

Traffic Feed Format (TraFF)

Version 0.8

Licensed under the Creative Commons Attribution-ShareAlike 4.0 International License.

Creative Commons Attribution-ShareAlike 4.0 International License

This is a human-readable summary of the full license below.

Under this license, you are free to:

- **Share** — copy and redistribute the material in any medium or format
- **Adapt** — remix, transform, and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

License terms:

- **Attribution** — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- **ShareAlike** — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.
- **No additional restrictions** — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

- You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.
- No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Creative Commons Corporation (“Creative Commons”) is not a law firm and does not provide legal services or legal advice. Distribution of Creative Commons public licenses does not create a lawyer-client or other relationship. Creative Commons makes its licenses and related information available on an “as-is” basis. Creative Commons gives no warranties regarding its licenses, any material licensed under their terms and conditions, or any related information. Creative Commons disclaims all liability for damages resulting from their use to the fullest extent possible.

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution-ShareAlike 4.0 International Public License (“Public License”). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution-ShareAlike 4.0 International Public License (“Public License”). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

1. Definitions

- Adapted Material** means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered, arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.
- Adapter's License** means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.
- BY-SA Compatible License** means a license listed at creativecommons.org/compatiblelicenses, approved by Creative Commons as essentially the equivalent of this Public License.
- Copyright and Similar Rights** means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section 2(b)(1)-(2) are not Copyright and Similar Rights.
- Effective Technological Measures** means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.
- Exceptions and Limitations** means fair use, fair dealing, and/or any other exception or limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.
- License Elements** means the license attributes listed in the name of a Creative Commons Public License. The License Elements of this Public License are Attribution and ShareAlike.
- Licensed Material** means the artistic or literary work, database, or other material to which the Licensor applied this Public License.
- Licensed Rights** means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.
- Licensor** means the individual(s) or entity(ies) granting rights under this Public License.

- k. **Share** means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the public including in ways that members of the public may access the material from a place and at a time individually chosen by them.
- l. **Sui Generis Database Rights** means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.
- m. **You** means the individual or entity exercising the Licensed Rights under this Public License. **Your** has a corresponding meaning.

2. Scope

a. License grant

1. Subject to the terms and conditions of this Public License, the Licensor hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:
 - A. reproduce and Share the Licensed Material, in whole or in part; and
 - B. produce, reproduce, and Share Adapted Material.
2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.
3. Term. The term of this Public License is specified in Section 6(a).
4. Media and formats; technical modifications allowed. The Licensor authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensor waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures. For purposes of this Public License, simply making modifications authorized by this Section 2(a)(4) never produces Adapted Material.
5. Downstream recipients.
 - A. Offer from the Licensor – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensor to exercise the Licensed Rights under the terms and conditions of this Public License.
 - B. Additional offer from the Licensor – Adapted Material. Every recipient of Adapted Material from You automatically receives an offer from the Licensor to exercise the Licensed Rights in the Adapted Material under the conditions of the Adapter’s License You apply.
 - C. No downstream restrictions. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.
6. No endorsement. Nothing in this Public License constitutes or may be construed as permission to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section 3(a)(1)(A)(i).

b. Other rights

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.
2. Patent and trademark rights are not licensed under this Public License.
3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

3. License Conditions

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution

1. If You Share the Licensed Material (including in modified form), You must:
 - A. retain the following if it is supplied by the Licensor with the Licensed Material:
 - i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);
 - ii. a copyright notice;
 - iii. a notice that refers to this Public License;
 - iv. a notice that refers to the disclaimer of warranties;
 - v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;
 - B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and
 - C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.
2. You may satisfy the conditions in Section 3(a)(1) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.

3. If requested by the Licensor, You must remove any of the information required by Section 3(a)(1)(A) to the extent reasonably practicable.
- b. **ShareAlike** In addition to the conditions in Section 3(a), if You Share Adapted Material You produce, the following conditions also apply.
- c. The Adapter's License You apply must be a Creative Commons license with the same License Elements, this version or later, or a BY-SA Compatible License.
- d. You must include the text of, or the URI or hyperlink to, the Adapter's License You apply. You may satisfy this condition in any reasonable manner based on the medium, means, and context in which You Share Adapted Material.
- e. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, Adapted Material that restrict exercise of the rights granted under the Adapter's License You apply.

4. Sui Generis Database Rights

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

- a. for the avoidance of doubt, Section 2(a)(1) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;
- b. if You include all or a substantial portion of the database contents in a database in which You have Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material, including for purposes of Section 3(b); and
- c. You must comply with the conditions in Section 3(a) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section 4 supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

5. Disclaimer of Warranties and Limitation of Liability

- a. Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors, whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.
- b. To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.
- c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

6. Term and Termination

- a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.
- b. Where Your right to use the Licensed Material has terminated under Section 6(a), it reinstates:
 - The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:
 1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or
 2. upon express reinstatement by the Licensor.For the avoidance of doubt, this Section 6(b) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.
- c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.
- d. Sections 1, 5, 6, 7, and 8 survive termination of this Public License.

7. Other Terms and Conditions

- a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.
- b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

8. Interpretation

- a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully be made without permission under this Public License.
- b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.
- c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.

- d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

Disclaimer

Creative Commons is not a party to its public licenses. Notwithstanding, Creative Commons may elect to apply one of its public licenses to material it publishes and in those instances will be considered the “Licensor.” The text of the Creative Commons public licenses is dedicated to the public domain under the CC0 Public Domain Dedication. Except for the limited purpose of indicating that material is shared under a Creative Commons public license or as otherwise permitted by the Creative Commons policies published at creativecommons.org/policies, Creative Commons does not authorize the use of the trademark “Creative Commons” or any other trademark or logo of Creative Commons without its prior written consent including, without limitation, in connection with any unauthorized modifications to any of its public licenses or any other arrangements, understandings, or agreements concerning use of licensed material. For the avoidance of doubt, this paragraph does not form part of the public licenses.

Creative Commons may be contacted at www.creativecommons.org

Version History

Version	Date	Author	Comment
0.1	2018-01-18	Michael von Glasow	Initial specification
0.2	2018-06-30	Michael von Glasow	Added platform specifics Revisited TMC conversion Added area-specific queries and subscriptions as features for future versions
0.3	2018-08-11	Michael von Glasow	Added official DNS domain name
0.4	2018-10-06	Michael von Glasow	Moved reserved events RESTRICTION_CONTRAFLOW, RESTRICTION_LANE_BLOCKED, RESTRICTION_LANE_CLOSED, RESTRICTION_REDUCED_LANES and reserved quantifier <code>q:ints</code> to official specification
0.5	2018-10-22	Michael von Glasow	Changed unit of <code>length</code> attribute (<code>event</code> element) to meters
0.6	2018-10-27	Michael von Glasow	Added CC-BY-SA 4.0 license
0.7	2018-12-19	Michael von Glasow	Added requirement for resilience against messages being missed by consumers Added Message Expiration Clarified semantics of <code>event.speed</code> Corrected example for <code>q_time</code> use Changed quantifier attributes from <code>q:*</code> to <code>q_*</code>
0.8	2021-01-11	Michael von Glasow	Split off event and supplementary information lists, both official and reserved, into separate documents Split off subscription management (previously mostly in the Platform Specifics chapter) into a separate document <code>destination</code> is now allowed for bidirectional locations and generally recommended; added <code>origin</code> to denote the opposite end Added <code>country</code> , <code>territory</code> and <code>town</code> attributes for locations Added value <code>MEDIUM_RES</code> for fuzziness, clarified meaning of <code>LOW_RES</code> Added <code>distance</code> attribute to location points; clarified that junction references are not to be used to indicate a nearby junction Moved reserved quantifiers into official list of quantifiers

Table of Contents

1	Introduction.....	8
2	Concepts.....	10
2.1	Source.....	10
2.2	Transport.....	10
2.3	Feed.....	10
2.4	Message.....	11
2.5	Event and Event Class.....	11
2.6	Supplementary Information.....	12
2.7	Location and Directionality.....	12
2.8	Fuzziness.....	15
2.9	Route Length.....	16
2.10	Updates, Cancellations and Mergers.....	16
2.11	Message Expiration.....	17
3	XML Structure.....	18
3.1	feed.....	18
3.2	message.....	18
3.3	merge.....	19
3.4	replaces.....	19
3.5	location.....	19
3.6	from.....	22
3.7	to.....	22
3.8	at.....	22
3.9	via.....	23
3.10	not_via.....	23
3.11	polyline.....	24
3.12	events.....	24
3.13	event.....	24
3.14	supplementary_info.....	25
4	Event Classes and Events.....	25
4.1	Events in the CONGESTION Class.....	25
4.2	Events in the RESTRICTION Class.....	26
5	Supplementary Information.....	26
6	Quantifiers.....	26
7	TMC Conversion.....	27
7.1	Identifiers.....	27
7.2	Events.....	29
7.3	Supplementary Information.....	29
7.4	Event and Message Nature.....	29
7.5	Location.....	30
8	Appendix.....	30
8.1	Features for Future Versions.....	30
8.1.1	Location Types.....	31
8.1.2	Linear Location Details.....	31

8.1.3	Diversions.....	32
8.2	Reserved Quantifier Types.....	32
8.3	Further Reading.....	32

1 Introduction

This specification was created out of the desire to have a universal “hub” for traffic news, capable of aggregating traffic reports from multiple sources and delivering them to applications. Such sources include TMC, various online sources or possibly crowdsourcing platforms. In the latter case, the data format is intended to be used both for reports sent by individual contributors to a central platform and for aggregated feeds delivered by that platform to its users. Uses may include:

- Offline navigation systems, which consider the current traffic situation when calculating routes
- Traffic maps with a color visualization of the traffic situation
- Long-term analysis to develop a forecast model for traffic density

Various formats were studied, but none was found to be sufficient for the intended purpose.

XML- or JSON-based formats, such as those used by Bingⁱ or Wazeⁱⁱ, have the advantage of relying upon well-established data formats. Parsers are widely available, and the raw data is easily parseable by humans. Feeds are self-contained, requiring no external data to interpret them. Locations can be encoded in a variety of formats, ultimately by means of coordinates. The data format is carrier-agnostic, allowing it to be transferred over a variety of channels. Bandwidth requirements are moderate, allowing their use over a narrow-band Internet connection (e.g. EDGE or dial-up). However, machine parseability of events is mediocre, as vital parts of the information are given in textual form: some aspects of the information cannot easily be transferred into a form which applications can utilize, and presentation to the user is limited to the language the announcements are written in.

TMC was designed to be used over unidirectional links with highly limited bandwidth, and a complete message can be encoded in as few as 21 bits. As there is no way to transmit prose descriptions, events and their locations are transmitted as identifiers, which the receiving end translates into human-readable descriptions using lookup tables. As a result, messages are completely machine-parseable. The downside is that lookup tables are required on the receiver side to decode both the events and their locations. For events this is less of an issue: the event table covers a wide range of events, including even codes for bull fights and air raids, currently using 1552 out of 2047 possible codes, is fairly static and detail can be added by combining multiple events in a single message.

Locations, however, are subject to serious constraints: each country has a separate location table; sometimes different services for the same country use different location tables. Availability of location tables varies greatly: some are freely downloadable, others require signing a contract imposing narrow restrictions on usage, while yet other services are commercial or semi-commercial and release their location tables only for a one-off or per-device royalty. Some countries have no

TMC service at all and therefore lack TMC location tables. Locations which are not in the location table cannot be represented, and location tables typically cover only the main road network.

Additionally, the data format itself is bound to RDS as the carrier medium, relying on certain items of information from RDS (such as the country code).

The encoding presents some anomalies and redundancies: for instance, each event in a multi-event message has its own urgency, directionality and potentially parameters for speed and queue length, requiring a set of rules to resolve contradictory information between multiple events in a multi-event message. The event table has many “convenience codes” which could easily be substituted by adding more events, a quantifier, supplementary information, queue length and/or a speed limit to the message. These were probably introduced in order fit some commonly-used combinations into a single RDS group, sacrificing orthogonality for transmission efficiency.

Bing, Waze	TMC
<ul style="list-style-type: none"> • Well-established data formats (XML, JSON), parsers widely available • Raw data easily parseable by humans • Self-contained feeds (no external data required to interpret) • Carrier-agnostic data format 	<ul style="list-style-type: none"> • Custom, complex binary format, requiring custom parser • Raw data difficult to parse by humans • Feeds require external lookup tables (events, locations) to interpret • Data format bound to carrier medium • Data format allows for anomalies • Redundant event codes
<ul style="list-style-type: none"> • Moderate bandwidth requirements (EDGE, dialup) • Any location can be encoded • Not all information can be parsed as details are in prose only • Few parseable event codes • Presentation to user bound to language in which announcement is written 	<ul style="list-style-type: none"> • Low bandwidth requirements (~100 bytes/sec) • Only predefined locations can be encoded • Fully machine-parseable (all information encoded) • Large number of event codes, multiple event codes per message allowed • Presentation can be in any language

TraFF aims to eliminate these shortcomings and take the best of both worlds:

- Well-established data format for which parsers are widely available and which is easily parseable by humans in raw form
- Minimal reliance on external data for interpretation; where such data is required, it shall be as static as possible
- Carrier-agnostic data format with moderate bandwidth requirements
- Fully parseable event information with high level of detail, allowing presentation in any language supported by the target application
- Support for encoding of arbitrary locations, without relying on a set of predefined locations

The approach chosen is essentially based upon the following principles:

- Use the TMC data model as a starting point
- Encode data in XML
- Resolve event codes into mnemonics

- Where multiple semantically equivalent codes exist that differ only in an explicit or implied quantifier, combine them into one and make the quantifier a parameter
- Instead of location codes, use locations based on coordinates
- Eliminate anomalies

2 Concepts

2.1 Source

A TraFF source gathers and aggregates traffic reports and broadcasts feeds of TraFF messages. Possible examples are:

- A TMC receiver which receives traffic reports from different radio stations and converts them into TraFF messages
- A client for a crowdsourced service which receives reports from other users (either directly or through an aggregator) and supplies aggregated TraFF messages to applications
- A client which receives traffic reports from an online source (city municipalities, highway operators) and forwards them to applications

Every source must have a unique identifier, with which it prefixes the identifiers of messages it emits. Nested source identifiers are possible: for example, if the source is a TMC receiver which picks up messages from different TMC services, it may prefix the resulting messages with its global identifier, followed by an identifier for the TMC service.

2.2 Transport

The TraFF data format aims to be transport-agnostic. Being an XML-based format with no explicit carrier dependencies, it is suitable for use over any connection which supports UTF encoded XML data. With an anticipated typical message size of around 800 bytes (the shortest messages being in the range of 400 bytes), message delivery rates of TMC can easily be exceeded with an analog modem or a 2.5G mobile data connection.

Possible transport mechanisms which have been considered are:

- The operating system's request broker architecture, such as D-Bus on Linux or Binder on Android
- A raw TCP connection
- HTTP (or HTTPS)

2.3 Feed

A feed is an aggregation of messages transmitted together.

Feeds can contain one or more messages. Single-message feeds are commonly used to pass individual messages as they are received, e.g. when messages are delivered in a continuous fashion via a broadcast-like channel. Multi-message feeds are used either when messages are received in bursts of two or more messages, or when querying the messages cached by an aggregator.

2.4 Message

A message is the atomic element of traffic information, referring to a particular condition at a given location. Examples:

- 1 km of stationary traffic on the A8/A9 before Barriera Milano Nord in northwestbound direction.
- L198 closed between Lech and Warth for at least the next day.
- Heavy snowfall in the area of Vilnius.

Every message in TraFF is identified by an alphanumeric identifier which incorporates the identifier of the source from which the message originated. No two messages with the same identifier can exist at the same time.

After a message expires, its identifier can be re-used for a new message. If a message has been explicitly canceled or replaced, its identifier should not be re-used until the last expiration time ever associated with this ID has elapsed. This is to avoid confusing receivers which may have missed some updates and may therefore assume the old ID to still be valid.

2.5 Event and Event Class

An event refers to a condition, its cause or its effect. Examples:

- Queuing traffic
- Road closure
- Accident

A message may contain multiple events. For example, if a road is closed due to an accident, the message would report events for an accident and a road closure.

Similar events may be grouped into event classes, for example:

- **Congestions:** This class comprises various events for different types of congestion, such as heavy traffic, slow traffic or stationary traffic.
- **Hazards:** This class includes events such as objects on the road, nearby fires, unprotected accident areas or black ice.
- **Restrictions:** This includes closed lanes, complete road closures as well as temporary weight or size limits.

2.6 Supplementary Information

Supplementary information is extra information which can be added to an event. Multiple supplementary information items can be added to the same event. They are qualifiers, not events in their own right. Examples:

- Vehicle types to which a restriction applies
- Refinements to the position (e.g. in a tunnel, in the left lane)
- Instructions to drivers (e.g. avoid the area, approach with care)

2.7 Location and Directionality

Locations can take many shapes in a traffic report. For example:

Type	Example
Point	S16, at exit Landeck-West
Pair of points	A12, between Zams and Imst-Pitztal
Area	Rosenheim

At the minimum, they are represented by one or more coordinates. However, coordinates come with several constraints:

- Low accuracy: Some formats use a single point to represent multiple adjacent points in reality. For example, TMC generally represents motorway junctions as a single point, whereas physically they are really four distinct points (entry and exit in both directions).
- Different map data: The encoding and decoding end may work with different maps, on which coordinates for the same feature may differ slightly. The same issue arises when the two sides use different versions of the same map: changed road layouts may be reflected in one map but not in the other.

TMC presents an additional challenge for ramps and connecting carriageways at motorway junctions: these are represented by a single coordinate pair (that of the junction). The fact that the message refers to the ramps rather than the motorway, as well as the ramps affected, must be inferred from event and/or supplementary information codes. The direction is expressed through the direction of the motorway.

For this reason, extra information should be included where applicable and available:

- Road class
- Road names and numbers
- Junction names and numbers
- Carriageways affected, if the message does not refer to the main carriageway

When matching a location to a map, data users should treat the above information as hints, not as hard requirements, as different map data on both ends may cause differences:

- Maps may disagree on the road class, or the road class may have changed
- Road names may be spelled differently (Via Rossi vs. Via Marco Rossi vs. Via Rossi Marco vs. Via privata Marco Rossi; Kowalskiego vs. ulica Kowalskiego; Jono Jonaičio gatvė vs. J. Jonaičio g.)
- Road numbers may differ in delimiters (A4 vs. A 4 vs. A-4), case differences (SS342DIR vs. SS342dir), use of prefixes (DK 8 vs. 8; FFB 11 vs. K FFB 11) or even different prefixes (A 8 vs. BAB 8); some sources even invent their own road numbers (TMC for local roads in Germany).

These possibilities need to be taken into account. Rather than opting for a “match or no match” decision, data users should establish criteria for ranking matches based on their quality.

Linear features, as well as points located on linear features, additionally have a directionality indicating the direction of travel to which the report applies. Directionality is expressed by an attribute marking the location as unidirectional or bidirectional. For unidirectional locations, the direction can be expressed in one of the following ways:

- Linear features have an inherent directionality, expressed by ordering the points in the order in which they are passed (i.e. from the end of the queue to the obstruction).
- For points located on linear features, the same can be achieved by two auxiliary points on the linear feature, one situated before, the other after the target location (as seen in the direction of travel). One of these points may coincide with the target location. If map data is available, this can be achieved by picking two points adjacent to the target location on the line feature. When translating TMC events into TraFF, an offset in either direction can be used.
- A compass direction, such as “northbound”. However, this may be ambiguous in situations as the one shown in the illustration: the point lies on a stretch of the road whose direction is contrary to the general direction of the road. What is southbound with respect to the whole road is northbound with respect to the stretch between the two adjacent turns, and vice versa.
- A destination on the road, such as “towards Chur” or “towards Bellinzona”. While this is easy to turn into a human-readable indication, it is difficult to translate into a direction on a map. This is only possible with map data to translate the destinations into coordinates, and unless the road in question is tagged with a matching destination, the destination indication suffers from similar ambiguities as a compass direction.

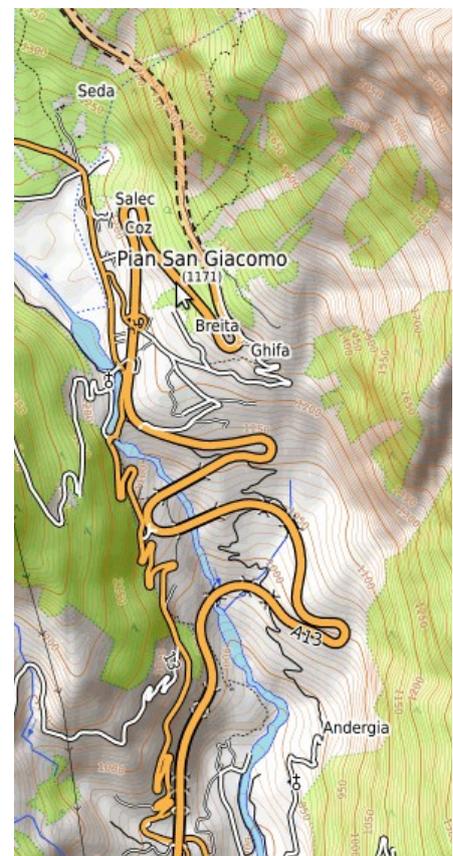


Illustration 1: A stretch of road running against the principal direction of the road

Ring roads present additional challenges here:

- The concept of a compass direction has only local relevance. Global designators would be “clockwise” and “counterclockwise”, although this may be similarly difficult to establish in practice if the road has sharp turns.
- Destinations are equally challenging, as any destination on a complete ring road can be reached in either direction. A destination on a ring road is only meaningful when coupled with the constraint to follow the shorter route, although establishing the shorter of the two possible routes may consume a significant amount of computation resources in practice, and may still be ambiguous if the location and the indicated destination are at near-opposite positions of the ring road.
- Directionality on ring roads can be expressed by an auxiliary point located in the middle of the condition, i.e. between the two end points and not coinciding with either end point. This results in three points and an order in which they are encountered, which unambiguously specifies the direction.
- The concept of a middle point will not work when converting from a service which relies on a location table (as TMC does) if the end points of the location are immediate neighbors, with no point in between which could serve as the middle point. This could be resolved by specifying a point located either before or after the location endpoints, resulting in three points to be passed in a particular order. In fact, it is not even necessary to specify whether the auxiliary point is encountered *before* or *after* the location, as the order of points will always be *auxiliary—start—end—auxiliary*. In other words, the auxiliary point is simply a point which does *not* lie within the location.
- An auxiliary point outside the location is subject to the same constraints as a middle point. However, one of the two will always be well-defined as long as the road has at least three points. Ring roads with two points would break direction in TMC as well (as each point would be the other point’s neighbor in either direction) and therefore need not be considered.
- For point locations on ring roads, the direction is unambiguous if two auxiliary points (see above) are specified, along with the order in which a driver would encounter them.

Ultimately, it is advisable to rely on coordinates for directionality wherever possible. A destination and origin (i.e. destination for the opposite direction) or a compass direction should be added, where available, for easier representation of the location in prose. In any case, these should be avoided where ambiguous.

A minimum of two points is necessary to unambiguously specify a direction on a road; three points (along with their order) are needed to specify a direction on a ring road.

A more comprehensive system for location referencing, based on similar concepts but more elaborate, is OpenLR.ⁱⁱⁱ Due to its complexity, it is not being considered for version 0.8 of this specification, although this may be revisited in future versions.

While the ability to match a location to a stretch of road on a map is one of the fundamental goals of TraFF, it may also be necessary to display locations in human-readable form. For example:

- Austria: S16 Landeck–Bludenz, between Landeck-Ost and Pians
- Poland: A4 Zgorzelec–Wrocław, between km 130 and km 135

The road reference number is a key attribute, but is guaranteed to be unique only within a certain area. For example, Austria also has a road named A4, and Poland has plans for a road named S16; therefore, the corresponding territory is necessary for disambiguation. In some cases the country is sufficient; in other cases a territorial entity within a country or even a town is needed. For example:

- “Switzerland: A11” is sufficient because the road number is unique within Switzerland.
- In the case of “Italy, Milano, viale Alcide de Gasperi” it is necessary to specify a town name because multiple towns may have streets with the same name.

The affected stretch of road can be indicated in different ways: enclosing junctions are a common representation; kilometric points are also widely used in some areas. This is the last part of each example above, and denotes the boundaries of the condition as precisely as possible.

Since drivers are not necessarily familiar with every junction or kilometric point, a second, more “global” reference is needed. Major towns along the route are frequently used for this purpose, either to indicate the stretch of road on which the affected stretch is located, the direction or both. It is recommended to choose places found on road signs for this purpose. Alternatively, or in addition, a direction indication such as “eastbound” or “westbound” can be specified.

2.8 Fuzziness

The characteristics of different traffic services have an impact on the resolution and accuracy with which they can represent locations:

- Services such as TMC are constrained by their use of location tables, which are mostly limited to junctions between major roads. Where the actual location does not coincide with one of these, the shortest pair of predefined locations which completely encloses the condition is used instead, which may cause the extent of the condition to appear longer than it actually is.
- Crowdsourced services which rely on reports by drivers inherently suffer from incomplete information: for example, a driver who has just approached the end of a traffic jam will not know how far away the obstruction is. Conversely, a driver who has just passed the obstruction has likely spent a significant amount of time in the queue and thus no longer knows how long the queue is, as it may have grown or shrunk while the driver was in it.

This is expressed by the fuzziness. The following constraints must be identifiable:

- Low resolution
- Start and/or end of condition unknown

2.9 Route Length

In prose, it is common to give reports like “8 km of stationary traffic between Kostomłoty and Kąty Wrocławskie”. TMC messages may also specify the length of the route affected, which allows for a higher granularity than the location database alone.

This typically conflicts with the inherent length of the location itself (i.e. the distance between its start and end point, or zero for a point location). TMC requires that the end points should fully enclose the location, thus their distance will always be larger than the route length.

Route length interpretation is based on two considerations:

- Be compatible with TMC: A pair of points with a length is to be interpreted in the same manner as in TMC.
- Reflect actual traffic development: In most cases, the location of the obstruction is fixed while the tail of the queue (where drivers would encounter the condition) can vary over time. Therefore, the former should be determined based on coordinates given and the latter calculated based on route length.

This results in the following rules:

- For a point location, the length indicates a stretch of road through which traffic would reach the location. A point location with a route length thus effectively becomes a point-to-point location whose end (the obstruction) coincides with the given coordinates; its start (where drivers would first encounter the condition) is implied by the route length.
- For a point-to-point location longer than the indicated length, the affected section may be located anywhere within that location.
- For a point-to-point location shorter than the indicated length, the start (where a driver would first encounter the condition) is to be inferred from the end coordinates and the route length; the start coordinates should be ignored.

2.10 Updates, Cancellations and Mergers

Messages may need to be updated as the traffic situation changes: accidents are cleared, the weather changes, queue length increases or decreases, to name just a few examples. Updates are regular messages, subject to the same format. They repeat all elements of the original message which are not being changed. Since messages are referred to by their ID, a new message replaces any previous message with the same ID.

A receiver has no way of telling if a particular message is an update or a new message other than by looking at previously stored messages. If an update is sent, one receiver may have the previous message in memory and treat the update as an update, while another receiver may have started listening after the previous message was sent and treat the message as a new message.

Cancellations are sent in a similar way: they use the same message format but use a special attribute to indicate cancellation; no location or events are included in the message. Upon receiving a

cancellation message, receivers shall delete the original message from memory and may perform additional steps (such as recalculating a route). If the original message is not found, receivers shall silently discard the cancellation message.

Crowdsourced services in particular have a need for a merge operation: Consider an accident on a motorway which causes a queue extending beyond the last junction before it. One driver approaches the end of the queue on the motorway, while another enters the motorway at the junction. Each driver sends a congestion report with a different location, and as it is not clear initially that both are referring to the same condition, two distinct messages exist. At a later time, a driver observing the congestion from the opposite carriageway sends a report with the full length of the congestion. Now up to three messages with distinct identifiers exist, which are known to refer to the same condition and thus need to be merged.

A merge operation is similar to an update: it is a full message (including a location and an event list) but additionally contains references to the messages it replaces. The merge message may have a new identifier or inherit the identifier of one of the messages to be merged together. The merge message may at the same time be a cancellation message, cancelling multiple messages at once.

The architecture of TraFF does not guarantee that every consumer will always receive the entire sequence of updates. There are many reasons why a consumer might miss individual updates to a message and keep the stale message in memory. The design of TraFF needs to take this into consideration. Specifically, consumers which have missed one or more updates to a message must have correct and up-to-date information after receiving the next update.

2.11 Message Expiration

Messages have multiple ways of indicating the period for which they are to be considered valid:

- Each message can specify an `expiration_time` attribute. This timestamp should be considered a checkpoint, i.e. the source must send an update before that time. The condition, however, may last longer than the expiration time, in which case the attribute needs to be increased through updates. After the expiration time has elapsed (unless other attributes indicate the message is still current), the message and the condition it indicates should be considered outdated. Therefore, TraFF sources should choose the duration for this attribute with care: ideally, it should be in the range of the expected duration of the condition.
- Messages can specify a `start_time` and/or `end_time` attribute, indicating when a condition is expected to begin or end. This is typically used with planned events, such as construction work or roads closed for mass events. If one or both of these are given, the message shall not be considered to have expired until the latest of these has elapsed.
- When a message updates or replaces another (including cancellations), the new message should not expire before any of the messages it supersedes would have expired. Unless enforced through other attributes, the `expiration_time` attribute should be used to indicate this.

3 XML Structure

An example of a TraFF feed may look like this:

```
<feed>
  <message id="tmc:5.1.1:5.1.1327.n.1" receive_time="2017-02-15T21:01:28+01:00"
update_time="2017-02-15T21:07:00+01:00" expiration_time="2017-02-
15T21:22:00+01:00">
    <events>
      <event class="CONGESTION" type="CONGESTION_SLOW_TRAFFIC"/>
    </events>
    <location road_class="MOTORWAY" road_ref="A4" fuzziness="LOW_RES">
      <from junction_name="Trezzo">+45.59612 +9.50253</from>
      <to junction_name="Dalmine">+45.64412 +9.62081</to>
    </location>
  </message>
</feed>
```

3.1 feed

Required	Yes, except for single-message feeds with no feed attributes, or for responses without a feed (it is required, though, for responses which include an empty feed)
Definition	This is the root element of a multi-message feed. It encapsulates all the messages in the feed.
Attributes	
Example	<feed>...</feed>
Subtags	message
Subtag of	response (see TraFF Subscription Management) if the feed is sent as part of a response in TraFF format, root tag otherwise

3.2 message

Required	Yes	
Definition	Encapsulates a single message. This is the root element of a single-message field.	
Attributes	cancellation	Boolean/Optional. <code>true</code> marks this message as a cancellation, indicating that existing messages with the same ID should be deleted or no longer considered current. All other attributes and subtags of a cancellation message should be ignored. Default is <code>false</code> .
	end_time	String/Optional. A timestamp in ISO8601 format indicating how long the condition is expected to last.
	expiration_time	String/Recommended if <code>end_time</code> is not specified. A timestamp in ISO8601 format indicating when the message will expire if not updated. An expired message should be deleted or no longer be considered current. If <code>end_time</code> is specified and is longer than <code>expiration_time</code> , <code>end_time</code> should be used to govern expiration and this attribute ignored. The expiration time governs how consumers will treat a message if they do not receive any further updates, e.g. because they lose the data connection. Sources should therefore choose the expiration time carefully: if a situation is not expected to resolve in a certain period of time, the expiration time should not be shorter. Conversely, if a situation is likely to have resolved within a certain period of time, the expiration time should not be longer than that.
	forecast	Boolean/Optional. If <code>false</code> , the message describes a current situation. If <code>true</code> , it describes an expected situation in the future. Default is <code>false</code> .

	id	String/Required. An identifier, which remains stable over the entire lifecycle of the message. The colon (:) is a reserved character to separate different levels of source identifiers from each other and from the local message identifier.
	receive_time	String/Required. A timestamp in ISO8601 format indicating at what time the message was first received by the source. Sources are expected to keep this attribute stable across all updates.
	start_time	String/Optional. A timestamp in ISO8601 format indicating when the condition is expected to begin.
	update_time	String/Required. A timestamp in ISO8601 format indicating at what time the last update to this message was received by the source.
	urgency	String/Recommended. One of NORMAL, URGENT or X_URGENT. This allows the consumer to decide how the message should be presented to the user. X_URGENT messages should be presented immediately, URGENT messages may be delayed and NORMAL messages may not need to be presented at all (e.g. for delays, it is sufficient to route the driver around them). If omitted, NORMAL shall be assumed.
Example	<code><message id="tmc:5.1.1:5.1.1327.n.1" receive_time="2017-02-15T21:01:28+01:00" update_time="2017-02-15T21:07:00+01:00" expiration_time="2017-02-15T21:22:00+01:00">...</message></code>	
Subtags	merge, events, location	
Subtag of	feed for multi-event feeds, root tag for single-event feeds	

3.3 merge

Required	No
Definition	Denotes this message as a merge message and lists all messages which it replaces.
Attributes	None
Example	<code><merge>...</merge></code>
Subtags	replaces
Subtag of	message

3.4 replaces

Required	No
Definition	A reference to another message which is replaced by this message. One element must be included for each message to be replaced. If a merge operation retains one of the original message identifiers, that message is not listed here.
Attributes	id String/Required. The identifier of a message which is replaced by the current message.
Example	<code><replaces id="foo:user5552878:45.64412N9.62081E"/></code>
Subtags	None
Subtag of	merge

3.5 location

Required	Yes, except for cancellation messages
-----------------	---------------------------------------

Definition	<p>Encapsulates the location data for the message.</p> <p>As of version 0.8, a location cannot span multiple roads, i.e. at least one of the following must be true:</p> <ul style="list-style-type: none"> • the road number remains unchanged throughout the location, • the road name remains unchanged throughout the location, • the stretch of road to which the location refers must not connect to any other road of the same or a higher category. <p>Some sections are shared by multiple roads and bear multiple road numbers. They are considered as belonging to either road and may be referenced by either number. A location may at the same time span a shared section and no more than one of the roads it belongs to, and reference them with the respective road number.</p>	
Attributes	country	String/Recommended. The ISO country code for the location.
	destination	String/Recommended. A destination, preferably the one given on road signs, which a driver would encounter after (or at) the <code>tO</code> location when following the road in the forward direction (e.g. “Milano”, “München”, “Vilnius”).
	direction	String/Optional for monodirectional locations, forbidden for bidirectional locations. A compass direction indicating the direction of travel to which this message applies (e.g. “N”, “SE”). Discouraged for ring roads (including partial ring roads) or for sections which significantly deviate from the principal direction of the main road.
	directionality	String/Recommended. One of <code>ONE_DIRECTION</code> or <code>BOTH_DIRECTIONS</code> . If omitted, <code>BOTH_DIRECTIONS</code> shall be assumed.
	fuzziness	String/Required where applicable. Permissible values: <ul style="list-style-type: none"> • <code>LOW_RES</code>: Locations are constrained to a predefined table of junctions and other waypoints (e.g. rest areas, bridges or tunnels). The actual extent of the condition may be shorter than indicated, or slightly longer if the location refers to a linear feature. When interpreting the location, matching junction attributes takes precedence over coordinates. • <code>MEDIUM_RES</code>: Locations are constrained to a predefined table of distance markers. The actual point may be located between this distance marker and it nearest neighbors in either direction. When interpreting the location, matching distance markers takes precedence over coordinates. • <code>END_UNKNOWN</code>: The end of the condition (where drivers leave the affected stretch) is unknown, as is typical for a report by a driver who has just encountered the end of a traffic jam. • <code>START_UNKNOWN</code>: The start of the condition (where drivers would first encounter it) is unknown, as is typical for a report by a driver who has just passed an obstruction. • <code>EXTENT_UNKNOWN</code>: It is unknown where the condition begins and ends, as is typical for a driver who is in the middle of a traffic jam.
	origin	String/Recommended. A destination, preferably the one given on road signs, which a driver would encounter before (or at) the <code>from</code> location when following the road in the forward direction (e.g. “Milano”, “München”, “Vilnius”). This is the counterpart of <code>destination</code> and is intended to be used for display purposes (“origin → destination”).

	ramps	String/Required where applicable. Permissible values: <ul style="list-style-type: none"> • ALL_RAMPS: Affects the entry and exit ramps only. • ENTRY_RAMP: Affects the entry ramp only. • EXIT_RAMP: Affects the exit ramp only. • NONE: No ramps are affected. Default is NONE, indicating that the message refers to the carriageways of the main road. Any of the other values imply that <i>only</i> the specified ramps are affected while the main road is not. Note that if any value other than NONE is used, the road_* attributes refer to the main road served by the ramp, not the ramp itself. This is mainly intended for compatibility with TMC, where junctions with all their ramps are represented by a single point. Other sources should use coordinate pairs instead.
	road_class	String/Recommended. The importance of the road within the road network, which roughly corresponds to the OpenStreetMap classification. One of: <ul style="list-style-type: none"> • MOTORWAY: Part of the motorway network, usually with segregated carriageways, no level crossings and restricted to motor vehicles. • TRUNK: The highest tier of non-motorway roads, may or may not be expressways with segregated carriageways. • PRIMARY: The next tier within the road network, often linking larger towns. This is often the lowest level of the national road network if TRUNK is used for expressways. • SECONDARY: The next tier, often linking towns or quarters of towns and often no longer part of the national road network but regionally administered. • TERTIARY: The next tier, often linking smaller towns and villages or neighborhoods within a town and often the lowest tier for numbered roads. • OTHER: Any road that does not fit into this classification.
	road_is_urban	Boolean/Recommended. Indicates whether the road is in a built-up area (this information is used by some data models).
	road_name	String/Recommended. A road name (e.g. “Savanorių prospektas”, “Viale Certosa”, “Wasserburger Landstraße”). Do not use this attribute if the road name is not consistent throughout the entire location.
	road_ref	String/Recommended. A road number (e.g. SP526). Do not use this attribute if the road number is not consistent throughout the entire location.
	territory	String/Recommended where applicable and needed for disambiguation. A territorial entity within a country (state, province, canton etc.), designated by its official shorthand where available and widely used (e.g. “ON” for Ontario, Canada; “NY” for New York, USA; “KL” for Kerala, India; “NSW” for New South Wales, Australia; “BE” for Bern, Switzerland), else the local name. Can be omitted if town is specified and the combination of country and town is unambiguous.
	town	String/Recommended where applicable and need for disambiguation. A town or city, designated by its local name (e.g. “Łódź”, “Rīga”, “München”). Generally needed for local street names.
Example	<location directionality="ONE_DIRECTION">...</location>	
Subtags	from, to, at, via, not_via, polyline	
Subtag of	message	

3.6 from

Required	Yes, for linear features and directional point locations on ring roads. For directional point locations on non-ring roads, either <code>from</code> or <code>to</code> must be specified along with <code>at</code> . Future versions of the specification may introduce other representations for geometries and relax the requirement for this element.	
Definition	The starting point of the location, i.e. the coordinates at which the driver would first encounter the condition reported. If <code>at</code> is specified, this element merely serves to indicate the direction, and the condition is limited to the point described by <code>at</code> .	
Attributes	<code>distance</code>	Float/Recommended where applicable. Distance marker corresponding to the point, in km (e.g 25 or 348.537).
	<code>junction_name</code>	String/Recommended where applicable. The name of a motorway junction (e.g. "Lainate"). Use only if the point refers to a junction, not to indicate the nearest junction.
	<code>junction_ref</code>	String/Recommended where applicable. The number of a motorway junction (e.g. 42). Use only if the point refers to a junction, not to indicate the nearest junction.
Example	<code><from junction_name="Trezzo">+45.59612 +9.50253</from></code>	
Subtags	None	
Subtag of	location	

3.7 to

Required	Yes, for linear features and directional point locations on ring roads. For directional point locations on non-ring roads, either <code>from</code> or <code>to</code> must be specified along with <code>at</code> . Future versions of the specification may introduce other representations for geometries and relax the requirement for this element.	
Definition	The end point of the location, i.e. the coordinates at which the driver would encounter the end of the condition reported. If <code>at</code> is specified, this element merely serves to indicate the direction, and the condition is limited to the point described by <code>at</code> .	
Attributes	<code>distance</code>	Float/Recommended where applicable. Distance marker corresponding to the point, in km (e.g 25 or 348.537).
	<code>junction_name</code>	String/Recommended where applicable. The name of a motorway junction (e.g. "Lainate"). Use only if the point refers to a junction, not to indicate the nearest junction.
	<code>junction_ref</code>	String/Recommended where applicable. The number of a motorway junction (e.g. 42). Use only if the point refers to a junction, not to indicate the nearest junction.
Example	<code><to junction_name="Dalmine">+45.64412 +9.62081</to></code>	
Subtags	None	
Subtag of	location	

3.8 at

Required	Yes, for point locations. For directional point locations on non-ring roads, either <code>from</code> or <code>to</code> must be specified along with <code>at</code> . For directional point locations on ring roads, all of <code>from</code> , <code>at</code> and <code>to</code> must be specified and no two points may coincide.	
Definition	The coordinates of the condition reported. If <code>from</code> and/or <code>to</code> are specified, they merely serve to indicate the direction, and the condition is limited to the point described by <code>at</code> .	
Attributes	<code>distance</code>	Float/Recommended where applicable. Distance marker corresponding to the point, in km (e.g 25 or 348.537).

	junction_name	String/Recommended where applicable. The name of a motorway junction (e.g. “Lainate”). Use only if the point refers to a junction, not to indicate the nearest junction.
	junction_ref	String/Recommended where applicable. The number of a motorway junction (e.g. 42). Use only if the point refers to a junction, not to indicate the nearest junction.
	radius	Not part of version 0.8 of the specification, reserved for future use.
Example	<code><at junction_name="Dalmine">+45.64412 +9.62081</at></code>	
Subtags	None	
Subtag of	location	

3.9 via

Required	Either <code>via</code> or <code>not_via</code> must be present for linear locations on ring roads. Both <code>from</code> and <code>to</code> must be specified along with <code>via</code> . Do not use for point locations.	
Definition	A location on the way between <code>from</code> and <code>to</code> , indicating the direction of a ring road to which the message applies.	
Attributes	distance	Float/Recommended where applicable. Distance marker corresponding to the point, in km (e.g 25 or 348.537).
	junction_name	String/Recommended where applicable. The name of a motorway junction (e.g. “Lainate”). Use only if the point refers to a junction, not to indicate the nearest junction.
	junction_ref	String/Recommended where applicable. The number of a motorway junction (e.g. 42). Use only if the point refers to a junction, not to indicate the nearest junction.
	radius	Not part of version 0.8 of the specification, reserved for future use.
Example	<code><via junction_name="Dalmine">+45.64412 +9.62081</at></code>	
Subtags	None	
Subtag of	location	

3.10 not_via

Required	Either <code>via</code> or <code>not_via</code> must be present for linear locations on ring roads. Both <code>from</code> and <code>to</code> must be specified along with <code>via</code> . Do not use for point locations.	
Definition	A location on the way between <code>from</code> and <code>to</code> , indicating the direction of a ring road to which the message applies.	
Attributes	distance	Float/Recommended where applicable. Distance marker corresponding to the point, in km (e.g 25 or 348.537).
	junction_name	String/Recommended where applicable. The name of a motorway junction (e.g. “Lainate”). Use only if the point refers to a junction, not to indicate the nearest junction.
	junction_ref	String/Recommended where applicable. The number of a motorway junction (e.g. 42). Use only if the point refers to a junction, not to indicate the nearest junction.
	radius	Not part of version 0.8 of the specification, reserved for future use.
Example	<code><not_via junction_name="Dalmine">+45.64412 +9.62081</at></code>	
Subtags	None	
Subtag of	location	

3.11 polyline

Not part of version 0.8 of the specification, reserved for future use.

Required	Yes, for bidirectional locations and linear features, unless <code>from</code> and <code>to</code> are used
Definition	The course of the location, from the coordinates at which the driver would first encounter the condition reported to the coordinates at which the driver would encounter its end. The location is assumed to approximately follow an uninterrupted line from one coordinate pair to the next. It is permissible to simplify geometries by omitting points in between, as long as it is still possible to unambiguously reconstruct the route on a map. It is recommended to: <ul style="list-style-type: none"> • Retain one point for each intersection with another road (except in cases where all intersecting roads are dead ends). • Retain points at noticeable turns in the road so that the geometry of the road is well represented. This ensures compatibility with later versions of the map in which intersections may have been added.
Attributes	None
Example	<code><polyline>+45.67781 +9.01643 +45.62333 +9.01344 +45.57642 +9.00929</polyline></code>
Subtags	None
Subtag of	location

3.12 events

Required	Yes, except for cancellation messages
Definition	Encapsulates the events for the message
Attributes	None
Example	<code><events>...</events></code>
Subtags	event
Subtag of	message

3.13 event

Required	Yes												
Definition	Describes a single event, such as the cause of a disruption or its impact on traffic, in machine-parseable form.												
Attributes	<table border="1"> <tr> <td><code>class</code></td> <td>String/Required. The event class (generic category).</td> </tr> <tr> <td><code>length</code></td> <td>Integer/Optional. The length of the affected route in meters.</td> </tr> <tr> <td><code>probability</code></td> <td>Not part of version 0.8 of the specification, reserved for future use. Integer/Optional for forecast events, forbidden otherwise. The probability in percent for the forecast (common in weather forecasts).</td> </tr> <tr> <td><code>q_*</code></td> <td>Variant/Optional. Additional quantifier for events allowing this. Permissible data types and their meanings depend on the event type.</td> </tr> <tr> <td><code>speed</code></td> <td>Integer/Optional. The speed at which vehicles can expect to pass through the affected stretch of road. This can either be a temporary speed limit or the average speed in practice: some events imply a particular interpretation; otherwise it refers to the lower of the two speeds.</td> </tr> <tr> <td><code>type</code></td> <td>String/Required. The event type, which uniquely identifies the event and can be mapped to a string to be displayed to the user.</td> </tr> </table>	<code>class</code>	String/Required. The event class (generic category).	<code>length</code>	Integer/Optional. The length of the affected route in meters.	<code>probability</code>	Not part of version 0.8 of the specification, reserved for future use. Integer/Optional for forecast events, forbidden otherwise. The probability in percent for the forecast (common in weather forecasts).	<code>q_*</code>	Variant/Optional. Additional quantifier for events allowing this. Permissible data types and their meanings depend on the event type.	<code>speed</code>	Integer/Optional. The speed at which vehicles can expect to pass through the affected stretch of road. This can either be a temporary speed limit or the average speed in practice: some events imply a particular interpretation; otherwise it refers to the lower of the two speeds.	<code>type</code>	String/Required. The event type, which uniquely identifies the event and can be mapped to a string to be displayed to the user.
<code>class</code>	String/Required. The event class (generic category).												
<code>length</code>	Integer/Optional. The length of the affected route in meters.												
<code>probability</code>	Not part of version 0.8 of the specification, reserved for future use. Integer/Optional for forecast events, forbidden otherwise. The probability in percent for the forecast (common in weather forecasts).												
<code>q_*</code>	Variant/Optional. Additional quantifier for events allowing this. Permissible data types and their meanings depend on the event type.												
<code>speed</code>	Integer/Optional. The speed at which vehicles can expect to pass through the affected stretch of road. This can either be a temporary speed limit or the average speed in practice: some events imply a particular interpretation; otherwise it refers to the lower of the two speeds.												
<code>type</code>	String/Required. The event type, which uniquely identifies the event and can be mapped to a string to be displayed to the user.												

Example	<code><event class="CONGESTION" type="CONGESTION_SLOW_TRAFFIC"> ...</event></code>
Subtags	supplementary_info
Subtag of	events

3.14 supplementary_info

Required	No	
Definition	Adds supplementary information to an event, such as vehicles or places to which the event applies, warnings or instructions.	
Attributes	class	String/Required. The supplementary information class (generic category).
	q_*	Variant/Optional. Additional quantifier for supplementary information types allowing this. Permissible data types and their meanings depend on the type.
	type	String/Required. The supplementary information type, which can be mapped to a string to be displayed to the user.
Example	<code><supplementary_info class="..." type="..."/></code>	
Subtags	None	
Subtag of	event	

4 Event Classes and Events

Beginning with version 0.8 of the TraFF specification, event classes and events are being maintained in a separate list. The main specification merely specifies general design principles.

\$Q in event descriptions refers to a quantifier which is commonly used with the event. Such values are always optional; if absent, the text in the `Text_Q0` column shall be used.

4.1 Events in the CONGESTION Class

This class has different event types for different degrees of congestion, giving an indication of average speed:

- Heavy traffic, traffic [much] heavier than normal, (ES: tráfico verde, DE: reger Verkehr, [sehr viel] dichter Verkehr als normal): Road is close to saturation. Ability to change lanes is somewhat restricted but average speed is close to the posted limit.
- Slow traffic (ES: tráfico amarillo, DE: dichter Verkehr): Capacity of the road is reached. Speed may vary rapidly and rarely reaches normal speed. Average speed is significantly lower than the posted limit but traffic is still moving.
- Queue (ES: tráfico rojo, DE: stockender Verkehr): Average speed 10–30 km/h.
- Stationary traffic (ES: tráfico rojo, DE: Stau): Average speed is less than 10 km/h.
- Long queues (ES: tráfico retenido, DE: lange Staus): Significant periods of standstill in which drivers turn off their engines and some leave their vehicles.

- Congestion (ES: congestión, DE: Verkehrsstörung): Generic category covering all of the previous four. Speed is significantly lower than the posted limit, anywhere between steadily moving traffic and long standstills. This event should be used when no accurate quantification is possible (as in the case of drivers observing a congestion without being in it, typically in the opposite direction of a dual-carriageway road).

If an explicit **speed** attribute is given, it takes precedence over the speed implied by the event. However, sources should keep the event consistent with the reported speed.

4.2 Events in the RESTRICTION Class

The restriction class contains events referring to closure or blockage of a *carriageway*. Unless otherwise indicated by a supplementary information item, these refer to the main carriageway.¹

Various events refer to a road, or parts of a road, being *blocked* or *closed*. The difference between these two is that a blockage refers to the presence of a physical obstacle which makes passage impossible. However, unless the road is also closed, attempts to use a blocked road will not normally result in a fine or other actions by authorities. Closures, on the contrary, refer to administrative restrictions: unless the road is also blocked, it can still be used by emergency services or with a special permit, but other traffic will be stopped by authorities or fined.

5 Supplementary Information

Beginning with version 0.8 of the TraFF specification, supplementary information classes and types are being maintained in a separate list.

6 Quantifiers

Quantifiers can be used with events and with supplementary information. The syntax is the same in both cases.

The quantifier type is specified in the attribute name itself. The following quantifier types are specified in version 0.8:

Name	Description	Default unit	Examples
q_dimension	A dimension (e.g. length/width/height restriction, visibility, amount of precipitation). Note that queue length or length of the route affected are expressed through the <code>length</code> attribute.	m	100 m 2.5 m 255 cm 0.5 cm 5 mm
q_duration	A duration (e.g. delay)	Minutes, hours/minutes	20 min 1:30 3 h

¹ In TMC, these events may be accompanied by a supplementary information item “use parallel carriageway”, indicating other carriageways are available which are not affected. Where this is not the case (and only one carriageway is present), this event indicates a blocked/closed road.

Name	Description	Default unit	Examples
q_int	An integer (e.g. number of distinct events, available spaces)	None	2 40
q_ints	A list of integer values	None	3, 2 2 (single integer value)
q_speed	A speed that does <i>not</i> refer to moving traffic (e.g. wind, queue length increase/ decrease). Note that the speed of moving traffic is expressed through the <code>speed</code> attribute.	km/h	40 km/h
q_temperature	A temperature	°C	-20 °C 40 °C
q_time	Time of day (e.g. ferry departure time)	ISO8601	2017-02-17T21:30+01:00
q_weight	A weight (e.g. weight beyond which a restriction applies)	t	20 t

Units other than the default unit may be used (e.g. imperial instead of metric units) but must be specified explicitly. The default unit is assumed if no unit is given.

7 TMC Conversion

7.1 Identifiers

Two levels of TraFF source identifiers are used: The first is the global identifier for TMC. The second level identifies the TMC service from which the message originated. It consists of the 4-bit country code, Location Table Number and Service Identifier.

The local message identifier contains the elements of the message which are not contained in the source identifier and, as per the TMC specification, are used to determine if a message overrides another: location code, direction, update classes and duration (the latter is for forecast messages only).

Because TMC messages may contain multiple events, each belonging to a different update class, and messages can override each other if any update class matches, it is impossible to design an identifier that is generated from the above identifiers, collision-free and at the same time guaranteed to be identical for any two messages which can override each other. Similar issues apply to the service identifier, as different services may be allowed to override each other's messages.

As identifiers need to be stable across updates, an updated message will always inherit the identifier of the message it replaces. This may result in discrepancies between the update classes of the updated message and those reported in the identifier: for example, if a message comprising events for update classes 1 and 2 is replaced with a message comprising events for update classes 2 and 3, the identifier will still report update classes 1 and 2. Similarly, if a message was first received by a service with SID 5 but later updated by a service with SID 6 (assuming the two services allow updates between each other), it will continue to report SID 5 in its ID.

Note that this also means that TMC converters on two different devices may assign different identifiers to the same message:

- Device A receives the first message in the sequence, as well as all updates to it. By the time device B picks up the first update to that message, the set of update classes has changed and device B will report different update classes in its identifier than device A does.
- Device A and B are listening to two different TMC services (different SIDs) which can update messages between them and share some or all of the messages they transmit. Both may receive the same message but assign it IDs reporting a different SID.

An example for a fully qualified TMC message ID would be `tmc:a.1.0:a.1.32908.n.3,5`. This can be broken down as follows:

- The message was received via TMC
- The Country Code of the service which sent this message is A (together with the Location Table Number, this identifies Austria)
- The Location Table Number of the service which sent this message is 1
- The Service Identifier is 0
- The Country Code and Location Table Number to resolve the Location Code are A and 1, respectively (same as for the sending service, unless it is an INTER-ROAD message)
- The Location Code is 32908
- The message refers to the negative direction of the road (but may be either unidirectional or bidirectional)
- The original message comprises events of update classes 3 and 5. (To minimize ambiguities, update classes should be deduplicated and sorted.)

Since the system has been designed to be primarily collision-free, messages with the same identifier can override each other in most cases. There are, however, scenarios in which this can break:

- A message with update classes 2 is received and assigned an ID that reports this update class.
- An update for the message is received, which now includes update class 2 and 3.
- Another update is received, which now includes only update class 3. Note that the ID still reports update class 2, and only 2.
- Now a different message is received, from the same service, referring to the same location and direction, and with update class 2. This message can coexist with our previous message in a TMC receiver. However, it would be assigned the same TraFF identifier as the previous message, therefore a TraFF consumer would treat it as an update to the earlier message.

This could be mitigated by the following approach:

- Generate the ID for each message (including updates) from scratch

- If the message updates an existing message and its ID is different from the ID of the old message, add a `replaces` element for the old ID. This element should be repeated in all updates until the expiration time for the last message with the old ID has elapsed.

Alternatively:

- When a new message is received and it does not replace an existing message, check if its generated ID is unique.
- If the generated ID is not unique, append a distinguishing element (e.g. a dot followed by a number which is incremented until the resulting ID is unique).

7.2 Events

Some TMC events have a direct equivalent in TraFF. However, many TMC events are “convenience” elements which assign a single event code to several items of information which frequently occur together. Since TraFF was designed for orthogonality rather than bandwidth economy, such TMC events may convert to several items of information in TraFF:

- One or more events
- A Supplementary Information item for each event
- A speed attribute or quantifier

7.3 Supplementary Information

Converting supplementary information from TMC into TraFF is straightforward in most cases. The two main points to be aware of are:

- Supplementary information in TraFF supports quantifiers just like events do. The quantifier format is the same.
- Some supplementary information items in TMC translate into events in TraFF.

7.4 Event and Message Nature

TMC defines three event *natures*: Info, Forecast and Silent. Info describes a current situation while Forecast describes a situation expected for the near future. Silent denotes messages which are to be processed internally without being presented to the user.

Each event has an assigned nature, which cannot be changed. Consequently, there are many event pairs, one referring to a current situation and the other referring to the same situation being expected in the future. The overall nature of the message is determined by the event which immediately precedes the duration field.

As only few events are of Silent nature and since TraFF leaves the decision how to process or present them to the processing application, TraFF lacks an explicit counterpart to Silent.

Whether a message describes a current situation or is a forecast message is expressed in an attribute of the `message` element, as TraFF has no mechanism to infer this information from events. TMC receivers are therefore expected to evaluate the nature of each received message and set the `forecast` attribute accordingly.

7.5 Location

A TMC message defines a *location*, a *direction* and an *extent*. A direction is even specified for bidirectional messages. The extent, together with the direction, implies a secondary location; however, the secondary location coincides with the primary location if the extent is zero. A zero extent thus implies a point location, a nonzero extent implies a linear location.

The road class implies whether the road is a ring road.

A TMC location can thus be transformed into a TraFF location in the following way:

If the extent is zero, the location is a point location and its coordinates must be stored in the `at` element. Unless the message is bidirectional, auxiliary points must be specified. `from` can be filled by obtaining the next point in the direction of the event; `to` can be filled by obtaining the next point in the opposite direction. One of these two can be omitted if no further points exist in the respective direction. On ring roads, two adjacent points will always be available.

If the extent is nonzero, the location is a linear location. The coordinates of the location must be stored in the `to` element. The `from` element is populated with the coordinates of the point at the offset indicated by the extent in the given direction. (Should no point be available at that offset, e.g. because the message is corrupt or an outdated location table is used, use the last point which can be retrieved.)

If the location is a linear location and the road class indicates a ring road, either the `via` or the `not_via` point must be populated. It can be obtained by picking an arbitrary offset that is greater than zero but less than the extent for `via` or greater than the extent (but less than the total number of points on the road) for `not_via`, and using the point at that offset in the given direction. The recommendation is to use `via` if the extent is greater than 1, and `not_via` with an extent of 2 otherwise.

8 Appendix

8.1 Features for Future Versions

- Dealing with mixed-nature TMC events such as “roadworks, delays expected”, whose TraFF equivalent is composed of two events, one of which is a forecast while the other is not.
- Parent messages: event in one location causing events in other locations
- Equivalent to TMC update class 30 (area and level of detail covered by this service)

8.1.1 Location Types

Locations may be extended to include:

- Areas (delimited by a polygon)
- Fuzzy areas (areas without a clear delimitation)
- Administrative areas

8.1.2 Linear Location Details

Currently, locations are identified by no more than three reference points (and some auxiliary attributes), which may not be sufficient for an unambiguous reference in all cases.

This could be solved by representing linear features as a series of points. Design criteria would be:

- **Compatibility:** The detailed geometry would be represented as a new child tag of `location` and used along with the existing point tags.
- **Orthogonality/anomaly prevention:** `from` and `to` would continue to be used as the start and end tag for the new geometry; the new tag would only list the points in between.
- **Economy:** The geometry comprises only the points necessary to unambiguously reconstruct the geometry with the aid of a map, but is not in itself an accurate representation of the geometry. Keeping the number of points low saves both processing power and computing resources for map matching.

The last criterion is intended as a recommendation rather than a hard requirement, given that compliance may be impractical to assess, and sources may not have the necessary information to encode a location in the recommended manner.

Unlike with other models, there will be no constraint on the maximum distance between two adjacent reference points, or on how far the actual road geometry may diverge from the straight line between two adjacent points.

The recommendation is to place one reference point on the middle of each *segment*. A segment is delimited by the start point, the end point or a *junction*. A junction is any point where a driver has a choice between at least two segments to continue their journey on, or which they can reach from at least two different segments (other than the segment on which they continue their journey). Therefore, any place where a carriageway splits in two or two carriageways merge are also considered junctions.

To account for rerouting of roads, the beginning and end of a rerouted section of road should also be treated as junctions, even if the old road is abandoned after the rerouted section is opened to traffic. This should also be applied to planned reroutings, and for a reasonable period of time after an abandoned road has been demolished. The same rule applies to planned junctions with a new road.

Map matching can then be performed by means of least-cost routing from the start point to the end point, using the length of segments as well as proximity to the nearest reference point to determine segment cost.

If a source does not have the necessary information to represent a location in the intended manner, it may need to fall back to an alternative strategy:

- If the source obtains its location from an Alert-C location table or similar, it may use reference points between the start and end point. If they are very close to each other and/or on a straight line, the source may choose to eliminate some of them.
- If the location was measured on the ground (e.g. with a GPS device), the source can identify turns and treat them as junctions. It should then eliminate points in the immediate vicinity of the junctions and finally reduce near-straight stretches to a small number of points.

8.1.3 Diversions

Messages may specify a recommended diversion. This is a feature already present in TMC, where a diversion is valid for the whole message.

TMC has two ways to specify a recommended diversion:

- A diversion flag. Setting this flag indicates that drivers are recommended to take a diversion route which is marked on the ground and valid for all destinations.
- An explicit diversion, given as a sequence of locations. Such diversions may be indicated as being valid for a particular destination only. Multiple explicit diversions may be specified.

Where multiple diversions are specified, all but one must refer to a particular destination.

8.2 Reserved Quantifier Types

The following quantifier types are reserved for use in future versions:

Name	Description	Default unit	Examples
q_frequency	A radio frequency (e.g. for a radio station broadcasting further information)	MHz	99.0 MHz 720 kHz
q_percent	A percentage	%	50 %

8.3 Further Reading

Additional documentation is available from the TraFF repository:

TraFF Events is a list of TraFF event classes and events, along with their specification status.

TraFF Supplementary Information is a list of TraFF supplementary information classes and supplementary information types, along with their specification status.

TraFF Subscription Management describes how TraFF feeds are transported between peers.

- i Traffic Incident Data. <https://msdn.microsoft.com/en-us/library/hh441730.aspx>, retrieved 2017-04-11.
- ii CIFS: Waze Closure and Incident Feed Specification v2.0, https://blog.waze.com/p/blog-page_19.html, retrieved 2017-04-11.
- iii OpenLR. <http://www.openlr.org/>, retrieved 2017-04-11.