



**International
Standard**

ISO/IEC 19987

Information technology — EPC Information Services (EPCIS)

*Technologies de l'information — Services d'information sur les
codes de produit électronique*

**Third edition
2024-03**

Reference number
ISO/IEC 19987:2024(en)

Full document details: ISO IEC 19987-2024

 **Documents**
© ISO/IEC 2024 Document Sharing Hub



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2024

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

© ISO/IEC 2024 – All rights reserved

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and <https://patents.iec.ch>. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by GS1 (as EPCIS Standard, Release 2.0) and drafted in accordance with its editorial rules. It was adopted, under the JTC 1 PAS procedure, by Joint Technical Committee ISO/IEC JTC 1, *Information technology*.

This third edition cancels and replaces the second edition (ISO/IEC 19987:2017), which has been technically revised.

The main changes are as follows:

- addition of JSON/SOJSON-LD syntax (alongside XML);
- addition of REST bindings (alongside SOAP/WSDL);
- complete overhaul of UML diagrams;
- clarification on distinction between standard vocabulary and user vocabulary;
- new AssociationEvent;
- new “How” event dimension;
- overview of EPCIS even “dimensions” with cross-references to relevant sections in EPCIS (this document) and CBV (ISO/IEC 19988);

ISO/IEC 19987:2024(en)

- new Persistent Disposition indicating non-transient business state of an object;
- new SensorElement to accommodate sensor data;
- addition of certificationInfo to core EPCISEvent;
- update of SimpleEventQuery parameters;
- removal of support for Simple Master Data Query and EPCIS Master Data Document.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Table of Contents

1 Introduction 1

2 Relationship to the GS1 System Architecture 2

2.1 Overview of GS1 standards 2

2.2 EPCIS in relation to the "Capture" and "Share" layers..... 3

2.3 EPCIS in Relation to trading partners 4

2.4 EPCIS in relation to other GS1 System Architecture components..... 5

3 EPCIS specification principles 8

4 Terminology and typographical conventions 9

5 EPCIS specification framework 10

5.1 Layers.....10

5.2 Extensibility 11

5.3 Modularity 11

6 Abstract data model layer 13

6.1 Event data and master data 13

6.1.1 Transmission of master data in EPCIS.....15

6.2 Standard vocabulary and user vocabulary 15

6.3 Extension mechanisms 17

6.4 Identifier representation 18

6.5 Hierarchical vocabularies 19

7 Data definition layer..... 20

7.1 General rules for specifying data definition layer modules 20

7.1.1 Content.....20

7.1.2 Notation.....21

7.1.3 Semantics 21

7.2 Core event types module – overview 23

7.2.1 UML Diagrams of EPCIS Event Types 24

7.2.2 Overview of EPCIS event "dimensions" (non-normative).....25

7.2.3 Table of vocabulary types 28

7.3 Core event types module – building blocks 29

7.3.1 Primitive types 29

7.3.2 Action type 30

7.3.3 The "What" dimension 31

7.3.4 The "When" dimension 32

7.3.5 The "Where" Dimension – read point and business location.....33

7.3.6 The "Why" dimension 36

7.3.7 The "How" dimension 40

7.3.8 Instance/Lot master data (ILMD) 47

7.4 Core event types module – events 48

7.4.1 EPCISEvent 48

7.4.2 ObjectEvent (subclass of EPCISEvent) 52

7.4.3 AggregationEvent (subclass of EPCISEvent).....56

7.4.4 TransactionEvent (subclass of EPCISEvent) 60

7.4.5 TransformationEvent (subclass of EPCISEvent) 65

7.4.6 AssociationEvent (subclass of EPCISEvent).....68

8	Service Layer	74
8.1	Core capture operations module.....	76
8.1.1	Authentication and authorisation.....	76
8.1.2	Capture service	77
8.2	Core Query operations module.....	78
8.2.1	Authentication.....	78
8.2.2	Authorisation and redaction.....	79
8.2.3	Queries for large amounts of data	79
8.2.4	Overly complex queries	80
8.2.5	Query framework (EPCIS query control interface).....	80
8.2.6	Error conditions.....	86
8.2.7	Predefined queries for EPCIS	88
8.2.8	Query callback interface	121
9	XML bindings for data definition modules	122
9.1	Extensibility mechanism	122
9.2	Standard business document header	125
9.3	EPCglobal Base schema	125
9.4	Master data in the XML binding	126
9.5	Schema for core event types	127
9.6	Core event types – examples (Non-Normative)	128
10	JSON/JSON-LD bindings for data definition	129
10.1	Brief introduction to JSON and JSON-LD in the context of EPCIS	129
10.1.1	JavaScript Object Notation (JSON)	130
10.1.2	JSON for Linked Data (JSON-LD).....	131
10.1.3	Features of the JSON-LD context resource.....	132
10.1.4	Compact URI Expressions (CURIEs).....	133
10.2	Expression and validation of EPCIS data structures in JSON and JSON-LD	134
10.2.1	Expressing data fields expecting simple values.....	134
10.2.2	Validating data fields expecting simple values.....	136
10.2.3	Validation of fields (e.g. 'action') that expect a string value from an enumerated list.....	138
10.2.4	Expressing simple lists of values	139
10.2.5	Validating lists of values	140
10.2.6	Expressing lists of elements with inline attributes expressing type	140
10.2.7	Modelling and validating subclasses of EPCIS event	143
10.2.8	Comparison of how validation rules are expressed in XSD, JSON Schema and SHACL	144
10.2.9	Mapping core SBDH fields to the JSON/JSON-LD data format for EPCIS	146
10.2.10	Online validation tools for JSON Schema and SHACL	146
10.2.11	Libraries and toolkits providing JSON-LD support.....	146
10.3	Validation schema (references to normative content)	146
10.4	Non-normative examples in JSON and JSON-LD	147
11	Bindings for core capture operations module	147
11.1	Message queue binding.....	147
11.2	HTTP binding.....	148
12	REST Bindings	149
12.1	Code conventions	149
12.2	Introduction to REST	149

12.3	Content negotiation, service discovery and custom headers for EPCIS	151
12.4	Authentication and Authorization.....	153
12.5	Pagination	154
12.6	Capturing EPCIS Events	154
12.6.1	Capture Interface	155
12.6.2	Capture Jobs Interface	156
12.7	Events interface.....	157
12.7.1	EPCIS events collections	157
12.7.2	EPCIS events endpoints.....	157
12.7.3	Event filtering with the EPCIS query language.....	158
12.7.4	Top-level resources.....	159
12.8	Query control interface	160
12.8.1	Creating and using named queries	162
12.8.2	Deleting named queries.....	162
12.8.3	Subscribing to named queries.....	162
12.8.4	EPCIS query language	167
12.8.5	EPCIS query in the URL	168
12.9	Backward Compatibility of REST bindings with EPCIS 1.2	169
12.10	EPCIS Error Conditions and HTTP Status Code Mapping	169
13	Bindings for core query operations module	172
13.1	XML schema for core query operations module.....	172
13.2	SOAP/HTTP binding for the query control interface.....	173
13.3	AS2 Binding for the query control interface.....	174
13.3.1	GS1 AS2 guidelines (Non-Normative)	175
13.4	Bindings for query callback interface.....	177
13.4.1	General Considerations for all XML-based bindings	178
13.4.2	HTTP binding of the query callback interface.....	178
13.4.3	HTTPS binding of the query callback interface	179
13.4.4	AS2 Binding of the query callback interface	179
14	Conformance	180
14.1	Conformance of EPCIS XML data	180
14.2	Conformance of EPCIS capture interface clients.....	180
14.3	Conformance of EPCIS capture interface servers	180
14.4	Conformance of EPCIS query interface clients	181
14.5	Conformance of EPCIS query interface servers	181
14.6	Conformance of EPCIS query callback interface implementations	181
14.7	Conformance of JSON/JSON-LD bindings.....	181
14.8	Conformance of REST Interface for EPCIS 2.0 Servers.....	182
15	UML Diagrams for SBDH	184
15.1	UML aligned with text of SBDH specification.....	185
15.2	UML aligned with XSD of SBDH specification	185
16	List of abbreviations (non-normative)	185
17	References	187

Index of figures

Figure 2-1 EPCIS in relation to the "Capture" and "Share" layers	3
Figure 2-2 EPCIS in relation to other GS1 System Architecture components	5
Figure 5-1 Layers of the EPCIS specification framework.....	10
Figure 6-1 Structure of event data and master data in EPCIS	14
Figure 7-1 EPCIS data definition notation	21
Figure 7-2 EPCIS UML with Ontology focus	24
Figure 7-3 EPCIS UML with Syntax focus.....	24
Figure 7-4 Example of the distinction between a read point and a business location	35
Figure 7-5 Coordinate reference systems	46
Figure 7-6 Association and Aggregation with returnable transport units (RTIs).....	69
Figure 7-7 Association and Aggregation with containers	70
Figure 7-8 Association and Aggregation in a room	70
Figure 8-1 EPCIS Service Layer	75
Figure 10-1 RDF Triple: Subject-Property-Value	130
Figure 10-2 Supporting multiple formats for EPCIS / CBV 2.0	134
Figure 12-1 Client first uses OPTIONS to discover which versions are supported and making GET request.....	153
Figure 12-2 Authentication and authorisation	153
Figure 12-3 Endpoint: Capture Interface workflow	155
Figure 12-4 EPCIS query as URL query parameters	158
Figure 12-5 Endpoint: Named queries workflow.....	161
Figure 12-6 Client creates a named query for EPCIS events and uses pagination to retrieve all EPCIS events	162
Figure 12-7 Scheduled query workflow	164
Figure 12-8 Event streaming query workflow.....	165
Figure 12-9 Query subscription with Webhook (HTTP Callback)	166
Figure 12-10 Query subscription with a WebSocket	167
Figure 15-1 UML aligned with text of SBDH	185
Figure 15-2 UML aligned with XSD of SBDH	185

1 Introduction

This document is a GS1 standard that defines Version 2.0 of EPC Information Services (EPCIS). The goal of EPCIS is to enable disparate applications to create and share visibility event data, both within and across enterprises. Ultimately, this sharing is aimed at enabling users to gain a shared view of physical or digital objects within a relevant business context.

"Objects" in the context of EPCIS typically refers to physical objects that are identified either at a class or instance level and which are handled in physical handling steps of an overall business process involving one or more organisations. Examples of such physical objects include trade items (products), logistic units, returnable assets, fixed assets, physical documents, etc. "Objects" may also refer to digital objects, also identified at either a class or instance level, which participate in comparable business process steps. Examples of such digital objects include digital trade items (music downloads, electronic books, etc.), digital documents (electronic coupons, etc.), and so forth. Throughout this document the word "object" is used to denote a physical or digital object, identified at a class or instance level, that is the subject of a business process step. EPCIS data consist of "visibility events," each of which is the record of the completion of a specific business process step acting upon one or more objects.

The EPCIS standard was originally conceived as part of a broader effort to enhance collaboration between trading partners by sharing of detailed information about physical or digital objects. The name EPCIS reflects the origins of this effort in the development of the Electronic Product Code (EPC). It should be noted, however, that EPCIS does not require the use of Electronic Product Codes, nor of Radio-Frequency Identification (RFID) data carriers, and does not even require instance-level identification (for which the Electronic Product Code was originally designed). The EPCIS standard applies to all situations in which visibility event data is to be captured and shared, and the presence of "EPC" within the name is of historical significance only.

EPCIS provides open, standardised interfaces that allow for seamless integration of well-defined services in inter-company environments as well as within companies. Standard interfaces are defined in the EPCIS standard to enable visibility event data to be captured and queried using a defined set of service operations and associated data standards, all combined with appropriate security mechanisms that satisfy the needs of user companies. In many or most cases, this will involve the use of one or more persistent databases of visibility event data, though elements of the Services approach could be used for direct application-to-application sharing without persistent databases.

With or without persistent databases, the EPCIS specification specifies only standard data sharing interfaces between applications that capture visibility event data and those that need access to it. *It does not specify how the service operations or databases themselves should be implemented.* This includes not defining how the EPCIS services should acquire and/or compute the data they need, except to the extent the data is captured using the standard EPCIS capture operations. The interfaces are needed for interoperability, while the implementations allow for competition among those providing the technology and implementing the standard.

EPCIS is intended to be used in conjunction with the GS1 Core Business Vocabulary (CBV) standard [CBV2.0]. EPCIS and the CBV are developed, maintained and published by GS1; EPCIS and the CBV are also published within ISO's PAS process as ISO/IEC 19987 and ISO/IEC 19988, respectively. The CBV standard provides definitions of data values that may be used to populate the data structures defined in the EPCIS standard. The use of the standardised vocabulary provided by the CBV standard is critical to interoperability and critical to provide for querying of data by reducing the variation in how different businesses express common intent. Therefore, applications should use the CBV standard to the greatest extent possible in constructing EPCIS data.

The companion EPCIS and CBV Implementation Guideline [EPCISGuideline] provides additional guidance for building visibility systems using EPCIS and CBV, including detailed discussion of how to model specific business situations using EPCIS/CBV data and methods for sharing such data between trading partners.

2 Relationship to the GS1 System Architecture

This section is largely quoted from [GS1Arch], and shows the relationship of EPCIS to other GS1 standards.

2.1 Overview of GS1 standards

GS1 standards support the information needs of end users interacting with each other in supply chains, specifically the information required to support the business processes through which supply chain participants interact. The subjects of such information are the real-world entities that are part of those business processes. Real-world entities include things traded between companies, such as products, parts, raw materials, packaging, and so on. Other real-world entities of relevance to trading partners include the equipment and material needed to carry out the business processes surrounding trade such as containers, transport, machinery; entities corresponding to physical locations in which the business processes are carried out; legal entities such as companies, divisions; service relationships; business transactions and documents; and others. Real-world entities may exist in the tangible world, or may be digital or conceptual. Examples of physical objects include a consumer electronics product, a transport container, and a manufacturing site (location entity). Examples of digital objects include an electronic music download, an eBook, and an electronic coupon. Examples of conceptual entities include a trade item class, a product category, and a legal entity.


GS1 standards may be divided into the following groups according to their role in supporting information needs related to real-world entities in supply chain business processes:

- Standards which provide the means to **identify** real-world entities so that they may be the subject of electronic information that is stored and/or communicated by end users. GS1 identification standards include standards that define unique identification codes (called GS1 identification keys).
- Standards which provide the means to automatically **capture** data that is carried directly on physical objects, bridging the world of physical things and the world of electronic information. GS1 data capture standards include definitions of barcode and radio-frequency identification (RFID) data carriers which allow identifiers to be affixed directly to a physical object, and standards that specify consistent interfaces to readers, printers, and other hardware and software components that connect the data carriers to business applications.
- Standards which provide the means to **Share** information, both between trading partners and internally, providing the foundation for electronic business transactions, electronic visibility of the physical or digital world, and other information applications. GS1 standards for information sharing include this EPCIS Standard which is a standard for visibility event data. Other standards in the "Share" group are standards for master data and for business transaction data, as well as discovery standards that help locate where relevant data resides across a supply chain and trust standards that help establish the conditions for sharing data with adequate security.

The EPCIS standard fits into the "Share" group, providing the data standard for visibility event data and the interface standards for capturing such information from data capture infrastructure (which employs standards from the "Capture" group) and for sharing such information with business applications and with trading partners.

2.2 EPCIS in relation to the "Capture" and "Share" layers

Figure 2-1 EPCIS in relation to the "Capture" and "Share" layers



The diagram above shows the relationship between EPCIS and other GS1 standards in the "Capture" and "Share" groups. (The "Identify" group of standards pervades the data at all levels of this architecture, and so is not explicitly shown.)

As depicted in the diagram above, the EPCIS Capture Interface exists as a bridge between the "Capture" and "Share" standards. The EPCIS Query Interface provides visibility event data both to internal applications and for sharing with trading partners.

At the centre of a data capture application is the data capture workflow that supervises the business process step within which data capture takes place. This is typically custom logic that is specific to the application. Beneath the data capture workflow in the diagram is the data path between the workflow and GS1 data carriers: barcodes and RFID. The green bars in the diagram denote GS1 standards that may be used as interfaces to the data carriers. At the top of the diagram are the interfaces between the data capture workflow and larger enterprise applications. Many of these interfaces are application- or enterprise-specific, though using GS1 data as building blocks; however, the EPCIS interface is a GS1 standard. Note that the interfaces at the top of the diagram, including EPCIS, are independent of the data carrier used at the bottom of the diagram.

The purpose of the interfaces and the reason for a multi-layer data capture architecture is to provide isolation between different levels of abstraction. Viewed from the perspective of an enterprise application (i.e., from the uppermost blue box in the figure), the entire data capture application shields the enterprise application from the details of exactly how data capture takes place. Through the application-level interfaces (uppermost green bars), an enterprise application interacts with the data capture workflow through data that is data carrier independent and in which all of the interaction between data capture components has been consolidated into that data. At a lower level, the data capture workflow is cognizant of whether it is interacting with

barcode scanners, RFID interrogators, human input, etc., but the transfer interfaces (green bars in the middle) shield the data capture workflow from low-level hardware details of exactly how the data carriers work. The lowest level interfaces (green bars on the bottom) embody those internal data carrier details. EPCIS and the “Share” layer in general differ from elements in the Capture layer in three key respects:

1. EPCIS deals explicitly with historical data (in addition to current data). The Capture layer, in contrast, is oriented exclusively towards real-time processing of captured data.
2. EPCIS often deals not just with raw data captured from data carriers such as barcodes and RFID tags, but also in contexts that imbue those observations with meaning relative to the physical or digital world and to specific steps in operational or analytical business processes. The Capture layers are more purely observational in nature. An EPCIS event, while containing much of the same “Identify” data as a Filtering & Collection event or a barcode scan, is at a semantically higher level because it incorporates an understanding of the business context in which the identifier data were obtained. Moreover, there is no requirement that an EPCIS event be directly related to a specific physical data carrier observation. For example, an EPCIS event may indicate that a perishable trade item has just crossed its expiration date; such an event may be generated purely by software.
3. EPCIS operates within enterprise IT environments at a level that is much more diverse and multi-purpose than exists at the Capture layer, where typically systems are self-contained and exist to serve a single business purpose. In part, and most importantly, this is due to the desire to share EPCIS data between enterprises which are likely to have different solutions deployed to perform similar tasks. In part, it is also due to the persistent nature of EPCIS data. And lastly, it is due to EPCIS being at the highest level of the overall architecture, and hence the natural point of entry into other enterprise systems, which vary widely from one enterprise to the next (or even within parts of the same enterprise).

2.3 EPCIS in Relation to trading partners

GS1 standards in the “Share” layer pertain to three categories of data that are shared between end users:

Data	Description	GS1 standards
Master data	Data, shared by one trading partner to many trading partners, that provide descriptive attributes of real-world entities identified by GS1 identification keys, including trade items, parties, and physical locations.	GDSN
Transaction data	Trade transactions triggering or confirming the execution of a function within a business process as defined by an explicit business agreement (e.g., a supply contract) or an implicit one (e.g., customs processing), from the start of the business process (e.g., ordering the product) to the end of it (e.g., final settlement), also making use of GS1 identification keys.	GS1 XML EANCOM
Visibility event data	Details about physical or digital activity in the supply chain of products and other assets, identified by keys, detailing where these objects are in time, and why; not just within one organisation’s four walls, but across organisations.	EPCIS

Transaction Data and Visibility Event Data have the characteristic that new documents of those types are continually created as more business is transacted in a supply chain in steady state, even if no new real-world entities are being created. Master data, in contrast, is more static: the master data for a given entity changes very slowly (if at all), and the quantity of master data only increases as new entities are created, not merely because existing entities participate in business processes. For example, as a given trade item instance moves through the supply chain, new transaction data and visibility event data are generated as that instance undergoes business transactions (such as purchase and sale) and physical handling processes (packing, picking, stocking, etc.). But new master data is only created when a new trade item or location is added to the supply chain.

The diagram illustrates the GS1 Networked Services architecture, divided into two main sections: **GS1 Networked Services** (top) and **Trading Partner 1 / Trading Partner 2** (bottom, separated by a dashed line).

GS1 Networked Services:

- Lookup services for data discovery** and **GDSN Global Registry** are at the top, each connected to a green box below it via a double-headed arrow.
- These green boxes are connected to a central **GS1 Networked Services** block via double-headed arrows.

Trading Partner 1 and Trading Partner 2:

The architecture is mirrored for both trading partners, connected by a central line labeled **EPCIS Query traffic - on-demand queries and standing queries**.

Trading Partner 1 Components:

- Top Layer:** **eCOM (GS1 XML / EANCOM) Interface** (green box) connects to **GDSN Interface and data pools** (green box) via a double-headed arrow.
- Accessing Layer:** **EPCIS Accessing Applications and other Enterprise-level Applications** (blue box) connects to the **GDSN Interface** and the **EPCIS Query Interface** via double-headed arrows.
- Repository Layer:** **EPCIS Repository** (blue box) connects to the **EPCIS Accessing Applications** and the **EPCIS Query Interface** via double-headed arrows.
- Interface Layer:** **EPCIS Query Interface (Control and Callback)** (green box) connects to the **EPCIS Repository** and the **EPCIS Accessing Applications** via double-headed arrows.
- Capture Layer:** **EPCIS Capture Interface** (green box) connects to the **EPCIS Accessing Applications** and the **App-specific interfaces** via double-headed arrows.
- Infrastructure:** **App-specific interfaces** (green box) and **Data Capture Infrastructure** (grey box) form the base layer.

Trading Partner 2 Components:

- Top Layer:** **eCOM (GS1 XML / EANCOM) Interface** (green box) connects to **GDSN Interface and data pools** (green box) via a double-headed arrow.
- Accessing Layer:** **EPCIS Accessing Applications and other Enterprise-level Applications** (blue box) connects to the **GDSN Interface** and the **EPCIS Query Interface** via double-headed arrows.
- Repository Layer:** **EPCIS Repository** (blue box) connects to the **EPCIS Accessing Applications** and the **EPCIS Query Interface** via double-headed arrows.
- Interface Layer:** **EPCIS Query Interface (Control and Callback)** (green box) connects to the **EPCIS Repository** and the **EPCIS Accessing Applications** via double-headed arrows.
- Capture Layer:** **EPCIS Capture Interface** (green box) connects to the **EPCIS Accessing Applications** and the **App-specific interfaces** via double-headed arrows.
- Infrastructure:** **App-specific interfaces** (green box) and **Data Capture Infrastructure** (grey box) form the base layer.

Central Flow:

- EPCIS Query traffic - on-demand queries and standing queries** flows from the **EPCIS Query Interface (Control and Callback)** of both trading partners to the **GS1 Networked Services** block.

In addition to the use of the EPCIS Query Interface as illustrated above, trading partners may by mutual agreement use the EPCIS Document structure defined in section 9.3 as a means to transport a collection of EPCIS events, optionally accompanied by relevant master data, as a single electronic document.

The following outlines the responsibilities of each element of the GS1 System Architecture as illustrated in the figures in the preceding sections. Further information may be found in [GS1Arch].

- *RFID and Barcode Readers* Make observations of RFID tags while they are in the read zone, and observations of barcodes when reading is triggered.
- *Low-Level [RFID] Reader Protocol (LLRP) Interface* Defines the control and delivery of raw RFID tag reads from RFID Readers to the Filtering & Collection role. Events at this interface say “Reader A saw EPC X at time T.”
- *Filtering & Collection* This role filters and collects raw RFID tag reads, over time intervals delimited by events defined by the EPCIS Capturing Application (e.g. tripping a motion detector). No comparable role typically exists for reading barcodes, because barcode readers typically only read a single barcode when triggered.
- *Filtering & Collection (ALE) Interface* Defines the control and delivery of filtered and collected RFID tag read data from the Filtering & Collection role to the Data Capture Workflow role. Events at this interface say “At Logical Reader L, between time T1 and T2, the following EPCs were observed,” where the list of EPCs has no

duplicates and has been filtered by criteria defined by the EPCIS Capturing Application. In the case of barcodes, comparable data is delivered to the Data Capture Workflow role directly from the barcode reader in the form of a GS1 Element String.

- *Data Capture Workflow* Supervises the operation of the lower-level architectural elements, and provides business context by coordinating with other sources of information involved in executing a particular step of a business process. The Data Capture Workflow may, for example, coordinate a conveyor system with Filtering & Collection events and barcode reads, may check for exceptional conditions and take corrective action (e.g., diverting a bad object into a rework area), may present information to a human operator, and so on. The Data Capture Workflow understands the business process step or steps during which EPCIS event data capture takes place. This role may be complex, involving the association of multiple Filtering & Collection events and/or barcode reads with one or more business events, as in the loading of a shipment. Or it may be straightforward, as in an inventory business process where there may be readers deployed that generate observations about objects that enter or leave the shelf. Here, the Filtering & Collection-level event or barcode read and the EPCIS-level event may be so similar that very little actual processing at the Data Capture Workflow level is necessary, and the Data Capture Workflow merely configures and routes events from the Filtering & Collection interface and/or barcode readers directly through the EPCIS Capture Interface to an EPCIS Repository or a business application. A Data Capture Workflow whose primary output consists of EPCIS events is called an “EPCIS Capturing Application” within this standard.
- *EPCIS Interfaces* The interfaces through which EPCIS data is delivered to enterprise-level roles, including EPCIS Repositories, EPCIS Accessing Applications, and data exchange with partners. Events at these interfaces say, for example, “At location X, at time T, the following contained objects (cases) were verified as being aggregated to the following containing object (pallet).” **There are three EPCIS Interfaces**, specified normatively in this document:
 - The **EPCIS Capture Interface** defines the delivery of EPCIS events from EPCIS Capturing Applications to other roles that consume the data in real time, including EPCIS Repositories, and real-time “push” to EPCIS Accessing Applications and trading partners.
 - The **EPCIS Query Control Interface** defines a means for EPCIS Accessing Applications and trading partners to obtain EPCIS data subsequent to capture, typically by interacting with an EPCIS Repository. The EPCIS Query Control Interface provides two modes of interaction. In “on-demand” or “synchronous” mode, a client makes a request through the EPCIS Query Control Interface and receives a response immediately. In “standing request” or “asynchronous” mode, a client establishes a subscription for a periodic query.
 - Each time the periodic query is executed, the results are delivered asynchronously (or “pushed”) to a recipient via the **EPCIS Query Callback Interface**. The EPCIS Query Callback Interface may also be used to deliver information immediately upon capture; this corresponds to the “possible bypass for real-time push” arrow in the diagram.
- *EPCIS Accessing Application*: Responsible for carrying out overall enterprise business processes, such as warehouse management, shipping and receiving, historical throughput analysis, and so forth, aided by EPCIS visibility event data.
- *EPCIS Repository*: Records EPCIS-level events generated by one or more EPCIS Capturing Applications and makes them available for later query by EPCIS Accessing Applications.
- *Partner Application*: Trading Partner systems that perform the same role as an EPCIS Accessing Application, though from outside the responding party’s network. Partner Applications may be granted access to a subset of the information that is available from an EPCIS Capturing Application or within an EPCIS Repository.

The interfaces within this stack are designed to insulate the higher levels of the architecture from unnecessary details of how the lower levels are implemented. One way to understand this is to consider what happens if certain changes are made:

- The *Low-Level [RFID] Reader Protocol (LLRP)* and *GS1 Element String* insulate the higher layers from knowing what RF protocols or barcode symbologies are in use, and what reader makes/models have been chosen. If a different reader is substituted, the information sent through these interfaces remains the same.
- In situations where RFID is used, the Filtering & Collection Interface insulates the higher layers from the physical design choices made regarding how RFID tags are sensed and accumulated, and how the time boundaries of events are triggered. If a single four-antenna RFID reader is replaced by a constellation of five single-antenna “smart antenna” readers, the events at the Filtering & Collection level remain the same. Likewise, if a different triggering mechanism is used to mark the start and end of the time interval over which reads are accumulated, the Filtering & Collection event remains the same.
- EPCIS insulates enterprise applications from understanding the details of how individual steps in a business process are carried out at a detailed level. For example, a typical EPCIS event is “At location X, at time T, the following cases were verified as being on the following pallet.” In a conveyor-based business implementation, this may correspond to a single Filtering & Collection event, in which reads are accumulated during a time interval whose start and end is triggered by the case crossing electric eyes surrounding a reader mounted on the conveyor. But another implementation could involve three strong people who move around the cases and use hand-held readers to read the tags. At the Filtering & Collection level, this looks very different (each triggering of the hand-held reader is likely a distinct Filtering & Collection event), and the processing done by the EPCIS Capturing Application is quite different (perhaps involving an interactive console that the people use to verify their work). But the EPCIS event is still the same for all these implementations.

In summary, EPCIS-level data differs from data employed at the Capture level in the GS1 System Architecture by incorporating semantic information about the business process in which data is collected, and providing historical observations. In doing so, EPCIS insulates applications that consume this information from knowing the low-level details of exactly how a given business process step is carried out.

3 EPCIS specification principles

The considerations in the previous two sections reveal that the requirements for standards at the EPCIS layer are considerably more complex than in the Capture layer of the GS1 System Architecture. The historical nature of EPCIS data implies that EPCIS interfaces need a richer set of access techniques than ALE or RFID and barcode reader interfaces. The incorporation of operational or business process context into EPCIS implies that EPCIS traffics in a richer set of data types, and moreover needs to be much more open to extension in order to accommodate the wide variety of business processes in the world. Finally, the diverse environment in which EPCIS operates implies that the EPCIS Standard be layered carefully so that even when EPCIS is used between external systems that differ widely in their details of operation, there is consistency and interoperability at the level of what the abstract structure of the data is and what the data means.

In response to these requirements, EPCIS is described by a framework specification and narrower, more detailed specifications that populate that framework. The framework is designed to be:

- *Layered*: In particular, the structure and meaning of data in an abstract sense is specified separately from the concrete details of data access services and bindings to particular interface protocols. This allows for variation in the concrete details over time and across enterprises while preserving a common meaning of the data itself. It also permits EPCIS data specifications to be reused in approaches other than the service-oriented approach of the present specification. For example, data definitions could be reused in an EDI framework.
- *Extensible*: The core specifications provide a core set of data types and operations, but also provide several means whereby the core set may be extended for purposes specific to a given industry or application area. Extensions not only provide for proprietary requirements to be addressed in a way that leverages as much of the standard framework as possible, but also provides a natural path for the standards to evolve and grow over time.
- *Modular*: The layering and extensibility mechanisms allow different parts of the complete EPCIS framework to be specified by different documents, while promoting coherence across the entire framework. This allows the process of standardisation (as well as of implementation) to scale.

The remainder of this document specifies the EPCIS framework. It also populates that framework with a core set of data types and data interfaces. The companion standard, the GS1 Core Business Vocabulary (CBV), provides additional data definitions that layer on top of what is provided by the EPCIS standard.

4 Terminology and typographical conventions

Within this specification, the terms SHALL, SHALL NOT, SHOULD, SHOULD NOT, MAY, CAN, and CANNOT are to be interpreted as specified in section 7 ("*Verbal forms for expressions of provisions*") of the ISO/IEC Directives, Part 2, 2018, 8th edition [ISODir2]. When used in this way, these terms will always be shown in ALL CAPS; when these words appear in ordinary typeface they are intended to have their ordinary English meaning.

All sections of this document, are normative, except where explicitly noted as non-normative.

The following typographical conventions are used throughout the document:

- ALL CAPS type is used for the special terms from [ISODir2] enumerated above.
- `Monospace` type is used to denote programming language, UML, XML, and JSON (or JSON-LD) identifiers, as well as for the text of XML and JSON (or JSON-LD) documents.


5 EPCIS specification framework

The EPCIS specification is designed to be layered, extensible, and modular.

5.1 Layers

The EPCIS specification framework is organised into several layers, as illustrated in the diagram and described below.

Figure 5-1 Layers of the EPCIS specification framework



- **Abstract Data Model Layer:** The Abstract Data Model Layer specifies the generic structure of EPCIS data. This is the only layer that is not extensible by mechanisms other than a revision to the EPCIS specification itself. The Abstract Data Model Layer specifies the general requirements for creating data definitions within the Data Definition Layer.
- **Data Definition Layer:** The Data Definition Layer specifies what data is exchanged through EPCIS, what its abstract structure is, and what it means. **One data definition module** is defined within the present specification, called **the Core Event Types Module**. Data definitions in the Data Definition Layer are specified abstractly, following rules defined by the Abstract Data Model Layer.
- **Service Layer:** The Service Layer defines service interfaces through which EPCIS clients interact. In the present specification, **two service layer modules** are defined.
 - The Core **Capture** Operations Module defines one service interface (the EPCIS **Capture Interface**) through which EPCIS Capturing Applications use to deliver Core Event Types to interested parties.
 - The Core **Query** Operations Module defines two service interfaces (the EPCIS **Query Control Interface** and the EPCIS **Query Callback Interface**) that

EPCIS Accessing Applications use to obtain data previously captured. Interface definitions in the Service Layer are specified abstractly using UML.

- *Bindings:* Bindings specify concrete realisations of the Data Definition Layer and the Service Layer. There may be many bindings defined for any given Data Definition or Service module. In this specification, a number of bindings are specified for the three modules defined in the Data Definition and Service Layers.
 - The data definitions in the **Core Event Types** data definition module are given bindings to XML schema, JSON schema and SHACL.
 - The EPCIS **Capture** Interface in the Core Capture Operations Module is given bindings for Message Queue, HTTP and REST.
 - The EPCIS **Query Control** Interface in the Core Query Operations Module is given a binding to SOAP over HTTP via a WSDL web services description, a binding for AS2, and a REST binding.
 - The EPCIS **Query Callback** Interface in the Core Query Operations Module is given bindings to HTTP, HTTPS, and AS2; the REST binding specifies the **WebSocket** to support subscriptions.
- *GS1 Core Business Vocabulary Standard:* The GS1 Core Business Vocabulary standard [CBV] is a companion to the EPCIS standard. It defines specific vocabulary elements that may be used to populate the data definitions specified in the Data Definition Layer of the EPCIS standard. While EPCIS may be used without CBV, by employing only private or proprietary data values, it is far more beneficial for EPCIS applications to make as much use of the CBV Standard as possible.

5.2 Extensibility

The layered technique for specification promotes extensibility, as one layer may be reused by more than one implementation in another layer. For example, while this specification includes an XML binding of the Core Event Types data definition module, another specification may define a binding of the same module to a different syntax, for example a CSV file.

Besides the extensibility inherent in layering, the EPCIS specification includes several specific mechanisms for extensibility:

- *Subclassing:* Data definitions in the Data Definition Layer are defined using UML, which allows a new data definition to be introduced by creating a subclass of an existing one. A subclass is a new type that includes all of the fields of an existing type, extending it with new fields. An instance of a subclass may be used in any context in which an instance of the parent class is expected.
- *Extension Points:* Data definitions and service specifications also include extension points, which vendors may use to provide extended functionality without creating subclasses.

5.3 Modularity

The EPCIS specification framework is designed to be modular. That is, it does not consist of a single specification, but rather a collection of individual specifications that are interrelated. This allows EPCIS to grow and evolve in a distributed fashion. The layered structure and the extension mechanisms provide the essential ingredients to achieving modularity, as does the grouping into modules.

While EPCIS specifications are modular, there is no requirement that the module boundaries of the specifications be visible or explicit within *implementations* of EPCIS. For example, there may be a particular software product that provides a SOAP/HTTP-based implementation of a case-to-pallet association service and a product catalogue service that traffics in data defined in the relevant data definition modules. This product may conform to as many as six different modules from the EPCIS standard: the data definition module that describes product catalogue data, the data definition module that defines case-to-pallet associations, the specifications for the respective services, and the respective SOAP/HTTP bindings. But the source code of the product may have no trace of these boundaries, and indeed the concrete database schema used by the product may denormalise the data so that product catalogue and case-to-

pallet association data are inextricably entwined. But as long as the net result conforms to the specifications, this implementation is permitted.

6 Abstract data model layer

This section gives a normative description of the abstract data model that underlies EPCIS.

6.1 Event data and master data


Generically, EPCIS deals in two kinds of data: event data and master data. Event data arises in the course of carrying out business processes, and is captured through the EPCIS Capture Interface and made available for query through the EPCIS Query Interfaces. Master data is additional data that provides the necessary context for interpreting the event data. It is available for query through the EPCIS Query Control Interface, but the means by which master data enters the system is not specified in the EPCIS standard.

The Abstract Data Model Layer does not attempt to define the meaning of the terms “event data” or “master data,” other than to provide precise definitions of the structure of the data as used by the EPCIS specification. The modelling of real-world business information as event data and master data is the responsibility of the Data Definition Layer, and of industry and end-user agreements that build on top of this specification.



Non-Normative: Explanation: While for the purposes of this specification the terms “event data” and “master data” mean nothing more than “data that fits the structure provided here,” the structures defined in the Abstract Data Model Layer are designed to provide an appropriate representation for data commonly requiring exchange through EPCIS. Informally, these two types of data may be understood as follows. Event data grows in quantity as more business is transacted, and refers to things that happen at specific moments in time. An example of event data is “At 1:23pm on 15 March 2004, EPC X was observed at Location L.” Master data does not generally grow merely because more business is transacted (though master data does tend to grow as organisations grow in size), is not typically tied to specific moments in time (though master data may change slowly over time), and provides interpretation for elements of event data. An example of master data is “Location L refers to the distribution centre located at 123 Elm Street, Anytown, US.” All of the data in the set of use cases considered in the creation of the EPCIS standard can be modelled as a combination of event data and master data of this kind.

Figure 6-1 Structure of event data and master data in EPCIS



The structure of event data and master data in EPCIS is illustrated above. (Note that this is an illustration only: the specific vocabulary elements and master data attribute names in this figure are not defined within this specification.)

The ingredients of the EPCIS Abstract Data Model are defined below:

- **Event Data:** A set of Events.
- **Event:** A structure consisting of an Event Type and one or more named Event Fields.
- **Event Type:** A namespace-qualified name (qname) that indicates to which of several possible Event structures (as defined by the Data Definition Layer) a given event conforms.
- **Event Field:** A named field within an Event. The name of the field is given by a qname, referring either to a field name specified by the Data Definition Layer or a field name defined as an extension to this specification. The value of the field may be a primitive type (such as an integer or timestamp), a Vocabulary Element, or a list of primitive types or Vocabulary Elements.
- **Master data:** A set of Vocabularies, together with master data attributes associated with elements of those Vocabularies.
- **Vocabulary:** A named set of identifiers. The name of a Vocabulary is a qname that may be used as a type name for an event field. The identifiers within a Vocabulary are called Vocabulary Elements. A Vocabulary represents a set of alternative values that may appear as the values of specific Event Fields. Vocabularies in EPCIS are used to model sets such as the set of available location names, the set of available business process step names, and so on.
- **Vocabulary Element:** An identifier that names one of the alternatives modelled by a Vocabulary. The value of an Event Field may be a Vocabulary Element. Vocabulary Elements are represented as Uniform Resource Identifiers (URIs). Each Vocabulary Element may have associated master data attributes.
- **Master data attributes:** An unordered set of name/value pairs associated with an individual Vocabulary Element. The name part of a pair is a qname. The value part of a pair may be a value of arbitrary type. A special attribute is a (possibly empty) list of children, each child being another vocabulary element from the same vocabulary. See [section 6.5](#).

New EPCIS Events are generated at the edge and delivered into EPCIS infrastructure through the EPCIS Capture Interface, where they can subsequently be delivered to interested applications through the EPCIS Query Interfaces. There is no mechanism provided in either interface by which an application can delete or modify an EPCIS Event. The only way to “retract” or “correct” an EPCIS Event is to generate a

subsequent event whose business meaning is to rescind or amend the effect of a prior event (section [7.4.1.2.1](#) discusses how this may be done).

While the EPCIS Capture Interface and EPCIS Query Interfaces provide no means for an application to explicitly request the deletion of an event, EPCIS Repositories MAY implement data retention policies that cause old EPCIS events to become inaccessible after some period of time.

Master data, in contrast, may change over time, though such changes are expected to be infrequent relative to the rate at which new event data is generated. The current version of this specification does not specify how master data changes (nor, as noted above, does it specify how master data is entered in the first place).

6.1.1 Transmission of master data in EPCIS

The EPCIS Capture and Query Interfaces are primarily concerned with the transmission of EPCIS Events. The means by which master data enters a system that implements these interfaces is not specified in the EPCIS standard. However, the EPCIS standard does provide mechanisms for transmission of master data, which an implementation may use to ensure that the recipient of EPCIS event data has access to the master data necessary to interpret that event data. Alternatively, master data may be transmitted by means entirely outside the EPCIS standard. The EPCIS standard does not impose any requirements on whether EPCIS event data is accompanied by master data or not, other than to require that master data accompanying event data be consistent with any master data in ILMD sections of those events.

The EPCIS standard provides two mechanisms for transmission of master data, summarised in the table below:

Table 6-1 Transmission of master data in EPCIS

Mechanism	Section	Description	Constraint
ILMD	7.3.8	An EPCIS event that marks the beginning of life for an instance-level or lot-level identifier may include corresponding master data directly in the event.	The master data in the event SHALL reflect the current values of master data attributes, as known to the event creator, as of the event time. Note that because this data is embedded directly in the event, it is permanently a part of that event and will always be included when this event is queried for (subject to redaction as specified in section 8.2.2).
Header of XML EPCIS document	9.5	An EPCIS document used for point-to-point transmission of a collection of EPCIS events outside of the EPCIS Query Interface may include relevant master data in the document header.	The master data in the document header SHALL reflect the current values of master data attributes, as known to the document creator, as of the time the document is created. Master data in the header of an EPCIS document SHALL NOT specify attribute values that conflict with the ILMD section of any event contained within the EPCIS document body.

6.2 Standard vocabulary and user vocabulary

Before version 2.0, EPCIS / CBV have previously used the terms 'vocabulary type' to refer to standard EPCIS event data fields whose value is expected to be populated either using one of the enumerated values of the corresponding 'standard vocabulary' defined within CBV or with a value 'user vocabulary' that might be defined by a single organisation, solution provider, industry sector or other consortium.

Starting with version 2.0, EPCIS and CBV add support for Linked Data (in particular, a JSON-LD representation and formal Linked Data ontologies for the data models). A Linked Data ontology is sometimes also referred to as a vocabulary. However, when specialists in ontologies, Linked Data or Semantic Web talk about a vocabulary, they are not only referring to enumerated code lists; a Linked Data vocabulary can also define classes, properties and identifiers in addition to code lists of enumerated

values. The table below provides a mapping of the historical terminology used within EPCIS/CBV versus the terminology used by ontologists, in which 'vocabulary' has a much broader meaning than 'enumeration'.

Table 6-2 Mapping of terminology used in EPCIS/CBV and by ontologists

Example	EPCIS/CBV describes as	Ontological description (A Linked Data Vocabulary or Web Vocabulary can define anything below)
ObjectEvent	Event Type	Class
eventTime field	A field	Datatype Property (expects a non-URI literal value)
epcList field	A field	Object Property (expecting a URI)
bizStep field	A 'Vocabulary Type'	Object Property (expecting a URI)
'shipping'	An individual value within CBV BizStep Vocabulary	Individual within a Class (value within a code list enumeration)
['shipping', 'receiving', ...]	a CBV Vocabulary, 'BizStep'	Class containing individuals, each individual having a definition, the class representing a code list

Vocabularies are used extensively within EPCIS to model physical, digital, and conceptual entities that exist in the real world. Examples of vocabularies defined in the core EPCIS Data Definition Layer are location names, object class names (an object class name is something like "Acme Deluxe Widget," as opposed to an EPC which names a specific instance of an Acme Deluxe Widget), and business step names. In each case, a vocabulary represents a finite (though open-ended) set of alternatives that may appear in specific fields of events.

It is useful to distinguish two kinds of vocabularies, which follow different patterns in the way they are defined and extended over time:

- **Standard Vocabulary:** A Standard Vocabulary represents a set of Vocabulary Elements whose definition and meaning must be agreed to in advance by trading partners who will exchange events using the vocabulary. For example, the EPCIS Core Data Definition Layer defines a vocabulary called "business step," whose elements are identifiers denoting such things as "shipping," "receiving," and so on. One trading partner may generate an event having a business step of "shipping," and another partner receiving that event through a query can interpret it because of a prior agreement as to what "shipping" means.

Standard Vocabulary elements are defined in the CBV, as well as in more application-specific GS1 application standards, having been developed and vetted by user-driven GS1 workgroups. The master data associated with Standard Vocabulary elements are defined by those same organisations and tend to be distributed to users as part of a specification or by some similar means. New vocabulary elements within a given Standard Vocabulary are introduced through a very deliberate process, such as the ratification of a new version of a standard or through a vote of a GS1 workgroup. While an individual end user organisation acting alone may introduce a new Standard Vocabulary element, such an element would have limited use in a data exchange setting and would probably only be used within an organisation's four walls.

- **User Vocabulary:** A User Vocabulary represents a set of Vocabulary Elements whose definition and meaning are under the control of a single organisation or consortium. For example, the EPCIS Core Data Definition Layer defines a vocabulary called "business location," whose elements are identifiers denoting such things as "Acme Corp. Distribution Centre #3." Acme Corp may generate an event having a business location of "Acme Corp. Distribution Centre #3," and another partner receiving that event through a query can interpret it either because it correlates it with other events naming the same location, or by looking at master data attributes associated with the location, or both.

User Vocabulary elements are primarily defined by users or consortia. New vocabulary elements within a given User Vocabulary are introduced at the sole discretion of these parties, and trading partners must be prepared to respond accordingly. Usually, however, the rules for constructing new User Vocabulary Elements are established by organisations of multiple end users, and in any case must follow the rules defined in section [6.4](#) below.

The lines between these two kinds of vocabularies are subjective. Because the mechanisms defined in the EPCIS specification make no distinction between the two vocabulary types, it is never necessary to identify a particular vocabulary as belonging to one type or the other. The terms “Standard Vocabulary” and “User Vocabulary” are introduced only because they are useful as a hint as to the way a given vocabulary is expected to be defined and extended.

The CBV provides standardised vocabulary elements for many of the vocabulary types used in EPCIS event types. In particular, the CBV defines vocabulary elements for the following EPCIS Standard Vocabulary types: Business Step, Disposition, Persistent Disposition, Business Transaction Type, Source/Destination Type, Error Reason, and Sensor Property Type. The CBV also defines templates for constructing vocabulary elements for the following EPCIS User Vocabulary types: Object (EPC), Object Class (EPCClass), Location (Read Point and Business Location), Business Transaction ID, Source/Destination ID, Transformation ID, Event ID, Chemical Substance ID, Microorganism ID, and Resource ID.

6.3 Extension mechanisms

A key feature of EPCIS is its ability to be extended by different organisations to adapt to particular business situations. In all, the Abstract Data Model Layer provides five methods by which the data processed by EPCIS may be extended (the Service Layer, in addition, provides mechanisms for adding additional services), enumerated here from the most invasive type of extension to the least invasive:

- *New Event Type:* A new Event Type may be added in the Data Definition Layer. Adding a new Event Type requires each of the Data Definition Bindings to be extended, and may also require extension to the Capture and Query Interfaces and their Bindings.
- *New Event Field:* A new field may be added to an existing Event Type in the Data Definition Layer. The bindings, capture interface, and query interfaces defined in this specification are designed to permit this type of extension without requiring changes to the specification itself. (The same may not be true of other bindings or query languages defined outside this specification.)
- *New Vocabulary Type:* A new Vocabulary Type may be added to the repertoire of available Vocabulary Types. No change to bindings or interfaces are required.
- *New master data attribute:* A new attribute name may be defined for an existing Vocabulary. No change to bindings or interfaces are required.
- *New Instance/Lot master data (ILMD) Attribute:* A new attribute name may be defined for use in Instance/Lot master data (ILMD); see section [7.3.8](#). No change to bindings or interfaces are required.
- *New Vocabulary Element* A new element may be added to an existing Vocabulary.

The Abstract Data Model Layer has been designed so that most extensions arising from adoption by different industries or increased understanding within a given industry can be accommodated by the latter methods in the above list, which do not require revision to the specification itself. The more invasive methods at the head of the list are available, however, in case a situation arises that cannot be accommodated by the latter methods.

It is expected that there will be several different ways to extend the EPCIS specification, as summarised below:

Table 6-3 Ways to extend the EPCIS specification

How extension is disseminated	Responsible organisation	Extension method				
		New Event Type	New Event Field	New Vocabulary Type	New master data or ILMD (section 7.3.8) Attribute	New Vocabulary Element
New Version of EPCIS standard	GS1 EPCIS/CBV MSWG	Yes	Yes	Yes	Occasionally	Rarely
New Version of CBV standard	GS1 EPCIS/CBV MSWG	No	No	No	Yes	Yes (Standard Vocabulary, User Vocabulary template)
GS1 Application Standard for a specific industry	GS1 Application Standard Working Group for a specific industry	Rarely	Rarely	Occasionally	Yes	Yes (Standard Vocabulary)
GS1 Member Organisation Local Recommendation Document for a specific industry within a specific geography	GS1 Member Organisation	Rarely	Rarely	Occasionally	Yes	Yes (Standard Vocabulary)
Private Group Interoperability Specification	Industry Consortium or Private End User Group outside GS1	Rarely	Rarely	Occasionally	Yes	Yes (Standard Vocabulary)
Updated master data via EPCIS Query or other data sync	Individual End User	Rarely	Rarely	Rarely	Rarely	Yes (User vocabulary)

6.4 Identifier representation

The Abstract Data Model Layer introduces several kinds of identifiers, including Event Type names, Event Field names, Vocabulary names, Vocabulary Elements, and master data Attribute Names. Because all of these namespaces are open to extension, this specification imposes some rules on the construction of these names so that independent organisations may create extensions without fear of name collision.

Vocabulary Elements are subject to the following rules. In all cases, a Vocabulary Element is represented as Uniform Resource Identifier (URI) whose general syntax is defined in [RFC3986]. The types of URIs admissible as Vocabulary Elements are those URIs for which there is an owning authority. This includes:

- URI representations for EPC codes [TDS]. The owning authority for a particular EPC URI is the organisation to whom the GS1 Company Prefix (or other issuing authority, depending on the EPC scheme) was assigned.
- Absolute Uniform Resource Locators (URLs) [RFC1738]. The owning authority for a particular URL is the organisation that owns the Internet domain name in the authority portion of the URL.

- Uniform Resource Names (URNs) [RFC8141] in the `oid` namespace that begin with a Private Enterprise Number (PEN). The owning authority for an OID-URN is the organisation to which the PEN was issued.
- Uniform Resource Names (URNs) [RFC8141] in the `epc` or `epcglobal` namespace, other than URIs used to represent EPCs [TDS]. The owning authority for these URNs is GS1.
- GS1 Digital Link URIs in the form normatively specified in the [GS1 Digital Link standard](#) [GS1DL], restricted to a highly constrained set of GS1 Digital Link URIs that corresponds to each of the EPC Pure Identity URI schemes defined in TDS. (See section 8 of CBV for details.) The canonical form of GS1 Digital Link URIs is recommended but optional.

Event Type names and Event Field names are represented as namespace-qualified names (qnames), consisting of a namespace URI and a name. This has a straightforward representation in XML bindings that is convenient for extension.

6.5 Hierarchical vocabularies

Some Vocabularies have a hierarchical or multi-hierarchical structure. For example, a vocabulary of location names may have an element that means "Acme Corp. Retail Store #3" as well others that mean "Acme Corp. Retail Store #3 Backroom" and "Acme Corp. Retail Store #3 Sales Floor." In this example, there is a natural hierarchical relationship in which the first identifier is the parent and the latter two identifiers are children.

Hierarchical relationships between vocabulary elements are represented through master data. Specifically, a parent identifier carries, in addition to its master data attributes, a list of its children identifiers. Each child identifier SHALL belong to the same Vocabulary as the parent. In the example above, the element meaning "Acme Corp. Distribution Centre #3" would have a children list including the element that means "Acme Corp. Distribution Centre #3 Door #5."

Elsewhere in this specification, the term "direct or indirect descendant" is used to refer to the set of vocabulary elements including the children of a given vocabulary element, the children of those children, etc. In other words, the "direct or indirect descendants" of a vocabulary element include both the vocabulary elements from the next lower hierarchical level, as well as vocabulary elements in subsequent lower hierarchical levels.

A given element MAY be the child of more than one parent. This allows for more than one way of grouping vocabulary elements; for example, locations could be grouped both by geography and by function. An element SHALL NOT, however, be a child of itself, either directly or indirectly.

- ✓ **Non-Normative:** Explanation: In the present version of this specification, only one hierarchical relationship is provided for, namely the relationship encoded in the special "children" list. Future versions of this specification may generalise this to allow more than one relationship, perhaps encoding each relationship via a different master data attribute.

Hierarchical relationships are given special treatment in queries (section 8.2), and may play a role in carrying out authorisation policies (section 8.2.2), but do not otherwise add any additional complexity or mechanism to the Abstract Data Model Layer.

The *GS1 GLN Allocation Rules Standard* [GLNAR3.0] provide a normative framework for identification of parties and locations.

7 Data definition layer

This section includes normative specifications of modules in the Data Definition Layer.

7.1 General rules for specifying data definition layer modules

The general rules for specifying modules in the Data Definition Layer are given here. These rules are then applied in section 7.2 to define the Core Event Types Module. These rules can also be applied by organisations wishing to layer a specification on top of this specification.

7.1.1 Content

In general, a Data Definition Module specification has these components, which populate the Abstract Data Model framework specified in section 6:


- **Value Types:** Definitions of data types that are used to describe the values of Event Fields and of master data attributes. The Core Event Types Module defines the primitive types that are available for use by all Data Definition Modules. Each Vocabulary that is defined is also implicitly a Value Type.
 - **Event Types:** Definitions of Event Types, each definition giving the name of the Event Type (which must be unique across all Event Types) and a list of standard Event Fields for that type. An Event Type may be defined as a subclass of an existing Event Type, meaning that the new Event Type includes all Event Fields of the existing Event Type plus any additional Event Fields provided as part of its specification.
 - **Event Fields:** Definitions of Event Fields within Event Types. Each Event Field definition specifies a name for the field (which must be unique across all fields of the enclosing Event Type) and the data type for values in that field. Event Field definitions within a Data Definition Module may be part of new Event Types introduced by that Module, or may extend Event Types defined in other Modules.
 - **Vocabulary Types:** Definitions of Vocabulary Types, each definition giving the name of the Vocabulary (which must be unique across all Vocabularies), a list of standard master data attributes for elements of that Vocabulary, and rules for constructing new Vocabulary Elements for that Vocabulary. (Any rules specified for constructing Vocabulary Elements in a Vocabulary Type must be consistent with the general rules given in section 6.4)
 - **Master data attributes:** Definitions of master data attributes for Vocabulary Types. Each master data attribute definition specifies a name for the Attribute (which must be unique across all attributes of the enclosing Vocabulary Type) and the data type for values of that attribute. Master data definitions within a Data Definition Module may belong to new Vocabulary Types introduced by that Module, or may extend Vocabulary Types defined in other Modules.
 - **Vocabulary Elements:** Definitions of Vocabulary Elements, each definition specifying a name (which must be unique across all elements within the Vocabulary, and conform to the general rules for Vocabulary Elements given in section 6.4 as well as any specific rules specified in the definition of the Vocabulary Type), and optionally specifying master data (specific attribute values) for that element.
- ✓ **Non-Normative:** Amplification: As explained in section 6.3, Data Definition Modules defined in this specification and by companion specifications developed by the EPCIS Working Group will tend to include definitions of Value Types, Event Types, Event Fields, and Vocabulary Types, while modules defined by other groups will tend to include definitions of Event Fields that extend existing Event Types, master data attributes that extend existing Vocabulary Types, and Vocabulary Elements that populate existing Vocabularies. Other groups may also occasionally define Vocabulary Types.

The word “Vocabulary” is used informally to refer to a Vocabulary Type and the set of all Vocabulary Elements that populate it.

7.1.2 Notation

In the sections below, Event Type and Event fields are specified using a restricted form of UML class diagram notation. UML class diagrams used for this purpose may contain classes that have attributes (fields) and associations, but not operations.

Figure 7-1 EPCIS data definition notation



The example UML diagram above shows a data definition for two Event Types, EventType1 and EventType2. These event types make use of four Value Types: Type1, Type2, DataClass3, and DataClass4. Type1 and Type2 are primitive types, while DataClass3 and DataClass4 are complex types whose structure is also specified in UML.

The Event Type 1 in this example has four fields. Field1 and Field2 are of primitive type Type1 and Type2 respectively. 1 has another field Field3 whose type is DataClass3. Finally, 1 has another field Field4 that contains a list of zero or more instances of type DataClass4 (the “0..*” notation indicates “zero or more”).

This diagram also shows a data definition for EventType2. The arrow with the open-triangle arrowhead indicates that EventType2 is a subclass of EventType1. This means that EventType2 actually has five fields: four fields inherited from EventType1 plus a fifth field5 of type Type1.

Within the UML descriptions, the notation <<extension point>> identifies a place where implementations SHALL provide for extensibility through the addition of new data members. (When one type has an extension point, and another type is defined as a subclass of the first type and also has an extension point, it does not mean the second type has two extension points; rather, it merely emphasises that the second type is also open to extension.) Extensibility mechanisms SHALL provide for both proprietary extensions by vendors of EPCIS-compliant products, and for extensions defined by GS1 through future versions of this specification or through new specifications.

In the case of the standard XML bindings, the extension points are implemented within the XML schema following the methodology described in section [9.1](#).

All definitions of Event Types SHALL include an extension point, to provide for the extensibility defined in section [6.3](#) (“New Event Fields”). Value Types MAY include an extension point.

7.1.3 Semantics

Each event (an instance of an Event Type) encodes several assertions which collectively define the semantics of the event. Some of these assertions say what was true at the time the event was captured. Other assertions say what is expected to be true following the event, until invalidated by a subsequent event. These are called, respectively, the *retrospective semantics* and the *prospective semantics* of the event.

For example, if widget #23 enters building #5 through door #6 at 11:23pm, then one retrospective assertion is that “widget #23 was observed at door #6 at 11:23pm,”, while a prospective assertion is that “widget #23 is in building #5.” The key difference is that the retrospective assertion refers to a specific time in the past (“widget #23 *was observed...*”), while the prospective assertion is a statement about the present condition of the object (“widget #23 *is in...*”). The prospective assertion presumes that if widget #23 ever leaves building #5, another EPCIS capture event will be recorded to supersede the prior one.

In general, retrospective semantics are things that were incontrovertibly known to be true at the time of event capture, and can usually be relied upon by EPCIS Accessing Applications as accurate statements of historical fact. Prospective semantics, since they attempt to say what is true after an event has taken place, must be considered at best to be statements of “what ought to be” rather than of “what is.” A prospective assertion may turn out not to be true if the capturing apparatus does not function perfectly, or if the business process or system architecture were not designed to capture EPCIS events in all circumstances. Moreover, in order to make use of a prospective assertion implicit in an event, an EPCIS Accessing Application must be sure that it has access to any subsequent event that might supersede the event in question.

The retrospective/prospective dichotomy plays an important role in EPCIS’s definition of location, in section [7.3.4](#).

In certain situations, an earlier event is subsequently discovered to be in error (the assertions its semantics makes are discovered to be incorrect), and the error cannot be corrected by recording a new event that adds additional assertions through its own semantics. For these cases, a mechanism is provided to record an event whose semantics assert that the assertions previously made by the erroneous event are in error. See section [7.4.1.2](#).

7.2 Core event types module – overview

The Core Event Types data definition module specifies the Event Types that represent EPCIS data capture events. These events are typically generated by an EPCIS Capturing Application and provided to EPCIS infrastructure using the data capture operations defined in section 8.1. These events are also returned in response to query operations that retrieve events according to query criteria.

The components of this module, following the outline given in section 7.1.1, are as follows:

- *Value Types*: Primitive types defined in section 7.3.1 and 7.3.2.
- *Event Types*: Event types as shown in the UML diagram in section 7.2.1, and defined in section 7.4.
- *Event Fields*: Included as part of the Event Types defined in section 7.4.
- *Vocabulary Types*: Types defined in section 7.3.3 through 7.3.8, and summarised in section 7.1.1.
- *Master data attributes*: Included as part of Vocabulary Types definitions. It is expected that industry vertical working groups will define additional master data attributes for the vocabularies defined here.
- *Vocabulary Elements*: None provided as part of this specification. It is expected that industry vertical working groups will define vocabulary elements for the *BusinessStep* vocabulary (section 7.3.6), the *Disposition* vocabulary (section 7.3.6.2), and the *BusinessTransactionType* vocabulary (section 7.3.6.3.1).

This module defines six event types, one very generic event and five subclasses that can represent events arising from supply chain activity across a wide variety of industries:


- *EPCISEvent* (section 7.4.1) is a generic base class for all event types in this module as well as others.
- *ObjectEvent* (section 7.4.2) represents an event that happened to one or more physical or digital objects.
- *AggregationEvent* (section 7.4.3) represents an event that happened to one or more objects that are physically aggregated together (physically constrained to be in the same place at the same time, as when cases are aggregated to a pallet).
- *TransactionEvent* (section 7.4.4) represents an event in which one or more objects become associated or disassociated with one or more identified business transactions.
- *TransformationEvent* (section 7.4.5) represents an event in which input objects are fully or partially consumed and output objects are produced, such that any of the input objects may have contributed to all of the output objects.
- *AssociationEvent* (section 7.4.6) is similar to an *AggregationEvent*, but allows for associations of objects with physical locations, and is especially suited to capture parent-child relationships that persist even after more temporarily linked children are disassociated from the parent.

A UML diagram showing these Event Types is depicted in the following section.

7.2.1 UML Diagrams of EPCIS Event Types


7.2.1.1 UML with Ontology focus

Figure 7-2 EPCIS UML with Ontology focus



7.2.1.2 UML with Syntax focus

Figure 7-3 EPCIS UML with Syntax focus



7.2.1.3 UML for SBDH



Note: The UML class diagrams for SBDH, which are not central to the EPCIS event data model, are included in section 15.

7.2.2 Overview of EPCIS event "dimensions" (non-normative)

Each of the core event types (not counting the generic `EPCISEvent`) has fields that represent five key dimensions of any EPCIS event. These five dimensions are: (1) the object(s) or other entities that are the subject of the event; (2) the date and time; (3) the location at which the event occurred and/or the whereabouts of the objects subsequent to the event; (4) the business context; (5) the condition of the objects that are the subject of the event. These five dimensions may be conveniently remembered as "**what, when, where, why** and **how**" (respectively). The "what" dimension varies depending on the event type (e.g., for an `ObjectEvent` the "what" dimension is one or more EPCs; for an `AggregationEvent` the "what" dimension is a parent ID and list of child EPCs). The "where" and "why" dimensions have both a retrospective aspect and a prospective aspect (see section 7.1.3), represented by different fields. The "how" dimension is populated on the basis of captured sensor data.

The following table summarises the fields of the event types that pertain to the five key dimensions;

in addition to the fields belonging to the five key dimensions, events may carry additional descriptive information in other fields, also indicated in the table below.

Dimension in EPCIS/CBV 1.x	Categorisation in EPCIS/CBV 2.0	Field	EPCIS section in which the field is defined	CBV section in which its value range is specified
WHAT	Objects in Focus (WHAT) 	I n s t a n c e	epcList	7.4.2 ObjectEvent 7.4.4 TransactionEvent <i>EPC Tag Data Standard (TDS) section 6, "EPC URI"</i>
			parentID	7.4.3 AggregationEvent 7.4.4 TransactionEvent 7.4.6 AssociationEvent 8.2 Physical or Digital Objects (Instance)
			childEPCs	7.4.3 AggregationEvent 7.4.6 AssociationEvent
			inputEPCList	7.4.5 TransformationEvent
			outputEPCList	
		C l a s s	quantityList	7.4.2 ObjectEvent 7.4.4 TransactionEvent <i>EPC Tag Data Standard (TDS) section 8, "URIs for EPC Pure Identity Patterns"</i>
			childQuantityList	7.4.3 AggregationEvent 7.4.6 AssociationEvent 8.3 Physical or Digital Objects (Class)
			inputQuantityList	7.4.5 TransformationEvent
			outputQuantityList	

Dimension in EPCIS/CBV 1.x	Categorisation in EPCIS/CBV 2.0	Field	EPCIS section in which the field is defined	CBV section in which its value range is specified
WHEN	Chronology (WHEN) 	eventTime	7.4.1 EPCISEvent	
		eventTimeZoneOffset		
		recordTime	7.4.1 EPCISEvent	
WHERE	Whereabouts (WHERE) 	readPoint	7.4.2 ObjectEvent	8.4 Locations
		bizLocation	7.4.3 AggregationEvent 7.4.4 TransactionEvent 7.4.5 TransformationEvent 7.4.6 AssociationEvent	
n/a	Condition (HOW) 	sensorElementList	7.4.2 ObjectEvent 7.4.3 AggregationEvent 7.4.4 TransactionEvent 7.4.5 TransformationEvent 7.4.6 AssociationEvent	7.6 Sensor Measurement Types 8.9 Chemical substance identifiers 8.10 Microorganism identifiers
WHY	Business Context (WHY) 	bizStep	7.4.2 ObjectEvent 7.4.3 AggregationEvent	7.1 Business Steps
		bizTransactionList	7.4.4 TransactionEvent	8.5 Business Transactions
		disposition	7.4.5 TransformationEvent	7.1 Dispositions

Dimension in EPCIS/CBV 1.x	Categorisation in EPCIS/CBV 2.0	Field	EPCIS section in which the field is defined	CBV section in which its value range is specified
		persistent Disposition	7.4.6 AssociationEvent	7.1 Dispositions
		sourceList		8.6 Source/Destination Identifiers
		destinationList		
	Other fields	ilmd	7.3.7 Instance/lot master data (ILMD)	9 Trade Item Master Data
(core field)		action	7.3.2 Action type	
(transformationID)		transformationID	7.4.5 TransformationEvent	8.7 Transformation Identifiers
(core field)		eventID	7.4.1 EPCISEvent	8.8 Event Identifiers
(core field)		errorDeclaration	7.4.1 EPCISEvent	7.5 Error Reason Identifiers 8.8 Event Identifiers

It is expected that the majority of additional descriptive information fields will be defined by industry-specific specifications layered on top of the core EPCIS and CBV standards.

7.2.3 Table of vocabulary types

The following table summarises the vocabulary types defined in this module. The URI column gives the formal name for the vocabulary used when the vocabulary must be referred to by name across the EPCIS interface.

Vocabulary type	Section	User or standard vocabulary	URI
ReadPointID	7.3.5	User	urn:epcglobal:epcis:vtype:ReadPoint
BusinessLocationID	7.3.5	User	urn:epcglobal:epcis:vtype:BusinessLocation
BusinessStepID	7.3.6.1	Standard	urn:epcglobal:epcis:vtype:BusinessStep
DispositionID	7.3.6.2	Standard	urn:epcglobal:epcis:vtype:Disposition
BusinessTransaction	7.3.6.3	User	urn:epcglobal:epcis:vtype:BusinessTransaction
BusinessTransactionTypeID	7.3.6.3.1	Standard	urn:epcglobal:epcis:vtype:BusinessTransactionType
EPCClass	7.3.3.1.2	User	urn:epcglobal:epcis:vtype:EPCClass
SourceDestTypeID	7.3.6.4.1	Standard	urn:epcglobal:epcis:vtype:SourceDestType
SourceDestID	7.3.6.4.2	User	urn:epcglobal:epcis:vtype:SourceDest
LocationID	7.3.5	User	urn:epcglobal:epcis:vtype:Location
PartyID	7.3.6.4.2	User	urn:epcglobal:epcis:vtype:Party
ErrorReasonID	7.4.1.2	Standard	urn:epcglobal:epcis:vtype:ErrorReason
SensorPropertyTypeID	7.3.7.1.4	Standard	urn:epcglobal:epcis:vtype:SensorPropertyType
MicroorganismID	7.3.7.1.6	User	urn:epcglobal:epcis:vtype:Microorganism
ChemicalSubstanceID	7.3.7.1.7	User	urn:epcglobal:epcis:vtype:ChemicalSubstance
ResourceID	7.3.7.1.8	User	urn:epcglobal:epcis:vtype:Resource

7.3 Core event types module – building blocks

This section specifies the building blocks for the event types defined in section [7.4](#).

7.3.1 Primitive types

The following primitive types are used within the Core Event Types Module.

Type	Description
xsd:int	An integer. Range restrictions are noted where applicable.

Type	Description
xsd:dateTimeStamp	A timestamp, giving the date and time in a time zone-independent manner. For bindings in which fields of this type are represented textually, an ISO 8601 [ISO8601] compliant representation SHOULD be used.
EPC	An Electronic Product Code Pure Identity URI, as defined in [TDS], or a GS1 Digital Link URI, as provided for in CBV section 8.
xsd:float	Patterned after the IEEE single-precision 32-bit floating point datatype; often used to approximate arbitrary real numbers. The <i>·value space·</i> of <i>float</i> contains the non-zero numbers $m \times 2^e$, where <i>m</i> is an integer whose absolute value is less than 224, and <i>e</i> is an integer between -149 and 104, inclusive.
xsd:double	The double-precision floating point datatype is patterned after the IEEE double-precision 64-bit floating point datatype, specified in IEEE 754-2008 [IEEE754]. Each double-precision floating point datatype has a value space that is a subset of the rational numbers. Floating point numbers are often used to approximate arbitrary real numbers.
xsd:boolean	Represents the values of two-valued logic; has the <i>·value space·</i> of two-valued logic: {true, false}.
xsd:hexBinary	Represents arbitrary hex-encoded binary data; the <i>·value space·</i> of <i>hexBinary</i> is the set of finite-length sequences of zero or more binary octets.
xsd:anyURI	Represents an Internationalized Resource Identifier Reference (IRI); an <i>anyURI</i> value can be absolute or relative.
xsd:decimal	Represents a subset of the real numbers, which can be represented by decimal numerals. The <i>·value space·</i> of <i>decimal</i> is the set of numbers that can be obtained by dividing an integer by a non-negative power of ten, i.e., expressible as $i / 10^n$ where <i>i</i> and <i>n</i> are integers and $n \geq 0$. Precision is not reflected in this value space; the number 2.0 is not distinct from the number 2.00.
xsd:string	Represents character strings.

The EPC type is defined as a primitive type for use in events when referring to EPCs that are not part of a Vocabulary Type. For example, an SGTIN EPC used to denote an instance of a trade item in the *epcList* field of an *ObjectEvent* is an instance of the EPC primitive type. But an SGLN EPC used as a read point identifier (section 7.3.4) in the *ReadPoint* field of an *ObjectEvent* is a Vocabulary Element, not an instance of the EPC primitive type.



Non-Normative: Explanation: This reflects a design decision not to consider individual trade item instances as Vocabulary Elements having master data, because trade item instances are constantly being created and hence new EPCs representing trade items are constantly being commissioned. In part, this design decision reflects consistent treatment of master data as excluding data that grows as more business is transacted (see comment in section 6.1), and in part reflects the pragmatic reality that data about trade item instances is likely to be managed more like event data than master data when it comes to aging, database design, etc.

7.3.2 Action type

The *Action* type says how an event relates to the lifecycle of the entity being described. For example, *AggregationEvent* (section 7.4.3) is used to capture events related to aggregations of objects, such as cases aggregated to a pallet. Throughout its life, the pallet load participates in many business process steps, each of which may generate an EPCIS event. The *action* field of each event says how the aggregation itself has changed during the event: have objects been added to the aggregation, have objects been removed from the aggregation, or has the aggregation simply been observed without change to its membership? The *action* is independent of the

`bizStep` (of type `BusinessStepID`) which identifies the specific business process step in which the action took place.

The `Action` type is an enumerated type having three possible values:

Action value	Meaning
ADD	The entity in question has been created or added to.
OBSERVE	The entity in question has not been changed: it has neither been created, added to, destroyed, or removed from.
DELETE	The entity in question has been removed from or destroyed altogether.

The description below for each event type that includes an `Action` value says more precisely what `Action` means in the context of that event.

Note that the values above are the only values possible for `Action`. Unlike other types defined below, `Action` is *not* a vocabulary type, and SHALL NOT be extended by industry groups.

7.3.3 The “What” dimension

Section 8.2 and section 8.3 of the CBV defines the data fields and the expected value types and data types used in the “What” dimension of EPCIS events, including discussion of class-level vs instance-level identification.

7.3.3.1 QuantityElement

A `QuantityElement` is a structure that identifies objects identified by a specific class-level identifier, either a specific quantity or an unspecified quantity. It has the following structure:

Field	Type	Description
<code>epcClass</code>	<code>EPCClass</code>	A class-level identifier for the class to which the specified quantity of objects belongs.
<code>quantity</code>	<code>xsd:double</code>	(Optional) A number that specifies how many or how much of the specified <code>EPCClass</code> is denoted by this <code>QuantityElement</code>. The <code>quantity</code> may be omitted to indicate that the quantity is unknown or not specified. If <code>quantity</code> is omitted, then <code>uom</code> SHALL be omitted as well. Otherwise, if <code>quantity</code> is specified: If the <code>QuantityElement</code> lacks a <code>uom</code> field (below), then the <code>quantity</code> SHALL have a positive integer value, and denotes a count of the number of instances of the specified <code>EPCClass</code> that are denoted by this <code>QuantityElement</code>. If the <code>QuantityElement</code> includes a <code>uom</code>, then the <code>quantity</code> SHALL have a positive value (but not necessarily an integer value), and denotes the magnitude of the physical measure that specifies how much of the specified <code>EPCClass</code> is denoted by this <code>QuantityElement</code>.
<code>uom</code>	<code>UOM</code>	(Optional) If present, specifies a unit of measure by which the specified quantity is to be interpreted as a physical measure, specifying how much of the specified <code>EPCClass</code> is denoted by this <code>QuantityElement</code>. The <code>uom</code> SHALL be omitted if <code>quantity</code> is omitted.

`EPCClass` is a Vocabulary whose elements denote classes of objects. `EPCClass` is a User Vocabulary as defined in section 6.2. Any EPC whose structure incorporates the concept of object class can be referenced as an `EPCClass`.

An `EPCClass` may refer to a class having fixed measure or variable measure. A fixed measure class has instances that may be counted; for example, a GTIN that refers to fixed-size cartons of a product. A variable measure class has instances that cannot be counted and so the quantity is specified as a physical measure; for example, a GTIN that refers to copper wire that is sold by length, carpeting that is sold by area, bulk oil

that is sold by volume, or fresh produce that is sold by weight. The following table summarises how the `quantity` and `uom` fields are used in each case:

EPCClass	quantity field	uom field	Meaning
Fixed measure	Positive integer	Omitted	The <code>quantity</code> field specifies the count of the specified class.
Variable measure	Positive number, not necessarily an integer	Present	The <code>quantity</code> field specifies the magnitude, and the <code>uom</code> field the physical unit, of a physical measure describing the amount of the specified class.
Fixed or Variable Measure	Omitted	Omitted	The quantity is unknown or not specified.

Master data attributes for the `EPCClass` vocabulary contain whatever master data is defined for the referenced objects independent of EPCIS (for example, product catalogue data); definitions of these are outside the scope of this specification.

7.3.3.1.1 UOM

As specified above, the `uom` field of a `QuantityElement` is present when the `QuantityElement` uses a physical measure to specify the quantity of the specified `EPCClass`. When a `uom` field is present, its value SHALL be the 2- or 3-character code for a physical unit specified in the "Common Code" column of UN/CEFACT Recommendation 20 [CEFACT20]. Moreover, the code SHALL be a code contained in a row of [CEFACT20] meeting all of the following criteria:

- The "Quantity" column contains one of the following quantities: *length*, *area*, *volume*, or *mass*.
- The "Status" column does *not* contain "X" (deleted) or "D" (deprecated).

For purposes of the first criterion, the quantity must appear as a complete phrase. Example: "metre" (MTR) is allowed, because the quantity includes *length* (among other quantities such as *breadth*, *height*, etc.). But "pound-force per foot" (F17) is *not* allowed, because the quantity is *force divided by length*, not just *length*.

7.3.3.1.2 Class-level identifiers

Class-level identifiers are specified -- as EPC URNs and GS1 Digital Link URIs -- in section 8.3 of CBV.

Implementations SHALL understand queries expressed in both EPC URN and GS1 Digital Link URI syntaxes, and MAY return query responses corresponding to either syntax, at the discretion of the responding implementation. The requesting client might need to translate the identifiers within the query response into its preferred syntax.

7.3.3.2 Identifier types (Non-Normative)

The normative specifications of identifiers are in the EPC Tag Data Standard [TDS], the EPC Core Business Vocabulary [CBV] and the GS1 Digital Link Standard [DL].

7.3.4 The "When" dimension

The "When" dimension of EPCIS includes data fields that provide a chronological context to visibility events.

7.3.4.1 The "When" dimension in the EPCISEvent common base type

The EPCISEvent common base type includes the `eventTime`, `recordTime` and `eventTimeZoneOffset` fields, as specified in section 7.4.1.

7.3.4.2 The "When" dimension in the Error Declaration

The EPCIS Error Declaration element, specified in section 7.4.1.2.1, includes -- beyond the reiterated fields of the original, erroneous event -- the `declarationTime` field.

7.3.4.3 The "When" dimension in Sensor Metadata

EPCIS Sensor Metadata, specified in section 7.3.7.1.1, includes the `time`, `startTime`, and `endTime` fields, which relate to the time of sensor data capture (i.e., rather than the time of the EPCIS event capture).

7.3.5 The "Where" Dimension – read point and business location

This section addresses the notion of location information as used in EPCIS.

EPCIS location types are defined as EPCIS Vocabulary Types, as follows:

Type	Description
<code>ReadPointID</code>	A Read Point is a discretely recorded location that is meant to identify the most specific place at which an EPCIS event took place. Read Points are determined by the EPCIS Capturing Application, perhaps inferred as a function of logical reader if stationary readers are used, perhaps determined overtly by reading a location tag if the reader is mobile, or in general determined by any other means the EPCIS Capturing Application chooses to use. Conceptually, the Read Point is designed to identify "where objects were at the time of the EPCIS event."
<code>BusinessLocationID</code>	A Business Location is a uniquely identified and discretely recorded location that is meant to designate the specific place where an object is assumed to be following an EPCIS event until it is reported to be at a different Business Location by a subsequent EPCIS event. As with the Read Point, the EPCIS Capturing Application determines the Business Location based on whatever means it chooses. Conceptually, the Business Location is designed to identify "where objects are following the EPCIS event."

`ReadPointID` and `BusinessLocationID` are User Vocabularies as defined in section 6.2. Some industries may wish to use EPCs as vocabulary elements, in which case pure identity URIs as defined in [TDS] SHALL be used.

Note:

The `LocationID` type is a supertype of `ReadPointID`, `BusinessLocationID`, and `SourceDestID`. In an EPCIS master data document (or master data header within an EPCIS document or EPCIS query document), the `urn:epcglobal:epcis:vtype:Location` URI may be used to specify a single vocabulary containing identifiers that may appear in the read point, business step, source, or destination field of associated EPCIS events.

- ✓ **Non-Normative:** Illustration: For example, in industries governed by GS1 General Specifications, `readPointID`, and `businessLocationID` may be SGLN-URIs.

Location vocabulary elements are not *required* to be EPCs.

- ✓ **Non-Normative:** Explanation: Allowing non-EPC URIs for locations gives organisations greater freedom to reuse existing ways of naming locations.

For all of the EPCIS Event Types defined in this section, capture events include separate fields for Read Point and Business Location. In most cases, both are optional, so that it is still possible for an EPCIS Capturing Application to include partial information if both are not known.




Non-Normative: Explanation: Logical Reader and Physical Reader are omitted from the definitions of EPCIS events in this specification. Physical Reader is generally not useful information for exchange between partners. For example, if a reader malfunctions and is replaced by another reader of identical make and model, the Physical Reader ID has changed. This information is of little interest to trading partners. Likewise, the Logical Reader ID may change if the capturing organisation makes a change in the way a particular business process is executed; again, not often of interest to a partner.

The distinction between Read Point and Business Location is very much related to the dichotomy between retrospective semantics and prospective semantics discussed above. In general, Read Points play a role in retrospective semantics, while Business Locations are involved in prospective statements. This is made explicit in the way each type of location enters the semantic descriptions given at the end of each section below that defines an EPCIS capture event.


7.3.5.1 Example of the distinction between a read point and a business location (Non-Normative)

Figure 7-4 Example of the distinction between a read point and a business location

Graph View:



Physical View:



Tag	Time	Read Point	Business Location	Comment
# 123	07:00	"RP-DC#88-A"	DC#88.Receive & Store	Product entered DC via DockDoor#R1
# 123	09:00	"RP-DC#88-K"	DC#88.Shipping	Product placed on conveyor for shipping
# 123	09:30	"RP-DC#88-N"	[omitted]	Product shipped via dock door#S2

The figure above shows a typical use case consisting of rooms with fixed doorways at the boundaries of the rooms. In such a case, Read Points correspond to the doorways (with RFID instrumentation) and Business Locations correspond to the rooms. Note that the Read Points and Business Locations are not in one-to-one correspondence; the only situation where Read Points and Business Locations could have a 1:1 relationship is the unusual case of a room with a single door, such a small storeroom.

Still considering the rooms-and-doors example, the Business Location is usually the location type of most interest to a business application, as it says which room an object is in. Thus, it is meaningful to ask the inventory of a Business Location such as the backroom. In contrast, the Read Point indicates the doorway through which the object entered the room. It is not meaningful to ask the inventory of a doorway. While sometimes not as relevant to a business application, the Read Point is nevertheless of significant interest to higher level software to understand the business process and the final status of the object, particularly in the presence of less than 100% read rates. Note that correct designation of the business location requires both that the tagged object be observed at the Read Point and that the direction of movement be correctly determined – again reporting the Read Point in the event will be very valuable for higher level software.

A supply chain like the rooms-and-doors example may be represented by a graph in which each node in the graph represents a room in which objects may be found, and each arc represents a doorway that connects two rooms. Business Locations, therefore, correspond to nodes of this graph, and Read Points correspond to the arcs. If the graph were a straight, unidirectional chain, the arcs traversed by a given object could be reconstructed from knowing the nodes; that is, Read Point information would

be redundant given the Business Location information. In more real-world situations, however, objects can take multiple paths and move “backwards” in the supply chain. In these real-world situations, providing Read Point information in addition to Business Location information is valuable for higher level software.

The key to balancing seemingly conflicting requirements is to define and relate various location types, and then to rely on the EPCIS Capturing Application to properly record them for a given capture event. This is why EPCIS events contain both a ReadPointID and a BusinessLocationID (the two primitive location types).

7.3.6 The “Why” dimension

This section defines the data fields and the expected value types and data types used in the “Why” dimension of the event types specified in section [7.3.6.4](#).

7.3.6.1 Business step

BusinessStepID is a vocabulary whose elements denote steps in business processes. An example is an identifier that denotes “shipping.” The business step field of an event specifies the business context of an event: what business process step was taking place that caused the event to be captured? BusinessStepID is an example of a Standard Vocabulary as defined in section [6.2](#).



Non-Normative: Explanation: Using an extensible vocabulary for business step identifiers allows GS1 standards (including and especially the GS1 Core Business Vocabulary) to define some common terms such as “shipping” or “receiving,” while allowing for industry groups and individual end-users to define their own terms. Master data provides additional information.

This specification defines no master data attributes for business step identifiers.

7.3.6.2 Disposition and Persistent Disposition

DispositionID is a vocabulary whose elements denote a business state of an object. An example is an identifier that denotes “recalled.”

7.3.6.2.1 Disposition

The disposition field of an event specifies the business condition of the event’s objects, subsequent to the event. The disposition is assumed to hold true until another event indicates a change of disposition. Intervening events that do not specify a disposition field have no effect on the presumed disposition of the object. DispositionID is an example of a Standard Vocabulary as defined in section [6.2](#).

Applicability of Disposition by event type

event type	Action value	Disposition applies to
ObjectEvent	ADD OBSERVE DELETE	epcList quantityList
AggregationEvent	ADD OBSERVE	parentID childEPCs childQuantityList
	DELETE	childEPCs childQuantityList
TransactionEvent	ADD OBSERVE DELETE	parentID epcList quantityList
TransformationEvent	n/a	outputEPCList outputQuantityList
AssociationEvent	ADD OBSERVE	parentID childEPCs childQuantityList
	DELETE	childEPCs childQuantityList

7.3.6.2.2 Persistent Disposition

The `persistentDisposition` field of an event specifies one or more business conditions of the event's objects, subsequent to the event. One or more `persistentDisposition` values can be `set` or `unset`, independently of each other. Values that are `set` are considered to remain valid until explicitly `unset`. Like `disposition`, `persistentDisposition` leverages the `DispositionID` Standard Vocabulary.

NOTE: `persistentDisposition` SHOULD only be used in conjunction with `ObjectEvents` with `Action` type `OBSERVE`, in order to avoid the potential for ambiguity with regard to inheritance of `persistentDisposition` by `child/input/output` objects.



Non-Normative: Explanation: Using an extensible vocabulary for disposition identifiers allows GS1 standards (including and especially the GS1 Core Business Vocabulary) to define some common terms such as "recalled" or "in transit," while allowing for industry groups and individual end-users to define their own terms. Master data may provide additional information.

This specification defines no master data attributes for disposition identifiers.

Applicability of persistentDisposition by event type

event type	Action value	set or unset of persistentDisposition applies to
ObjectEvent	ADD OBSERVE DELETE	epcList quantityList
AggregationEvent	ADD OBSERVE DELETE	n/a
TransactionEvent	ADD OBSERVE DELETE	n/a
TransformationEvent	n/a	outputEPCList outputQuantityList
AssociationEvent	ADD OBSERVE DELETE	n/a

7.3.6.3 Business transaction

A `BusinessTransaction` identifies a particular business transaction. An example of a business transaction is a specific purchase order. Business Transaction information may be included in EPCIS events to record an event's participation in particular business transactions.

A business transaction is described in EPCIS by a structured type consisting of a pair of identifiers, as follows.

Field	Type	Description
type	<code>BusinessTransactionTypeID</code>	(Optional) An identifier that indicates what kind of business transaction this <code>BusinessTransaction</code> denotes. If omitted, no information is available about the type of business transaction apart from what is implied by the value of the <code>bizTransaction</code> field itself.
bizTransaction	<code>BusinessTransactionID</code>	An identifier that denotes a specific business transaction.

The two vocabulary types `BusinessTransactionTypeID` and `BusinessTransactionID` are defined in the sections below.

7.3.6.3.1 Business transaction type

`BusinessTransactionTypeID` is a vocabulary whose elements denote a specific type of business transaction. An example is an identifier that denotes "purchase order." `BusinessTransactionTypeID` is an example of a Standard Vocabulary as defined in section 6.2.



Non-Normative: Explanation: Using an extensible vocabulary for business transaction type identifiers allows GS1 standards to define some common terms such as "purchase order" while allowing for industry groups and individual end-users to define their own terms. Master data may provide additional information.

This specification defines no master data attributes for business transaction type identifiers.

7.3.6.3.2 Business transaction ID

`BusinessTransactionID` is a vocabulary whose elements denote specific business transactions. An example is an identifier that denotes “Acme Corp purchase order number 12345678.” `BusinessTransactionID` is a User Vocabulary as defined in section 6.2.

 **Non-Normative:** Explanation: URIs are used to provide extensibility and a convenient way for organisations to distinguish one kind of transaction identifier from another. For example, if Acme Corporation has purchase orders (one kind of business transaction) identified with an 8-digit number as well as shipments (another kind of business transaction) identified by a 6-character string, the following business transaction IDs might be associated with a particular EPC over time:

<https://transaction.acme.com/po/12345678>
<https://transaction.acme.com/shipment/34ABC8>

An EPCIS Accessing Application might query EPCIS and discover both of the transaction IDs; using URIs gives the application a way to understand which ID is of interest to it.

7.3.6.4 Source and destination

A `Source` or `Destination` is used to provide additional business context when an EPCIS event is part of a business transfer; that is, a process in which there is a transfer of ownership, responsibility, and/or custody of physical or digital objects.

In many cases, a business transfer requires several individual business steps (and therefore several EPCIS events) to execute; for example, shipping followed by receiving, or a more complex sequence such as loading → departing → transporting → arriving → unloading → accepting. The `ReadPoint` and `BusinessLocation` in the “where” dimension of these EPCIS events indicate the known physical location at each step of the process. `Source` and `Destination`, in contrast, may be used to indicate the parties and/or location that are the intended endpoints of the business transfer. In a multi-step business transfer, some or all of the EPCIS events may carry `Source` and `Destination`, and the information would be the same for all events in a given transfer.

`Source` and `Destination` provide a standardised way to indicate the parties and/or physical locations involved in the transfer, complementing the business transaction information (e.g., purchase orders, invoices, etc.) that may be referred to by `BusinessTransaction` elements.

A source or destination is described in EPCIS by a structured type consisting of a pair of identifiers, as follows.

Field	Type	Description
type	SourceDestTypeID	(Optional in <code>BizTransaction</code> , required in <code>SourceOrDestination</code> , and optional in <code>SensorReport</code>). Indicates the kind of <code>BizTransaction</code> document (e.g., <code>po</code>), role of <code>SourceOrDestination</code> (e.g., <code>owning_party</code>), or kind of measurement in <code>SensorReport</code> (e.g., <code>Mass</code>).
source or destination	SourceDestID	An identifier that denotes a specific source or destination.

The two vocabulary types `SourceDestTypeID`, and `SourceDestID` are defined in the sections below.

7.3.6.4.1 Source/Destination type

`SourceDestTypeID` is a vocabulary whose elements denote a specific type of business transfer source or destination. An example is an identifier that denotes "owning party." `SourceDestTypeID` is an example of a Standard Vocabulary as defined in section 6.2.



Non-Normative: Explanation: Using an extensible vocabulary for source/destination type identifiers allows GS1 standards to define some common terms such as "owning party" while allowing for industry groups and individual end-users to define their own terms. Master data may provide additional information.

This specification defines no master data attributes for source/destination type identifiers.

7.3.6.4.2 Source/Destination ID

`SourceDestID` is a vocabulary whose elements denote specific sources and destinations. An example is an identifier that denotes "Acme Corporation (an owning party)." `SourceDestID` is a User Vocabulary as defined in section 6.2.

Note:

The `PartyID` type is a supertype of `SourceDestID`. The `urn:epcglobal:epcis:vtype:Location` URI may be used to specify a single vocabulary containing identifiers that may populate the source or destination field(s) of associated EPCIS events.



Non-Normative: Explanation: URIs are used to provide extensibility and a convenient way for organisations to distinguish one kind of source or destination identifier from another.

7.3.7 The "How" dimension

This section defines the data fields and the expected value types and data types used in the "How" dimension.

7.3.7.1 SensorElement

A `SensorElement` is a structure that contains an optional `sensorMetadata` element and one or several `sensorReport` elements, described as follows:

Field	Type	Description
<code>sensorMetadata</code>	<code>SensorMetadata</code>	(Optional) An element containing one or several metadata attributes, which are applicable to all <code>sensorReport</code> elements that are part of the same <code>sensorElement</code> .
<code>sensorReport</code>	<code>SensorReport</code>	An element containing one or several attributes that pertain to a specific sensor observation.

For the sake of compactness, all elements contained in either the `sensorMetadata` or `sensorReport` element are inline attributes.

The following rules apply:

1. A `sensorElement` parent element MAY contain exactly one `sensorMetadata` element and SHALL contain one or more `sensorReport` elements.
2. The values of any inline attributes specified within the `sensorMetadata` element are assumed to apply for all `sensorReport` elements contained within that same `sensorElement`.
3. The following inline attributes are permitted to appear both in a `sensorMetadata` as well as in a `sensorReport` element: `dataProcessingMethod`, `deviceId`, `deviceMetadata`, `rawData`, and `time` (which are specified in section 7.3.7.1.1). If they are present in a `sensorElement` container, they SHALL either be specified in the `sensorMetadata` element or in the `sensorReport` element(s) and SHALL NOT appear in both.

Non-normative: Explanation: Even though it would technically be feasible, EPCIS SHOULD NOT be used to accommodate **raw** sensor data unless there is a strong reason to do so. The added value of the `sensorElement` in EPCIS consists in the abstraction from raw sensor data and provisioning of aggregated, business-oriented data to accessing applications. For instance, instead of capturing thousands of time-stamped datasets, it is often far more appropriate and efficient to only indicate the range of values of a given sensor property within a given period of time. For that purpose, an EPCIS capturing application would only need to populate four fields: `minValue`, `maxValue`, `startTime`, and `endTime`. Even if there is a business need to have the ability to access the underlying raw sensor data, it is neither required nor advisable to include raw data in the EPCIS event. Instead, it is advisable to include a Web URI in the `rawData` element, pointing to a resource through which clients can access the underlying raw sensor data.

7.3.7.1.1 SensorMetadata

A `SensorMetadata` element contains a number of inline attributes, defined as follows.

Field	Type	Description
time	xsd:dateTimeStamp	(Optional) The actual point in time of an observation as transmitted by a sensor device. If present, the <code>time</code> value SHALL be less (earlier) than or equal to the <code>eventTime</code> value. It SHALL also contain a time zone offset value. The coordination and integrity of distributed computing requires time synchronisation of EPCIS events conveying sensor data. Therefore, when populating the <code>time</code>, <code>startTime</code>, and <code>endTime</code> field, EPCIS capture applications SHOULD apply established time synchronisation protocols such as IEEE 1588-2008, which provides a standard method to synchronise device clocks in a network. <u>Note:</u> The <code>eventTime</code> applies to the completion of a business step, not a sensor observation. For instance, for a receiving event accommodating a sensor element, the <code>eventTime</code> indicates when goods were received – it gives no information when certain conditions (e.g. a specific temperature value) held true. In some circumstances, event and sensor observation times may correspond though (e.g. if a quality inspector checks certain properties of goods). In such cases, indicating the <code>eventTime</code> may be sufficient.
startTime	xsd:dateTimeStamp	(Optional) The lowest (earliest) value of a given observation period as transmitted by a sensor device. If present, the <code>startTime</code> SHALL be less (earlier) than the <code>eventTime</code> value and the <code>endTime</code> value.
endTime	xsd:dateTimeStamp	(Optional) The highest (most recent) value of a given observation period, as transmitted by a sensor device. If present, the <code>endTime</code> SHALL be less (earlier) than or equal to the <code>eventTime</code> value.
deviceID	EPC	(Optional) Device from which the sensor data originates.
deviceMetadata	ResourceID	(Optional) Storage location of an electronic document accommodating metadata of the device from which the sensor data originates.
rawData	ResourceID	(Optional) Storage/service location of the raw sensor data on which the aggregated/business-oriented data contained in the <code>sensorElement</code> is based.
dataProcessingMethod	ResourceID	(Optional) Storage location of an electronic document accommodating the data processing method of the contained sensor data, if applicable. For instance, before sensor data is captured in an EPCIS event, the latter might be redacted or refined by means of specific algorithms.