI) path string that is meaningful in context of the server. CoAP defines a core set of methods, GET, POST, PUT, DELETE that can be applied to the resources on a server. Resource representation returned by a server provides the context and links to other resources. It is up-to to the application architects to define the representations and the resource organization, to drive the application state.</p> <p>CoAP servers can expose resource representations in a variety of formats using a variety of data models. CoAP servers can use the <i>CoRE link format</i> for the resource representations. The CoRE link format identifies the paths to the resources in the CoAP server and provides annotations to the resources, which includes items like the content-type, interface and the resource type. CoAP clients can use these annotations to better understand the semantics of the resource.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>The CoAP resource URI path string provides a way of identifying and addressing a data object within a server. The server itself is addressed as a network IP address and port number. The result of combing the URI with the network endpoint is a Uniform Resource Locator (URL), expressed as “coap://” or “coaps:// (when DTLS is used for security).</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>CoAP does not dictate a specific data type system to define resources. CoAP simply transports resources as a payload on CoAP messages. However, the <i>CoRE link format</i> does provide target attributes for CoAP servers to report the content type of the resource.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object’s lifecycle? Summarize the salient aspects.</i></p> <p>CoAP provides a means for managing data object lifecycles. A client can use the POST or the PUT method to create a data object; the GET method to retrieve the data object’s representation; the PUT method to update a data object’s representation; and the DELETE method to delete its representation. The server controls which methods are applicable on a data object, via the response for each operation.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>CoAP does provide cache mechanisms for proxies and CoAP clients to maintain responses of previous requests. These responses may contain the representations of the resources or links of where the resources are located. The CoAP specification defines mechanisms for determining the freshness of the caches and provides rules for access control and applicability of a cached response.</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>CoAP provides a way to <i>observe</i> (subscribe) resources and get notifications when their state changes. The <i>OBSERVE</i> is an extension of the GET method with an additional option that requests the server to keep the representation updated over time. It is up-to to the server to determine how and when to notify the client of updates to resources. There is an IETF effort in progress to develop a <i>publish-subscribe broker</i> on top of this mechanism.</p>

E.6.4 Functional Viewpoint	
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>Yes, CoAP provides a means to request the state of data objects. This is the fundamental means of communicating using CoAP. All CoAP requests have an associated response.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>Yes, data objects can be discovered via the embedded resource links in the response from a server. The <i>CoRE link format</i> defines a simple format for exposing the resources offered by a CoAP server, and forms the basis for a resource directory.</p>
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>Yes, CoAP provides limited exception handling that resolves around the timeout when a CoAP request does not receive a response.</p>
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>CoAP supports two levels of QoS: Confirmable and Non-confirmable (best efforts). The CoAP messaging layer provides implementations for confirmable and non-confirmable message delivery over unreliable network layer protocols (e.g., UDP).</p>
Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>An IETF effort is in progress to define Object Security of CoAP.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>CoAP's API is a generic REST API. It does not provide a standardized programming API. However, libraries are available in most popular programming languages and provide user-friendly APIs.</p>
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>CoAP does not define a standardized way to configure, administer, and manage a server. Configuration, administration, and monitoring of CoAP servers is implementation specific.</p>

E.6.4 Functional Viewpoint	
E.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>The CoAP messaging protocol is a IP based protocol. It supports multiple bindings that are based on IP (i.e., UDP, TCP, SMS, Web Sockets). By default, it works over UDP. CoAP messages size is based on a 32-bit integer but CoAP messages work best without fragmentation, as such implementations tend to keep message sizes less than the underlying network transport layer payload sizes. A CoAP message, appropriately encapsulated, should fit within a single IP packet to packet to avoid IP fragmentation. When necessary CoAP does provide a mechanism to fragment and reassemble larger messages sizes.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>CoAP supports both unicast (default) and multicast (when available by the underlying transport).</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>A transport endpoint is a server IP address and a port number. There is no inherent limitation on the number of endpoints.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>CoAP does not require a connected circuit between a client and server.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>CoAP itself does not provide a way to prioritize messages.</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, CoAP does not provide a way to synchronize time between clients and servers.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>CoAP can use Datagram Transport Layer Security (DTLS) over UDP to provide message security. CoAP's default choice of DTLS parameters is equivalent to 3072-bit RSA keys, yet still it runs well on the smallest nodes.</p>

E.6.5 Implementation Viewpoint	
E.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<p><i>Does the connectivity framework require running a special process or broker?</i></p> <p>No brokers are required.</p> <p>Communication occurs between endpoints acting as clients or servers, and so is peer-to-peer oriented.</p>
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<p><i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i></p> <p>Clients do not have to be aware of the server behavior to participate in a data exchange. Clients need to have mechanisms for finding and operating on resources much as on the web. Servers can dynamically provide their interfaces to the clients.</p> <p>In practice, depending on the use case, it is feasible to build data-centric (RESTful, dynamic APIs) or device-centric (fixed API) architectures.</p>
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<p><i>Is the governance explicit and shareable?</i></p> <p>The governance is implicit, embedded in the request and response headers and data exchanged between a client and a server.</p>
E.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<p><i>Can a content-filter specify the data subset of interest?</i></p> <p>No, CoAP does not provide a content filtering mechanism to specify a data subset of interest. However, it does support the concept of “content negotiation” between a client and a server. A client can express interest in only a subset of the data via the query parameters on a resource. It is left up-to the server to define the results of the content negotiation.</p>
Time-Based Selection (Section 4.2.2.2)	<p><i>Can sub-sampling specify the data subset of interest?</i></p> <p>No, CoAP does not provide a sub-sampling mechanism to specify a data subset of interest. A client can express interest in only a subset of the data via the query parameters on a resource being observed. It is left up-to the server to define the results of the content negotiation.</p>
E.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<p><i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i></p> <p>CoAP is not aimed at real-time applications, but rather at resource constrained applications. Similar to TCP/IP connectivity, the exponential back-off and retry algorithm for confirmed reliability is not deterministic. CoAP does not provide mechanisms to ensure timeliness of data; that is left to the connectivity framework.</p>
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<p><i>How does the latency and jitter change with throughput? What limits the throughput?</i></p> <p>Compared to HTTP, CoAP endpoints should not experience more latency due to the use of CoAP, as it is very constrained and avoids fragmentation at multiple layers. CoAP should have smaller latency and jitter, compared to HTTP when used over UDP.</p>

E.6.5 Implementation Viewpoint	
E.6.5.4 Scalability Considerations	
Data Objects (Section 4.2.4.1)	<p><i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i></p> <p>Yes, CoAP can handle an effectively increasing number of data objects without constraints, as the limit of the resource identifier is the string size.</p>
Apps (Section 4.2.4.2)	<p><i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i></p> <p>CoAP is designed with evolution and long-term robustness in mind; it can support future changes or extensions to the servers or clients, much like HTTP.</p>
E.6.5.5 Availability Considerations	
Redundancy (Section 4.2.5.1)	<p><i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i></p> <p>Data availability will be determined by the availability of the CoAP Server, if the server is down a cached version of the resource can be accessed but it won't be "fresh".</p> <p>CoAP doesn't provide functionality for redundancy as part of the protocol. This is considered part of the application layer.</p>
Recovery (Section 4.2.5.2)	<p><i>Can the connectivity framework support recovery when fault conditions occur?</i></p> <p>CoAP doesn't provide functionality for recovery of fault conditions. This is considered part of the application layer.</p>
E.6.5.6 Deployment Considerations	
Platforms Constraints (Section 4.2.6.1)	<p><i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i></p> <p>CoAP is supported by several platforms, OSs and hardware. CoAP has been architected to work with constrained devices and networks in mind, and so can be made available on the smallest of platforms.</p>
Incremental Upgrades (Section 4.2.6.2)	<p><i>Does the connectivity framework facilitate incremental upgrades?</i></p> <p>CoAP, like HTTP, is designed with incremental updates and long-lasting client and server lifecycles. As a protocol CoAP does not place constraints on the upgradability of CoAP clients or servers. CoAP clients and server can use standard techniques for upgrading the CoAP client or server (load-balancers, clusters, virtualized environments).</p>

E.6.5 Implementation Viewpoint**E.6.5.7 Network Layer Considerations**

Topology <i>(Section 5.2.1.1)</i>	<i>What network topologies are allowed?</i> CoAP is agnostic to network topologies. CoAP has been used on low power networks that have a central point of connectivity to the outside (sink).
Span <i>(Section 5.2.1.2)</i>	<i>What is the span of the transport: LAN vs. WAN?</i> CoAP can be used within the LAN or across the WAN. CoAP clients and servers can be located in either the LAN or WAN. CoAP implementations provide proxy mechanisms to deal with firewalls and other restrictions encountered when going across the WAN.
Segmentation <i>(Section 5.2.1.3)</i>	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Yes, CoAP is an IP based communication protocol. As such it can support multiple independent and isolated communication paths between the same network endpoints using the underlying IP network's mechanisms for path redundancy and isolation.

Annex F ASSESSMENT TEMPLATE: MQTT

This Annex contains the assessment template for MQTT (formerly MQ Telemetry Transport).

F.6.1 General Info (Section 6.1)	
Name	<i>Common and formal name of the connectivity technology.</i> MQTT (formerly MQ Telemetry Transport)
Contacts	<i>Responsible standards development organization (SDO), task group or author(s), respective companies and email addresses.</i> OASIS
Description	<i>Short synopsis of the technology.</i> MQTT is a connectivity transport for lightweight machine-to-machine (M2M) messaging. MQTT uses a centralized broker and supports publish-subscribe communications pattern running on top of TCP/IP.
Application Domain(s)	<i>Application domains targeted by the connectivity technology.</i> Telemetry. Connecting remote sensors to the cloud. IoT scenarios where small code footprint is required and/or network bandwidth is at a premium.
Dependencies	<i>Possible commonalities with or reliance on other connectivity elements.</i> <ul style="list-style-type: none"> • TCP/IP • Recent addition MQTT-SN supports UDP/IP • TLS or DTLS for security
References	<i>Website¹ and other useful links to the technology.</i>

¹ See [MQTT]

F.6.2 Business Viewpoint	
F.6.2.1 Purpose (Section 6.2.1)	<p><i>Give the general motivation and expectation for the Connectivity Technology. This section provides the business rationale. It communicates the fundamental "why and what" for the project.</i></p> <p>Provide connectivity to M2M applications where small code footprint is required or network bandwidth is at a premium.</p> <p>MQTT may be considered for applications that exhibit high-cost connections, high latency, variable availability and negotiated delivery guarantees.</p>
F.6.2.2 Pedigree (Section 6.2.2)	<p><i>Describe the derivation, origin or history of the system. The objective is to understand the brief evolutionary context of this technology.</i></p> <p>The protocol was created by IBM in 1999 as the MQ Telemetry Protocol (MQTT).</p> <p>In 2010 IBM published the protocol under royalty-free terms.</p> <p>In 2011 IBM contributed the MQTT standard to OASIS and in 2012 the source code to Eclipse.</p> <p>The first OASIS standard version of MQTT (version 3.1.1) was approved in 2014.</p>
F.6.2.3 Variants (Section 6.2.3)	<p><i>Describe the options and variants from the original generic description of the technology.</i></p> <p>MQTT-SN is a variation aimed at embedded devices on non-TCP/IP networks.</p>
F.6.2.4 Maturity (Section 6.2.4)	<p><i>Estimate the technology maturity, state of development and condition relative to perfection. How refined are the connectivity concepts, requirements and demonstrated capabilities? Is the technology consistent and uniform?</i></p> <p>A website¹ maintains a list of notable projects that use MQTT.</p>
F.6.2.5 Stability (Section 6.2.5)	<p><i>Describe whether the connectivity technology has been in use for long enough that most of its initial faults and inherent problems have been removed or reduced; how easy is it to use for both non-experts and professionals? Has there been a reduction in the rate of new breakthrough advances related to it?</i></p> <p>The baseline MQTT protocol is stable and has been available for a long time. The more recent MQTT-SN protocol is not as mature.</p>
F.6.2.6 Standards Body (Section 6.2.6)	<p><i>List the relevant organizational bodies developing, coordinating, promulgating, revising, amending, reissuing, interpreting or otherwise producing technical standards and guidelines intended to address the needs of the base of affected adopters.</i></p> <p>OASIS²</p>
F.6.2.7 Openness (Section 6.2.7)	<p><i>Is it an open standard? Who can participate? Are the specifications freely available? Are open source implementations available? Does it require any single component from any single vendor?</i></p> <p>Yes, it is an open standard. OASIS members can participate in its development. The specifications are freely available. Open source implementations are available. It does not require any single component from any single vendor.</p>

¹ See [MQTT-P]

² See [OASIS]

F.6.3 Usage Viewpoint	
F.6.3.1 Architecture (Section 6.3.1)	<p><i>Summarize the main concepts, and high-level architecture, and terminology. Describe the end-to-end information exchange path.</i></p> <p>MQTT consists of multiple MQTT-Clients connected to a MQTT-Server (or broker). MQTT-Clients publish and subscribe to messages on one or more MQTT-Topics. A message published at a client is sent to the MQTT-Server, which sends it to all the subscribed MQTT-Clients. An MQTT message is an opaque vector of bytes.</p>
F.6.3.2 Technology Options (Section 6.3.2)	<p><i>List the choices to be made for using the connectivity technology in a system.</i></p> <ul style="list-style-type: none"> • Selection of MQTT versus MQTT-SN • Selection of the MQTT broker. This is one for a segment of connected applications. • Selection of client libraries (can be different for each client application).
F.6.3.3 Applications (Section 6.3.3)	<p><i>A general statement of the typical applications that rely on this connectivity technology and the reason for using the connectivity technology.</i></p> <p>According to the OASIS MQTT Technical Committee, target applications are sensors communicating to a broker via satellite links, occasional medical device dial-up connections with healthcare providers, home automation and small device scenarios. MQTT also targets mobile applications.</p>
F.6.3.4 Typical Usage (Section 2.2)	<p><i>What function or where in the system this technology is typically used?</i></p> <p>Centralized data collection.</p>
F.6.3.5 Operations (Section 2.3.8)	<p><i>Can one monitor, manage, and dynamically replace elements of the connectivity function?</i></p> <p>No, MQTT does not provide standardized mechanisms to monitor and manage a MQTT-Server. However, MQTT-clients can be replaced at any time.</p> <p>The broker routes all messages in the system. To avoid becoming a bottleneck it is deployed so that there is high-bandwidth connectivity to all critical clients.</p> <p>The broker should be specially protected against security breaches and denial of service attacks.</p>
F.6.3.6 Security (Section 2.3.5)	<p><i>What are the system security implications of this connectivity technology?</i></p> <p>Security is provided only at the transport level between each client and the broker. There is no end-to-end (client to client security). Therefore, if the broker is compromised, all data in the system will be compromised.</p> <p>The broker introduces a potential target to denial-service-attacks.</p>
F.6.3.7 Safety (Section 2.3.9)	<p><i>For systems that need it, are certifiable implementations available?</i></p> <p>There are currently no safety-certified client libraries or brokers.</p>
F.6.3.8 Gateways (Section 3.3)	<p><i>List of gateways to core connectivity standards and other relevant connectivity technologies.</i></p> <p>Custom application gateways have been developed for MQTT to DDS and HTTP to meet the needs of specific applications.</p>

F.6.4 Functional Viewpoint	
F.6.4.1 Core Framework Layer Functions	
Data Resource Model (Section 4.1.1)	<p><i>Does it provide a data resource model? Summarize the salient aspects.</i></p> <p>No, MQTT does not provide a data resource model.</p> <p>Messages are directed to Topics, which represent logical flows or streams.</p> <p>There is no explicit data resource model in MQTT. A single Topic might be used to represent data from multiple resources and the association will be encoded in the data and maintained by the clients.</p>
ID & Addressing (Section 4.1.2)	<p><i>Does it provide a way to identifying and addressing data objects? Summarize the identification and addressing scheme.</i></p> <p>No, it does not provide a way of identifying and addressing data objects. Addressing of individual resources within the streams is left to the application code.</p>
Data Type System (Section 4.1.3)	<p><i>Does it provide a data type system? Summarize the salient aspects.</i></p> <p>No, MQTT does not define a data type system. It transmits opaque data to be interpreted by the applications.</p>
Data Resource Lifecycle (CRUD) (Section 4.1.4)	<p><i>Does it provide a means of managing a data object's lifecycle? Summarize the salient aspects.</i></p> <p>No, it does not provide a means of managing a data object's lifecycle. There is no explicit resource management. This would be implemented by the client applications.</p>
State Management (Section 4.1.5)	<p><i>Does it provide a means to manage the recent history of data objects? Summarize the salient aspects.</i></p> <p>No, it does not provide a means to manage the recent history of data objects. Given that there is no resource management there is also no state management provided by MQTT.</p> <p>However, the MQTT broker can retain messages to be delivered to late joining applications and clients could use this to build state management at the application level.</p>
Publish-Subscribe (Section 4.1.6)	<p><i>Does it provide a means to publish and subscribe the state of data objects? Summarize the salient aspects.</i></p> <p>It provides a means to publish and subscribe messages on topics, but since there is no data-resource model, applications have to maintain the mapping of messages to state updates on data objects.</p>
Request-Reply (Section 4.1.7)	<p><i>Does it provide a means to request the state of data objects? Summarize the salient aspects.</i></p> <p>No, it does not provide a means to request the state of data objects.</p>
Discovery (Section 4.1.8)	<p><i>Does it provide a means to discover the data objects? Summarize the salient aspects.</i></p> <p>No, it does not provide a means to discover the data objects. Discovery is implicit by the fact that applications communicate via a broker. All client applications must connect to the same broker that has full knowledge of the topology of the system.</p>
Exception Handling (Section 4.1.9)	<p><i>Does it provide a means to handle exceptions when quality of service or connectivity violations happen? Summarize the salient aspects.</i></p> <p>No, it does not provide a means to handle exceptions when quality of service or connectivity violations happen.</p>
Data Quality of Service (QoS) (Section 4.1.10)	<p><i>Does it support data QoS? Summarize the scope and coverage. Highlight the salient aspects.</i></p> <p>MQTT provides limited QoS support. It includes best efforts and reliable delivery, and a minimal level of durability so that subscribers can receive a special message after a publisher goes offline.</p>

F.6.4 Functional Viewpoint	
Data Security (Section 4.1.11)	<p><i>Does it provide a data object security model? Summarize the salient aspects.</i></p> <p>No, it does not provide a data object security model.</p> <p>Only user name and password authentication is provided by the protocol. Security model is implemented by the broker and is not part of the MQTT standard.</p>
API (Section 4.1.12)	<p><i>Is there a standard API? Which programming languages is it available for?</i></p> <p>No, there is no standard programming API. It is implementation dependent.</p>
Governance (Section 4.1.13)	<p><i>Does it standardize the mechanisms for configuration, administration, and monitoring? Summarize the salient aspects.</i></p> <p>MQTT does not define a standardized way to configure, administer, and manage a broker. Configuration, administration, and monitoring of MQTT brokers is implementation specific.</p>
F.6.4.2 Core Transport Layer Functions	
Messaging Protocol (Section 5.1.1)	<p><i>Does it require UDP or TCP? What are the salient aspects of the messaging protocol? What are the message size limitations? What are the usage assumptions? Is it optimized for certain message requirements?</i></p> <p>The MQTT standard is the messaging protocol.</p> <p>Applications are responsible for building the communication framework on top of the MQTT transport protocol. There are no standards for this.</p> <p>MQTT requires TCP. MQTT-SN works over UDP.</p>
Communication Modes (Section 5.1.2)	<p><i>Which communication modes does it support?</i></p> <p>MQTT relies on unicast.</p> <p>MQTT-SN can use multicast but not with security.</p>
Endpoint Addressing (Section 5.1.3)	<p><i>Describe the transport endpoints. How are the endpoints addressed? What are the limitations, if any, on the number of endpoints?</i></p> <p>MQTT endpoints are the MQTT-Client and the MQTT-Server. MQTT uses standard IP host and port number addressing combined with the name of the Topic to direct messages.</p> <p>The number of TCP connections on the server host and the memory limits the number of endpoints.</p>
Connectedness (Section 5.1.4)	<p><i>Does it require a connected circuit between the endpoints? Summarize the salient aspects.</i></p> <p>MQTT is a connection-oriented transport on top of TCP.</p> <p>MQTT-SN is a connectionless transport on top of UDP.</p>
Prioritization (Section 5.1.5)	<p><i>Does it provide a means to prioritize messages? Summarize the salient aspects.</i></p> <p>No, it does not provide a means to prioritize messages.</p>
Timing & Synchronization (Section 5.1.6)	<p><i>Does it provide the ability to synchronize time? Summarize the salient aspects.</i></p> <p>No, it does not provide the ability to synchronize time.</p>
Message Security (Section 5.1.7)	<p><i>Does it provide mechanisms for message security? Summarize the salient aspects.</i></p> <p>No, it does not provide any mechanism for message security. Instead, it relies on transport-level security to authenticate the broker and provide integrity and confidentiality of the information: Transport Level Security (TLS) for MQTT and Datagram Transport Level Security (DTLS) for MQTT-SN.</p>

F.6.5 Implementation Viewpoint	
F.6.5.1 System Architecture Considerations	
Peer-to-Peer vs. Broker: (Section 4.2.1.1)	<i>Does the connectivity framework require running a special process or broker?</i> Broker-based. It requires running MQTT-Server, a special broker process, for clients to communicate.
Data-Centric vs. Device/App-Centric: (Section 4.2.1.2)	<i>Does the application code (or business logic) have to be aware of the other endpoints in order to participate in information exchange?</i> The application code (or business logic) does not have to be aware of the other endpoints to participate in data exchange. MQTT is only a connectivity transport. Data-centric or device-centric connectivity frameworks can be built with application-specific code.
Explicit vs. Implicit Governance: (Section 4.2.1.3)	<i>Is the governance explicit and shareable?</i> No, the governance is not explicit and shareable. Governance is enforced by the server implementation. The means for this are not standardized.
F.6.5.2 Data Considerations	
Content-Based Selection (Section 4.2.2.1)	<i>Can a content-filter specify the data subset of interest?</i> No, the data subset of interest cannot be specified by content. The content is opaque to MQTT.
Time-Based Selection (Section 4.2.2.2)	<i>Can sub-sampling specify the data subset of interest?</i> No, one cannot subscribe to a sub-sampled data subset of interest. MQTT attempts to deliver the published messages to all the subscribers on a topic.
F.6.5.3 Performance Considerations	
Real-Time (Section 4.2.3.1)	<i>Does the connectivity technology support real-time data distribution? Is the latency deterministic (smaller jitter is better)?</i> No, MQTT is a TCP and broker-based protocol, and is not intended for real-time. Binary protocol offers low overhead, but the use of TCP and relay via a broker provides non-deterministic latency.
Latency and Jitter vs. Throughput (Section 4.2.3.2)	<i>How does the latency and jitter change with throughput? What limits the throughput?</i> Implementation dependent, but use of a broker is likely to make latency highly dependent on throughput. Small protocol overhead benefits throughput, but the use of broker limits this to what a single broker can relay.

F.6.5 Implementation Viewpoint

F.6.5.4 Scalability Considerations

Data Objects (Section 4.2.4.1)	<p><i>Can the connectivity framework effectively handle an increasing number of data objects? What limits data object size?</i></p> <p>MQTT does not provide a notion of data objects or data object caching; just the notion of a Topic. The number of topics supported by a MQTT-Server will depend on the server memory and the number of clients. The number of clients will be limited by the capabilities of the MQTT broker and the number of connections it can sustain.</p> <p>There are no explicit message-size limits in MQTT, since it runs over TCP. MQTT-SN messages are over UDP will limit message size to what it can fit in a network datagram (64KB).</p>
Apps (Section 4.2.4.2)	<p><i>Can the connectivity framework effectively support interface evolution for an increasing number of distributed application components?</i></p> <p>Yes, MQTT can effectively support interface evolution for an increasing number of distributed application components, since the message are opaque. However, the applications will need to manage the message versioning and evolution.</p>

F.6.5.5 Availability Considerations

Redundancy (Section 4.2.5.1)	<p><i>Can the connectivity framework support continuous availability over a defined system-relevant time period?</i></p> <p>No, MQTT-Server does not support continuous availability over a defined system-relevant time period. The single point of failure introduced by the MQTT-Server will impact availability.</p>
Recovery (Section 4.2.5.2)	<p><i>Can the connectivity framework support recovery when fault conditions occur?</i></p> <p>No, MQTT does not support recovery when fault conditions occur.</p> <p>Broker health should be monitored to ensure system availability. There should be mechanisms in place to re-start the broker in case of malfunction or failure.</p>

F.6.5.6 Deployment Considerations

Platforms Constraints (Section 4.2.6.1)	<p><i>Does the connectivity framework support the operating system (OS), the CPU and the resource constraints on the platform(s) being used?</i></p> <p>MQTT is available for a number of platforms. Open-source implementations are available and could be built for target platforms.</p>
Incremental Upgrades (Section 4.2.6.2)	<p><i>Does the connectivity framework facilitate incremental upgrades?</i></p> <p>Yes, MQTT can facilitate incremental upgrades, since it is built on the publish-subscribe data exchange pattern.</p> <p>During deployment, the main requirement is to configure all clients to connect to the same broker.</p> <p>The centralization of the configuration on the broker simplifies deployment but it requires provisioning and maintenance of a service that is separate from all client applications that is common to all.</p> <p>Integrating separate applications developed using different brokers requires consolidation of the brokers.</p>

F.6.5 Implementation Viewpoint**F.6.5.7 Network Layer Considerations**

Topology (Section 5.2.1.1)	<i>What network topologies are allowed?</i> Hub and Spoke. The broker (MQTT-Server) is the hub. All messages flow via the broker.
Span (Section 5.2.1.2)	<i>What is the span of the transport: LAN vs. WAN?</i> MQTT can span globally over the WAN, as long as the broker is accessible via TCP/IP.
Segmentation (Section 5.2.1.3)	<i>Can the transport support multiple independent and isolated communication paths between the same network endpoints?</i> Segmentation is tied to the MQTT-Servers. Clients connected to different Servers are segmented from each other.

Annex G REVISION HISTORY

Revision	Date	Editor	Changes Made
V1.00	2017-02-28	Joshi, Mellor, Didier	Initial Release

Annex H ACRONYMS

API	Application Programming Interface
CoAP	Constrained Application Protocol
DDS	Data Distribution Service
DDSI-RTPS	Data Distribution Service Interoperability Wire Protocol (DDSI)– Real-Time Publish-Subscribe Protocol (RTPS)
DHCP	Dynamic Host Configuration Protocol
DNP	Distributed Network Protocol
DTLS	Datagram Transport Layer Security
HTTP	Hyper Text Transfer Protocol
ICMP	Internet Control Message Protocol
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IIC	Industrial Internet Consortium
IIoT	Industrial Internet of Things
IoT	Internet of Things
IP	Internet Protocol
ISO	International Organization for Standardization
IT	Information Technology
JSON	JavaScript Object Notation
kHz	Kilohertz
MAC	Media Access Control layer
MIB	Management Information Base
MTU	Maximum Transmission Unit
NAN	Neighborhood Area Network
ND	Neighbor Discovery
OMG	Object Management Group
OSI	Open Systems Interconnection
OT	Operational Technology
PHY	Physical Communications Layer
QoS	Quality of Service
REST	Representational State Transfer
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
UI	User Interface
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
XML	Extensible Markup Language
X-Types	Extensible Data Types

Annex I GLOSSARY

This document uses specific words and phrases which are defined in the Industrial Internet Vocabulary¹.

¹ See [IIC-IIV2015]

Annex J REFERENCES

- [CoAP] CoAP: RFC 7252 Constrained Application Protocol, retrieved 2017-02-16
<http://www.coap.technology>
- [CoAP-Impl] CoAP: Implementations, retrieved 2017-02-16
<http://www.coap.technology/impls.html>
- [Fielding-2000] Fielding, Roy Thomas: Chapter 5: Representational State Transfer (REST) Architectural Styles and the Design of Network-based Software Architectures (Ph.D.). University of California, Irvine, 2000, retrieved at 2017-01-29 download at
http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm
- [GOO-PB] Google Developers: Protocol Buffers, retrieved 2017-01-10
<https://developers.google.com/protocol-buffers/>
- [HTTPWG] IETF HTTP Working Group home page, retrieved 2017-02-16
<http://httpwg.org>
- [IETF] Internet Engineering Taskforce (IETF), retrieved 2017-01-10
<https://www.ietf.org/>
- [IETF-RFC768] Internet Engineering Task Force (IETF), Postel J.: RFC 768, User Datagram Protocol, 1980, retrieved 2017-01-29
<https://tools.ietf.org/html/rfc768>
- [IETF-RFC793] Internet Engineering Task Force (IETF), Postel J.: RFC 793, Transmission Control Protocol, 1981, retrieved 2017-01-29
<https://tools.ietf.org/html/rfc793>
- [IETF-RFC1122] Internet Engineering Task Force (IETF), Braden R.: RFC 1122, Requirements for Internet Hosts -- Communication Layers, 1989, retrieved 2017-01-29
<https://tools.ietf.org/html/rfc1122>
- [IETF-RFC4279] Internet Engineering Task Force (IETF), Kronen P., Tschofenig, H.: RFC 4279, Pre-Shared Key Ciphersuites for Transport Layer Security (TLS), 2005, retrieved 2017-01-10
<https://tools.ietf.org/html/rfc4279>
- [IETF-RFC7252] Internet Engineering Task Force (IETF), Shelby Z., Hartke K., Bormann C.: The Constrained Application Protocol (CoAP), 2014, retrieved 2017-02-16
<https://tools.ietf.org/html/rfc7252>
- [IIC-IIRA2015] IIC: The Industrial Internet, Volume G1: Reference Architecture Technical Report, version 1.7, 2015-June-04, retrieved 2017-01-10
<http://www.iiconsortium.org/IIRA.htm>

- [IIC-IISF2016] IIC: The Industrial Internet, Volume G4: Security Framework Technical Report, version 1.0, 2016-Sep-26, retrieved 2017-01-10
<http://www.iiconsortium.org/IISF.htm>
- [IIC-IIV2015] IIC: The Industrial Internet, Volume G8: Vocabulary Technical Report, version 1.0, 2015-May-07, retrieved 2017-01-10
<http://www.iiconsortium.org/vocab/index.htm>
- [ISO-7498-1] ISO/IEC standard 7498-1:1994
download at
[http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994(E).zip)
- [MQTT] MQTT protocol, retrieved 2017-02-16
<http://www.mqtt.org>
- [MQTT-P] MQTT projects, retrieved 2017-02-16
<http://www.mqtt.org/projects>
- [OASIS] OASIS: Advancing open standards for the information society, retrieved 2017-02-16
<https://www.oasis-open.org>
- [OMG-DDS] Object Management Group: DDS Portal – Data Distribution Service, retrieved 2017-01-10
<http://portals.omg.org/dds/>
- [OMG-DDSI-RTPS] Object Management Group: The Real-Time Publish-Subscribe Wire Protocol DDS Interoperability Wire Protocol (DDSI-RTPS), version 2.2, 2014 September, retrieved 2017-01-10
<http://www.omg.org/spec/DDSI-RTPS/>
- [OMG-DDSRPC] Object Management Group: Remote Procedure Call over DDS, version 1.0, 2016 June, retrieved 2017-01-10
<http://www.omg.org/spec/DDS-RPC/>
- [OMG-DDSSTD] Object Management Group: What’s in the DDS Standard? Open International Data-Centric Connectivity Standard, retrieved 2017-02-16
<http://portals.omg.org/dds/omg-dds-standard/>
- [OMG-DDSWEB] Object Management Group: Web-Enables DDS, version 1.0, 2016 December, retrieved 2017-01-10
<http://www.omg.org/spec/DDS-WEB/>
- [ONEM2M] OneM2M: Standards for M2M and the Internet of Things, retrieved 2017-02-16
<http://www.oneM2M.org>

- [ONEM2M-27] OneM2M: TR-0027, DDS usage in oneM2M system, Draft Technical Reports, version 0.1.0, 2016-08-07, retrieved 2017-01-10
http://www.onem2m.org/component/rsfiles/download-file/files?path=Draft_TR%255CTR-0027-DDS_usage_in_oneM2M-V0_1_0.DOC&Itemid=238
from
<http://www.onem2m.org/technical/latest-drafts>
- [ONEM2M-PS] OneM2M: Published Specifications, retrieved 2017-02-16
<http://www.onem2m.org/technical/published-documents>
- [OPC-CS] OPC Foundation: Case Studies, retrieved 2017-02-16
<https://opcfoundation.org/resources/case-studies/>
- [OPC-DDS] OPC Foundation: OPC Foundation and Object Management Group (OMG) Announce Collaborative Strategy for the OPC UA and DDS Connectivity Standards, 2016-04-06, retrieved 2017-01-10
<https://opcfoundation.org/wp-content/uploads/2016/04/OPCF-OPCUA-OMG-DDS-Positionspaper-final-v1.pdf>
from
<https://opcfoundation.org/news/press-releases/opc-foundation-and-object-management-group-omg-announce-collaborative-strategy-for-the-opc-ua-and-dds-connectivity-standards/>
- [OPC-MEM] OPC Foundation: Members, retrieved 2017-02-16
<https://opcfoundation.org/members>
- [OPC-UA] OPC Foundation: OPC Unified Architecture, retrieved 2016-09-05
<https://opcfoundation.org/about/opc-technologies/opc-ua/>
- [Tolk-2007] Tolk, Andreas and Diallo Y. Saikou, and Turnitsa, D. Charles: Applying the Levels of Conceptual Interoperability Model in Support of Integrability, Interoperability, and Composability for System-of-Systems Engineering, Journal of Systems, Cybernetics and Informatics, 2007
Download at [http://www.iiscisci.org/journal/cv\\$/sci/pdfs/p468106.pdf](http://www.iiscisci.org/journal/cv$/sci/pdfs/p468106.pdf)
- [W3C] World Wide Web Consortium (W3C)
<https://www.w3.org/>
- [W3C-WSA] World Wide Web Consortium (W3C): Web Services Architecture, W3C Working Group Note 11 February 2004, retrieved 2017-01-29
<https://www.w3.org/TR/ws-arch/>
- [WKPD-CI] Wikipedia: Conceptual Interoperability, retrieved 2017-01-10
http://en.wikipedia.org/wiki/Conceptual_interoperability
- [WKPD-IPS] Wikipedia: Internet Protocol Suite.
https://en.wikipedia.org/wiki/Internet_protocol_suite

- [WKPD-OSI] Wikipedia: OSI-Model, retrieved 2017-01-10
https://en.wikipedia.org/wiki/OSI_model
- [WKPD-REST] Wikipedia: Representational state transfer (REST), retrieved 2017-01-10
https://en.wikipedia.org/wiki/Representational_state_transfer
- [WKPD-WS] Wikipedia: Web Service, retrieved 2017-10-10
https://en.wikipedia.org/wiki/Web_service

USE OF INFORMATION—TERMS, CONDITIONS AND NOTICES

This is an Industrial Internet Consortium document (the “Document”) and is to be used in accordance with the terms, conditions and notices set forth below. This Document does not represent a commitment by any person to implement any portion or recommendation contained in it in any products or services. The information contained in this Document is subject to change without notice.

LICENSES

The companies listed above have granted to the Object Management Group, Inc. (OMG) and its Industrial Internet Consortium (the “IIC”) a nonexclusive, irrevocable, royalty-free, paid up, worldwide license to copy and distribute this Document and to modify this Document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having copied, distributed or used such material set forth herein.

Subject to all of the terms and conditions below, the owners of the copyright in this Document hereby grant you a fully-paid up, non-exclusive, nontransferable, perpetual, worldwide license (without the right to sublicense) to use, copy and distribute this Document (the “Permission”), provided that: (1) both the copyright notice above, and a copy of this Permission paragraph, appear on any copies of this Document made by you or by those acting on your behalf; (2) the use of the Document is only for informational purposes in connection with the IIC’s mission, purposes and activities; (3) the Document is not copied or posted on any network computer, publicly performed or displayed, or broadcast in any media and will not be otherwise resold or transferred for commercial purposes; and (4) no modifications are made to this Document.

This limited Permission is effective until terminated. You may terminate it at any time by ceasing all use of the Document and destroying all copies. The IIC may terminate it at any time by notice to you. This Permission automatically terminates without notice if you breach any of these terms or conditions. Upon termination, or at any time upon the IIC’s express written request, you will destroy immediately any copies of this Document in your possession or control.

The Licenses and Permission relate only to copyrights and do not convey rights in any patents (see below).

PATENTS

Compliance with or adoption of any advice, guidance or recommendations contained in any IIC reports or other IIC documents may require use of an invention covered by patent rights. *OMG and the IIC are not responsible for identifying patents for which a license may be required to comply with any IIC document or advice, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. IIC documents are informational and advisory only. Readers of this Document are responsible for protecting themselves against*

liability for infringement of patents and other intellectual property that may arise from following any IIC recommendations or advice. OMG disclaims all responsibility for such infringement.

GENERAL USE RESTRICTIONS

This Document contains content that is protected by copyright. Any unauthorized use of this Document may violate copyright laws, trademark laws and communications regulations and statutes. Except as provided by the above Licenses, no part of this work covered by copyright may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping or information storage and retrieval systems—without permission of the copyright owner(s).

DISCLAIMER OF WARRANTY

WHILE THIS DOCUMENT IS BELIEVED TO BE ACCURATE, IT IS PROVIDED “AS IS” AND MAY CONTAIN ERRORS OR MISPRINTS. THE OBJECT MANAGEMENT GROUP, INC. (INCLUDING THE IIC) AND THE COPYRIGHT OWNERS LISTED ABOVE MAKE NO WARRANTY, REPRESENTATION OR CONDITIONS OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS DOCUMENT, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, ANY IMPLIED WARRANTY OR MERCHANTABILITY OR ANY WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE OR USE. IN NO EVENT SHALL THE OBJECT MANAGEMENT GROUP, INC. (INCLUDING THE IIC) OR ANY OF THE COPYRIGHT OWNERS BE LIABLE FOR ERRORS CONTAINED HEREIN OR FOR DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, RELIANCE OR COVER DAMAGES, INCLUDING LOSS OF PROFITS, REVENUE, DATA OR USE, INCURRED BY ANY USER OR ANY THIRD PARTY IN CONNECTION WITH THE FURNISHING, PERFORMANCE, REPRODUCTION, DISTRIBUTION OR USE OF THIS MATERIAL, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The entire risk as to the quality and performance of any software or technology developed using this Document is borne by you. This disclaimer of warranty constitutes an essential part of the Licenses granted to you to use this Document.

LIMITED RIGHTS NOTICE

This Document contains technical data that was developed at private expense and (i) embodies trade secrets, or (ii) is confidential and either commercial or financial. This document was not produced in the performance of a government contract and is not in the public domain. The use, duplication or disclosure of this Document by the U.S. Government is subject to the restrictions set forth in 48 C.F.R. 52.227-14—Rights in Data “Limited Rights Notice (Dec. 2007) (a) and (b),” or as specified in 48 C.F.R. 12.211 of the Federal Acquisition Regulations and its successors, as applicable. This data may only be reproduced and used by the U.S. Government with the express limitation that it will not, without written permission of the copyright owners, be used for purposes of manufacture nor disclosed outside the Government. The copyright owners are as indicated above and may be contacted through the Object Management Group, Inc., 109 Highland Avenue, Needham, MA 02494, U.S.A.

TRADEMARKS

The trademarks, service marks, trade names and other special designations that appear on and within the Document are the marks of OMG, the copyright holders listed above and possibly other manufacturers and suppliers identified in the Document and may not be used or reproduced without the express written permission of the owner, except as necessary to reproduce, distribute and refer to this Document as authorized herein.