



OpenLCB Working Note	
Location Services	
Jul 2, 2024	Preliminary

1 Introduction

A Working Note is an intermediate step in the documentation process. It gathers together the content from various informal development documents, discussions, etc into a single place. One or more Working Notes form the basic for the next step, which is one or more Standard/TechNote pairs.

This working note is the start of effort toward a general location services protocol which can be used by e.g. RFID and QR scanners and other detectors to publish information about the detected location of rolling stock. It uses event transfer with payload to globally distribute information about location-specific information from the layout.

A “location” is a point on the layout, a block of track, or an area of a layout for which a location report may be generated by a scanner or another type of node. There may be multiple scanners or nodes that generate or modify location reports for the same or overlapping locations.

See also the DCC Detection Protocol Working Note for a different approach to some of these use cases.

1.1 Served Use Cases

- 1) A train enters a current-detected block. A detector recognizes this, successfully queries the decoder for its address, and reports that information on the OpenLCB bus. That information results in a speed command such as Slow or Stop being sent to that particular train.

There is a simpler use case for detectors without the ability to read the address:

A signal needs to determine occupancy in the block it protects. It is not concerned with identity, but it does need to know location and entry/exit.

This can be handled efficiently with standard PCER messages carrying just the EventID.

- 2) A train enters a current-detected block. A detector recognizes this, and attempts unsuccessfully to query the decoder for its address. The detector emits a report without decoder address information.
- 3) A train exits a current-detected block. A detector recognizes this, and reports that interaction on the OpenLCB bus.
- 4) A train crosses an RFID reader which determines the train’s DCC address and reports that address on the OpenLCB bus.

- 5) A train-announcement unit needs to determine the identity and location of approaching trains so it can announce their arrival.
- 6) A train display needs to gather all train locations for display to the users.
- 35 7) A train display needs to gather the location and identity of all trains within a specified area.
- 8) An application needs to gather the location and identity of all trains.
- 9) A node powers up and emits a query which generates messages carrying all available location information.
- 40 10) An RFID reader reads a tag and uses this protocol to distribute the raw tag contents. A separate database node receives this content, does an internal lookup to locate the rolling stock's Unique ID, and then uses this protocol to distribute a human-readable version of the original distribution that also carries the rolling stock's Unique ID.
- 11) A detector observes that a piece of rolling stock has been added to the layout. It reports this. A throttle or command station provides that rolling stock in its roster display.
- 45 12) A detector observes that a piece of rolling stock has been removed from the layout. A throttle or command station removes that rolling stock from its roster display.
- 13) The user has a throttle, which has a train on it. There is a location report coming in for that train. The user would like to display on the throttle the name of the location where that train was seen. (This has to be a user-visible name, so coming out of some configuration.) This raises two indirection issues that need to be solvable:
 - 50 A) Indirection of the OpenLCB train and the detector address. Say this is a DCC train from a command station, then there is a known DCC address, which is recognized by RailCom detectors, and reported. The harder case of OpenLCB trains needs to be handled too.
 - B) Indirection from the detector address to a user-visible (configured) block name.
- 55 14) Cab signaling: display on the throttle what the next signal aspect is for the selected train. This raises indirection issues:
 - A) Going from train to block as above
 - B) Going from physical orientation plus reverser setting on the logical locomotive to an absolute direction inside the block
 - 60 C) Going from block + absolute direction to a next signal (which is a multi-state state machine)
 - D) Aspect interpretation (state to human readable text)

1.2 Unserved Use Cases

(To be determined as the proposal is developed)

1.3 Requirements

- 65 1) The protocol should cover multiple hardware types, including RFID and QR scanners, RailCom and other DCC detectors, and follow-on hardware that can detect and identify rolling stock at a point, on a track section, or in an area.
- 2) To the extent possible, the protocol should identify standard forms for common information.
- 70 A) The Unique ID of the identified rolling stock should be in a common location in the protocol's message(s) so that it can be extracted regardless of the form of the sensor that generated it: a RFID detector and a QR detector on the same layout should both provide the same information on the identity of the detected rolling stock.
- B) The Unique ID of the scanner should be in a common location in the protocol's message(s) so that it can be extracted regardless of the form of the sensor that generated it.
- 75 3) To the extent possible, the protocol should identify methods for its message(s) to carry detector-specific information. For example, a RFID detector and a QR detector on the same layout should both be able to publish their tag-specific information.
- 4) The identification of the scanning hardware should be guaranteed unique. This lets hardware be added to a layout without having to resolve conflicts.
- 80 5) The messages should not rely on the transmitting-node-address field of the message for decoding. I.e. the protocol doesn't say "You identify who made the scan by looking at who sent the message". This allows replacing the physical hardware that is broken with another piece of physical hardware carrying the same configuration; we expect that the entire network continues to operate as it did before. It is also useful to allow messages to be created by e.g. a debugging node during development.
- 85 6) It should be possible to look at the message(s) carrying a specific scan and know that they contain a scan without additional information. I.e. you don't have to configure a layout monitor with all possible scanner IDs before it can start processing or displaying scanning results.
- 90 7) The identification of which scanner is reporting should be early in the message(s). This allows nodes to drop the receipt of the message(s) early, without having to buffer them completely.
- 8) The identification of the detected rolling stock should be early in the message(s). This allows nodes to drop the receipt of the message(s) early, without having to buffer them completely.
- 9) There should be a standard algorithm to find the UniqueID of the scanned piece of rolling stock within the message(s), when the scanner can provide one. In the common case of DCC, this is the UniqueID mapped from the DCC address. For some scanners, e.g. some types of RFID scanner, the scanner might not be able to map the scanned value to a specific piece of rolling stock, in which case this information field would be blank.
- 95 10) Note that the item directly above is referring to a full Unique ID (node ID), not just a DCC address. Someday, there will be LCC-native trains and this protocol should be ready for them. The alternative is inevitably a messy migration at some point in the future.
- 100

11) Ideally, information about short circuit and current overload status could be conveyed by this protocol. These are often generated by the same device as the transponder messages e.g. LocoNet devices. These are providing location-specific information. It would be good to have some commonality between all location based messages. Another example would be current or voltage measurements at a point.

105

1.4 Not Requirements; Unmet Requirements

1. The messages for this protocol should be a super-set of those from the DCC Detection Protocol.

2 Specified Sections

110

This is the usual section organization for a Technical Note, to accumulate the Standard and Technical Note content in its eventual order.

In this section, deeply indented content is intended for the eventual Technical Note.

2.1 Introduction

Note that this section of the Standard is informative, not normative.

115

A “scanner” is a device that can gather information from the railroad, typically when something happens. RFID readers, QR code readers, point optical detectors and track current detectors are all examples of scanners. There may be many scanners associated with an individual node.

A scanner forwards one instance of the gathered information as a “report”.

RailCom is a trademark of Lenz GMBH. Transponding is a trademark of Digitrax, Inc. ~~StarRF is a trademark of 292 Design UK, Ltd.~~

120

2.2 Intended Use

Note that this section of the Standard is informative, not normative.

This protocol is intended to provide a standardized way for scanners to distribute report information across the OpenLCB network.

2.3 Reference and Context

125

See also the DCC Detection working note for a simpler form specific to DCC.

This working note assumes adoption of the Event With Payload working note as a Standard.

2.4 Message Formats

2.4.1 Report Message

130

The proposal is to carry reports within a Producer Consumer Event Report (PCER) as an event with payload defined by:

Note bytes 0-7 are the Event ID of the message, 8 through the end of the message are the payload of the message.

- Bytes 0-1: The defined prefix 0x0102 is assigned to this.

This unambiguously identifies this PCER message as part of this protocol.

135 That enables nodes not interested in this protocol to discard this event-with-payload transmission immediately.

Using a unique prefix followed by a UniqueID (see immediately below) is a mechanism for creating a fully-unique Event ID. No other Event ID allocation mechanism can duplicate one of these IDs.

140 As an alternative, using multiple prefixes to separate out various forms of content was considered. This makes the interactions below more complicated, without any particular benefit, so a single prefix is used.

- Bytes 2-7: A Unique ID that specifies the device that read the information e.g. a scanner.

145 Having this in the 1st eight bytes, i.e. the Event ID, allows using the various identify message exchanges to locate scanners.

Putting this near the start of the message (1st CAN frame) allows early rejection of uninteresting scanner sources.

150 Note that each scanner has a Unique ID that may be separate from the Unique ID of the node supporting it. In the case of multiple scanners associated with a node, each will have its own Unique ID, and only one at most may be the same as that node.

Note that this Unique ID identifies this device to other nodes on the OpenLCB bus. As such, it should be a configurable value so that e.g. a faulty hardware node can be replaced, with the replacement configured to continue to send the same Unique ID.

- Bytes 8-9: Flag bits, defined below

- 155 • Bytes 10-15: A Unique ID that specifies the item being scanned, e.g. a train node ID, DCC address mapped to a Unique ID, or zeros if the information is not available.

This will fit in the 2nd CAN frame, allowing early rejection of uninteresting items when available.

160 Placing this in the content and having the scanner ID in the first eight bytes was preferred because it simplified the defined interactions.

The process for mapping a DCC address to a Unique ID is defined elsewhere.

- Bytes 16-NN: Any read (scanned) content that's available.

165 Note that there is no restriction on the length of the message in the Standard.—When designing something using this protocol, it's important to consider the impact of the data length on network bandwidth and latency, particularly in light of the section on “Retrieving most recent scan reports” below.

2.4.1.1 Flag bits

bits 0-1: entry/exit

- 170 | ~~0b00: unknown – This message may or may not be due to an item appearing or disappearing.~~
 | ~~_____ E.g. this might refer to an item that's remaining within the scanned region.~~
 | ~~0b00: exit – The report refers to an item disappearing from the detected area.~~
 | 0b01: entry without state – The report refers to an item appearing in the detected area; no report
 | of the item eventually disappearing from the detected area will follow.
 | 0b10: entry with state – The report refers to an item appearing in the detected area; a report of
 175 | the item eventually disappearing from the detected area will follow.
 | ~~0b11: exit – The report refers to an item disappearing from the detected area.~~
 | ~~0b11: unknown – This message may or may not be due to an item appearing or disappearing.~~
 | ~~_____ E.g. this might refer to an item that's remaining within the scanned region.~~

180 | bits 2-3: direction of relative motion – Defined with respect to the designated front of the piece of
 rolling stock.

- | ~~_____ 0b00: stopped~~
 | ~~_____ 0b00: unknown~~
 | 0b01: forward
 | 0b10: reverse
 185 | ~~_____ 0b11: stopped~~
 | ~~_____ 0b11: unknown~~

bits 4-5: direction of absolute motion – The geographical meanings of these are defined during set up
 of the scanner. Note that this is the direction of motion, not direction the front of the item is pointing.

- | ~~_____ 0b00: stopped~~
 190 | ~~_____ 0b00: unknown~~
 | 0b01: east/north
 | 0b10: west/south
 | ~~_____ 0b11: unknown~~
 | ~~_____ 0b11: stopped~~

195 | bits 6-7: facing direction – The geographical direction that the front of the item is facing. The
 | geographical meanings of these are defined during the set up of the scanners.

| ~~0b00: reserved – Do not send, ignore on receipt.~~

~~0b01: east/north~~

~~0b10: west/south~~

200 | ~~0b00H: unknown – The scanner did not determine a value for this item.~~

~~There could be scanners that can detect this but can't detect forward/reverse motion and east/west motion.~~

~~0b01: east/north~~

~~0b10: west/south~~

205 | ~~0b11: reserved – Do not send, ignore on receipt.~~

~~bits 8-9: flags content that spans multiple EWP messages~~

~~0b00 – only EWP message of the content~~

~~0b01 – first EWP message of the content, will be followed by more EWP messages~~

~~0b11 – middle EWP message of the content, will be followed by more EWP messages~~

210 | ~~0b10 – last EWP message of the content, must be preceded by an EWP marked “first EWP”.~~

~~It is strongly recommended that nodes sending Location Services messages define their content to be as short as possible, such that multiple-EWP messages are not necessary. For example, nodes should not repeat previous Location Services message content when they send an additional Location Services message with additional content. A node that's adding human-readable information to a prior Location Services report that already contains several optional fields, e.g. position information, should not repeat those optional fields.~~

215 | ~~Nodes that are expecting several types of optional Location Services information for a single reporting event should expect that information to arrive in multiple Location Services messages.~~

bits ~~108~~-11: reserved, send as zero and ignore on receipt

bits 12-15: content format – see below

The content format field is used to indicate the format of bytes 16-NN

~~0b000x0~~ reserved, do not send; ignore message upon receipt ~~of 0b00~~

225 | ~~0b010x1~~ (initial) standard content form, see below

2.4.1.2 Standard Content Form

The standard content form consists of one or more blocks formatted as

230 | byte 0: length of content N, including type byte and data but not including the length byte. The value of 0xFF is reserved for expansion to longer blocks and should not be used until that is defined.

235

Numerology: Although the length of the data field is not limited by Standards Although the length of EWP messages is limited to a total of 256 bytes, it is possible for the Location Services content to be longer than that through the use of the first/middle/last flags. Messages with over 256 bytes of content in a single ~~item~~ block are expected to be rare and are strongly recommended against at this time. Should that be required at some point, that data can be placed in multiple blocks defined with their own first, middle, end typeflags, or a system for identifying longer blocks can be defined. There is not enough information available to define this method at this time, although not doing so will result in some backward-compatibility issues.

240

byte 1: type of content – see below

byte 2-N: content defined by type byte

There can be more than one block, in which case the blocks are concatenated with no intervening bytes. Blocks of length 0 (i.e. a single zero byte) can optionally appear at the end of the message.

2.4.1.3 Standard Content Block Types

245

The type-byte values are allocated by this document. It's expected that new types will be routinely added. The current set of types are:

0: Reserved, do not send, ignore the associated block on receipt

This is meant as a convenient flag for software while creating and parsing messages.

250

1: Human-readable format. Not intended for machine parsing, this is meant to carry messages for human display.

An example would be “Train 123 arriving at station ABC”.

2: Raw-eContent of a read RFID tag

255

In the future, it may be desirable to define types for specific types of tag, e.g. ISO14443, ISO15693, etc. That's left for a future version once more usage experience is available.

3: Raw-eContent of a QR or bar code read

4: Raw-eContent of a RailCom read

260

Is there a document we can refer to about the format and content of this?

The scanner is expected to parse the contents of the read data and fill in as many of the flag fields as possible.

5: Raw-eContent of a Digitrax Transponding read

Is there a document we can refer to about the format and content of this?

Two bytes of Transponding zone, followed by two bytes of locomotive address. These hold the values reported by the Transponding reader, without further processing. Any additional

265 | information in the Transponding report will appear after the defined four bytes. Its presence will be indicated by the specified length of the block.

The scanner is expected to parse the contents of the read data and fill in as many of the flag fields as possible.

270 | 6: Position information: Three 16-bit IEEE floating point numbers representing the X, Y and Z position of the detected object on the layout. The units are meters.

The origin and direction of axes is specified when the scanner is configured by the user, but conventionally X is -left/+right, Y is -front/+back and Z is -down/+up from some location on the layout. This corresponds to putting the origin at the left-front of the layout on the lowest level.

275 | Note that 8-bit IEEE does not have a suitable resolution at a suitable maximum range. 16-bit IEEE provides roughly 16mm steps at +/-25m

7. DCC address: Two byte short or long DCC address, defined with the leading 0xC000 for long addresses.

8. Set speed: Float16 value of the current set speed; NaN if not available

280 | 9. Commanded speed: Float16 value of the current commanded speed; NaN if not available

10. Actual speed: Float16 value of the current actual speed; NaN if not available

117: Analog information: For reporting analog information (voltage, current, temperature, water level, ...) as an IEEE-16 value in SI units, followed by a byte for units, followed by an optional human-readable string. The defined units are

285 | 0: Unknown, no unit specified.

1: Volts

2: Amperes

3: Watts

4: Ohms

290 | 5: Degrees C

6: Seconds

7: Meters

8: Meters²

9: Meters³

295 | 10: Meters/second

11: Meters/second/second

12: Kilograms

13: Newtons

All others: Reserved, do not send, treat as unknown on receipt

300 There needs to be something in the block so that there can be multiple independent analog information readings in the block. Note that there can be multiple readings of a single type in the same message e.g. “Ambient temperature” and “Motor temperature”.

All others reserved, do not send; ignore the block on receipt.

305 Definitions for JSON and XML formats were considered. They were not included because that would define yet another level of syntax that would require standardization, and it was felt that the standard format could be defined to carry the same data as needed.

2.4.2 Consumed and Produced Event ID

310 Every node that produceessends information via this protocol must also consume a specific well-known Event ID consisting of the well-known 0x0102 prefix assigned followed by six 0x00 bytes. When asked to Identify Consumers for this event, the reply should include the Valid modifier because the consuming nodes are always prepared to respond to this event.

The six 0x00 bytes is an invalid Unique ID, so no scanner will allocate this as its own.

315 Nodes that will inquire for the current values of reports produce a specific well-known Event ID consisting of the 0x0102 prefix followed by six 0x00 bytes. When asked to Identify Producers for this event, the reply should include the Unknown modifier because there’s no state associated with this inside the producer.

2.5 States

This protocol has no states.

320 Remembering that a report for entry to a detected area has been sent, so that an exit message can later be sent, is not considered a public state of the protocol.

2.6 Interactions

325 Note that the Unique ID that specified the scanner doing the reading (bytes 2-7) does not need to be the Node ID of the scanning node. It just needs to be uniquely assigned by the node creator from some assigned space to indicate the device doing the read.

330 For example, consider a board that supports four RFID readers that operate independently. Each of those readers will have a distinct Unique ID that is used when its detections are reported. In addition, the board will have a single Unique ID that it uses as its node ID. This may, but does not have to be, the same as the Unique ID of one of the readers.

When requested to Identify Producers for one of the Report Messages defined above, the Producer Identified response should carry as modifiers:

- Valid – the last Report Message’s contents are still valid, and would be produced now

- Invalid – the last Report Message’s contents are no long valid, and would not be produced now
- Unknown – if neither of the above apply

These values are chosen to be consistent with the general definitions in the Event Transport protocol definition.

2.6.1 Reporting new information

When a scanner gathers new information, it may emit a PCER message with the above format. PCER messages are global, so all interested nodes will receive this message. The interested nodes can then process them as desired.

Note this says “may emit”. That choice of whether or not to report every time new information is received is up to the device manufacturer.

Note that a single node implementing multiple readers will send PCER event-with-payload messages from those separate readers but must not interleave the messages making up the payload of those separate PCER messages. The event with payload protocol requires that the sending node only sends one PCER event-with-payload sequence at a time.

Example: An RFID reader reads the fixed information from a tag. This does not directly contain a Unique ID for the scanned item. The reader produces a report with a zero value for bytes 10-15, the identity of the scanned item, and a block of RFID information. A second node with a RFID → Node ID mapping database consumes that message, looks up the appropriate Unique ID of the scanned item, and produces a report message that contains it and the original RFID information block.

Example: A scanner produces a message that carries only the minimal information. A second node with other sources of data consumes that message. Using the Unique IDs for the scanner and the scanned item, this second node creates suitable human-readable content such as “Train 123 arriving on track 4 five minutes late”. It then adds that content as a new block and produces a new message with otherwise the same content. Meanwhile, a third node driving a station display consumes the first message, sees it’s from its associated scanner but there’s no human readable content and discards it. It then consumes the second message, sees it’s from its appropriate scanner and contains a human-readable block, extracts the human readable content, and displays it.

Example: A control node is tasked with controlling the speed of each train that approaches a signal. A scanner is configured to report the arrival and departure of each train on the track approaching the signal, including Unique ID and directional information. The control node recognizes reports from that scanner, extracts the train’s unique ID, and determines whether the train is approaching the signal or receding from it. The control node then has enough information to control the train via the traction protocol.

Example: A node facilitating cab signaling will consume all messages from this protocol. It consumes all reports reporting arrival of all trains on the track approaching each signal. For each report, it determines if the train is approaching the signal,

375 (somehow) determines the current aspect of the relevant signal, and then transmits that aspect to the relevant train throttle(s) using some protocol not defined here.

This could also be done on the throttle by checking for its own scanned ID in reports, but that requires substantial context information that seems better centralized.

380 Or there could be a mechanism for directly inquiring about the current aspect of the signal associated with a particular reader's Unique ID, which would allow a distributed mechanism.

385 Example: A video system is directly tracking trains. For each train, it periodically produces a report that contains the X, Y, Z position of the train. Each of these comes from the same scanner ID, the Unique ID associated with the video system. For this to work with a signal system from another vendor, a node is provided that can compute occupancy in virtual blocks: When the train is inside a specific volume, that's considered to be occupancy in a specific virtual block associated with a specific virtual scanner's Unique ID. Upon finding a train in one of these, this node will produce a report that includes the Unique ID of the specific virtual scanner. This in turn will be consumed in the usual way by the signal system node(s) that are looking for occupancy on that associated section of track.

2.6.2 Identification of scanners

Sending the defined 0x01 02 00 00 00 00 00 00 Event ID in an Identify Consumer message will result in Consumer Identified messages that indicate which nodes can take part in this protocol.

395 This removes the need to include this protocol in the Protocol Identification Protocol bits, *although that can also be considered for ease of access.*

Note that this refers to nodes. The individual scanners are not nodes. They're not expected to each implement the full OpenLCB protocol stack. Instead, this protocol has a single node providing that service for multiple scanners.

400 Once the nodes that use this protocol have been identified, an Identify Events directed to each of them will return Producer Identified messages indicating the Unique IDs of all the scanners.

2.6.3 Identification of scannable objects

405 There is currently no way to efficiently request the unique IDs of all scannable objects. Some of them might not be currently visible to the system. You can retrieve all presently-visible objects by retrieving the most recent scan reports, see section below.

2.6.4 Retrieving most recent scan reports

410 When a new node joins the network, it may want to catch up on previous scan reports. In that case, the new node can produce (in a PCER message) the 0x01 02 00 00 00 00 00 00 Event ID. This is consumed by all nodes that implement this protocol. When consumed, each of those nodes will re-emit the most recent report from each of its scanner implementations.

How to handle multiple overlapping reports from the same scanner? Just send them all? In what order? Option for the implementor?

This might be a large load on the network. ~~PCER messages are high priority, and t~~This request will result ~~on~~in one or more PCER messages from (most? all?) scanners.

415 There is no standard interaction defined to get a report from just one scanned device, i.e. there is no standard interaction defined that does “Any scanner that currently sees device A.B.C.D.E.F, please resend your last report”.

2.6.5 Identifying users of scanned information

420 A node can identify nodes that use inquire for scanned information by emitting an Identify Producer message for the 0x01 02 00 00 00 00 00 00 Event Id.

The Standard doesn't actually require that a using node be able to produce that event if it's not going to inquire about the current reports. So this is not “identify all nodes that use scanned information”.

2.6.6 Traffic limiting

425 Producing nodes may not produce the same event with payload content more than twice per second.

The goal of this section is to reduce the bandwidth consumed by repeating detection messages. Repeating the same message is a bad practice, and should be limited to the extent possible.

430 Another approach would have been to prohibit sending the same content twice in a row. However, there may be situations where that's a valid thing to do.

3 Background Information

RFID readers and tags come in multiple forms. With 13.56MHz RFID, be it ISO14443 (as used in the RC522 readers), or ISO15693 ~~as used in StaRFHshrail~~, the following points apply:

- 435 • Up to 8 bytes can be read "on the fly", i.e. as the tag on a train passes over a reader at speed. Depending on how the reader is set up, these 8 bytes can either be the UID of the tag, programmed in at time of manufacture, or 8 bytes of EPROM, which can be programmed into the tag by the user. The usage of these 8 bytes would be determined by the user, but could, say, be 2 bytes for Stock ID, 2 bytes for DCC address, and 4 bytes for other purposes.
- 440 • In addition, there at least another 96 bytes of user programmable EPROM, but these need to be read when the tag is essentially static over a reader, due to the time taken to read all the bytes.

The tags used by the MERG CAN RFID readers provide five bytes of programmable payload. Multiple conventions have been proposed for the contents of those bytes.

445 QR codes come in multiple forms. Depending on version, size and error correction level, they can contain thousands of characters, many more than can reasonably be carried across OpenLCB CAN links.

Some scanners can report the direction of the detected train. For some, this is whether the train is moving in forward or reverse. For others, it's whether the train is moving in one direction on the track vs the other (North vs South, East vs West), e.g. ProTrak Grapevine systems can report moving North/East, not moving, and moving South/West.

450 RPS detectors reported absolute X, Y and Z coordinates of the detected rolling stock. The report was emitted at specified intervals. Although these are no longer being produced, something similar may emerge in the future with developments in e.g. video monitoring of a layout.

Traditional block detectors and point detectors are an interesting special case. They're simple, low-cost and common on layouts. They can't sense any address information from the detected rolling stock.

455 Although their needs are already met by the Event Transfer protocol itself, they can also use this protocol to convey their information too.

DRAFT

Table of Contents

1	Introduction.....	1
1.1	Served Use Cases.....	1
1.2	Unserved Use Cases.....	2
1.3	Requirements.....	2
1.4	Not Requirements; Unmet Requirements.....	3
2	Specified Sections.....	4
2.1	Introduction.....	4
2.2	Intended Use.....	4
2.3	Reference and Context.....	4
2.4	Message Formats.....	4
2.4.1	Report Message.....	4
2.4.1.1	Flag bits.....	5
2.4.1.2	Standard Content Form.....	6
2.4.1.3	Standard Content Block Types.....	7
2.4.2	Consumed and Produced Event ID.....	8
2.5	States.....	9
2.6	Interactions.....	9
2.6.1	Reporting new information.....	9
2.6.2	Identification of scanners.....	10
2.6.3	Identification of scannable objects.....	11
2.6.4	Retrieving most recent scan reports.....	11
2.6.5	Identifying users of scanned information.....	11
2.6.6	Traffic limiting.....	11
3	Background Information.....	11