



OpenLCB Technical Note	
Protocol Identification Protocol	
May 12, 2012	Preliminary

1 Introduction

OpenLCB defines a number of specific protocols for interacting with nodes, including event exchange, datagrams, streams, configuration, etc. These protocols are optional, in the sense that not every node will implement every one.

- 5 To determine which protocols a node implements, a Protocol Support Inquiry message is sent to the specific node. It will reply with a Protocol Support Reply message that contains six bytes of data. A specific bit position has been reserved for each defined protocol. If the bit is zero or not present, the protocol is not supported and requests to use it will result in a error. If present and 1, the protocol is supported.
- 10 It is not necessary to check whether an addressed protocol is supported by a node before attempting to use the protocol. If it's not, the standard error handling mechanism will indicate that.

This protocol provides a way to check, without errors, whether the protocol is supported. Avoiding errors provides a cleaner system. Further, this protocol can check support for protocols that use global (non-addressed messages); nodes are not permitted to return errors for global messages.

2 Annotations to the Standard

2.1 Introduction

2.2 Intended Use

- 20 Although this protocol isn't mandatory, it's a very good idea to provide it. High-function nodes such as PC programs are likely to use it to customize their operation. Programmers are more likely to use this protocol rather than write special error-handling code for each of the many different protocols. Use of this protocol also avoids "intentional errors", which might confuse or annoy users.

2.3 Reference and Context

We don't explicitly reference all the protocol definitions that are associated with specific bits. Implementors can find them from the protocol names, and we don't want to clog up this section of the Standard with a set of references that we'll have to update continually.

2.4 Messages

- 30 The CAN messages were defined to be part of the “simple” subset. Low-end nodes may want to implement this protocol so that higher-function nodes can easily learn their limitations.

Generally, the node designer will just provide a simple fixed value for the reply. On CAN, the entire frame is fixed except for the source and destination aliases. The destination alias can be taken directly from the request packet.

35 2.5 Interactions

Other nodes can snoop on these interactions to learn the protocols supported by a node, without having to send their own inquiry. It's not expected that the protocols supported by a specific node will change with time, although the Standard does not require that they be immutable.

- A node supporting this protocol must promptly reply to the Inquiry message with the Reply message.
- 40 On the other hand, the general message standard requires that a node not supporting this protocol reply to the Inquiry message with a Optional Interaction Rejected message. Either way, a node sending a Protocol Id Inquiry message can count on getting a rapid reply.

2.6 Protocol Identification Values

The length of the field is constrained by the CAN packet length.

- 45 OpenLCB is big-endian, so these have been assigned from the MSB of the 1st byte.

The MSB for Protocol Identification Protocol is always 1, because it has to be present for the reply to be generated.

0x00 00 00 00 00 08 through 0x00 00 00 00 00 01 are reserved for control of future expansion. They'll eventually be needed when there are more protocols defined than can fit in a single CAN frame.

- 50 In general, nodes using this protocol don't need to know the meaning of bits defined after the node was created. A node is looking to see whether a particular protocol is present or not so that it can select what actions to take. A newly-defined protocol isn't among the things that the node will try to reason about. This allows us to eventually extend the length of the reply in any one of several ways.

- 55 We define the bits here, rather than in the individual protocol definitions, to reduce the risk of duplicate assignments. Duplicate assignments would be obvious here, but not so much when spread across separate documents.

Table of Contents

1 Introduction.....	1
2 Annotations to the Standard.....	1
2.1 Introduction.....	1
2.2 Intended Use.....	1

2.3 Reference and Context.....	1
2.4 Messages.....	2
2.5 Interactions.....	2
2.6 Protocol Identification Values.....	2

DRAFT