



FIX Simple Binary Encoding Technical Specification

Version 2.0 Release Candidate 1 August 16,2018

THIS DOCUMENT IS A RELEASE CANDIDATE FOR A PROPOSED FIX TECHNICAL STANDARD. A RELEASE CANDIDATE HAS BEEN APPROVED BY THE GLOBAL TECHNICAL COMMITTEE AS AN INITIAL STEP IN CREATING A NEW FIX TECHNICAL STANDARD. POTENTIAL ADOPTERS ARE STRONGLY ENCOURAGED TO BEGIN WORKING WITH THE RELEASE CANDIDATE AND TO PROVIDE FEEDBACK TO THE GLOBAL TECHNICAL COMMITTEE AND THE WORKING GROUP THAT SUBMITTED THE PROPOSAL. THE FEEDBACK TO THE RELEASE CANDIDATE WILL DETERMINE IF ANOTHER REVISION AND RELEASE CANDIDATE IS NECESSARY OR IF THE RELEASE CANDIDATE CAN BE PROMOTED TO BECOME A FIX TECHNICAL STANDARD DRAFT.

DISCLAIMER

THE INFORMATION CONTAINED HEREIN AND THE FINANCIAL INFORMATION EXCHANGE PROTOCOL (COLLECTIVELY, THE "FIX PROTOCOL") ARE PROVIDED "AS IS" AND NO PERSON OR ENTITY ASSOCIATED WITH THE FIX PROTOCOL MAKES ANY REPRESENTATION OR WARRANTY, EXPRESS OR IMPLIED, AS TO THE FIX PROTOCOL (OR THE RESULTS TO BE OBTAINED BY THE USE THEREOF) OR ANY OTHER MATTER AND EACH SUCH PERSON AND ENTITY SPECIFICALLY DISCLAIMS ANY WARRANTY OF ORIGINALITY, ACCURACY, COMPLETENESS, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SUCH PERSONS AND ENTITIES DO NOT WARRANT THAT THE FIX PROTOCOL WILL CONFORM TO ANY DESCRIPTION THEREOF OR BE FREE OF ERRORS. THE ENTIRE RISK OF ANY USE OF THE FIX PROTOCOL IS ASSUMED BY THE USER.

NO PERSON OR ENTITY ASSOCIATED WITH THE FIX PROTOCOL SHALL HAVE ANY LIABILITY FOR DAMAGES OF ANY KIND ARISING IN ANY MANNER OUT OF OR IN CONNECTION WITH ANY USER'S USE OF (OR ANY INABILITY TO USE) THE FIX PROTOCOL, WHETHER DIRECT, INDIRECT, INCIDENTAL, SPECIAL OR CONSEQUENTIAL (INCLUDING, WITHOUT LIMITATION, LOSS OF DATA, LOSS OF USE, CLAIMS OF THIRD PARTIES OR LOST PROFITS OR REVENUES OR OTHER ECONOMIC LOSS), WHETHER IN TORT (INCLUDING NEGLIGENCE AND STRICT LIABILITY), CONTRACT OR OTHERWISE, WHETHER OR NOT ANY SUCH PERSON OR ENTITY HAS BEEN ADVISED OF, OR OTHERWISE MIGHT HAVE ANTICIPATED THE POSSIBILITY OF, SUCH DAMAGES.

DRAFT OR NOT RATIFIED PROPOSALS (REFER TO PROPOSAL STATUS AND/OR SUBMISSION STATUS ON COVER PAGE) ARE PROVIDED "AS IS" TO INTERESTED PARTIES FOR DISCUSSION ONLY. PARTIES THAT CHOOSE TO IMPLEMENT THIS DRAFT PROPOSAL DO SO AT THEIR OWN RISK. IT IS A DRAFT DOCUMENT AND MAY BE UPDATED, REPLACED, OR MADE OBSOLETE BY OTHER DOCUMENTS AT ANY TIME. THE FIX GLOBAL TECHNICAL COMMITTEE WILL NOT ALLOW EARLY IMPLEMENTATION TO CONSTRAIN ITS ABILITY TO MAKE CHANGES TO THIS SPECIFICATION PRIOR TO FINAL RELEASE. IT IS INAPPROPRIATE TO USE FIX WORKING DRAFTS AS REFERENCE MATERIAL OR TO CITE THEM AS OTHER THAN "WORKS IN PROGRESS". THE FIX GLOBAL TECHNICAL COMMITTEE WILL ISSUE, UPON COMPLETION OF REVIEW AND RATIFICATION, AN OFFICIAL STATUS ("APPROVED") OF/FOR THE PROPOSAL AND A RELEASE NUMBER.

No proprietary or ownership interest of any kind is granted with respect to the FIX Protocol (or any rights therein).

Copyright 2013-2018 FIX Protocol Ltd., all rights reserved.



FIX Performance Session Layer specification by [FIX Protocol Ltd.](#) is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](#).

Based on a work at <https://github.com/FIXTradingCommunity/fixp-specification>.

Contents

1	Introduction	8
1.1	Binary type system.....	8
1.2	Design principles.....	8
1.3	Message schema.....	8
1.4	Glossary.....	9
1.5	Documentation	9
1.5.1	Specification terms	9
1.5.2	Document format.....	10
1.6	References	10
1.6.1	Related FIX Standards	10
1.6.2	Dependencies on other standards.....	10
2	Field Encoding.....	11
2.1	Field aspects.....	11
2.1.1	Semantic data type	11
2.1.2	Encoding.....	11
2.1.3	Metadata.....	11
2.1.4	Field presence	12
2.1.5	Default value.....	12
2.2	Data type summary.....	12
2.3	Common field schema attributes	15
2.3.1	Non-FIX types.....	15
2.4	Integer encoding.....	15
2.4.1	Primitive type encodings.....	16
2.4.2	Range attributes for integer fields.....	16
2.4.3	Byte order	16
2.4.4	Integer encoding specifications	17
2.4.5	Examples of integer fields.....	17
2.5	Decimal encoding.....	18
2.5.1	Composite encodings.....	18
2.5.2	Range attributes for decimal fields.....	18
2.5.3	Encoding specifications for decimal types	19
2.5.4	Composite encoding padding	19
2.5.5	Examples of decimal fields.....	19
2.6	Float encoding.....	20
2.6.1	Primitive types	20
2.6.2	Null values.....	20
2.6.3	Byte order	20

2.6.4	Float encoding specifications.....	21
2.6.5	Examples of floating point fields.....	21
2.7	String encodings.....	21
2.7.1	Character.....	21
2.7.2	Fixed-length character array.....	22
2.7.3	Variable-length string encoding.....	23
2.7.4	Range attributes for string Length.....	23
2.7.5	Encoding specifications for variable-length string.....	24
2.7.6	Example of a variable-length string field.....	24
2.8	Data encodings.....	24
2.8.1	Fixed-length data.....	25
2.8.2	Variable-length data encoding.....	25
2.8.3	Range attributes for variable-length data Length.....	26
2.8.4	Encoding specifications for variable-length data.....	26
2.8.5	Example of a data field.....	26
2.9	MonthYear encoding.....	26
2.9.1	Composite encoding padding.....	27
2.9.2	Encoding specifications for MonthYear.....	27
2.10	Date and time encoding.....	27
2.10.1	Epoch.....	27
2.10.2	Time unit.....	28
2.10.3	Encoding specifications for date and time.....	28
2.10.4	Examples of date/time fields.....	29
2.11	Local date encoding.....	30
2.12	Local time encoding.....	30
2.12.1	TZTimestamp encoding.....	30
2.12.2	Composite encoding padding.....	31
2.12.3	TZTimeOnly encoding.....	31
2.12.4	Composite encoding padding.....	32
2.13	Enumeration encoding.....	32
2.13.1	Primitive type encodings.....	32
2.13.2	Value encoding.....	32
2.13.3	Encoding specification of enumeration.....	33
2.13.4	Enumeration examples.....	33
2.13.5	Constant field of an enumeration value.....	34
2.13.6	Boolean encoding.....	34
2.14	Multi-value choice encoding.....	34
2.14.1	Primitive type encodings.....	35

2.14.2	Value encoding.....	35
2.14.3	Encoding specification of multi-value choice.....	35
2.14.4	Multi-value example	35
2.15	Field value validation	36
3	Message Structure	37
3.1	Message Framing	37
3.1.1	Simple Open Framing Header	37
3.2	SBE Message Encoding Header	37
3.2.1	Message header schema.....	38
3.2.2	Root block length	39
3.2.3	Template ID.....	39
3.2.4	Schema ID	39
3.2.5	Schema version	39
3.2.6	Number of repeating groups	39
3.2.7	Number of variable-length fields	39
3.3	Message Body	39
3.3.1	Data only on the wire.....	39
3.3.2	Direct access	39
3.3.3	Field position and padding.....	40
3.4	Repeating Groups	42
3.4.1	Schema specification of a group	42
3.4.2	Group block length.....	42
3.4.3	Byte alignment of a group entry	42
3.4.4	Padding at end of a group entry	42
3.4.5	Entry counter	43
3.4.6	Empty group.....	43
3.4.7	Multiple repeating groups	43
3.4.8	Nested repeating group specification.....	43
3.4.9	Nested repeating group wire format.....	44
3.4.10	Empty group means nested group is empty	44
3.4.11	Group dimension encoding.....	44
3.5	Sequence of message body elements.....	45
3.5.1	Root level elements	45
3.5.2	Repeating group elements.....	46
3.6	Message structure validation.....	46
4	Message Schema.....	47
4.1	XML schema for SBE message schemas.....	47
4.2	XML namespace	47

4.3	Name convention.....	47
4.3.1	Capitalization	47
4.4	Root element	47
4.4.1	<messageSchema> attributes.....	47
4.4.2	Schema versioning	48
4.5	Data encodings.....	48
4.5.1	Encoding sets	48
4.5.2	Encoding name.....	49
4.5.3	Simple encodings	49
4.5.4	Composite encodings.....	50
4.5.5	Reference to reusable types	52
4.5.6	Enumeration encodings	53
4.5.7	Multi-value choice encodings (bitset).....	55
4.6	Message template.....	57
4.6.1	Byte alignment	57
4.6.2	Reserved space	57
4.6.3	Message members	57
4.6.4	Member order.....	57
4.7	Field attributes.....	58
4.8	Repeating group schema	60
4.9	Schema validation.....	61
4.9.1	Message with a repeating group	62
4.9.2	Message with raw data fields	62
4.10	Reserved element names.....	63
4.10.1	Composite types	63
4.10.2	Composite type elements	63
5	Schema Extension Mechanism	64
5.1	Objective	64
5.1.1	Constraints	64
5.2	Message schema features for extension	64
5.2.1	Schema version	64
5.2.2	Since version	65
5.2.3	Block length.....	65
5.2.4	Deprecated elements.....	65
5.3	Wire format features for extension	65
5.3.1	Block size	65
5.3.2	Number of repeating groups and variable data.....	65
5.4	Comaptibility strategy.....	66

5.5	Message schema extension example	66
6	Usage Guidelines.....	68
6.1	Identifier encodings	68
7	Examples	69
7.1	SBE Message Schema.....	69
7.2	Flat, fixed-length message	72
7.2.1	Wire format of an order message.....	73
7.3	Message with a repeating group	73
7.3.1	Wire format of an execution message.....	74
7.3.2	Interpretation.....	74
7.4	Message with a variable-length field	75
7.4.1	Wire format of a business reject message.....	75
7.4.2	Interpretation.....	75
8	Release Notes.....	76
8.1	SBE version 2.0 Release Candidate 1	76
8.1.1	Compatibility.....	76

1 Introduction

FIX Simple Binary Encoding (SBE) targets high performance trading systems. It is optimized for low latency of encoding and decoding while keeping bandwidth utilization reasonably small. For compatibility, it is intended to represent all FIX semantics.

This encoding specification describes the wire protocol for messages. Thus, it provides a standard for interoperability between communicating parties. Users are free to implement the standard in a way that best suits their needs. Implementers may use the most appropriate techniques and idioms of programming languages to access data off the wire and manipulate it in applications.

The encoding standard is complimentary to other FIX standards for session protocol and application level behavior.

1.1 Binary type system

In order to support traditional FIX semantics, all the documented field types are supported. However, instead of printable character representations of tag-value encoding, the type system binds to native binary data types, and defines derived types as needed.

The binary type system has been enhanced in these ways:

- Provides a means to specify precision of decimal numbers and timestamps, as well as valid ranges of numbers.
- Differentiates fixed-length character arrays from variable-length strings. Allows a way to specify the minimum and maximum length of strings that an application can accept.
- Provides a consistent system of enumerations, Boolean switches and multiple-choice fields.

1.2 Design principles

The message design strives for direct data access without complex transformations or conditional logic. This is achieved by:

- Usage of native binary data types and simple types derived from native binaries, such as prices and timestamps.
- Preference for fixed positions and fixed length fields, supporting direct access to data and avoiding the need for management of heaps of variable-length elements which must be sequentially processed.

1.3 Message schema

This standard describes how fields are encoded and the general structure of messages. The content of a message type is specified by a message schema. A message schema tells which fields belong to a message and their location within a message. Additionally, the metadata describes valid value ranges and information that need not be sent on the wire, such as constant values.

Message schemas may be based on standard FIX message specifications, or may be customized as needed by agreement between counterparties.

1.4 Glossary

Data type - A field type with its associated encoding attributes, including backing primitive types and valid values or range. Some types have additional attributes, e.g. epoch of a date.

Encoding - a message format for interchange. The term is commonly used to mean the conversion of one data format to another, such as text to binary. However, Simple Binary Encoding strives to use native binary data types in order to make conversion unnecessary, or at least trivial. Encoding also refers to the act of formatting a message, as opposed to decoding.

Message schema - metadata that specifies messages and their data types and identifiers. Message schemas may be disseminated out of band. For Simple Binary Encoding, message schemas are expressed as an XML document that conforms to an XML schema that is published as part of this standard.

Message template - metadata that specifies the fields that belong to one particular message type. A message template is contained by a message schema.

Session protocol - a protocol concerned with the reliable delivery of messages over a transport. FIX protocol makes a distinction between session protocol and the encoding of a message payload, as described by this document. See the specifications section of FIX protocol web site for supported protocols. The original FIX session protocol is known as FIXT.

XML schema - defines the elements and attributes that may appear in an XML document. The SBE message schema is defined in W3C (XSD) schema language since it is the most widely adopted format for XML schemas.

1.5 Documentation

This document explains:

- The binary type system for field encoding
- Message structure, including field arrangement, repeating groups, and relationship to a message header that may be provided by a session protocol.
- The Simple Binary Encoding message schema.

1.5.1 Specification terms

These key words in this document are to be interpreted as described in [Internet Engineering Task Force RFC2119](#). These terms indicate an absolute requirement for implementations of the standard: "**must**", or "**required**".

This term indicates an absolute prohibition: "**must not**".

These terms indicate that a feature is allowed by the standard but not required: "**may**", "**optional**". An implementation that does not provide an optional feature must be prepared to interoperate with one that does.

These terms give guidance, recommendation or best practices: "**should**" or "**recommended**". A recommended choice among alternatives is described as "**preferred**".

These terms give guidance that a practice is not recommended: "**should not**" or "**not recommended**".

1.5.2 Document format

In this document, these formats are used for technical specifications and data examples.

This is a sample encoding specification

```
<type name="short" primitiveType="int16"/>
```

This is sample data as it would be transmitted on the wire

```
10270000
```

1.6 References

1.6.1 Related FIX Standards

Simple Open Framing Header, FIX Protocol, Limited. Version 1.0 Draft Standard specification has been published at <http://www.fixtradingcommunity.org/>

For FIX semantics, see the current FIX message specification, which is currently [FIX 5.0 Service Pack 2](#) with Extension Packs.

1.6.2 Dependencies on other standards

SBE is dependent on several industry standards. Implementations must conform to these standards to interoperate. Therefore, they are normative for SBE.

[IEEE 754-2008](#) A Standard for Binary Floating-Point Arithmetic

[ISO 639-1:2002](#) Codes for the representation of names of languages - Part 1: Alpha-2 code

[ISO 3166-1:2013](#) Codes for the representation of names of countries and their subdivisions - Part 1: Country codes

[ISO 4217:2015](#) Codes for the representation of currencies and funds

[ISO 8601:2004](#) Data elements and interchange formats - Information interchange - Representation of dates and times

[ISO 10383:2012](#) Securities and related financial instruments - Codes for exchanges and market identification (MIC)

[W3C XML Schema version 1.1](#)

2 Field Encoding

2.1 Field aspects

A field is a unit of data contained by a FIX message. Every field has the following aspects: semantic data type, encoding, and metadata. They will be specified in more detail in the sections on data type encoding and message schema but are introduced here as an overview.

2.1.1 Semantic data type

The FIX semantic data type of a field tells a data domain interpreted at the application layer, for example, whether it is numeric or character data, or whether it represents a time or price. Simple Binary Encoding represents most of the semantic data types that FIX protocol has defined across all encodings. In field specifications, FIX data type is declared with attribute `semanticType`.

See the "Data type summary" below for a listing of FIX types with the usual mapping to SBE encodings. Implementations should not restrict a semantic type to a specific encoding. Message designers may wish to imply a conversion or specialized interpretation at the application layer.

2.1.2 Encoding

Encoding tells how a data element is encoded on the wire. Encoding belongs strictly to the presentation layer. It is context-free and carries no business semantics.

An encoding maps a FIX data type to either a simple, primitive data type, such as a 32 bit signed integer, or to a composite type. A composite type is composed of two or more simple primitive types. For example, the FIX data type Price is encoded as a decimal, a composite type containing a mantissa and an exponent. Note that many fields may share a data type and an encoding. The sections that follow explain the valid encodings for each data type.

2.1.3 Metadata

Field metadata, part of a message schema, describes a field to application developers. Elements of field metadata are:

- Field ID, also known as FIX tag, is a unique identifier of a field for semantic purposes. For example, tag 55 identifies the Symbol field of an instrument.
- Field name, as it is known in FIX specifications
- The FIX semantic data type and encoding type that it maps to
- Valid values or data range accepted
- Documentation

Metadata is normally *not* sent on the wire with Simple Binary Encoding messages. It is necessary to possess the message schema that was used to encode a message in order to decode it. In other words, Simple Binary Encoding messages are not self-describing. Rather, message schemas are typically exchanged out-of-band between counterparties.

See section 4 below for a detailed message schema specification.

2.1.4 Field presence

By default, fields are assumed to be required in a message. However, fields may be specified as optional. To indicate that a value is not set, a special null indicator value is sent on the wire. The null value varies according to data type and encoding. Global defaults for null value may be overridden in a message schema by explicitly specifying the value that indicates nullness.

Alternatively, fields may be specified as constant. In which case, the data is not sent on the wire, but may be treated as constants by applications.

2.1.5 Default value

Default value handling is not specified by the encoding layer. A null value of an optional field does not necessarily imply that a default value should be applied. Rather, default handling is left to application layer specifications.

2.2 Data type summary

FIX semantic types are typically mapped to binary field encodings as follows. See sections below for more detail about each type.

Schema attributes may restrict the range of valid values for a field. See Common field schema attributes below.

FIX semantic type	Binary type	Section	Description
int	Integer encoding	Integer encoding	An integer number
Length	Integer encoding	Integer encoding	Field length in octets. Value must be non-negative.
TagNum	Integer encoding	Integer encoding	A field's tag number. Value must be positive.
SeqNum	Integer encoding	Integer encoding	A field representing a message sequence number. Value must be positive
NumInGroup	Group dimension encoding	Group dimension encoding	A counter representing the number of entries in a repeating group. Value must be positive.
DayOfMonth	Integer encoding	Integer encoding	A field representing a day during a particular month (values 1 to 31).
Qty	Decimal encoding	Decimal encoding	A number representing quantity of a security, such as shares. The encoding may constrain values to integers, if desired.

FIX semantic type	Binary type	Section	Description
float	Float or Decimal	Float encoding Decimal encoding	A real number with binary representation of specified precision. Alternatively, Decimal encoding is preferred for exact decimal numbers.
Price	Decimal encoding	Decimal encoding	A decimal number representing a price
PriceOffset	Decimal encoding	Decimal encoding	A decimal number representing a price offset, which can be mathematically added to a Price.
Amt	Decimal encoding	Decimal encoding	A field typically representing a Price times a Qty.
Percentage	Decimal encoding	Decimal encoding	A field representing a percentage (e.g. 0.05 represents 5% and 0.9525 represents 95.25%).
char	Character	Character encoding	Single US-ASCII character value. Can include any alphanumeric character or punctuation. All char fields are case sensitive (i.e. m != M).
String	Fixed-length character array	Fixed-length character	A fixed-length character array of ASCII encoding
String	Variable-length data encoding	Variable-length string	Alpha-numeric free format strings can include any character or punctuation. All String fields are case sensitive (i.e. morstatt != Morstatt). ASCII encoding.
String—EncodedText	String encoding	Variable-length string	Non-ASCII string. The character encoding may be specified by a schema attribute.
XMLData	String encoding	Variable-length string	Variable-length XML. Must be paired with a Length field.
data	Fixed-length data	Fixed-length data	Fixed-length non-character data
data	Variable-length data encoding	Variable-length data	Variable-length data. Must be paired with a Length field.
Country	Fixed-length character array; size = 2 or a subset of values may use Enumeration encoding	Fixed-length character	ISO 3166-1:2013 Country code

FIX semantic type	Binary type	Section	Description
Currency	Fixed-length character array; size = 3 or a subset of values may use Enumeration encoding	Fixed-length character	ISO 4217:2015 Currency code (3 character)
Exchange	Fixed-length character array; size = 4 or a subset of values may use Enumeration encoding	Fixed-length character	ISO 10383:2012 Market Identifier Code (MIC)
Language	Fixed-length character array; size = 2 or a subset of values may use Enumeration encoding	Fixed-length character	National language - uses ISO 639-1:2002 standard
Implicit enumeration—char or int	Enumeration encoding	Enumeration encoding	A single choice of alternative values
Boolean	Boolean encoding	Boolean encoding	Values true or false
MultipleCharValue	Multi-value choice encoding	Multi-value choice encoding	Multiple choice of a set of values
MultipleStringValue	Multi-value choice encoding. String choices must be mapped to int values.	Multi-value choice encoding	Multiple choice of a set of values
MonthYear	MonthYear encoding	MonthYear encoding	A flexible date format that must include month and year at least, but may also include day or week.
UTCTimestamp	Date and time encoding	Date and time encoding	Time/date combination represented in UTC (Universal Time Coordinated, also known as "GMT")
UTCTimeOnly	Date and time encoding	Date and time encoding	Time-only represented in UTC (Universal Time Coordinated, also known as "GMT")
UTCDateOnly	Date and time encoding	Date and time encoding	Date represented in UTC (Universal Time Coordinated, also known as "GMT")
LocalMktDate	Local date encoding	Date and time encoding	Local date (as oppose to UTC)

FIX semantic type	Binary type	Section	Description
TZTimeOnly	TZTimeOnly	TZTimeOnly encoding	Time of day
TZTimestamp	TZTimestamp	TZTimestamp encoding	Time/date combination representing local time with an offset to UTC to allow identification of local time and timezone offset of that time. The representation is based on ISO 8601:2004

The FIX semantic types listed above are spelled and capitalized exactly as they are in the FIX repository from which official FIX documents and references are derived.

2.3 Common field schema attributes

Schema attributes alter the range of valid values for a field. Attributes are optional unless specified otherwise.

Schema attribute	Description
presence=required	The field must always be set. This is the default presence. Mutually exclusive with nullValue.
presence=constant	The field has a constant value that need not be transmitted on the wire. Mutually exclusive with nullValue, minValue, and maxValue attributes.
presence=optional	The field need not be populated. A special null value indicates that a field is not set. The presence attribute may be specified on either on a field or its encoding.
nullValue	A special value that indicates that an optional value is not set. See encodings below for default nullValue for each type. Mutually exclusive with presence=required and constant.
minValue	The lowest valid value of a range. Applies to scalar data types, but not to String or data types.
maxValue	The highest valid value of a range (inclusive unless specified otherwise). Applies to scalar data types, but not to String or data types.
semanticType	Tells the semantic interpretation of a field in the FIX data type taxonomy.

2.3.1 Non-FIX types

Encodings may be added to SBE messages that do not correspond to listed FIX data types. In that case, the fields that use the encoding will not have a semanticType attribute.

2.4 Integer encoding

Integer encodings should be used for cardinal or ordinal number fields. Signed integers are encoded in a two's complement binary format.

2.4.1 Primitive type encodings

Numeric data types may be specified by range and signed or unsigned attribute. Integer types are intended to convey common platform primitive data types as they reside in memory. An integer type should be selected to hold the maximum range of values that a field is expected to hold.

Primitive type	Description	Length (octets)
int8	Signed byte	1
uint8	Unsigned byte / single-byte character	1
int16	16-bit signed integer	2
uint16	16-bit unsigned integer	2
int32	32-bit signed integer	4
uint32	32-bit unsigned integer	4
int64	64-bit signed integer	8
uint64	64-bit unsigned integer	8

2.4.2 Range attributes for integer fields

The default data ranges and null indicator are listed below for each integer encoding.

A message schema may optionally specify a more restricted range of valid values for a field.

For optional fields, a special null value is used to indicate that a field value is not set. The default null indicator may also be overridden by a message schema.

Required and optional fields of the same primitive type have the same data range. The null value must not be set for a required field.

Schema attribute	int8	uint8	int16	uint16	int32	uint32	int64	uint64
minValue	-127	0	-32767	0	$-2^{31} + 1$	0	$-2^{63} + 1$	0
maxValue	127	254	32767	65534	$2^{31} - 1$	$2^{32} - 2$	$2^{63} - 1$	$2^{64} - 2$
nullValue	-128	255	-32768	65535	-2^{31}	$2^{32} - 1$	-2^{63}	$2^{64} - 1$

2.4.3 Byte order

The byte order of integer fields, and for derived types that use integer components, is specified globally in a message schema. Little-Endian order is the default encoding, meaning that the least significant byte is serialized first on the wire.

See section 4 for specification of message schema attributes, including byteOrder. Message schema designers should specify the byte order most appropriate to their system architecture and that of their counterparties.

2.4.4 Integer encoding specifications

By nature, integers map to simple encodings. These are valid encoding specifications for each of the integer primitive types.

```
<type name="int8" primitiveType="int8" />
<type name="int16" primitiveType="int16" />
<type name="int32" primitiveType="int32" />
<type name="int64" primitiveType="int64" />
<type name="uint8" primitiveType="uint8" />
<type name="uint16" primitiveType="uint16" />
<type name="uint32" primitiveType="uint32" />
<type name="uint64" primitiveType="uint64" />
```

2.4.5 Examples of integer fields

Examples show example schemas and encoded bytes on the wire as hexadecimal digits in Little-Endian byte order.

Example integer field specification

```
<field type="uint32" name="ListSeqNo" id="67" semanticType="int"
  description="Order number within the list" />
```

Value on the wire - uint32 value decimal 10,000, hexadecimal 2710.

10270000

Optional field with a valid range 0-6

```
<type name="range06" primitiveType="uint8" maxValue="6"
  presence="optional" nullValue="255" />
<field type="range06" name="MaxPriceLevels" id="1090"
  semanticType="int"/>
```

Wire format of uint8 value decimal 3.

03

Sequence number field with integer encoding

```
<field type="uint64" name="MsgSeqNum" id="34"
  semanticType="SeqNum" />
```

Wire format of uint64 value decimal 100,000,000,000, hexadecimal 174876E800.

00e8764817000000

Wire format of uint16 value decimal 10000, hexadecimal 2710.

1027

Wire format of uint32 null value $2^{32} - 1$

ffffffff

2.5 Decimal encoding

Decimal encodings should be used for prices and related monetary data types like PriceOffset and Amt.

FIX specifies Qty as a float type to support fractional quantities. However, decimal encoding may be constrained to integer values if that is appropriate to the application or market.

2.5.1 Composite encodings

Prices are encoded as a scaled decimal, consisting of a signed integer mantissa and signed exponent. For example, a mantissa of 123456 and exponent of -4 represents the decimal number 12.3456.

2.5.1.1 Mantissa

Mantissa represents the significant digits of a decimal number. Mantissa is a commonly used term in computing, but it is properly known in mathematics as significand or coefficient.

2.5.1.2 Exponent

Exponent represents scale of a decimal number as a power of 10.

2.5.1.3 Floating point and fixed point encodings

A floating-point decimal transmits the exponent on the wire while a fixed-point decimal specifies a fixed exponent in a message schema. A constant negative exponent specifies a number of assumed decimal places to the right of the decimal point.

Implementations should support both 32 bit and 64 bit mantissa. The usage depends on the data range that must be represented for a particular application. It is expected that an 8 bit exponent should be sufficient for all FIX uses.

Encoding type	Description	Backing primitives	Length (octets)
decimal	Floating-point decimal	Composite: int64 mantiss, int8 exponent	9
decimal64	Fixed-point decimal	int64 mantissa, constant exponent	8
decimal32	Fixed-point decimal	int32 mantissa, constant exponent	4

Optionally, implementations may support any other signed integer types for mantissa and exponent.

2.5.2 Range attributes for decimal fields

The default data ranges and null indicator are listed below for each decimal encoding.

A message schema may optionally specify a more restricted range of valid values for a field.

For all optional decimal fields, a special mantissa value is used to indicate that a field value is null. For floating-point decimals, the exponent should also be set to its null indicator.

Schema attribute	decimal	decimal64	decimal32
exponent range	-128 to 127	-128 to 127	-128 to 127
mantissa range	$-2^{63} + 1$ to $2^{63} - 1$	$-2^{63} + 1$ to $2^{63} - 1$	$-2^{31} + 1$ to $2^{31} - 1$
minValue	$(-2^{63} + 1) * 10^{127}$	$(-2^{63} + 1) * 10^{127}$	$(-2^{31} + 1) * 10^{127}$
maxValue	$(2^{63} - 1) * 10^{127}$	$(2^{63} - 1) * 10^{127}$	$(2^{31} - 1) * 10^{127}$
nullValue	mantissa=-263, exponent=-128	mantissa =-263	mantissa =-231

2.5.3 Encoding specifications for decimal types

Decimal encodings are composite types, consisting of two members, named mantissa and exponent. They must be listed in that order in the composite type. The exponent may either be serialized on the wire or may be set to constant. A constant exponent is a way to specify an assumed number of decimal places.

Decimal encoding specifications that an implementation must support

```
<composite name="decimal" >
  <type name="mantissa" primitiveType="int64" />
  <type name="exponent" primitiveType="int8" />
</composite>
```

```
<composite name="decimal32" >
  <type name="mantissa" primitiveType="int32" />
  <type name="exponent" primitiveType="int8"
    presence="constant">-2</type>
</composite>
```

```
<composite name="decimal64">
  <type name="mantissa" primitiveType="int64" />
  <type name="exponent" primitiveType="int8"
    presence="constant">-2</type>
</composite>
```

2.5.4 Composite encoding padding

When both mantissa and exponent are sent on the wire for a decimal, the elements are packed by default. However, byte alignment may be controlled by specifying offset of the exponent within the composite encoding. See section 4 below.

2.5.5 Examples of decimal fields

Examples show encoded bytes on the wire as hexadecimal digits, little-endian.

FIX Qty data type is a float type, but a decimal may be constrained to integer values by setting exponent to zero.

```
<composite name="intQty32">
  <type name="mantissa" primitiveType="int32" />
```

```

    <type name="exponent" primitiveType="int8"
      presence="constant">0</type>
  </composite>

```

A field uses the encoding and adds semanticType.

```

<field type="intQty32" name="OrderQty" id="38" semanticType="Qty"
  description="Total number of shares" />

```

Wire format of decimal 123.45 with 2 significant decimal places.

```
3930000000000000fe
```

Wire format of null floating-point decimal.

```
ffffffffffffffffffff
```

Wire format of decimal64 123.45 with 2 significant decimal places. Schema attribute exponent = -2

```
3930000000000000
```

Wire format of decimal32 123.45 with 2 significant decimal places. Schema attribute exponent = -2

```
39300000
```

2.6 Float encoding

Binary floating point encodings are compatible with IEEE Standard for Floating-Point Arithmetic (IEEE 754-2008). They should be used for floating point numeric fields that do not represent prices or monetary amounts. Examples include interest rates, volatility and dimensionless quantities such as ratios. On the other hand, decimal prices should be encoded as decimals; see [Decimal encoding](#) above.

2.6.1 Primitive types

Both single and double precision encodings are supported as primitive data types. See the IEEE 754-2008 standard for ranges and details of the encodings.

Primitive type	Description	IEEE 754-2008 format	Length (octets)
float	Single precision floating point	binary32	4
double	Double precision floating point	binary64	8

2.6.2 Null values

For both float and double precision encodings, null value of an optional field is represented by the Not-a-Number format (NaN) of the standard encoding. Technically, it is indicated by the so-called quiet NaN.

2.6.3 Byte order

Like integer encodings, floating point encodings follow the byte order specified by message schema. See section 4 for specification of message schema attributes, including byteOrder.

2.6.4 Float encoding specifications

These are valid encoding specifications for each of the floating point primitive types.

```
<type name="float" primitiveType="float" />
<type name="double" primitiveType="double" />
```

Designers should determine whether a field of FIX data type float should be represented by a binary floating point number or an exact decimal number. If it is the latter, then see "Decimal encoding" above.

2.6.5 Examples of floating point fields

Examples show encoded bytes on the wire as hexadecimal digits, little-endian.

A single precision ratio

```
<type name="ratio" primitiveType="float" />

<field type="ratio" name="CurrencyRatio" id="1382"
  semanticType="float"/>
```

Wire format of float 255.678

91ad7f43

Wire format of double 255.678

04560e2db2f56f40

2.7 String encodings

Character data may either be of fixed size or variable size. In Simple Binary Encoding, fixed-length fields are recommended in order to support direct access to data. Variable-length encoding should be reserved for character strings that cannot be constrained to a specific size. It may also be used for non-ASCII encoded strings.

2.7.1 Character

Character fields hold a single character. They are most commonly used for field with character code enumerations. See [Enumeration encoding](#) below for discussion of enum fields.

FIX data type	Description	Backing primitive	Length (octet)
char	A single US-ASCII character	char	1

2.7.1.1 Range attributes for char fields

Valid values of a char field are printable characters of the US-ASCII character set (codes 20 to 7E hex.) The implicit nullValue is the NUL control character (code 0).

Schema attribute	char
minValue	hex 20

Schema attribute	char
maxValue	hex 7e
nullValue	0

2.7.1.2 Encoding of char type

This is the standard encoding for char type. Note that the length attribute defaults to 1, producing a single character rather than a character array.

```
<type name="charType" primitiveType="char"/>
```

A field may be specified with a constant character value.

```
<field type="charType" name="OptAttribute" id="206"
presence="constant">P</field>
```

Wire format of char encoding of "A" (ASCII value 65, hexadecimal 41)

41

2.7.2 Fixed-length character array

Character arrays are allocated a fixed space in a message, supporting direct access to fields. A fixed size character array is distinguished from a variable length string by the presence of a length schema attribute or a constant attribute.

FIX data type	Description	Backing primitives	Length (octets)	Required schema attribute
String	character array	Array of char of specified length, delimited by NUL character if a string is shorter than the length specified for a field.	Specified by length attribute	length

A length attribute set to zero indicates variable length. See [Variable-length string](#).

2.7.2.1 Encoding specifications for fixed-length character array

A fixed-length character array encoding must specify primitiveType="char" and a length attribute is required.

Range attributes minValue and maxValue do not apply to fixed-length character arrays.

US-ASCII is the default encoding of character arrays to conform to usual FIX values. The characterEncoding attribute may be specified to override encoding.

2.7.2.2 Examples of fixed-length character arrays

A typical string encoding specification

```
<type name="string6" primitiveType="char" length="6" />
```

```
<field type="string6" name="Symbol" id="55" semanticType="String"/>
```

Wire format of a character array in character and hexadecimal formats

M S F T

4d5346540000

A character array constant specification. As for a non-constant value, if the constant is shorter than the specified length, it should be padded with NUL characters to fill the fixed-length array.

```
<type name="MarketID" primitiveType="char" length="4" description="MIC code"/>
```

```
<field name="EurexMarketID" type="MarketID" id="1301" semanticType="Exchange" presence="constant">XEUR</field>
```

2.7.3 Variable-length string encoding

Variable-length string encoding is used for variable length ASCII strings or embedded non-ASCII character data (like EncodedText field). A length member conveys the size of the string that follows.

On the wire, length immediately precedes the data.

The length member may not be null, but may be set to zero for an empty string. In that case, no space is reserved for the data. No distinction is made at an encoding layer between an empty string and a null string. Semantics of an empty variable-length string should be specified at an application layer.

FIX data type	Description	Backing primitives	Length (octets)
Length	The length of variable data in octets	primitiveType="uint8" or "uint16" May not hold null value.	1 or 2
data	Raw data	Array of octet of size specified in associated Length field. The data field itself should be specified as variable length. primitiveType="uint8" length="0" indicates variable length	variable

Optionally, implementations may support any other unsigned integer types for length.

2.7.4 Range attributes for string Length

Schema attribute	length uint8	length uint16	data
minValue	0	0	N/A
maxValue	254	65534	N/A

If the Length element has minValue and maxValue attributes, it specifies the minimum and maximum *length* of the variable-length data.

Range attributes `minValue`, `maxValue`, and `nullValue` do not apply to the data element.

If a field is required, both the `Length` and data fields must be set to a "required" attribute.

2.7.5 Encoding specifications for variable-length string

Variable length string is encoded as a composite type, consisting of a `length` member and `varData` member. In the composite specification, they must be listed with those names and in that order. The `length` attribute of the `varData` member is set to zero in the XML message schema as special value to indicate that the character data is of variable length.

To map an SBE data field specification to traditional FIX, the field ID of a data field is used. Its associated length is implicitly contained by the composite type rather than specified as a separate field.

Encoding specification for variable length data up to 65535 octets

```
<composite name="varString" description="Variable-length string">
  <type name="length" primitiveType="uint16"/>
  <type name="varData" length="0" primitiveType="uint8"
    characterEncoding="UTF-16"/>
</composite>
```

```
<data name="SecurityDesc" id="107" type="varString"/>
```

The `characterEncoding` attribute tells which variable-sized encoding is used if the data field represents encoded text. UTF-8 is the recommended encoding, but there is no default in the XML schema

2.7.6 Example of a variable-length string field

Example shows encoded bytes on the wire.

Wire format of variable-length String in character and hexadecimal formats, preceded by `uint16` length of 4 octets in little-endian byte order

M S F T

04004d534654

2.8 Data encodings

Raw data is opaque to SBE. In other words, it is not constrained by any value range or structure known to the messaging layer other than length. Data fields simply convey arrays of octets.

Data may either be of fixed-length or variable-length. In Simple Binary Encoding, fixed-length data encoding may be used for data of predetermined length, even though it does not represent a FIX data type. Variable-length encoding should be reserved for raw data when its length is not known until run-time.

2.8.1 Fixed-length data

Data arrays are allocated as a fixed space in a message, supporting direct access to fields. A fixed size array is distinguished from a variable length data by the presence of a length schema attribute rather than sending length on the wire.

FIX data type	Description	Backing primitives	Length (octets)	Required schema attribute
data	octet array	Array of uint8 of specified length.	Specified by length attribute	length

2.8.1.1 Encoding specifications for fixed-length data

A fixed-length octet array encoding should specify `primitiveType="uint8"` and a length attribute is required.

Data range attributes `minValue` and `maxValue` do not apply.

Since raw data is not constrained to a character set, `characterEncoding` attribute should not be specified.

2.8.1.2 Example of fixed-length data encoding

A fixed-length data encoding specification for a binary user ID

```
<type name="uuid" primitiveType="uint8" length="16" description="RFC 4122 compliant UUID"/>
```

```
<field type="uuid" name="Username" id="553" />
```

2.8.2 Variable-length data encoding

Variable-length data is used for variable length non-character data (such as `RawData`). A length member conveys the size of the data that follows. On the wire, length immediately precedes the data.

The length member may not be null, but it may be set to zero. In that case, no space is reserved for the data. Semantics of an empty variable-length data element should be specified at an application layer.

FIX data type	Description	Backing primitives	Length (octets)
Length	The length of variable data in octets	<code>primitiveType="uint8"</code> or <code>"uint16"</code> May not hold null value.	1 or 2
data	Raw data	Array of octet of size specified in associated Length member. The <code>varData</code> member should be specified as variable length. <code>primitiveType="uint8"</code>	variable

Optionally, implementations may support any other unsigned integer types for length.

2.8.3 Range attributes for variable-length data Length

Schema attribute	length uint8	length uint16	data
minValue	0	0	N/A
maxValue	254	65534	N/A

If the Length field has min**Value** and max**Value** attributes, it specifies the minimum and maximum *length* of the variable-length data. Data range attributes min**Value** , max**Value**, and null**Value** do not apply to a data field.

If a field is required, both the Length and data fields must be set to a "required" attribute.

2.8.4 Encoding specifications for variable-length data

Variable length data is encoded as composite type, consisting of a length member and varData member. They must be listed with those names and in that order in the composite specification.

To map an SBE data field specification to traditional FIX, the field ID of a data field is used. Its associated length is implicitly contained by the composite type rather than specified as a separate field.

Encoding specification for variable length data up to 65535 octets

```
<composite name="DATA" description="Variable-length data">
  <type name="length" primitiveType="uint16"/>
  <type name="varData" length="0" primitiveType="uint8"/>
</composite>
```

```
<data name="RawData" id="96" type="DATA"/>
```

2.8.5 Example of a data field

Example shows encoded bytes on the wire.

Wire format of data in character and hexadecimal formats, preceded by uint16 length of 4 octets in little-endian byte order

M S F T

04004d534654

2.9 MonthYear encoding

MonthYear encoding contains four members representing respectively year, month, and optionally day or week. A field of this type is not constrained to one date format. One message may contain only year and month while another contains year, month and day in the same field, for example.

Values are distinguished by position in the field. Year and month must always be populated for a non-null field. Day and week are set to special value indicating null if not present. If Year is set to the null value, then the entire field is considered null.

Member	Primitive type	Length (octets)	Null value
Year	uint16	2	65535
Month (1-12)	uint8	1	—
Day of the month(1-31) optional	uint8	1	255
Week of the month (1-5) optional	uint8	1	255

2.9.1 Composite encoding padding

The four members of MonthYear are packed at an octet level by default. However, byte alignment may be controlled by specifying offset of the elements within the composite encoding. See section 4.4.4.3 below.

2.9.2 Encoding specifications for MonthYear

MonthYear data type is based on a composite encoding that carries its required and optional elements. The members must listed with the names and in the order shown.

The standard encoding specification for MonthYear

```
<composite name="monthYear">
  <type name="year" primitiveType="uint16" presence="optional"
    nullValue="65536" />
  <type name="month" primitiveType="uint8" minValue="1" maxValue="12" />
  <type name="day" primitiveType="uint8" minValue="1" maxValue="31"
    presence="optional" nullValue="255" />
  <type name="week" description="week of month" primitiveType="uint8"
    minValue="1" maxValue="5" presence="optional" nullValue="255" />
</composite>
```

Example MonthYear field specification

```
<field type="monthYear" name="MaturityMonthYear" id="200"
semanticType="MonthYear"/>
```

Wire format of MonthYear 2014 June week 3 as hexadecimal

de0706ff03

2.10 Date and time encoding

Dates and times represent Coordinated Universal Time (UTC). This is the preferred date/time format, except where regulations require local time with time zone to be reported (see time zone encoding below).

2.10.1 Epoch

Each time type has an epoch, or start of a time period to count values. For timestamp and date, the standard epoch is the UNIX epoch, midnight January 1, 1970 UTC.

A time-only value may be thought of as a time with an epoch of midnight of the current day. Like current time, the epoch is also referenced as UTC.

2.10.2 Time unit

Time unit tells the precision at which times can be collected. Time unit may be serialized on the wire if timestamps are of mixed precision. On the other hand, if all timestamps have the same precision, then time unit may be set to a constant in the message schema. Then it need not be sent on the wire.

FIX data type	Description	Backing primitives	Length (octets)	Schema attributes
UTCTimestamp	UTC date/time Default: nanoseconds since Unix epoch Range Jan. 1, 1970 - July 21, 2554	uint64 time	8	epoch="unix" (default)
	timeUnit = second or millisecond or microsecond or nanosecond May be constant	uint8 unit	1	
UTCTimeOnly	UTC time of day only Default: nanoseconds since midnight today	uint64 time	8	
	timeUnit = second or millisecond or microsecond or nanosecond May be constant	uint8 unit	1	
UTCDateOnly	UTC calendar date Default: days since Unix epoch. Range: Jan. 1, 1970 - June 7, 2149	uint16	2	epoch="unix" (default)

2.10.3 Encoding specifications for date and time

Time specifications use an enumeration of time units. See [Enumeration encoding](#) below for a fuller explanation of enumerations.

Enumeration of time units:

```
<enum name="TimeUnit" encodingType="uint8">
  <validValue name="second">0</validValue>
  <validValue name="millisecond">3</validValue>
  <validValue name="microsecond">6</validValue>
  <validValue name="nanosecond">9</validValue>
</enum>
```

Composite type members must listed with the names and in the order shown.

Timestamp with variable time units:

```
<composite name="UTCTimestamp" description="UTC timestamp with precision on the wire">
  <type name="time" primitiveType="uint64" />
```

```

    <type name="unit" primitiveType="uint8" />
</composite>

```

Timestamp with constant time unit:

```

<composite name="UTCTimestampNanos" description="UTC timestamp with
nanosecond precision">
    <type name="time" primitiveType="uint64" />
    <type name="unit" primitiveType="uint8" presence="constant"
valueRef="TimeUnit.nanosecond" />
</composite>

```

Time only with variable time units:

```

<composite name="UTCTime" description="Time of day with precision on the
wire">
    <type name="time" primitiveType="uint64" />
    <type name="unit" primitiveType="uint8" />
</composite>

```

Time only with constant time unit:

```

<composite name="UTCTimeNanos" description="Time of day with millisecond
precision">
    <type name="time" primitiveType="uint64" />
    <type name="unit" primitiveType="uint8" presence="constant"
valueRef="TimeUnit.millisecond" />
</composite>

```

Date only specification:

```

<type name="date" primitiveType="uint16"/>

```

2.10.4 Examples of date/time fields

timestamp 14:17:22 Friday, October 4, 2024 UTC (20,000 days and 14 hours, 17 minutes and 22 seconds since the UNIX epoch) with default schema attributes

```

<composite name="UTCTimestampNanos" description="UTC timestamp with
nanosecond precision">
<type name="time" primitiveType="uint64" />
<type name="unit" primitiveType="uint8" presence="constant"
valueRef="TimeUnit.nanosecond" />
</composite>

```

Wire format of UTCTimestamp with constant time unit in little-Endian byte order

4047baa145fb17

time 10:24:39.123456000 (37,479 seconds and 123456000 nanoseconds since midnight UTC) with default schema attributes

```
<composite name="UTCTimeOnlyNanos" description="UTC time of day with
nanosecond precision">
  <type name="time" primitiveType="uint64" />
  <type name="unit" primitiveType="uint8" presence="constant"
valueRef="TimeUnit.nanosecond" />
</composite>
```

Wire format of UTCTimeOnly

10d74916220000

date Friday, October 4, 2024 (20,000 days since UNIX epoch) with default schema attributes

```
<type name="date" primitiveType="uint16"/>
```

Wire format of UTCTimeOnly

204e

2.11 Local date encoding

Local date is encoded the same as UTCTimeOnly, but it represents local time at the market instead of UTC time.

FIX data type	Description	Backing primitives	Length (octets)	Schema attributes
LocalMktDate	Local calendar date Default: days since Unix epoch. Range: Jan. 1, 1970 - June 7, 2149 local time	uint16	2	epoch="unix" (default)

The standard encoding specification for LocalMktDate

```
<type name="localMktDate" primitiveType="uint16"/>
```

2.12 Local time encoding

Time with time zone encoding should only be used when required by market regulations. Otherwise, use UTC time encoding (see above).

Time zone is represented as an offset from UTC in the ISO 8601:2004 format \pm hhmm.

2.12.1 TZTimestamp encoding

A binary UTCTimestamp followed by a number representing the time zone indicator as defined in ISO 8601:2004.

FIX data type	Description	Backing primitives	Length (octets)	Schema attributes
TZTimestamp	date/time with timezone Default: nanoseconds since Unix epoch Range Jan. 1, 1970 - July 21, 2554	uint64	8	epoch="unix" (default) Represents Jan. 1, 1970 local time
	timeUnit = second or millisecond or microsecond or nanosecond May be constant	uint8	1	
	Time zone hour offset	int8	1	None
	Time zone minute offset	uint8	1	None

2.12.2 Composite encoding padding

The members of TZTimestamp are packed at an octet level by default. However, byte alignment may be controlled by specifying offset of the elements within the composite encoding. See section 4 below.

The members must listed with the names and in the order shown.

Standard TZTimestamp encoding specification

```
<composite name="tzTimestamp">
  <type name="time" primitiveType="uint64" />
  <type name="unit" primitiveType="uint8" />
  <!-- Sign of timezone offset is on the hour member -->
  <type name="timezoneHour" primitiveType="int8" minValue="-12"
maxValue="14" />
  <type name="timezoneMinute" primitiveType="uint8" maxValue="59" />
</composite>
```

Wire format of TZTimestamp 8:30 17 September 2013 with Chicago time zone offset (-6:00)

0050d489fea22413fa00

2.12.3 TZTimeOnly encoding

A binary UTCTimeOnly followed by a number representing the time zone indicator as defined in ISO 8601:2004.

The time zone hour offset tells the number of hours different to UTC time. The time zone minute tells the number of minutes different to UTC. The sign telling ahead or behind UTC is on the hour member.

FIX data type	Description	Backing primitives	Length (octets)	Schema attributes
TZTimeOnly	Time of day only with time zone Default: nanoseconds since midnight today, local time	uint64	8	None

FIX data type	Description	Backing primitives	Length (octets)	Schema attributes
	timeUnit = second or millisecond or microsecond or nanosecond May be constant	uint8	1	None
	Time zone hour offset	int8	1	None
	Time zone minute offset	uint8	1	None

2.12.4 Composite encoding padding

The members of TZTimeOnly are packed at an octet level by default. However, byte alignment may be controlled by specifying offset of the elements within the composite encoding. See section 4 below.

Standard TZTimeOnly encoding specification

```
<composite name="tzTimeOnly">
  <type name="time" primitiveType="uint64" />
  <type name="unit" primitiveType="uint8" />
  <!-- Sign of timezone offset is on the hour member -->
  <type name="timezoneHour" primitiveType="int8"
    minValue="-12" maxValue="14" />
  <type name="timezoneMinute" primitiveType="uint8" minValue="0"
    maxValue="59" />
</composite>
```

Wire format of TZTimeOnly 8:30 with Chicago time zone offset (-6:00)

006c5ebe76000000fa00

2.13 Enumeration encoding

An enumeration conveys a single choice of mutually exclusive valid values.

2.13.1 Primitive type encodings

An unsigned integer or character primitive type is selected to contain the number of choices. Implementations must support char and uint8 types. They may additionally support other unsigned integer types to allow more choices.

Primitive type	Description	Length (octets)	Maximum number of choices
char	character	1	95
uint8	8-bit unsigned integer	1	255

2.13.2 Value encoding

If a field is of FIX data type char, then its valid values are restricted to US-ASCII printable characters. See [Character encoding](#) above.

If the field is of FIX data type int, then a primitive integer data type should be selected that can contain the number of choices. For most cases, an 8 bit integer will be sufficient, allowing 255 possible values.

Enumerations of other data types, such as String valid values specified in FIX, should be mapped to an integer wire format in SBE.

2.13.3 Encoding specification of enumeration

In a message schema, the choices are specified as `<validValue>` members of an `<enum>`. An `<enum>` specification must contain at least one `<validValue>`.

The name and value of a `validValue` element must be unique within an enumeration.

An `<enum>` element must have an `encodingType` attribute to specify the type of its values. Two formats of `encodingType` are acceptable:

- In-line style: the value of `encodingType` is its primitive data type.
- Reference style: the value of `encodingType` is the name of a `<type>` element that specifies the wire format.

The length of a `<type>` associated to an enumeration must be 1. That is, enumerations should only be backed by scalar types, not arrays.

2.13.4 Enumeration examples

These examples use a char field for enumerated code values.

Example enum lists acceptable values and gives the underlying encoding, which in this case is char (in-line style)

```
<enum name="SideEnum" encodingType="char">
  <validValue name="Buy">1</validValue>
  <validValue name="Sell">2</validValue>
  <validValue name="SellShort">5</validValue>
  <validValue name="SellShortExempt">6</validValue>
  <!-- not all FIX values shown -->
</enum>
```

Reference to type: This specification is equivalent to the one above.

```
<type name="charEnumType" primitiveType="char"/>
  <enum name="SideEnum" encodingType="charEnumType">
    <!-- valid values as above -->
  </enum>
```

Side field specification references the enumeration type

```
<field name="Side" type="SideEnum" id="54" />
```

Wire format of Side "Buy" code as hexadecimal

01

2.13.5 Constant field of an enumeration value

A constant field may be specified as a value of an enumeration. The attribute valueRef is a cross-reference to validValue entry by symbolic name.

Example of a char field using a constant enum value

```
<enum name="PartyIDSourceEnum" encodingType="char">
  <validValue name="BIC">B</validValue>
  <validValue name="GeneralIdentifier">C</validValue>
  <validValue name="Proprietary">D</validValue>
</enum>

<field type="PartyIDSourceEnum" name="PartyIDSource" id="447"
  description="Party ID source is fixed" presence="constant"
  valueRef="PartyIDSourceEnum.GeneralIdentifier" />
```

2.13.6 Boolean encoding

A Boolean field is a special enumeration with predefined valid values: true and false. Like a standard enumeration, an optional Boolean field may have nullValue that indicates that the field is null (or not applicable).

Standard encoding specifications for required and optional Boolean fields

```
<enum name="booleanEnum" encodingType="uint8">
  <validValue name="false">0</validValue>
  <validValue name="true">1</validValue>
</enum>
```

Example optional Boolean field

```
<field type="requiredBoolean" name="SolicitedFlag" id="377"
  presence="required" semanticType="Boolean" />

<field type="optionalBoolean" name="SolicitedFlag" id="377"
  presence="optional" nullValue="255" semanticType="Boolean" />
```

Wire format of true value as hexadecimal

01

Wire format of false value as hexadecimal

00

Wire format of null Boolean (or N/A) value as hexadecimal

ff

2.14 Multi-value choice encoding

A multi-value field conveys a choice of zero or more non-exclusive valid values.

2.14.1 Primitive type encodings

The binary encoding uses a bitset (a fixed-size sequence of bits, also known as bitmap, bit array or bit vector) to represent up to 64 possible choices. The encoding is backed by an unsigned integer. The smallest unsigned primitive type should be selected that can contain the number of valid choices.

Primitive type	Description	Length (octets)	Maximum number of choices
uint8	8-bit unsigned integer	1	8
uint16	16-bit unsigned integer	2	16
uint32	32-bit unsigned integer	4	32
uint64	64-bit unsigned integer	8	64

Like other integer-backed encodings, multi-value encodings follow the byte order specified by message schema when serializing to the wire. See section 4 for specification of message schema attributes, including byteOrder.

2.14.2 Value encoding

Each choice is assigned a bit of the primitive integer encoding, starting with the least significant bit. For each choice the value is selected or not, depending on whether its corresponding bit is set or cleared.

Any remaining unassigned bits in an octet should be cleared.

There is no explicit null value for multi-value choice encoding other than to set all bits off when no choices are selected.

2.14.3 Encoding specification of multi-value choice

In a message schema, the choices are specified as <choice> members of an <set> element. Choices are assigned values as an ordinal of bits in the bit set. The first Choice "0" is assigned the least significant bit; choice "1" is the second bit, and so forth.

The name and value (bit position) must be unique for element of a set.

A <set> element must have an encodingType attribute to specify the wire format of its values. Two formats of encodingType are recognized :

- In-line style: the value of encodingType is its primitive data type.
- Reference style: the value of encodingType is the name of a <type> element that specifies the wire format.

The length of a <type> associated to an bitset must be 1. That is, bitsets should not be specified as arrays.

2.14.4 Multi-value example

Example of a multi-value choice (was MultipleCharValue in tag-value encoding) Encoding type is in-line style.

```
<set name="FinancialStatusEnum" encodingType="uint8">
  <choice name="Bankrupt">0</choice>
  <choice name="Pending delisting">1</choice>
  <choice name="Restricted">2</choice>
</set>
```

Reference to type. This is equivalent to the example above.

```
<type name="u8Bitset" primitiveType="uint8"/>

<set name="FinancialStatusEnum" encodingType="u8Bitset">
<!--choices as above -->
</set>
```

A field using the multi-choice encoding

```
<field type="FinancialStatus" name="FinancialStatusEnum"
  id="291" semanticType="MultipleCharValue"/>
```

Wire format of choices "Bankrupt" + "Pending delisting" (first and second bits set)

03

2.15 Field value validation

These validations apply to message field values.

If a value violation is detected on a received message, the message should be rejected back to the counterparty in a way appropriate to the session protocol.

Error condition	Error description
Field value less than minValue	The encoded value falls below the specified valid range.
Field value greater than maxValue	The encoded value exceeds the specified valid range.
Null value set for required field	The null value of a data type is invalid for a required field.
String contains invalid characters	A String contains non-US-ASCII printable characters or other invalid sequence if a different characterEncoding is specified.
Required members not populated in MonthYear	Year and month must be populated with non-null values, and the month must be in the range 1-12.
UTCTimeOnly exceeds day range	The value must not exceed the number of time units in a day, e.g. greater than 86400 seconds.
TZTimestamp and TZTimeOnly has missing or invalid time zone	The time zone hour and minute offset members must correspond to an actual time zone recognized by international standards.
Value must match valid value of an enumeration field	A value is invalid if it does not match one of the explicitly listed valid values.

3 Message Structure

3.1 Message Framing

SBE messages have no defined message delimiter. SBE makes it possible to walk the elements of a message to determine its limit, even when the message has been extended. Nevertheless, since internal framing depends on a correct starting point and not encountering malformed messages, it may be desirable to use an external framing protocol when used with transports that do not preserve message boundaries, such as when they are transmitted on a streaming session protocol or when persisting messages in storage.

3.1.1 Simple Open Framing Header

FIX Protocol Ltd. offers the Simple Open Framing Header standard for framing messages encoded with binary wire formats, such as Simple Binary Encoding.

The framing header provides two features:

- An overall message size including headers to support framing
- An identifier of the encoding used in the message payload. This supports selecting the correct decoder in the case where multiple message encodings are used on a session. It also aids tooling such as protocol analyzers to identify message protocols contained in network packets.

While the Simple Open Framing Header specification is normative, the following is an interpretation of that standard as an SBE encoding. Note that the framing standard specifies that the framing header will always be encoded in big-endian byte order, also known as network byte order.

Simple Open Framing Header as an SBE composite encoding (big-endian)

```
<composite name="framingHeader"/>
  <type name="messageLength" primitiveType="uint32" />
  <type name="encodingType" primitiveType="uint16" />
</composite>
```

The values of encodingType used to indicate SBE payloads are currently defined as:

Encoding	encodingType value
SBE version 1.0 big-endian	0x5BE0
SBE version 1.0 little-endian	0xEB50

The Simple Open Framing Header specification also lists values for other wire formats.

3.2 SBE Message Encoding Header

The purpose of the message encoding header is to tell which message template was used to encode the message and to give information about the size of the message body to aid in decoding, even when a message template has been extended in a later version. See section 5 below for an explanation of the schema extension mechanism.

The fields of the SBE message header are:

- **Block length of the message root** - the total space reserved for the root level of the message not counting any repeating groups or variable-length fields.
- **Template ID** - identifier of the message template
- **Schema ID** - identifier of the message schema that contains the template
- **Schema version** - the version of the message schema in which the message is defined
- **Group count** - the number of repeating groups in the root level of the message
- **Variable-length field count** - the number of variable-length fields in the root level of the message

Block length is specified in a message schema, but it is also serialized on the wire. By default, block length is set to the sum of the sizes of body fields in the message. However, it may be increased to force padding at the end of block. See [Padding](#).

3.2.1 Message header schema

The header fields precede the message body of every message in a fixed position as shown below. Each of these fields must be encoded as an unsigned integer type. The encoding must carry the name "messageHeader".

The message header is encoded in the same byte order as the message body, as specified in a message schema. See section 4.

Recommended message header encoding

```
<composite name="messageHeader" description="Template ID and length of
message root">
  <type name="blockLength" primitiveType="uint16"/>
  <type name="templateId" primitiveType="uint16"/>
  <type name="schemaId" primitiveType="uint16"/>
  <type name="version" primitiveType="uint16"/>
  <type name="numGroups" primitiveType="uint16" />
  <type name="numVarDataFields" primitiveType="uint16" />
</composite>
```

The recommended header encoding is 12 octets.

Element	Description	Primitive type	Length (octets)	Offset
blockLength	Root block length	uint16	2	0
templateId	Template ID	uint16	2	2
schemald	Schema ID	uint16	2	4
version	Schema Version	uint16	2	6
numGroups	Number of repeating groups	uint16	2	8
numVarDataFields	Number of variable-length fields	uint16	2	10

Optionally, implementations may support any other unsigned integer types for blockLength.

3.2.2 Root block length

The total space reserved for the root level of the message not counting any repeating groups or variable-length fields. (Repeating groups have their own block length; see [Group block length](#) below. Length of a variable-length Data field is given by its corresponding Length field; see section 2.) Block length only represents message body fields; it does not include the length of the message header itself, which is a fixed size.

The block size must be at least the sum of lengths of all fields at the root level of the message, and that is its default value. However, it may be set larger to reserve more space to effect alignment of blocks. This is specified by setting the `blockLength` attribute in a message schema.

3.2.3 Template ID

The identifier of a message type in a message schema. See section 4 for schema attributes of a message.

3.2.4 Schema ID

The identifier of a message schema. See section 4 for schema attributes.

3.2.5 Schema version

The version number of the message schema that was used to encode a message. See section 4 for schema attributes.

3.2.6 Number of repeating groups

A count of repeating groups at the root level of the message. The count does not include nested repeating groups.

3.2.7 Number of variable-length fields

A count of the variable-length fields at the root level of the message. The count does not include variable-length fields within repeating groups.

3.3 Message Body

The message body conveys the business information of the message.

3.3.1 Data only on the wire

In SBE, fields of a message occupy proximate space without delimiters or metadata, such as tags.

3.3.2 Direct access

Access to data is positional, guided by a message schema that specifies a message type.

Data fields in the message body correspond to message schema fields; they are arranged in the same sequence. The first data field has the type and size specified by the first message schema field, the second data field is described by the second message schema field, and so forth. Since a message decoder follows the field descriptions in the schema for position, it is not necessary to send field tags on the wire.

In the simplest case, a message is flat record with a fixed length. Based on the sequence of field data types, the offset to a given data field is constant for a message type. This offset may be computed in advance, based on a message schema. Decoding a field consists of accessing the data at this fixed location.

3.3.3 Field position and padding

3.3.3.1 No padding by default

By default, there is no padding between fields. In other words, a field value is packed against values of its preceding and following fields. No consideration is given to byte boundary alignment.

By default, the position of a field in a message is determined by the sum of the sizes of prior fields, as they are defined by the message schema.

```
<field name="ClOrdID" id="11" type="string14"
  semanticType="String"/>
<field name="Side" id="54" type="char" semanticType="char"/>
<field name="OrderQty" id="38" type="intQty32"
  semanticType="Qty"/>
<field name="Symbol" id="55" type="string8" semanticType="String"/>
```

Field	Size	Offset
ClOrdID	14	0
Side	1	14
OrderQty	4	15
Symbol	8	19

3.3.3.2 Field offset specified by message schema

If a message designer wishes to introduce padding or control byte boundary alignment or map to an existing data structure, field offset may optionally be specified in a message schema. Field offset is the number of octets from the start of the message body or group to the first octet of the field. Offset is a zero-based index.

If specified, field offset must be greater than or equal to the sum of the sizes of prior fields. In other words, an offset is invalid if it would cause fields to overlap.

Extra octets specified for padding should never be interpreted as business data. They should be filled with binary zeros.

The offset attribute is mutually exclusive with alignment; see below.

Example of fields with specified offsets

```
<field name="ClOrdID" id="11" type="string14" offset="0"
  semanticType="String"/>
<field name="Side" id="54" type="char" offset="14"
  semanticType="char"/>
<field name="OrderQty" id="38" type="intQty32" offset="16"
```



```

    semanticType="Qty"/>
  <field name="Symbol" id="55" type="string8" offset="20"
    semanticType="String"/>

```

Field	Size	Padding preceding field	Offset
ClOrdID	14	0	0
Side	1	0	14
OrderQty	4	1	16
Symbol	8	0	20

3.3.3.3 Field byte alignment

As an alternative to specifying a field's offset from the start of a block, its byte alignment may be specified. If specified, the start of a field should reside at an offset from the start of the buffer that is the next multiple of the alignment attribute. If the offset without alignment would not be a multiple of the alignment attribute, then the buffer is padded with octets to that multiple. The padding should be set to binary zeroes.

The alignment attribute is mutually exclusive with offset.

If neither offset nor alignment attributes are set, then a field is positioned at the next available octet in a buffer without padding.

Computation of padding and effective offset is as follows, where effective-offset is the resulting position in the buffer:

```

padding = (alignment - (unaligned-offset modulo alignment)) modulo alignment
effective-offset = unaligned-offset + padding

```

Example of fields with alignment that is equivalent to the offset example above:

```

<field name="ClOrdID" id="11" type="string14" semanticType="String"/>
<field name="Side" id="54" type="char" semanticType="char"/>
<field name="OrderQty" id="38" type="intQty32" alignment="4"
  semanticType="Qty"/>
<field name="Symbol" id="55" type="string8" alignment="4"
  semanticType="String"/>

```

3.3.3.4 Padding at end of a message or group

In order to force messages or groups to align on byte boundaries or map to an existing data structure, they may optionally be specified to occupy a certain space with a `blockLength` attribute in the message schema. The extra space is padded at the end of the message or group. If specified, `blockLength` must be greater than or equal to the sum of the sizes of all fields in the message or group.

The `blockLength` attribute applies only to the portion of message that contains fix-length fields; it does not apply to variable-length data elements of a message.

Extra octets specified for padding should be filled with binary zeros.

Example of blockLength specification for 24 octets

```
<message name="ListOrder" id="2" blockLength="24">
```

3.4 Repeating Groups

A repeating group is a message structure that contains a variable number of entries. Each entry contains fields specified by a message schema.

The order and data types of the fields are the same for each entry in a group. That is, the entries are homogeneous. Position of a given field within any entry is fixed, with the exception of variable-length fields.

A message may have no groups or an unlimited number of repeating groups specified in its schema.

3.4.1 Schema specification of a group

A repeating group is defined in a message schema by adding a <group> element to a message template. An unlimited number of <field> elements may be added to a group, but a group must contain at least one field.

Example repeating group encoding specification

```
<group name="Parties" id="1012" blockLength="16">
  <field name="PartyID" id="448" type="string14"
    semanticType="String"/>
  <field name="PartyIDSource" id="447" type="char"
    semanticType="char"/>
  <field name="PartyRole" id="452" type="uint8" semanticType="int"/>
</group>
```

3.4.2 Group block length

The blockLength part of a group dimension represents total space reserved for each group entry, not counting any nested repeating groups or variable-length fields. (Length of a variable-length Data field is given by its corresponding Length field.) Block length only represents message body fields; it does not include the length of the group dimension itself, which is a fixed size.

3.4.3 Byte alignment of a group entry

If byte alignment of a repeating group is specified, the start of the first field of each entry should reside at an offset from the start of the buffer that is the next multiple of the alignment attribute. If the offset without alignment would not be a multiple of the alignment attribute, then the buffer is padded with octets to that multiple. The padding should be set to binary zeroes.

If byte alignment is not specified for a group, then no padding is introduced to align entries.

3.4.4 Padding at end of a group entry

By default, the space reserved for an entry is the sum of a group's field lengths, as defined by a message schema, without regard to byte alignment.

The space reserved for an entry may optionally be increased to effect alignment of entries or to plan for future growth. This is specified by adding the group attribute `blockLength` to reserve a specified number of octets per entry. If specified, the extra space is padded at the end of each entry and should be set to zeroes by encoders. The `blockLength` value does not include the group dimensions itself.

Note that padding will only result in deterministic alignment if the repeating group contains no variable-length fields. Therefore, the alignment attribute is a more direct solution to that need.

3.4.5 Entry counter

Each group is associated with a required counter field of semantic data type `NumInGroup` to tell how many entries are contained by a message. The value of the counter is a non-negative integer. See "Encoding of repeating group dimensions" section below for encoding of that counter.

3.4.6 Empty group

The space reserved for all entries of a group is the product of the space reserved for each entry times the value of the associated `NumInGroup` counter. If the counter field is set to zero, then no entries are sent in the message, and no space is reserved for entries. The group dimensions including the zero-value counter is still transmitted, however.

3.4.7 Multiple repeating groups

A message may contain multiple repeating groups at the same level.

Example of encoding specification with multiple repeating groups

```
<message name="ExecutionReport" id="8">
  <group name="ContraGrp" id="2012">
    <!-- ContraGrp group fields -->
  </group>
  <group name="PreAllocGrp" id="2026">
    <!-- PreAllocGrp group fields -->
  </group>
</message>
```

3.4.8 Nested repeating group specification

Repeating groups may be nested to an arbitrary depth. That is, a `<group>` in a message schema may contain one or more `<group>` child elements, each associated with their own counter fields.

The encoding specification of nested repeating groups is in the same format as groups at the root level of a message in a recursive procedure.

Example of nested repeating group specification

```
<group name="ListOrdGrp" id="2030">
  <field name="ClOrdID" id="11" type="string14" semanticType="String"/>
  <field name="ListSeqNo" id="67" type="uint32" semanticType="int"/>
  <field name="Symbol" id="55" type="string8" semanticType="String"/>
  <field name="Side" id="54" type="char" semanticType="char"/>
  <field name="OrderQty" id="38" type="intQty32" semanticType="Qty"/>
```

```

<group name="Parties" id="1012">
  <field name="PartyID" id="448" type="string14"
semanticType="String"/>
  <field name="PartyRole" id="452" type="int" semanticType="int"/>
</group>
</group>

```

3.4.9 Nested repeating group wire format

Nested repeating groups are encoded on the wire by a depth-first walk of the data hierarchy. For example, all inner entries under the first outer entry must be encoded before encoding outer entry 2. (This is the same element order as FIX tag=value encoding.)

On decoding, nested repeating groups do not support direct access to fields. It is necessary to walk all elements in sequence to discover the number of entries in each repeating group.

3.4.10 Empty group means nested group is empty

If a group contains nested repeating groups, then a NumInGroup counter of zero implies that both that group and its child groups are empty. In that case, no NumInGroup is encoded on the wire for the child groups.

3.4.11 Group dimension encoding

Every repeating group must be immediately preceded on the wire by its dimensions. The two dimensions are the count of entries in a repeating group and the space reserved for each entry of the group.

3.4.11.1 Range of group entry count

Implementations should support uint8 and uint16 types for repeating group entry counts. Optionally, implementations may support any other unsigned integer types.

By default, the minimum number of entries is zero, and the maximum number is the largest value of the primitiveType of the counter.

Primitive type	Description	Length (octets)	Maximum number of entries
uint8	8-bit unsigned integer	1	255
uint16	16-bit unsigned integer	2	65535

The number of entries may be restricted to a specific range; see "Restricting repeating group entries" below.

3.4.11.2 Encoding of repeating group dimensions

Conventionally in FIX, a NumInGroup field conveys the number of entries in a repeating group. In SBE, the encoding conveys two dimensions: the number of entries and the length of each entry in number octets. Therefore, the encoding is a composite of those two elements. Block length and entry count subfields must be encoded as unsigned integer types.

By default, the name of the group dimension encoding is `groupSizeEncoding`. This name may be overridden by setting the `dimensionType` attribute of a `<group>` element.

Recommended encoding of repeating group dimensions

```
<composite name="groupSizeEncoding">
  <type name="blockLength" primitiveType="uint16"/>
  <type name="numInGroup" primitiveType="uint16"/>
  <type name="numGroups" primitiveType="uint16" />
  <type name="numVarDataFields" primitiveType="uint16" />
</composite>
```

3.4.11.3 Block length

The total space reserved for the fixed-length fields of this repeating group, not counting any repeating groups or variable-length fields.

3.4.11.4 Number of entries

The number of entries in this repeating group, called `NumInGroup` in FIX.

3.4.11.5 Number of repeating groups

A count nested repeating groups in this repeating group.

3.4.11.6 Number of variable-length fields

A count of the variable-length fields in this repeating group.

Wire format of `NumInGroup` with block length 55 octets by 3 entries, containing one nested group and two variable-length fields.

```
3700030001000200
```

3.4.11.7 Restricting repeating group entries

The occurrences of a repeating group may be restricted to a specific range by modifying the `numInGroup` member of the group dimension encoding. The `minValue` attribute controls the minimum number of entries, overriding the default of zero, and the `maxValue` attribute restricts the maximum entry count to something less than the maximum corresponding to its `primitiveType`. Either or both attributes may be specified.

Example of a restricted group encoding

```
<type name="numInGroup" primitiveType="uint16" minValue="1" maxValue="10" />
```

3.5 Sequence of message body elements

3.5.1 Root level elements

To maximize deterministic field positions, message schemas must be specified with this sequence of message body elements:

- i. Fixed-length fields that reside at the root level of the message (that is, not members of repeating groups), including any of the following, in the order specified by the message schema::
 - a. Fixed-length scalar fields, such as integers
 - b. Fixed-length character arrays
 - c. Fixed-length composite types, such as MonthYear
- ii. Repeating groups, if any.
- iii. Data fields, including raw data and variable-length strings, if any.

3.5.2 Repeating group elements

Repeating group entries are recursively organized in the same fashion as the root level: fixed-length fields, then nested repeating groups, and finally, variable-length data fields.

3.6 Message structure validation

Aside from message schema validations (see section 4), these validations apply to message structure.

If a message structure violation is detected on a received message, the message should be rejected back to the counterparty in a way appropriate to the session protocol.

Error condition	Error description
Wrong message size in header	A message size value smaller than the actual message may cause a message to be truncated.
Wrong or unknown template ID in header	A mismatch of message schema would likely render a message unintelligible or cause fields to be misinterpreted.
Fixed-length field after repeating group or variable-length field	All fixed-length fields in the root of a message or in a repeating group entry must be listed before any (nested) repeating group or variable-length field.
Repeating group after variable-length field	All repeating groups at the root level or in a nested repeating group must be listed before any variable length field at the same level.

4 Message Schema

4.1 XML schema for SBE message schemas

See [sbe.xsd](#) for the normative XML Schema Definition (XSD) for SBE.

The SBE schema is defined with W3C XML Schema Definition Language (XSD) version 1.1.

4.2 XML namespace

The Simple Binary Encoding XML schema is identified by this URL [*tentative*]:

```
xmlns:sbe=http://fixprotocol.io/2017/sbe
```

Conventionally, the URI of the XML schema is aliased by the prefix "sbe".

Caution: Users should treat the SBE XML namespace as a URI (unique identifier), not as a URL (physical resource locator). Firms should not depend on access to the FIX Trading Community web site to validate XML schemas at run-time

4.3 Name convention

All symbolic names in a message schema are restricted to alphanumeric characters plus underscore without spaces. This is the same restriction applied to all names in FIX specifications.

4.3.1 Capitalization

The value of a field's `semanticType` attribute is a FIX data type. In this document, FIX types are capitalized exactly as in the FIX repository, from which all official FIX documentation and references are derived. Since the capitalization is somewhat inconsistent, however, it is recommended that matching of type names should be case insensitive in schema parsers.

4.4 Root element

The root element of the XML document is `<messageSchema>`.

4.4.1 `<messageSchema>` attributes

The root element provides basic identification of a schema.

The `byteOrder` attribute controls the byte order of integer encodings within the schema. It is a global setting for all specified messages and their encodings.

Schema attribute	Description	XML type	Usage	Valid values
package	Name or category of a schema	string	optional	Should be unique between counterparties but no naming convention is imposed.
id	Unique identifier of a schema	unsignedInt		Should be unique between counterparties
version	Version of this schema	nonnegativeInteger		Initial version is zero and is incremented for each version
semanticVersion	Version of FIX semantics	string	optional	FIX versions, such as "FIX.5.0_SP2"
byteOrder	Byte order of encoding	token	default = littleEndian	littleEndian bigEndian
description	Documentation of the schema	string	optional	
headerType	Name of the encoding type of the message header, which is the same for all messages in a schema.	string	default=messageHeader	An encoding with this name must be contained by `.

4.4.2 Schema versioning

Changes to a message schema may be tracked by its `version` attribute. A version of a schema is a snapshot in time. All elements in a given generation of the schema share the same version number. That is, elements are not versioned individually. By convention, the initial version of a schema is version zero, and subsequent changes increment the version number.

The `package` attribute should remain constant between versions, if it is supplied.

4.5 Data encodings

4.5.1 Encoding sets

The `<types>` element contains one or more sets of data encodings used for messages within the schema.

Within each set, an unbound number of encodings will be listed in any sequence:

- Element `<type>` defines a simple encoding
- Element `<composite>` defines a composite encoding
- Element `<enum>` defines an enumeration
- Element `<set>` defines a multi-value choice bitset encoding

4.5.2 Encoding name

The namespace for encoding names is global across all encodings included in a schema, including simple, composite and enumeration types. That is, the name must be unique among all encoding instances.

All symbolic names should be alphanumeric without spaces.

4.5.2.1 Importing encodings

A suggested usage is to import common encodings that are used across message schemas as one set while defining custom encodings that are particular to a schema in another set.

Example of XML include usage to import common encoding types

```
<!-- included XML contains a <types> element -->
<xi:include href="sbe-builtins.xml"/>
```

4.5.3 Simple encodings

A simple encoding is backed by either a scalar type or an array of scalars, such as a character array. One or more simple encodings may be defined, each specified by a `<type>` element.

The `<type>` XML element is also used for a scalar or array member of a composite encoding; see below.

4.5.3.1 `<type>` element content

If the element has a value, it is used to indicate a special value of the encoding.

4.5.3.1.1 Constant value

The element value represents a constant if attribute `presence="constant"`. In this case, the value is required.

4.5.3.2 `<type>` attributes

<code><type></code> attribute	Description	XML type	Usage	Valid values
name	Name of encoding	symbolicName_t	required	Must be unique among all encoding types in a schema.
description	Documentation of the type	string	optional	
presence	Presence of this element within a composite type	token		required optional constant

<type> attribute	Description	XML type	Usage	Valid values
length	Number of elements of the primitive data type	nonnegativeInteger	default = 1	Value "0" represents variable length.
offset	If a member of a composite type, tells the offset from the beginning of the composite. By default, the offset is the sum of preceding element sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	See section 4.4.4.3 below
primitiveType	The primitive data type that backs the encoding	token	required	char int8 int16 int32 int64 uint8 uint16 uint32 uint64 float double
sinceVersion	Documents the version of a schema in which a type was added. Not valid on a member of a composite type.	nonnegativeInteger	default = 0	Must be less than or equal to the version of the message schema.
deprecated	Documents the version of a schema in which a type was deprecated. It should no longer be used in new messages. Not valid on a member of a composite type.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

Simple type examples

```
<type name="FLOAT" primitiveType="double"/>
<type name="idString" length="8" primitiveType="char"/>
```

4.5.4 Composite encodings

Composite encoding types are composed of two or more simple types.

4.5.4.1 <composite> attributes

<composite> attribute	Description	XML type	Usage	Valid values
name	Name of encoding	symbolicName_t	required	Must be unique among all encoding types.

<composite> attribute	Description	XML type	Usage	Valid values
offset	The offset from the beginning of the composite. By default, the offset is the sum of preceding element sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	
description	Documentation of the type	string	optional	
sinceVersion	Documents the version of a schema in which a type was added	nonnegativeInteger	default = 0	Must be less than or equal to the version of the message schema.
deprecated	Documents the version of a schema in which a type was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

4.5.4.2 Composite type elements

A <composite> composite encoding element may be composed of any combination of types, including <type> simple encoding, <enum> enumeration, <set> bitset, and nested composite type. The elements that compose a composite type carry the same XML attributes as stand-alone types.

Composite type example

In this example, a Price is encoded as 32 bit integer mantissa and a constant exponent, which is not sent on the wire.

```
<composite name="decimal32">
  <type name="mantissa" primitiveType="int32" />
  <type name="exponent" primitiveType="int8"
    presence="constant">-4</type>
</composite>
```

4.5.4.3 Element offset within a composite type

If a message designer wishes to control byte boundary alignment or map to an existing data structure, element offset may optionally be specified on a simple type, enum or bitset within a composite type. Offset is the number of octets from the start of the composite; it is a zero-based index.

If specified, offset must be greater than or equal to the sum of the sizes of prior elements. In other words, an offset is invalid if it would cause elements to overlap.

4.5.4.4 Null value of a composite type

For a composite type, nullness is indicated by the value of its first element. For example, if a price field is optional, a null value in its mantissa element indicates that the price is null.

4.5.5 Reference to reusable types

A composite type often has its elements defined in-line within the <composite> XML element as shown in the example above. Alternatively, a common type may be defined once on its own, and then referred to by name with the composite type using a <ref> element.

4.5.5.1 <ref> attributes

<ref> attribute	Description	XML type	Usage	Valid values
name	Usage of the type in this composite	symbolicName_t	required	
type	Name of referenced encoding	symbolicName_t	required	Must match a defined type, enum or set or composite name attribute.
description	Documentation of how the referenced type is used	string	optional	
offset	The offset from the beginning of the composite. By default, the offset is the sum of preceding element sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	

4.5.5.2 Type reference examples

Reference to an enum

In this example, a futuresPrice is encoded as 64 bit integer mantissa, 8 bit exponent, and a reused enum type.

```
<enum name="booleanEnum" encodingType="uint8">
  <validValue name="false">0</validValue>
  <validValue name="true">1</validValue>
</enum>

<composite name="futuresPrice">
  <type name="mantissa" primitiveType="int64" />
  <type name="exponent" primitiveType="int8" />
  <ref name="isSettlement" type="boolEnum" />
</composite>
```

Reference to a composite type

In this example, a nested composite is formed by using a reference to another composite type. It supports the expression of a monetary amount with its currency, such as USD150.45. Note that a reference may carry an offset within the composite encoding that contains it.

```
<composite name="price">
  <type name="mantissa" primitiveType="int64" />
  <type name="exponent" primitiveType="int8" />
</composite>

<composite name="money">
  <type name="currencyCode" primitiveType="char" length="3"
semanticType="Currency" />
  <ref name="amount" type="price" offset="3" />
</composite>
```

4.5.6 Enumeration encodings

An enumeration explicitly lists the valid values of a data domain. Any number of fields may share the same enumeration.

4.5.6.1 <enum> element

Each enumeration is represented by an <enum> element. It contains any number of <validValue> elements.

The encodingType attribute refers to a simple encoding of scalar type. The encoding of an enumeration may be char or any unsigned integer type.

<enum> attribute	Description	XML type	Usage	Valid values
name	Name of encoding	symbolicName_t	required	Must be unique among all encoding types.
description	Documentation of the type	string	optional	
encodingType	Name of a simple encoding type	symbolicName_t	required	Must match the name attribute of a scalar <type> element <i>or</i> a primitive type: char uint8 uint16 uint32 uint64
sinceVersion	Documents the version of a schema in which a type was added	nonnegativeInteger	default = 0	Must be less than or equal to the version of the message schema.

<enum> attribute	Description	XML type	Usage	Valid values
deprecated	Documents the version of a schema in which a type was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.
offset	If a member of a composite type, tells the offset from the beginning of the composite. By default, the offset is the sum of preceding element sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	

4.5.6.2 <validValue> element attributes

The name attribute of the <validValue> uniquely identifies it.

<validValue> attribute	Description	XML type	Usage	Valid values
name	Symbolic name of value	symbolicName_t	required	Must be unique among valid values in the enumeration.
description	Documentation of the value	string	optional	
sinceVersion	Documents the version of a schema in which a value was added	nonNegativeInteger	default = 0	
deprecated	Documents the version of a schema in which a value was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

4.5.6.3 <validValue> element content

The element is required to carry a value, which is the valid value as a string. The string value in XML must be convertible to the data type of the encoding, such as an integer.

<enum> and <validValue> elements

Enumeration example (not all valid values listed)

This enumeration is encoded as an 8 bit unsigned integer value. Others are encoded as char codes.

```
<type name="intEnum" primitiveType="uint8" />
```

```
<enum name="PartyRole" encodingType="intEnum">
```

```

<validValue name="ExecutingFirm">1</validValue>
<validValue name="BrokerOfCredit">2</validValue>
<validValue name="ClientID">3</validValue>
<validValue name="ClearingFirm">4</validValue>
</enum>

```

4.5.7 Multi-value choice encodings (bitset)

An enumeration explicitly lists the valid values of a data domain. Any number of fields may share the same set of choices.

4.5.7.1 <set> element

Each multi-value choice is represented by a <set> element. It may contain a number of <choice> elements up to the number of bits in the primitive encoding type. The largest number possible is 64 choices in a uint64 encoding.

The encodingType attribute refers to a simple encoding of scalar type. The encoding of a bitset should be an unsigned integer type.

<set> attribute	Description	XML type	Usage	Valid values
name	Name of encoding	symbolicName_t	required	Must be unique among all encoding types.
description	Documentation of the type	string	optional	
encodingType	Name of a simple encoding type	string	required	Must match the name attribute of a scalar <type> element <i>or</i> a primitive type: uint8 uint16 uint32 uint64
sinceVersion	Documents the version of a schema in which a type was added	nonnegativeInteger	default = 0	Must be less than or equal to the version of the message schema.
deprecated	Documents the version of a schema in which a type was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

<set> attribute	Description	XML type	Usage	Valid values
offset	If a member of a composite type, tells the offset from the beginning of the composite. By default, the offset is the sum of preceding element sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	

4.5.7.2 <choice> element attributes

The name attribute of the <choice> uniquely identifies it.

<choice> attribute	Description	XML type	Usage	Valid values
name	Symbolic name of value	symbolicName_t	required	Must be unique among choices in the set.
description	Documentation of the value	string	optional	
sinceVersion	Documents the version of a schema in which a choice was added	nonNegativeInteger	default = 0	
deprecated	Documents the version of a schema in which a choice was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

4.5.7.3 < choice > element content

The element is required to carry a value, which is an unsigned integer representing a zero-based index to a bit within a bitset. Zero is the least significant bit.

<set> and <choice> XML elements

Multi-value choice example, The choice is encoded as a bitset.

```
<type name="bitset" primitiveType="uint8" />
```

```
<set name="Scope" encodingType="bitset" >
  <choice name="LocalMarket">0</choice>
  <choice name="National">1</choice>
  <choice name="Global">2</choice>
</set>
```


4.6 Message template

To define a message type, add a `<message>` element to the root element of the XML document, `<messageSchema>`.

The name and `id` attributes are required. The first is a display name for a message, while the latter is a unique numeric identifier, commonly called template ID.

4.6.1 Byte alignment

By default, a message starts at the first available byte of buffer. However, if desired, it may be aligned to a boundary by specifying the `alignment` attribute to improve performance on certain platforms. If specified, the start of the message should reside in the buffer at the next multiple of the alignment attribute. For example, if `alignment=4` then the message should start at the next address that is a multiple of 4.

4.6.2 Reserved space

By default, message size is the sum of its field lengths. However, a larger size may be reserved by setting `blockLength`, either to allow for future growth or for desired byte alignment. If so, the extra reserved space should be filled with zeros by message encoders.

4.6.3 Message members

A `<message>` element contains its field definitions in three categories, which must appear in this sequence:

- i. Element `<field>` defines a fixed-length field
- ii. Element `<group>` defines a repeating group
- iii. Element `<data>` defines a variable-length field, such as raw data

The number of members of each type is unbound.

4.6.4 Member order

The order that fields are listed in the message schema governs the order that they are encoded on the wire.

`<message>` element attributes

<code><message></code> attribute	Description	XML type	Usage	Valid values
<code>name</code>	Name of a message	<code>symbolicName_t</code>	required	Must be unique among all messages in a schema
<code>id</code>	Unique message template identifier	<code>unsignedInt</code>	required	Must be unique within a schema
<code>description</code>	Documentation	<code>string</code>	optional	

<message> attribute	Description	XML type	Usage	Valid values
blockLength	Reserved size in number of octets for root level of message body	unsignedInt	optional	If specified, must be greater than or equal to the sum of field lengths.
semanticType	Documents value of FIX MsgType for a message	token	optional	Listed in FIX specifications
sinceVersion	Documents the version of a schema in which a message was added	nonNegativeInteger	default = 0	
deprecated	Documents the version of a schema in which a message was deprecated. It should no longer be sent but is documented for back-compatibility.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.
alignment	Controls byte alignment of the start of a message in its buffer	positiveInteger	optional	

Note that there need not be a one-to-one relationship between message template (identified by `id` attribute) and `semanticType` attribute. You might design multiple templates for the same FIX MsgType to optimize different scenarios.

Example <message> element

```
<sbe:message name="NewOrderSingle" id="2" semanticType="D">
```

4.7 Field attributes

Fields are added to a <message> element as child elements. See Field Encoding section above for a listing of all field types.

These are the common attributes of all field types.

Schema attribute	Description	XML type	Usage	Valid values
name	Name of a field	symbolicName_t	required	Name and id must uniquely identify a field type within a message schema.
id	Unique field identifier (FIX tag)	unsignedShort	required	
description	Documentation	string	optional	

Schema attribute	Description	XML type	Usage	Valid values
type	Encoding type name, one of simple type, composite type or enumeration.	string	required	Must match the name attribute of a simple <type>, <composite> encoding type, <enum> or <set>.
offset	Offset to the start of the field within a message or repeating group entry. By default, the offset is the sum of preceding field sizes, but it may be increased to effect byte alignment.	unsignedInt	optional	Must be greater than or equal to the sum of preceding field sizes. Mutually exclusive with alignment
alignment	Controls byte alignment of the start of a field in its buffer	positiveInteger	optional	Mutually exclusive with offset
presence	Field presence	enumeration	Default = required	required = field value is required; not tested for null. optional = field value may be null. constant = constant value not sent on wire.
valueRef	Constant value of a field as a valid value of an enumeration	qualifiedName_t	optional Valid only if presence="constant"	If provided, the qualified name must match the name attribute of a <validValue> within an <enum>
sinceVersion	The version of a message schema in which this field was added.	nonnegativeInteger	default=0	Must not be greater than version attribute of <messageSchema> element.
deprecated	Documents the version of a schema in which a field was deprecated. It should no longer be used in new messages.	nonnegativeInteger	optional	Must be less than or equal to the version of the message schema.

Example field schemas

Field that uses a composite encoding

```

<composite name="intQty32">
  <type name="mantissa" primitiveType="int32" />
  <type name="exponent" primitiveType="int8"
    presence="constant">0\</type>
</composite>

<field type="intQty32" name="OrderQty" id="38" offset="16"
  description="Shares: Total number of shares" />

```

4.8 Repeating group schema

A <group> has the same attributes as a <message> element since they both inherit attributes from the blockType XML type. A group has the same child members as a message, and they must appear in the same order:

- i. Element <field> defines a fixed-length field
- ii. Element <group> defines a repeating group. Groups may be nested to any level.
- iii. Element <data> defines a variable-length field, such as raw data

The number of members of each type is unbound.

<group> attribute	Description	XML type	Usage	Valid values
name	Name of a group	symbolicName_t	required	Name and id must uniquely identify a group type within a message schema.
id	Unique group identifier	unsignedShort	required	
description	Documentation	string	optional	
dimensionType	Dimensions of the repeating group	symbolicName_t	default = groupSizeEncoding	If specified, must be greater than or equal to the sum of field lengths.
alignment	Controls byte alignment of the start of each group instance in its buffer	positiveInteger	optional	
<group> element inherits attributes of blockType. See <message>				
above.				

Example group schema with default dimension encoding

```
<composite name="groupSizeEncoding">
  <type name="blockLength" primitiveType="uint16"/>
  <type name="numInGroup" primitiveType="uint16"/>
  <type name="numGroups" primitiveType="uint16" />
  <type name="numVarDataFields" primitiveType="uint16" />
</composite>

<group name="Parties" id="1012" >
  <field type="string14" name="PartyID" id="448" />
  <field type="partyRoleEnum" name="PartyRole" id="452" />
</group>
```

4.9 Schema validation

The first level of schema validation is enforced by XML schema validation tools to make sure that a schema is well-formed according to XSD schema rules. Well-formed XML is necessary but insufficient to prove that a schema is correct according to FIX Simple Binary Encoding rules.

Additional conditions that render a schema invalid include the following.

Error condition	Error description
Missing field encoding	A field or <enum> references a type name that is undefined.
Missing message header encoding	Missing encoding type for headerType specified in <messageSchema>. Default name is "messageHeader".
Duplicate encoding name	An encoding name is non-unique, rendering a reference ambiguous.
nullValue specified for non-null encoding	Attribute nullValue is inconsistent with presence=required or constant
Attributes nullValue, minValue or maxValue of wrong data range	The specified values must be convertible to a scalar value consistent with the encoding. For example, if the primitive type is uint8, then the value must be in the range 0 through 255.
semanticType mismatch	If the attribute is specified on both a field and the encoding that it references, the values must be identical.
presence mismatch	If the attribute is specified on both a field and the encoding that it references, the values must be identical.
Missing constant value	If presence=constant is specified for a field or encoding, the element value must contain the constant value.
Missing validValue content	A <validValue> element is required to carry its value.
Incompatible offset and blockLength	A field offset greater than message or group blockLength is invalid

Error condition	Error description
Duplicate ID or name of field or group	Attributes id and name must uniquely identify a type within a message schema. This applies to fields and groups. To be clear, the same field or group ID may be used in multiple messages, but each instance must represent the same type. Each of those instances must match on both id and name attributes.

4.9.1 Message with a repeating group

```
<message name="ListOrder" id="2" description="Simplified
NewOrderList. Demonstrates repeating group">
  <field name="ListID" id="66" type="string14" semanticType="String"/>
  <field name="BidType" id="394" type="uint8" semanticType="int"/>
  <group name="ListOrdGrp" id="2030" >
    <field name="ClOrdID" id="11" type="string14" semanticType="String"/>
    <field name="ListSeqNo" id="67" type="uint32" semanticType="int"/>
    <field name="Symbol" id="55" type="string8" semanticType="String"/>
    <field name="Side" id="54" type="char" semanticType="char"/>
    <field name="OrderQty" id="38" type="intQty32" semanticType="Qty"/>
  </group>
</message>
```

4.9.2 Message with raw data fields

```
<message name="UserRequest" id="4" description="Demonstrates raw data usage">
  <field name="UserRequestId" id="923" type="string14"
semanticType="String"/>
  <field name="UserRequestType" id="924" type="uint8" semanticType="int"/>
  <field name="UserName" id="553" type="string14" semanticType="String"/>
  <field name="Password" id="554" type="string14" semanticType="String"/>
  <field name="NewPassword" id="925" type="string14"
semanticType="String"/>
  <field name="EncryptedPasswordMethod" id="1400" type="uint8"
description="This should be an enum but values undefined."
semanticType="int"/>
  <field name="EncryptedPasswordLen" id="1401" type="uint8"
semanticType="Length"/>
  <field name="EncryptedNewPasswordLen" id="1403" type="uint8"
semanticType="Length"/>
  <field name="RawDataLength" id="95" type="uint8" semanticType="Length"/>
  <data name="EncryptedPassword" id="1402" type="rawData"
semanticType="data"/>
  <data name="EncryptedNewPassword" id="1404" type="rawData"
semanticType="data"/>
  <data name="RawData" id="96" type="rawData" semanticType="data"/>
</message>
```

4.10 Reserved element names

4.10.1 Composite types

Encoding type name (default names)
messageHeader
groupSizeEncoding

4.10.2 Composite type elements

Type name	Composite type
blockLength	messageHeader and groupSize
day	MonthYear
exponent	decimal
length	variable-length data
mantissa	decimal
month	MonthYear
numInGroup	groupSize
templateId	messageHeader
time	timestamp, TZ time
timezoneHour	TZ time
timezoneMinute	TZ time
unit	timestamp, TZ time
varData	variable-length data
version	messageHeader
week	MonthYear
year	MonthYear

5 Schema Extension Mechanism

5.1 Objective

It is not always practical to update all message publishers and consumers simultaneously. Within certain constraints, messages and repeating groups can be extended in a controlled way. Consumers using an older version of a schema should be compatible if interpretation of added fields or messages is not required for business processing.

Message templates and repeating groups may be extended with new fields. However, the extension mechanism does not support extension of composite types that back existing fields.

This specification only details compatibility at the presentation layer. It does not relieve application developers of any responsibility for carefully planning a migration strategy and for handling exceptions at the application layer.

5.1.1 Constraints

Compatibility is only ensured under these conditions:

- Fields may be added to either the root of a message or to a repeating group, but in each case, they must be appended to end of a block.
- Existing fields cannot change data type or move within a message.
- Message and repeating group byte alignment may not change.
- A repeating group may be added after existing groups at the root level or nested within another repeating group.
- A variable-length data field may be added after existing variable-length data at the root level or within a repeating group.
- Message header encoding cannot change.
- In general, metadata changes such as name or description corrections do not break compatibility so long as wire format does not change.

Changes that break those constraints require consumers to update to the current schema used by publishers. An message template that has changed in an incompatible way must be assigned a new template "id" attribute.

5.2 Message schema features for extension

5.2.1 Schema version

The <messageSchema> root element contains a version number attribute. By default, version is zero, the initial version of a message schema. Each time a message schema is changed, the version number is incremented.

Version applies to the schema as a whole, not to individual elements. Version is sent in the message header so the consumer can determine which version of the message schema was used to encode the message.

See section 4.3.1 above for schema attributes.

5.2.2 Since version

When a new field, enumeration value, group or message is added to a message schema, the extension may be documented by adding a `sinceVersion` attribute to the element. The `sinceVersion` attribute tells in which schema version the element was added. This attribute remains the same for that element for the lifetime of the schema. This attribute is for documentation purposes only, it is not sent on the wire.

Over time, multiple extensions may be added to a message schema. New fields must be appended following earlier extensions. By documenting when each element was added, it possible to verify that extensions were appended in proper order.

5.2.3 Block length

The length of the root level of the message may optionally be documented on a `<message>` element in the schema using the `blockLength` attribute. See section 4.5.3 above for message attributes. If not set in the schema, block length of the message root is the sum of its field lengths. Whether it is set in the schema or not, the block length is sent on the wire to consumers.

Likewise, a repeating group has a `blockLength` attribute to tell how much space is reserved for group entries, and the value is sent on the wire. It is encoded in the schema as part of the `NumInGroup` field encoding. See section 3.4.8.2 above.

5.2.4 Deprecated elements

A message schema may document obsolete elements, such as messages, fields, and valid values of enumerations with `deprecated` attribute. Updated applications should not publish deprecated messages or values, but declarations may remain in the message schema during a staged migration to replacement message layouts.

5.3 Wire format features for extension

5.3.1 Block size

The length of the root level of the message is sent on the wire in the SBE message header. See section 3.2.2 above. Therefore, if new fields were appended in a later version of the schema, the consumer would still know how many octets to consume to find the next message element, such as repeating group or variable-length Data field. Without the current schema version, the consumer cannot interpret the new fields, but it does not break parsing of earlier fields.

Likewise, block size of a repeating group is conveyed in the `NumInGroup` encoding.

5.3.2 Number of repeating groups and variable data

Message headers and repeating group dimensions carry a count of the number of repeating groups and a count of variable-length data fields on the wire. This supports a walk by a decoder of all the elements of a message, even when the decoder was built with an older version of a schema. As for added fixed-length fields, new repeating groups cannot be interpreted by the decoder, but it still can process the ones it knows, and it can correctly reach the end of a message.

5.4 Compatibility strategy

This suggested strategy is non-normative.

A message decoder compares the schema version in a received message header to the version that the decoder was built with.

If the *received version is equal to the decoder's version*, then all fields known to the decoder may be parsed, and no further analysis is required.

If the *received version is greater than the decoder's version* (that is, the producer's encoder is newer than the consumer's decoder), then all fields known to the decoder may be parsed but it will be unable to parse added fields.

Also, an old decoder may encounter unexpected enumeration values. The application layer determines whether an unexpected value is a fatal error. Probably so for a required field since the business meaning is unknown, but it may choose to allow an unknown value of an optional field to pass through. For example, if OrdType value J="Market If Touched" is added to a schema, and the consumer does not recognize it, then the application returns an order rejection with reason "order type not supported", even if it does not know what "J" represents. Note that this is not strictly a versioning problem, however. This exception handling is indistinguishable from the case where "J" was never added to the enum but was simply sent in error.

If the *received version is less than the decoder's version* (that is, the producer's encoder is older than the consumer's decoder), then only the fields of the older version may be parsed. This information is available through metadata as "sinceVersion" attribute of a field. If sinceVersion is greater than received schema version, then the field is not available. How a decoder signals an application that a field is unavailable is an implementation detail. One strategy is for an application to provide a default value for unavailable fields.

5.5 Message schema extension example

Initial version of a message schema

```
<messageSchema package="FIXBinaryTest" byteOrder="littleEndian">
  <types>
    <type name="int8" primitiveType="int8"/>
  </types>

  <message name="FIX Binary Message1" id="1" blockLength="4">
    <field name="Field1" id="1" type="int8" semanticType="int"/>
  </message>
</messageSchema>
```

Second version - a new message is added

```
<messageSchema package="FIXBinaryTest" byteOrder="littleEndian"
  version="1">
```

```
<types>
  <type name="int8" primitiveType="int8"/>
  <type name="int16" primitiveType="int16"
    sinceVersion="1"/>
</types>

<message name="FIX Binary Message1" id="1" blockLength="4">
  <field name="Field1" id="1" type="int8" semanticType="int"/>
</message>

<!-- New message added in this version-->
<message name="FIX Binary Message2" id="2" blockLength="4"
  sinceVersion="1">
  <field name="Field2" id="2" type="int16" semanticType="int"/>
</message>
</messageSchema>
```

Third version - a field is added

```
<messageSchema package="FIXBinaryTest" byteOrder="littleEndian"
  version="2">

<types>
  <type name="int8" primitiveType="int8"/>
  <type name="int16" primitiveType="int16"
    sinceVersion="1"/>
  <type name="int32" primitiveType="int32"
    sinceVersion="2"/>
</types>

<message name="FIX Binary Message1" id="1" blockLength="8">
  <field name="Field1" id="1" type="int8" semanticType="int"/>
  <field name="Field11" id="11" type="int32" semanticType="int"
    sinceVersion="2"/>
</message>

<message name="FIX Binary Message2" id="2" blockLength="4"
  sinceVersion="1">
  <field name="Field2" id="2" type="int16" semanticType="int"/>
</message>
</messageSchema>
```

6 Usage Guidelines

6.1 Identifier encodings

FIX specifies request and entity identifiers as String type. Common practice is to specify an identifier field as fixed-length character of a certain size.

Optionally, a message schema may restrict such identifiers to numeric encodings.

Example of an identifier field with character encoding

```
<type name="idString" primitiveType="char" length="16" />
```

```
<field name="QuoteReqId" id="131" type="idString"
  semanticType="String"/>
```

Example of an identifier field with numeric encoding

```
<type name="uint64" primitiveType="uint64" />
```

```
<field name="QuoteReqId" id="131" type="uint64"
  semanticType="String"/>
```

7 Examples

The example messages are preceded by Simple Open Framing Header. Note that SOFH encoding is always big-endian, regardless of the byte order of the SBE message body. See that FIX standard for details.

Not all FIX enumeration values are listed in the samples.

7.1 SBE Message Schema

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<sbe:messageSchema xmlns:sbe="http://fixprotocol.io/2017/sbe"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" package="Examples"
id="91" version="0" byteOrder="littleEndian">
  <types>
    <type name="date" primitiveType="uint16"/>
    <type name="enumEncoding" primitiveType="char"/>
    <type name="idString" length="8" primitiveType="char"/>
    <type name="intEnumEncoding" primitiveType="uint8"/>
    <type name="currency" length="3" primitiveType="char"
description="ISO 4217"/>
    <composite name="DATA">
      <type name="length" primitiveType="uint16"/>
      <type name="varData" length="0" primitiveType="uint8"/>
    </composite>
    <composite name="MONTH_YEAR">
      <type name="year" primitiveType="uint16"/>
      <type name="month" primitiveType="uint8"/>
      <type name="day" primitiveType="uint8"/>
      <type name="week" primitiveType="uint8"/>
    </composite>
    <composite name="groupSizeEncoding">
      <type name="blockLength" primitiveType="uint16"/>
      <type name="numInGroup" primitiveType="uint16"/>
      <type name="numGroups" primitiveType="uint16"/>
      <type name="numVarDataFields" primitiveType="uint16"/>
    </composite>
    <composite name="messageHeader">
      <type name="blockLength" primitiveType="uint16"/>
      <type name="templateId" primitiveType="uint16"/>
      <type name="schemaId" primitiveType="uint16"/>
      <type name="version" primitiveType="uint16"/>
      <type name="numGroups" primitiveType="uint16"/>
      <type name="numVarDataFields" primitiveType="uint16"/>
    </composite>
    <composite name="decimalEncoding">
      <type name="mantissa" presence="optional" primitiveType="int64"/>
      <type name="exponent" presence="constant" primitiveType="int8">-
3</type>
```

```

    </composite>
    <composite name="qtyEncoding">
      <type name="mantissa" primitiveType="int32"/>
      <type name="exponent" presence="constant"
primitiveType="int8">0</type>
    </composite>
    <composite name="timestampEncoding" description="UTC timestamp with
nanosecond precision">
      <type name="time" primitiveType="uint64"/>
      <type name="unit" primitiveType="uint8" presence="constant"
valueRef="TimeUnit.nanosecond"/>
    </composite>
    <enum name="TimeUnit" encodingType="uint8">
      <validValue name="second">0</validValue>
      <validValue name="millisecond">3</validValue>
      <validValue name="microsecond">6</validValue>
      <validValue name="nanosecond">9</validValue>
    </enum>
    <enum name="businessRejectReasonEnum" encodingType="intEnumEncoding">
      <validValue name="Other">0</validValue>
      <validValue name="UnknownID">1</validValue>
      <validValue name="UnknownSecurity">2</validValue>
      <validValue name="ApplicationNotAvailable">4</validValue>
      <validValue name="NotAuthorized">6</validValue>
    </enum>
    <enum name="execTypeEnum" encodingType="enumEncoding">
      <validValue name="New">0</validValue>
      <validValue name="DoneForDay">3</validValue>
      <validValue name="Canceled">4</validValue>
      <validValue name="Replaced">5</validValue>
      <validValue name="PendingCancel">6</validValue>
      <validValue name="Rejected">8</validValue>
      <validValue name="PendingNew">A</validValue>
      <validValue name="Trade">F</validValue>
    </enum>
    <enum name="ordStatusEnum" encodingType="enumEncoding">
      <validValue name="New">0</validValue>
      <validValue name="PartialFilled">1</validValue>
      <validValue name="Filled">2</validValue>
      <validValue name="DoneForDay">3</validValue>
      <validValue name="Canceled">4</validValue>
      <validValue name="PendingCancel">6</validValue>
      <validValue name="Rejected">8</validValue>
      <validValue name="PendingNew">A</validValue>
      <validValue name="PendingReplace">E</validValue>
    </enum>
    <enum name="ordTypeEnum" encodingType="enumEncoding">
      <validValue name="Market">1</validValue>
      <validValue name="Limit">2</validValue>

```

```

        <validValue name="Stop">3</validValue>
        <validValue name="StopLimit">4</validValue>
    </enum>
    <enum name="sideEnum" encodingType="enumEncoding">
        <validValue name="Buy">1</validValue>
        <validValue name="Sell">2</validValue>
    </enum>
</types>
<sbe:message name="BusinessMessageReject" id="97" blockLength="9"
semanticType="j">
    <field name="BusinesRejectRefId" id="379" type="idString" offset="0"
semanticType="String"/>
    <field name="BusinessRejectReason" id="380"
type="businessRejectReasonEnum" offset="8" semanticType="int"/>
    <data name="Text" id="58" type="DATA" semanticType="data"/>
</sbe:message>
<sbe:message name="ExecutionReport" id="98" blockLength="42"
semanticType="8">
    <field name="OrderID" id="37" type="idString" offset="0"
semanticType="String"/>
    <field name="ExecID" id="17" type="idString" offset="8"
semanticType="String"/>
    <field name="ExecType" id="150" type="execTypeEnum" offset="16"
semanticType="char"/>
    <field name="OrdStatus" id="39" type="ordStatusEnum" offset="17"
semanticType="char"/>
    <field name="Symbol" id="55" type="idString" offset="18"
semanticType="String"/>
    <field name="MaturityMonthYear" id="200" type="MONTH_YEAR"
offset="26" semanticType="MonthYear"/>
    <field name="Side" id="54" type="sideEnum" offset="31"
semanticType="char"/>
    <field name="LeavesQty" id="151" type="qtyEncoding" offset="32"
semanticType="Qty"/>
    <field name="CumQty" id="14" type="qtyEncoding" offset="36"
semanticType="Qty"/>
    <field name="TradeDate" id="75" type="date" offset="40"
semanticType="LocalMktDate"/>
    <group name="FillsGrp" id="2112" blockLength="12"
dimensionType="groupSizeEncoding">
        <field name="FillPx" id="1364" type="decimalEncoding" offset="0"
semanticType="Price"/>
        <field name="FillQty" id="1365" type="qtyEncoding" offset="8"
semanticType="Qty"/>
    </group>
</sbe:message>
<sbe:message name="NewOrderSingle" id="99" blockLength="54"
semanticType="D">
    <field name="ClOrdId" id="11" type="idString" offset="0"

```

```

semanticType="String"/>
  <field name="Account" id="1" type="idString" offset="8"
semanticType="String"/>
  <field name="Symbol" id="55" type="idString" offset="16"
semanticType="String"/>
  <field name="Side" id="54" type="sideEnum" offset="24"
semanticType="char"/>
  <field name="TransactTime" id="60" type="timestampEncoding"
offset="25" semanticType="UTCTimestamp"/>
  <field name="OrderQty" id="38" type="qtyEncoding" offset="33"
semanticType="Qty"/>
  <field name="OrdType" id="40" type="ordTypeEnum" offset="37"
semanticType="char"/>
  <field name="Price" id="44" type="decimalEncoding" offset="38"
semanticType="Price" presence="optional"/>
  <field name="StopPx" id="99" type="decimalEncoding" offset="46"
semanticType="Price" presence="optional"/>
  </sbe:message>
</sbe:messageSchema>

```

7.2 Flat, fixed-length message

This is an example of a simple, flat order message without repeating groups or variable-length data.

Notes on the message schema

In this case, there is a lot of verbiage for one message, but in practice, a schema would define a set of messages. The same encodings within the <types> element would be used for a whole collection of messages. For example, a price encoding need only be defined once but can be used in any number of messages in a schema. Many of the attributes, such as description, offset, and semanticType, are optional but are shown here for a full illustration.

All character fields in the message are fixed-length. Values may be shorter than the specified field length, but not longer. Since all fields are fixed-length, they are always in a fixed position, supporting direct access to data.

An enumeration gives the valid values of a field. Both enumerations in the example use character encoding, but note that some enumerations in FIX are of integer type.

There are two decimal encodings. The one used for quantity sets the exponent to constant zero. In effect there is no fractional part and only the mantissa is sent on the wire, acting as an integer. However, FIX defines Qty as a float type since certain asset classes may use fractional shares.

The other decimal encoding is used for prices. The exponent is constant -3. In essence, each price is transmitted as an integer on the wire with assumed three decimal places. Each of the prices in the message is conditionally required. If OrdType=Limit, then Price field required. If OrdType=Stop then StopPx is required. Otherwise, if OrdType=Market, then neither price is required. Therefore, the price takes an optional encoding. To indicate that it is null, a special value is sent on the wire. See the table in section 2.4.2 above for the null value of the int64 mantissa.

In this example, all fields are packed without special byte alignment. Performance testing may prove better results with a different arrangement of the fields or adjustments to field offsets. However, those sorts of optimizations are platform dependent.

7.2.1 Wire format of an order message

Hexadecimal and ASCII representations (little-endian byte order):

```
00 00 00 48 eb 50 36 00 63 00 5b 00 00 00 00 00 : H P6 c [
00 00 4f 52 44 30 30 30 30 31 41 43 43 54 30 31 : ORD00001ACCT01
00 00 47 45 4d 34 00 00 00 00 31 00 a0 37 3a ce : GEM4 1 7:
10 40 15 07 00 00 00 32 1a 85 01 00 00 00 00 00 : @ 2
00 00 00 00 00 00 00 80 :
```

Interpretation

Wire format	Field ID	Name	Offset	Length	Interpreted value
00000048		SOFH message length	0	4	72
eb50		SOFH encoding	4	2	SBE little-endian
3600		SBE root block length	0	2	54
6300		SBE template ID	2	2	99
5b00		SBE schema ID	4	2	91
0000		SBE schema version	6	2	0
0000		No. of groups	8	2	0
0000		No. of var data	10	2	0
4f52443030303031	11	ClOrdId	0	8	ORD00001
4143435430310000	1	Account	8	8	ACCT01
47454d3400000000	55	Symbol	16	8	GEM4
31	54	Side	24	1	Buy
00a0373ace104015	60	TransactTime	25	8	2018-07-10T17:05:51.232Z
07000000	38	OrderQty	33	4	7
32	40	OrdType	37	1	Limit
1a85010000000000	44	Price	38	8	99.610
0000000000000080	99	StopPx	46	8	null

7.3 Message with a repeating group

This is an example of a message with a repeating group.

Notes on the message schema

The message contains a MonthYear field. It is encoded as a composite type with year, month, day and week subfields.

This message layout contains a repeating group containing a collection of partial fills for an execution report. The <group> XML tag enclosed the fields within a group entry. The dimensions of the repeating group are encoding as a composite type called groupSizeEncoding.

7.3.1 Wire format of an execution message

Hexadecimal and ASCII representations (little-endian byte order):

```
00 00 00 5c eb 50 2a 00 62 00 5b 00 00 00 01 00 : \ P* b [
00 00 4f 30 30 30 30 30 31 45 58 45 43 30 30 : 00000001EXEC00
30 30 46 31 47 45 4d 34 00 00 00 00 de 07 06 ff :00F1GEM4
ff 31 01 00 00 00 06 00 00 00 75 3e 0c 00 02 00 : 1          u>
00 00 00 00 1a 85 01 00 00 00 00 00 02 00 00 00 :
24 85 01 00 00 00 00 00 04 00 00 00 :$
```

7.3.2 Interpretation

Offset is from beginning of block.

Wire format	Field ID	Name	Offset	Length	Interpreted value
0000005c		SOFH message length	0	4	92
eb50		SOFH encoding	4	2	SBE little-endian
2a00		SBE root block length	0	2	42
6200		SBE template ID	2	2	98
5b00		SBE schema ID	4	2	91
0000		SBE schema version	6	2	0
0100		No. of groups	8	2	1
0000		No. of var data	10	2	0
4f30303030303031	37	OrderID	8	8	O0000001
4558454330303030	17	ExecID	8	8	EXEC0000
31	39	OrdStatus	1	1	PartialFilled
47454d3400000000	55	Symbol	18	8	GEM4
de0706ffff	200	MaturityMonthYear	26	5	year=2014 month=6
31	54	Side	1	1	Buy
01000000	151	LeavesQty	32	4	1
06000000	14	CumQty	36	4	6
753e	75	TradeDate	40	2	2013-10-11
0c00		Group block length	0	2	12
0000		NumInGroup	4	2	2
0000		No. of groups	4	2	0
0000		No. of var data	6	2	0

Wire format	Field ID	Name	Offset	Length	Interpreted value
1a85010000000000	1364	FillPx	0	8	99.610
02000000	1365	FillQty	8	4	2
2485010000000000	1364	FillPx	0	8	99.620
04000000	1365	FillQty	8	4	4

7.4 Message with a variable-length field

7.4.1 Wire format of a business reject message

Hexadecimal and ASCII representations (little-endian byte order):

```

00 00 00 44 eb 50 09 00 61 00 5b 00 00 00 00 00 :  D P a [
01 00 4f 52 44 30 30 30 31 06 27 00 4e 6f 74 :  ORD00001 ' Not
20 61 75 74 68 6f 72 69 7a 65 64 20 74 6f 20 74 : authorized to t
72 61 64 65 20 74 68 61 74 20 69 6e 73 74 72 75 :rade that instru
6d 65 6e 74                                     :ment

```

7.4.2 Interpretation

Wire format	Field ID	Name	Offset	Length	Interpreted value
00000044		SOFH message length	0	4	68
eb50		SOFH encoding	4	2	SBE little-endian
0900		SBE root block length	0	2	9
6100		SBE template ID	2	2	97
5b00		SBE schema ID	4	2	91
0000		SBE schema version	6	2	0
0000		No. of groups	8	2	0
0100		No. of var data	10	2	1
4f52443030303031	379	BusinesRejectRefId	0	8	ORD00001
06	380	BusinessRejectReason	8	1	NotAuthorized
4e6f74206175...	58	Text	0	39	Not authorized to trade that instrument

8 Release Notes

These are the changes made since SBE version 1.0.

8.1 SBE version 2.0 Release Candidate 1

These issues were resolved and accepted for version 2.0 Release Candidate 1. See [issues](#) and [pull requests](#) in GitHub for details and changes.

Issue	Description
26	Adding new variable length field in a repeating group
31	Presence attribute belongs to a field
35	Field attributes
36	Document cross-references
37	Extend message with repeating groups and vardata
39	sinceVersion on nested types within a composite
40	Clarification on constant char encodings
44	Remove deprecated attributes
45	Update examples for changed headers
48	Member names for data encoding type
51	Semantic type with Decimal encoding
59	refType is missing semanticAttributes
62	Decimal representation for optional mantissa with non-constant exponent
64	Support for easy constant value setting in fields
65	Alignment of var data and repeating groups
87	update examples for version 2.0

8.1.1 Compatibility

Version 2.0 is not back-compatible with SBE version 1.0, either in wire format or XML schema.