

**Draft Recommendation for
Space Data System Standards**

**SCHEDULE-AWARE
BUNDLE ROUTING**

DRAFT RECOMMENDED STANDARD

CCSDS 734.3-R-1

RED BOOK
July 2018



CCSDS

The Consultative Committee for Space Data Systems

**Draft Recommendation for
Space Data System Standards**

**SCHEDULE-AWARE
BUNDLE ROUTING**

DRAFT RECOMMENDED STANDARD

CCSDS 734.3-R-1

RED BOOK

July 2018

AUTHORITY

Issue:	Red Book, Issue 1
Date:	July 2018
Location:	Not Applicable

(WHEN THIS RECOMMENDED STANDARD IS FINALIZED, IT WILL CONTAIN THE FOLLOWING STATEMENT OF AUTHORITY:)

This document has been approved for publication by the Management Council of the Consultative Committee for Space Data Systems (CCSDS) and represents the consensus technical agreement of the participating CCSDS Member Agencies. The procedure for review and authorization of CCSDS documents is detailed in *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-4), and the record of Agency participation in the authorization of this document can be obtained from the CCSDS Secretariat at the e-mail address below.

This document is published and maintained by:

CCSDS Secretariat
National Aeronautics and Space Administration
Washington, DC, USA
E-mail: secretariat@mailman.ccsds.org

STATEMENT OF INTENT

(WHEN THIS RECOMMENDED STANDARD IS FINALIZED, IT WILL CONTAIN THE FOLLOWING STATEMENT OF INTENT:)

The Consultative Committee for Space Data Systems (CCSDS) is an organization officially established by the management of its members. The Committee meets periodically to address data systems problems that are common to all participants, and to formulate sound technical solutions to these problems. Inasmuch as participation in the CCSDS is completely voluntary, the results of Committee actions are termed **Recommended Standards** and are not considered binding on any Agency.

This **Recommended Standard** is issued by, and represents the consensus of, the CCSDS members. Endorsement of this **Recommendation** is entirely voluntary. Endorsement, however, indicates the following understandings:

- o Whenever a member establishes a CCSDS-related **standard**, this **standard** will be in accord with the relevant **Recommended Standard**. Establishing such a **standard** does not preclude other provisions which a member may develop.
- o Whenever a member establishes a CCSDS-related **standard**, that member will provide other CCSDS members with the following information:
 - The **standard** itself.
 - The anticipated date of initial operational capability.
 - The anticipated duration of operational service.
- o Specific service arrangements shall be made via memoranda of agreement. Neither this **Recommended Standard** nor any ensuing **standard** is a substitute for a memorandum of agreement.

No later than five years from its date of issuance, this **Recommended Standard** will be reviewed by the CCSDS to determine whether it should: (1) remain in effect without change; (2) be changed to reflect the impact of new technologies, new requirements, or new directions; or (3) be retired or canceled.

In those instances when a new version of a **Recommended Standard** is issued, existing CCSDS-related member standards and implementations are not negated or deemed to be non-CCSDS compatible. It is the responsibility of each member to determine when such standards or implementations are to be modified. Each member is, however, strongly encouraged to direct planning for its new standards and implementations towards the later version of the Recommended Standard.

FOREWORD

Through the process of normal evolution, it is expected that expansion, deletion, or modification of this document may occur. This Recommended Standard is therefore subject to CCSDS document management and change control procedures, which are defined in the *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-4). Current versions of CCSDS documents are maintained at the CCSDS Web site:

<http://www.ccsds.org/>

Questions relating to the contents or status of this document should be sent to the CCSDS Secretariat at the e-mail address indicated on page i.

At time of publication, the active Member and Observer Agencies of the CCSDS were:

Member Agencies

- Agenzia Spaziale Italiana (ASI)/Italy.
- Canadian Space Agency (CSA)/Canada.
- Centre National d'Etudes Spatiales (CNES)/France.
- China National Space Administration (CNSA)/People's Republic of China.
- Deutsches Zentrum für Luft- und Raumfahrt (DLR)/Germany.
- European Space Agency (ESA)/Europe.
- Federal Space Agency (FSA)/Russian Federation.
- Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
- Japan Aerospace Exploration Agency (JAXA)/Japan.
- National Aeronautics and Space Administration (NASA)/USA.
- UK Space Agency/United Kingdom.

Observer Agencies

- Austrian Space Agency (ASA)/Austria.
- Belgian Federal Science Policy Office (BFSP0)/Belgium.
- Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
- China Satellite Launch and Tracking Control General, Beijing Institute of Tracking and Telecommunications Technology (CLTC/BITTT)/China.
- Chinese Academy of Sciences (CAS)/China.
- Chinese Academy of Space Technology (CAST)/China.
- Commonwealth Scientific and Industrial Research Organization (CSIRO)/Australia.
- Danish National Space Center (DNSC)/Denmark.
- Departamento de Ciência e Tecnologia Aeroespacial (DCTA)/Brazil.
- Electronics and Telecommunications Research Institute (ETRI)/Korea.
- European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
- European Telecommunications Satellite Organization (EUTELSAT)/Europe.
- Geo-Informatics and Space Technology Development Agency (GISTDA)/Thailand.
- Hellenic National Space Committee (HNSC)/Greece.
- Indian Space Research Organization (ISRO)/India.
- Institute of Space Research (IKI)/Russian Federation.
- Korea Aerospace Research Institute (KARI)/Korea.
- Ministry of Communications (MOC)/Israel.
- Mohammed Bin Rashid Space Centre (MBRSC)/United Arab Emirates.
- National Institute of Information and Communications Technology (NICT)/Japan.
- National Oceanic and Atmospheric Administration (NOAA)/USA.
- National Space Agency of the Republic of Kazakhstan (NSARK)/Kazakhstan.
- National Space Organization (NSPO)/Chinese Taipei.
- Naval Center for Space Technology (NCST)/USA.
- Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
- Scientific and Technological Research Council of Turkey (TUBITAK)/Turkey.
- South African National Space Agency (SANSA)/Republic of South Africa.
- Space and Upper Atmosphere Research Commission (SUPARCO)/Pakistan.
- Swedish Space Corporation (SSC)/Sweden.
- Swiss Space Office (SSO)/Switzerland.
- United States Geological Survey (USGS)/USA.

PREFACE

This document is a draft CCSDS Recommended Standard. Its 'Red Book' status indicates that the CCSDS believes the document to be technically mature and has released it for formal review by appropriate technical organizations. As such, its technical contents are not stable, and several iterations of it may occur in response to comments received during the review process.

Implementers are cautioned **not** to fabricate any final equipment in accordance with this document's technical content.

Recipients of this draft are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 734.3-R-1	Schedule-Aware Bundle Routing, Draft Recommended Standard, Issue 1	July 2018	Current draft

CONTENTS

<u>Section</u>	<u>Page</u>
1 INTRODUCTION	1-1
1.1 PURPOSE.....	1-1
1.2 SCOPE.....	1-1
1.3 NOMENCLATURE	1-1
1.4 DEFINITIONS.....	1-2
1.5 REFERENCES	1-3
2 OVERVIEW	2-1
2.1 GENERAL.....	2-1
2.2 ARCHITECTURAL CONSIDERATIONS.....	2-3
2.3 DATA STRUCTURES.....	2-4
2.4 KEY CONCEPTS.....	2-6
3 ROUTE DETERMINATION PROCEDURES	3-1
3.1 GENERAL.....	3-1
3.2 CONTACT GRAPH ROUTING.....	3-1
3.3 ALTERNATIVE ROUTE DETERMINATION PROCEDURES	3-10
3.4 EXCEPTION HANDLING	3-11
ANNEX A PROTOCOL IMPLEMENTATION CONFORMANCE STATEMENT PROFORMA (NORMATIVE)	A-1
ANNEX B SECURITY, SANA, AND PATENT CONSIDERATIONS (INFORMATIVE)	B-1
ANNEX C INFORMATIVE REFERENCES (INFORMATIVE)	C-1
ANNEX D ABBREVIATIONS (INFORMATIVE)	D-1

Figure

3-1 Network Topology Example.....	3-2
3-2 Contact Plan Example: Contacts	3-2
3-3 Contact Plan Example: Range Intervals	3-3
3-4 Node A’s Contact Graph for Node D, Given This Contact Plan.....	3-3

1 INTRODUCTION

1.1 PURPOSE

This document defines a Recommended Standard for Schedule-Aware Bundle Routing (SABR) in the forwarding of Delay-Tolerant Networking (DTN) bundles in the space environment. SABR provides dynamic route computation in an environment of stable topology but time-varying connectivity when instances of connectivity are scheduled rather than opportunistic.

1.2 SCOPE

SABR is intended for use in operation of the Bundle Protocol (BP) in the Solar System Internet.

1.3 NOMENCLATURE

1.3.1 NORMATIVE TEXT

The following conventions apply for the normative specifications in this Recommended Standard:

- a) the words ‘shall’ and ‘must’ imply a binding and verifiable specification;
- b) the word ‘should’ implies an optional, but desirable, specification;
- c) the word ‘may’ implies an optional specification;
- d) the words ‘is’, ‘are’, and ‘will’ imply statements of fact.

These conventions do not imply constraints on diction in text that is clearly informative in nature.

1.3.2 INFORMATIVE TEXT

In the normative sections of this document, informative text is set off from the normative specifications either in notes or under one of the following subsection headings:

- Overview;
- Background;
- Rationale;
- Discussion.

1.4 DEFINITIONS

1.4.1 BUNDLE PROTOCOL TERMS

Within the context of this document the following definitions apply:

bundle: A unit of data transmitted via the DTN bundle protocol from one DTN node (termed the bundle's *source*) to another (termed the bundle's *destination*). Each bundle comprises a primary block, zero or more extension blocks, and a payload block.

class of service: An indicator of the importance of a bundle as assessed by the application that provided the content of the bundle's payload block.

contact plan: A plan made up of contacts and range intervals.

contact: An interval during which it is expected that data will be transmitted by a sending node and most or all of the transmitted data will be received by a receiving node.

critical bundle: A bundle that has the critical quality-of-service flag set.

custody: Retention of a copy of a forwarded bundle until conditions that allow deletion of that bundle are satisfied.

duration: Of a contact, end time minus start time.

entry node: The receiving node for the first contact of a route.

estimated convergence-layer overhead: For a bundle whose header is of size M and whose payload is of size N , three percent of $(M + N)$, or 100 bytes, whichever is larger.

estimated volume consumption, EVC: For a bundle, the sum of the sizes of the bundle's payload and header and the estimated convergence-layer overhead.

excluded neighbor: A neighboring node that refuses custody of a bundle destined for some remote node.

expiration time: Of a bundle, creation time plus Time To Live (TTL).

header: Of a bundle, the concatenation of the primary block, all extension blocks that precede the payload block, and the block header of the payload block itself.

one-way-light-time margin, OWLT margin: The maximum delta by which the one-way light time between any pair of nodes can change during the time a bundle is in transit between them.

payload: Of a bundle, the content of the bundle's payload block.

range interval: A period of time during which the displacement between two nodes A and B is expected to vary by less than one light second from a stated anticipated distance.

route list: A list of one or more routes to a destination.

route: For a bundle whose current location is node X and whose destination is node D, a sequence of contacts such that (a) the sending node for the first contact is X, (b) the receiving node for the last contact is D, (c) the receiving node for contact *i* is the sending node for contact *i*+1, and (d) the time at which contact *i*+1 ends is no earlier than the time at which contact *i* begins.

routing table: A list, constructed locally by each node in the network, of route lists.

termination time: Of a route, the earliest end time among all contacts in the route.

volume: Of a contact, the product of its duration and its data transmission rate.

1.5 REFERENCES

The following publications contain provisions which, through reference in this text, constitute provisions of this document. At the time of publication, the editions indicated were valid. All publications are subject to revision, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the publications indicated below. The CCSDS Secretariat maintains a register of currently valid CCSDS publications.

- [1] *CCSDS Bundle Protocol Specification*. Issue 1. Recommendation for Space Data System Standards (Blue Book), CCSDS 734.2-B-1. Washington, D.C.: CCSDS, September 2015.
- [2] S. Burleigh. *Compressed Bundle Header Encoding (CBHE)*. RFC 6260. Reston, Virginia: ISOC, May 2011.
- [3] T. Berners-Lee, R. Fielding, and R. Fielding. *Uniform Resource Identifier (URI): Generic Syntax*. STD 66. Reston, Virginia: ISOC, January 2005.
- [4] E. W. Dijkstra. "A Note on Two Problems in Connexion with Graphs." *Numerische Mathematik* 1, no. 1 (1959): 269–271.
- [5] Jin Y. Yen. "Finding the K Shortest Loopless Paths in a Network." *Management Science* 17, no. 11 (July 1971): 712–716.

2 OVERVIEW

2.1 GENERAL

The Bundle Protocol specification (reference [1]) defines procedures for forwarding data *bundles* through a delay-tolerant network such as the Solar System Internet. However, it intentionally omits definition of procedures for determining which nodes of the network to forward a given bundle through in order to ensure that it reaches its destination, leaving those specifications to other documents that are targeted for various network environments requiring a variety of specialized algorithms. One such environment is the infrastructure used to conduct space flight missions. SABR is designed for that environment.

In the Internet, protocol operations can be largely driven by currently effective information that is discovered opportunistically and immediately, at the time it is needed, because the latency in communicating this information over the network is negligible: distances between communicating entities are small and connectivity is continuous. In a space flight mission environment, however, ad-hoc information discovery would in many cases take so much time that it could not be completed before the information lost currency and effectiveness. Instead, protocol operations must be largely driven by information that is preplaced at the network nodes and tagged with the dates and times at which it becomes effective.

More specifically, the forwarding of bundles through a DTN-based flight mission network differs in several ways from the forwarding of packets through an IP-based network. In an IP-based network:

- Connectivity, the ability of topologically adjacent (‘neighboring’) network nodes to exchange packets, is typically continuous throughout the network. Lapses in connectivity are anomalous and may be interpreted as changes in topology.
- Signal propagation delays are very small.
- Together, these characteristics ensure that the rate at which information on changes in connectivity may be propagated through the network far exceeds the rate at which those changes occur.

A flight mission network based on DTN is different:

- There is no expectation of continuous connectivity throughout the network. Lapses in connectivity may be routine, lengthy, and recurring; they should not be interpreted as changes in topology.
- Signal propagation delays may be large.
- Together, these characteristics imply that the rate at which information on changes in connectivity may be propagated through the network may be far lower than the rate at which those changes occur.

Because of these differences, the constraints within which forwarding routes are computed in a DTN-based network are different from those within which IP routes are computed, so route computation procedures must be different.

In particular, IP routing at each router can be based on a local understanding of current connectivity in the network that may be assumed to be generally accurate and generally stable over time. The route to a given destination host, once computed, may be stored in a routing table for future reference and will only need to be changed upon the arrival of new connectivity information, conveyed by routing protocol messages generated immediately in response to detected changes in connectivity, that invalidates that route.

DTN routing enjoys no such advantages. The potential delay in the arrival of information regarding connectivity changes makes all such information potentially obsolete; a Bundle Protocol Agent (BPA) that relied solely on this flow of information might never have a fully accurate understanding of current connectivity in the network.

Yet BPAs that must compute routes in a DTN-based network have no alternative but to rely on that understanding, imperfect as it may be. Each BPA must therefore augment its model of connectivity in the network by other means. Some elements of the model may simply be asserted by network management, i.e., as static routes. Some changes in proximate DTN network connectivity may be discovered in real time. Other connectivity changes may be predicted on a probabilistic basis.

Schedule-Aware Bundle Routing is designed for use in networks where changes in connectivity are planned and scheduled, rather than predicted, discovered, or contemporaneously asserted.

Scheduled changes in connectivity characterize a number of potential DTN application environments:

- Episodes of communication between robotic spacecraft in interplanetary space and ground tracking stations on Earth are typically scheduled weeks or months before they occur.
- The beginning and end of each communication opportunity between an orbiting spacecraft and a communication asset on a planetary surface, either Earth or another planet, can readily be computed from known orbital elements.
- Power-conserving motes of sensor webs may communicate on infrequent, fixed intervals established by network configuration.

Where changes in connectivity are scheduled, a global ‘contact plan’ of all such events may be distributed in advance to all BPAs, enabling each BPA to have a theoretically accurate understanding of connectivity in the network at any specified moment. The Schedule-Aware Bundle Routing procedures compute bundle forwarding decisions from this time-varying model of network connectivity, using a technique called Contact Graph Routing (CGR).

2.2 ARCHITECTURAL CONSIDERATIONS

As discussed in the BP specification (reference [1]), the source and destination of each bundle are BP endpoints, identified by BP Endpoint ID (EID) strings that are Uniform Record Identifiers (URIs) (reference [3]).

However, the actual agents of bundle origination, forwarding, and delivery are instances of Bundle Protocol procedures implementation (bundle protocol agents) that are installed at physical computational entities termed ‘nodes’.

For bundle forwarding purposes, a BP endpoint only exists so long as at least one node is ‘registered’ in that endpoint: only the operation of the BPA at a node can cause a bundle to be delivered to an endpoint, and a BPA can only deliver a bundle to an endpoint within which that BPA’s host node is registered.

An endpoint in which only a single node may be registered at any time is termed a ‘singleton’ endpoint. The forwarding of a bundle to a singleton endpoint is functionally equivalent to ‘unicast’ transmission in the Internet and is the most familiar and widely implemented mode of network communications.

No specifications yet exist that would govern the forwarding of a bundle to a non-singleton endpoint, e.g., ‘multicast’ transmission, so for the purposes of this document only bundle ‘unicast’ transmission is considered; SABR is not applicable to the forwarding of bundles to non-singleton endpoints. So for the purposes of this document, the existence of a node may be regarded as a precondition for the existence of an endpoint, and arrival of a bundle at some node’s BPA is a precondition for the delivery of that bundle to an endpoint.

Moreover, it is not unusual for a single node to be registered in multiple endpoints, each serving the needs of a different DTN application operating at that node. When this is the case, the arrival of a bundle at some single BPA is a precondition for the delivery of that bundle to any of a potentially large number of (singleton) endpoints.

For these reasons, the design of SABR is based on the concept of forwarding each bundle to the sole node that is registered in the bundle’s destination endpoint (rather than directly to the destination endpoint), leaving to the node’s BPA the task of final delivery.

Execution of this concept requires that nodes be recognized as first-class BP architectural elements, which must be uniquely identified in order to ensure accurate bundle delivery to the correct destination endpoints. SABR assumes that all nodes in the network are identified by unique ‘node numbers’ as discussed in the specification for Compressed Bundle Header Encoding (CBHE) (reference [2]). When SABR is used to forward a bundle to an endpoint identified by a CBHE-conformant EID, the destination node number can simply be extracted from the EID. SABR can also be used to forward bundles to an endpoint identified by a non-CBHE-conformant EID, but only if that EID can somehow be mapped to the appropriate destination node number; mechanisms for accomplishing this mapping are beyond the scope of this specification.

NOTE – The design of SABR precludes its use for routing a bundle to a ‘multicast’ endpoint: by definition, a multicast EID cannot be mapped to the number identifying a single node. SABR can only be used for forwarding bundles to ‘singleton’ endpoints.

2.3 DATA STRUCTURES

2.3.1 CONTACT PLANS

The basic strategy of SABR is to take advantage of the fact that, since communication operations are planned in detail, the communication routes between any pair of bundle protocol agents in a population of DTN nodes that have all been informed of one another’s plans can be inferred from those plans rather than discovered via dialogue (which is impractical over long-one-way-light-time space links). This information takes the form of *contact plans*.

A contact plan comprises two types of information items: *contacts* and *range intervals*.

A *contact* is here defined as an interval during which it is expected that data will be transmitted by DTN node A (the contact’s sending node) and most or all of the transmitted data will be received by node B (the contact’s receiving node). Implicitly, the sending node will utilize some ‘convergence-layer’ protocol underneath the Bundle Protocol to effect this transmission of data directly (that is, without any relay of data by any other DTN node) to the receiving node. Each contact is characterized by its start time, its end time, the identities of the sending and receiving nodes, and the mean rate at which data are expected to be transmitted by the sending node throughout the indicated time period.

A *terminated contact* is a contact for which the end time is not later than the current time.

The *duration* of a contact is given by subtracting the contact’s start time from its end time.

The *volume* of a contact is the product of its duration and its expected mean data transmission rate.

NOTE – A contact is specifically not an episode of activity on a link. Episodes of activity on different links, e.g., different radio transponders operating on the same spacecraft, may well overlap, but contacts by definition cannot; they are bounded time intervals and as such are innately ‘tiled’. For example, given the case in which transmission on link X from node A to node B, at data rate R_X , begins at time T_1 and ends at time T_2 ; and transmission on link Y from node A to node B, at data rate R_Y , begins at time T_3 and ends at time T_4 ; if $T_1 = T_3$ and $T_2 = T_4$, then there is a single contact from time T_1 to time T_2 at data rate $R_X + R_Y$; if $T_1 < T_3$ and $T_2 = T_4$, then there are two contiguous contacts: one from T_1 to T_3 at data rate R_X , and one from T_3 to T_2 at data rate $R_X + R_Y$; if $T_1 < T_3$ and $T_3 < T_2 < T_4$, then there are three contiguous contacts: one from T_1 to T_3 at data rate R_X , one from T_3 to T_2 at data rate $R_X + R_Y$, and one from T_2 to T_4 at data rate R_Y ; and so on.

A *range interval* is a period of time during which the displacement between two nodes A and B is expected to vary by less than one light second from a stated anticipated distance. (It is expected that this information will be readily computable from the known orbital elements of all nodes.) Each range interval is characterized by its start time, its end time, the identities of the two nodes to which it pertains, and the anticipated approximate distance between those nodes throughout the indicated time period, to the nearest light second.

Protocols for distributing contact plan information to bundle protocol agents are beyond the scope of this specification.

2.3.2 ROUTING LISTS

2.3.2.1 The *termination time* of a route is the earliest end time among all contacts in the route. A *terminated route* is a route whose termination time is not greater than the current time.

2.3.2.2 The list of all routes through the contact graph from node X to destination node D that (a) have been computed and (b) are not terminated is termed the *route list* for destination node D.

2.3.2.3 Any node that is the receiving node for at least one non-terminated contact whose sending node is X is termed a *neighbor* of node X. All, and only, the neighbors of node X are adjacent to X in the time-varying network topology described by the contact plan.

2.3.2.4 Each route to destination node D, from the local node X, whose entry node is G is referred to as a route to D ‘through’ G.

2.4 KEY CONCEPTS

2.4.1 EXPIRATION TIME

The *expiration time* of a bundle is computed as its creation time plus its Time To Live (TTL).

NOTE – Every bundle transmitted via DTN has a TTL, the length of time after which the bundle is subject to destruction if it has not yet been delivered to its destination. When computing the next-hop destination for a bundle that the local bundle agent is required to forward, there is no point in selecting a route that cannot get the bundle to its final destination prior to the bundle's expiration time.

2.4.2 OWLT MARGIN

The One-Way Light Time (OWLT) margin (*OWLT margin*) is defined as the maximum delta by which the OWLT between any pair of nodes can change during the time a bundle is in transit between them.

OWLT, that is, distance, is obviously a factor in delivering a bundle to a node prior to a given time. OWLT can actually change during the time a bundle is en route, but route computation becomes intractably complex if it cannot be assumed that an OWLT 'safety margin', a maximum delta by which OWLT between any pair of nodes can change during the time a bundle is in transit between them.

The OWLT delta is necessarily mission-specific, but in practice it may be simplest to assume a worst-case constant. For example, as of the date of publication of this document it might be posited that the maximum rate of change in distance between any two nodes in the network is 450,000 miles (720,000 km) per hour, which is 125 miles (200 km) per second. (This is the projected maximum speed of the Solar Probe Plus spacecraft, planned for launch in 2018.)

At this speed, the distance between any two nodes that are initially separated by a distance of N light seconds will increase by a maximum of 125 miles (200 km) per second of transit. This will result in data arrival no later than roughly $(N + Q)$ seconds after transmission, where the OWLT margin value Q is $(125 * N)$ divided by 186,000, rather than just N seconds after transmission, as would be the case if the two nodes were stationary relative to each other. When computing the expected time of arrival of a transmitted bundle, the most pessimistic case, $N + Q$, is used as the anticipated total in-transit time.

2.4.3 ESTIMATED VOLUME CONSUMPTION

The size of a bundle is the sum of the sizes of its payload and its header, but bundle size is not the only lien on the volume of a contact. The total Estimated Volume Consumption (EVC) for a bundle is the sum of the sizes of the bundle's payload and header and the estimated convergence-layer overhead. For a bundle whose header is of size M and whose

payload is of size N , the estimated convergence-layer overhead is defined as three percent of $(M + N)$, or 100 bytes, whichever is larger.

2.4.4 EXCLUDED NEIGHBORS

A neighboring node C that refuses ‘custody’ (as described in reference [1]) of a bundle destined for some remote node D is termed an *excluded neighbor* for (that is, with respect to computing routes to) D . So long as C remains an excluded neighbor for D , no bundles destined for D will be forwarded to C , except that occasionally (once per lapse of the Round-Trip Time [RTT] between the local node and C) a custodial bundle destined for D may be forwarded to C to ‘probe’ the link, i.e., to test whether or not the neighbor must remain excluded. (Bundles that are forwarded under such circumstances are termed *probe bundles*.) C ceases to be an excluded neighbor for D as soon as it accepts custody of a bundle destined for D .

2.4.5 CRITICAL BUNDLES

A *critical bundle* is one that has the critical quality-of-service flag set, notionally because it absolutely has got to reach its destination.

For an ordinary non-critical bundle, the SABR dynamic route computation algorithm uses the routing table to select a single neighboring node to forward the bundle through. It is possible, though, that as a result of some unforeseen delay the selected neighbor may prove to be a suboptimal forwarder: the bundle might arrive later than it would have if another neighbor had been selected, or it might not even arrive at all.

For critical bundles, the SABR dynamic route computation algorithm causes the bundle to be inserted into the outbound transmission queues for transmission to all neighboring nodes that can plausibly forward the bundle to its final destination. The bundle is therefore guaranteed to travel over the most successful route, as well as over all other plausible routes. It should be noted that this may result in multiple copies of a critical bundle arriving at the final destination.

3 ROUTE DETERMINATION PROCEDURES

3.1 GENERAL

The neighboring node(s) to which a bundle shall be forwarded from node X in order to arrive at node D shall be determined as follows.

3.2 CONTACT GRAPH ROUTING

3.2.1 CONTACT GRAPH

The contact graph for node D at node X is a conceptual directed acyclic graph that shall comprise:

- a) at the root vertex, a notional contact from node X (the local node) to itself;
- b) a terminal vertex, a notional contact from node D to itself;
- c) one additional vertex for each contact in the contact plan that signifies transmission either directly 'to' node D or indirectly to node D (i.e., to the 'from' node of some other contact that signifies transmission directly or indirectly to node D) **and** either directly 'from' node X or indirectly from node X (i.e., from the 'to' node of some other contact that signifies transmission directly or indirectly from node X);
- d) an edge between two vertices wherever one vertex corresponds to a contact signifying transmission 'to' some node (the origin of the edge) and the other vertex corresponds to a contact signifying transmission 'from' that same node (the termination of the edge).

NOTE – The structure of the contact graph may seem somewhat counterintuitive; it bears almost no relation to the topology of the network, a more familiar graph. The vertices of the graph correspond to **contacts**, not to nodes. Contacts are intervals of opportunity for data transmission from one node to another. The root vertex may be thought of as corresponding to bundle creation, the conveyance of bundle content 'from' the application 'to' the BPA, both at the source node. The terminal vertex may be thought of as corresponding to bundle delivery, the conveyance of bundle content 'from' the BPA 'to' the application, both at the destination node. The edges are episodes of data retention at a node. That is, the vertex for contact 1 (transmission of data from node A to node B) is connected to the vertex for contact 2 (transmission of data from node B to node C) by an edge indicating a period of data retention at node B while node B waits for contact 2 to start. (See figures 3-1, 3-2, 3-3, and 3-4 for an illustration.)




Figure 3-1: Network Topology Example

Contact	Sender	Recvr	From	Until	Rate
1	A	B	1000	1100	1000
2	B	A	1000	1100	1000
3	B	D	1100	1200	1000
4	D	B	1100	1200	1000
5	A	C	1100	1200	1000
6	C	A	1100	1200	1000
7	A	B	1300	1400	1000
8	B	A	1300	1400	1000
9	B	D	1400	1500	1000
10	D	B	1400	1500	1000
11	C	D	1500	1600	1000
12	D	D	1500	1600	1000

Figure 3-2: Contact Plan Example: Contacts

Sender	Recvr	From	Until	Range (light seconds)
A	B	1000	1100	1
A	C	1100	1200	30
B	D	1400	1500	120
C	D	1500	1600	90

Figure 3-3: Contact Plan Example: Range Intervals




Figure 3-4: Node A's Contact Graph for Node D, Given This Contact Plan

3.2.2 CONTACT PLAN CHECK

If no contacts in the contact plan identify transmission to node D, then the Contact Graph Routing (CGR) procedures shall not be used to select the neighboring node(s) to which this bundle shall be forwarded. The alternative route determination procedures defined in 3.3 shall be used instead.

3.2.3 ROUTE PRUNING

3.2.3.1 If the contact plan has been modified in any way since the route computation procedures were most recently performed, all route lists shall be discarded.

NOTE – Contact plan changes may invalidate any or all earlier route computations.

3.2.3.2 Every contact whose end time is in the past shall be deleted from the contact plan (and therefore, implicitly, from the contact graphs for all destination nodes).

3.2.4 ROUTE COMPUTATION

3.2.4.1 The following definitions shall constrain SABR route computation:

3.2.4.1.1 The *earliest transmission time* for a contact from the local node to one of its neighbors is defined as the start time of that contact or the current time, whichever is later. The earliest transmission time for any other contact is defined as the start time of that contact or the earliest arrival time (defined below) for the immediately preceding contact in the route, whichever is later. No contact whose end time is before its earliest transmission time (i.e., before the earliest arrival time for the preceding contact in the route under consideration) shall be included in a route.

3.2.4.1.2 The *earliest arrival time* for a contact is defined as the sum of the earliest transmission time for that contact plus the range in light seconds from the contact's sending node to its receiving node, plus the applicable OWLT margin.

3.2.4.2 The *best-case delivery time* characterizing a route is defined as the earliest arrival time for the contact that immediately precedes the terminal vertex contact in this route.

3.2.5 CGR PREPARATION

3.2.5.1 The conceptual list of nodes to which the bundle may be forwarded, here termed the *proximate nodes list*, shall initially be empty.

3.2.5.1.1 This list shall contain one entry for each node that is the entry node of a candidate route (selected as discussed below) that could result in arrival of the bundle at node D.

3.2.5.1.2 Each entry in this list shall indicate the following properties of the first (i.e., lowest-cost) route in the route list of the indicated node: the ID for the proximate node, the projected bundle arrival time, the number of included contacts, and the termination time.

3.2.5.2 The conceptual list of nodes to which the bundle must **not** be forwarded, here termed the *excluded nodes list*, shall be populated as follows:

- a) If the bundle is a non-critical bundle that was previously forwarded to a node that refused custody (and is now being re-forwarded due to that custody refusal), then *backward propagation* of the bundle (that is, transmission of the bundle back to the node from which it was directly received) is authorized; otherwise, backward propagation of the bundle is not authorized. If backward propagation of the bundle is **not** authorized, then the node from which the bundle was directly received shall be added to this list.
- b) Every excluded neighbor for node D, for which this bundle would **not** serve as a probe bundle if forwarded to that neighbor, shall be added to this list.

3.2.6 POPULATING THE PROXIMATE NODES LIST

3.2.6.1 Consideration of routes on which the bundle might be forwarded shall be subject to the following definitions.

3.2.6.2 The *earliest transmission opportunity* for a route shall be computed as follows:

- a) The *adjusted start time* for a contact is defined as the contact's start time or the current time, whichever is later.
- b) The *applicable backlog* for the route is the sum of the EVCs of all bundles currently queued for transmission to the route's entry node whose priority is greater than or equal to that of the bundle that is to be forwarded.
- c) An *applicable prior contact* for the route is any contact that has end time later than the current time and has the same sending and receiving nodes as the route's initial contact but an earlier start time.
- d) The *applicable duration* of an applicable prior contact is given by the contact's end time minus its adjusted start time.
- e) An *applicable prior contact volume* is defined as the product of data transmission rate and applicable duration for some applicable prior contact.
- f) The *applicable backlog relief* for the route is the sum of all of its applicable prior contact volumes.
- g) The *residual backlog* for the route is the applicable backlog minus the applicable backlog relief, or zero, whichever is greater.

NOTE – This is the projected backlog at the adjusted start time of the route's initial contact.

- h) The *backlog lien* on the route's initial contact is given by the residual backlog divided by the contact's data transmission rate.
- i) The *earliest transmission opportunity* for the route is given by the adjusted start time of the route's initial contact plus the backlog lien on that contact.

3.2.6.3 The *first byte transmission time* for the initial contact of a route shall be the bundle's earliest transmission opportunity on this route. The first byte transmission time for each subsequent contact on that route is defined as the start time of the contact or the last byte arrival time (as defined below) for the immediately preceding contact in that route, whichever is later.

3.2.6.4 The *last byte transmission time* for a contact shall be the contact's first byte transmission time plus the *applicable radiation latency*, which is given by the EVC of the bundle divided by the contact's data transmission rate, or the contact's end time, whichever is earlier.

3.2.6.5 The *first byte arrival time* for a contact shall be the first byte transmission time for that contact plus the range in light seconds from the contact's sending node to its receiving node, plus the applicable OWLT margin.

3.2.6.6 The *last byte arrival time* for a contact shall be the last byte transmission time for that contact plus the range in light seconds from the contact's sending node to its receiving node, plus the applicable OWLT margin.

3.2.6.7 The *projected bundle arrival time* for the route, then, shall be the computed last byte arrival time for the contact immediately preceding the terminal vertex contact.

3.2.6.8 Route selection considerations are further constrained by the following definitions pertaining to transmission volumes.

3.2.6.8.1 The *priority* of a bundle is defined as an indicator of the required order of precedence in the transmission of this bundle among other bundles; for the purposes of SABR, it is assumed that a bundle may be transmitted to a neighbor only after every bundle with higher priority that is queued for transmission to that same neighbor has been transmitted. Priority is distinct from class of service, but priority may be determined from class of service in various ways; selecting the manner in which the priority of a bundle is determined is an implementation matter. Every bundle shall be characterized by one of N levels of priority where $N > 0$, the levels of priority are monotonically increasing in indicated required order of precedence, one level of priority is designated *top priority* (first in order of precedence), and one level of priority is designated *bottom priority* (last in order of precedence).

NOTE – One possibly useful priority determination method is to utilize three levels of priority where bundles of service class 0 are assigned bottom priority and bundles of service class 2 are assigned top priority; this scale can help optimize contact utilization. Another method is to utilize just one level of priority, such that all bundles are assigned both top priority and bottom priority; this scale can help minimize computing resource requirements.

3.2.6.8.2 For a given route and a given level of priority P , the sum of the estimated volume consumptions of all bundles of priority P for which that route has been selected for forwarding (as discussed later), since instantiation of the route, is termed the volume that is currently *reserved* for transmission of bundles of priority P or higher on this route.

3.2.6.8.3 The volume that is currently reserved for transmission of bundles of priority P or higher on some route is additionally similarly reserved on every contact in that route.

3.2.6.8.4 A contact is considered *depleted* with regard to priority level P if the volume reserved for transmission of bundles of priority P or higher on this contact is not less than the total volume of the contact.

3.2.6.8.5 The *effective start time* of one of the contacts in a route is the contact's first byte transmission time.

3.2.6.8.6 The *effective stop time* of one of the contacts in a route is whichever is less, the stop time of that contact or the smallest value of stop time among all successor contacts in the route.

3.2.6.8.7 The *effective duration* of one of the contacts in a route is the contact's effective stop time minus its effective start time.

3.2.6.8.8 The *maximum transmission volume (MTV)*, for priority level P, of a contact is defined as the portion of the volume of this contact that is not currently reserved for transmission of bundles of priority P or higher. The initial value of MTV of a given contact, for all levels of priority, is the volume of the contact.

3.2.6.8.9 The *effective volume limit (EVL)*, for priority level P, of any one of the contacts in a route is the portion of the volume of this contact that was nominally available for transmission of bundles of priority P or higher at the time this route was computed, as constrained by the preceding and succeeding contacts. EVL is defined as whichever is less, (a) the contact's MTV for this level of priority or (b) the product of the contact's data rate and the contact's effective duration.

3.2.6.8.10 The *transmission potential (TP)*, for priority level P, of a route is the number of bytes that are currently available to be reserved for the transmission of bundles of priority P or higher on this route. The initial value of TP for a given route, for a given level of priority, is the smallest value of EVL for that level of priority among all contacts in the route.

3.2.6.8.11 A route is considered *depleted* with regard to priority level P when the route's value of TP for P is not greater than zero.

3.2.6.9 Subject to these considerations, the route(s) in node D's route list at node X that shall be selected for the forwarding of a given bundle that is destined for node D shall be determined in the following manner:

- a) Each route whose best-case delivery time is after the bundle's expiration time shall be ignored.
- b) Each route whose entry node is a member of the excluded nodes list shall be ignored.
- c) Each route that includes any contact indicating transmission to node X shall be ignored unless node D and node X are identical ('loopback' transmission).
- d) Each route for which the earliest transmission opportunity is after the end time of the initial contact shall be ignored.
- e) Each route for which projected bundle arrival time is after the bundle's expiration time shall be ignored.
- f) Each route that is depleted with regard to the bundle's level of priority shall be ignored.
- g) All other routes shall be deemed *candidate routes*.

- h) So long as the route list contains no candidate routes, either (a) the next best route from X to D through the contact graph for node D shall be computed as described below and added to the route list or (b) route selection shall be deemed to have concluded. Identification of the conditions under which the computing of additional routes must cease is an implementation matter.

If the bundle is flagged as a critical bundle, then so long as the route list contains no candidate routes through one or more of the neighbors of node X, either (a) the next best route from X to D through the contact graph for node D shall be computed as described below and added to the route list or (b) route selection shall be deemed to have concluded. Identification of the conditions under which the computing of additional routes must cease is an implementation matter.

NOTE – A candidate route can include a contact whose MTV, for the bundle's level of priority, is less than the bundle's estimated volume consumption. In general, inability to transmit a given bundle during any contact with any node other than the entry node cannot ever be accurately anticipated, because of the possibility of transmission backlog (of unknown and unknowable size) at any node. Anticipatory fragmentation, described later, can mitigate this risk somewhat, but in practice ad-hoc recovery mechanisms are used to address such routing failures.

3.2.6.10 The next best route from X to D through the contact graph for node D shall be computed by identifying the shortest path from X to D – that is, beginning at the root of the graph and ending at the terminal vertex – excluding all previously identified paths. For this purpose, the cost of edge N shall be the earliest arrival time of contact K.

NOTE – Yen's K Shortest Path algorithm (reference [5]) may be utilized as a means of identifying the shortest path from X to D excluding all previously identified paths. The details of this procedure are beyond the scope of this specification.

3.2.6.11 The entry node of each candidate route shall be added to the proximate nodes list.

3.2.7 PROXIMATE NODES LIST CHECK

If the proximate nodes list contains no nodes, the CGR procedures shall not be used to select the neighboring node(s) to which this bundle shall be forwarded. The alternative route determination procedures defined in 3.3 shall be used instead.

3.2.8 CGR FORWARDING

3.2.8.1 The best candidate route for a neighbor is defined as the route with the smallest value of projected arrival time among all candidate routes for which this neighbor is the entry node.

3.2.8.2 Whenever a bundle is enqueued for transmission to a neighbor, the best candidate route's TP, for that bundle's level of priority and every lower level of priority, shall be

reduced by the EVC of that bundle, as shall the corresponding MTVs of all contacts in that route.

3.2.8.3 If the bundle is flagged as a critical bundle, then a copy of this bundle shall be enqueued for transmission to every node in the proximate nodes list that is the entry node for at least one of the candidate routes.

3.2.8.4 Otherwise:

- a) For each neighbor, the projected bundle arrival time associated with that neighbor is the projected arrival time of the best candidate route for that neighbor.
- b) The most preferred neighbor in the proximate nodes list shall be selected as follows:
 - 1) If one of the nodes in this list is associated with a projected bundle arrival time that is earlier than that of all other nodes in the list, then it shall be the most preferred neighbor.
 - 2) Otherwise, if one of the nodes with the earliest projected bundle arrival time is associated with a smaller number of contacts than every other node with the same projected bundle arrival time, then it shall be the most preferred neighbor.
 - 3) Otherwise, if one of the nodes with the earliest projected bundle arrival time and smallest number of contacts is associated with a later termination time than every other node with the same projected bundle arrival time and number of contacts, then it shall be the most preferred neighbor.
 - 4) Otherwise, the node with the smallest node number among all nodes with the earliest projected bundle arrival time and smallest number of contacts and latest termination time shall arbitrarily be chosen as the most preferred neighbor.
- c) If Z is the bundle's estimated volume consumption, and S is the TP, for the bundle's level of priority, of the best candidate route for the most preferred neighbor, then:
 - 1) if S is greater than or equal to Z , or if the bundle cannot be fragmented, or if anticipatory fragmentation (an implementation option, described below) is not desired, the bundle shall simply be enqueued for transmission to the most preferred neighbor;
 - 2) otherwise the anticipatory fragmentation procedures shall be performed.

NOTE – An implementation might find it advantageous to note each successfully forwarded bundle in a history list, removing an item from the history list whenever the TTL of the corresponding bundle expires. This will enable a newly received bundle to be discarded immediately if it was previously forwarded, minimizing the incidence of routing loops. Great care would need to be used with such a mechanism, however. A bundle that reaches a downstream node from which it cannot be forwarded toward the destination might be forwarded back to an earlier forwarding point from which an alternate route might be taken; upon receiving such a bundle, if the TTL of this bundle is not expired, the earlier forwarder will find that bundle in the history list but even so must not discard it.

3.2.8.5 Anticipatory Fragmentation

NOTE – Transmission opportunity utilization might in some cases be improved by fragmenting a bundle at the time the CGR procedures have computed the bundle's route. If the best candidate route associated with the most preferred neighbor includes a contact whose MTV, for the bundle's level of priority, is less than the size of the bundle, then it can be assumed that the bundle would need to be forwarded in multiple episodes from that contact's sending node. Since contacts enabling those episodes might or might not be available, while contacts on less theoretically optimal routes from the current forwarding node might be undersubscribed, it may instead be desirable to fragment the bundle immediately and forward the fragments on different routes.

Anticipatory fragmentation shall be performed as follows:

- a) The bundle shall be fragmented into two fragmentary bundles, bundle A containing the first S octets of the original bundle's payload and bundle B containing the last (Z-S) octets of that payload.
- b) Bundle A shall be enqueued for transmission to the most preferred neighbor.
- c) Route determination shall be performed for bundle B as detailed in this specification, except that the best candidate route for the neighbor for which bundle A was enqueued shall **not** be deemed a candidate route for bundle B or for any fragment of bundle B.

3.3 ALTERNATIVE ROUTE DETERMINATION PROCEDURES

3.3.1 ROUTING TO NEIGHBOR

If it is known (e.g., by management) that bundles can be transmitted directly to node D, then the bundle shall be enqueued for transmission to that neighboring node and no other alternative route determination procedure shall be performed.

NOTE – Contact graph routing is attempted first, before checking for a direct transmission option, because the direct next contact with node D might be far in the future; it is possible that an indirect route through some other neighbor might enable the bundle to be delivered earlier.

3.3.2 STATIC ROUTING

3.3.2.1 If at least one static route has been defined, by management, that indicates that all bundles destined for node D may be directed to some gateway node, then the bundle shall be forwarded to the gateway node associated with the most narrowly defined of all such static routes and no other alternative route determination procedure shall be performed.

NOTE – The details of static route declaration syntax are an implementation matter. As an example of what is intended, in a case in which each static route declaration associates a gateway node's ID (expressed as a BP endpoint ID) with a range of destination node numbers, and the range of nodes numbered 10 through 30 are associated with the gateway node identified by the EID 'ipn:901.0', while the range of nodes numbered 16 through 19 are associated with the gateway node identified by the EID 'ipn:816.0', a bundle destined for node 27 would be forwarded to ipn:901.0, but a bundle destined for node 17 would be forwarded to ipn:816.0. That is because, while 17 is between 10 and 30, it is also between 16 and 19, the more narrowly defined static route. In effect, the route to ipn:816.0 overrides the more general static route to ipn:901.0, but only for bundles whose destination nodes are in the range 16 to 19.

3.3.2.2 For this purpose, the route to the gateway node shall be determined by exercise of the route determination procedures defined in this Recommended Standard.

3.3.3 ROUTING FAILURE

If none of the route determination procedures defined in this Recommended Standard result in the enqueueing of the bundle for transmission on some convergence-layer transmission channel, then route determination shall be deemed failed.

NOTE – Procedures for responding to failure in route determination shall be an implementation matter.

3.4 EXCEPTION HANDLING

3.4.1 OVERBOOKING MANAGEMENT

Enqueueing a bundle for transmission to some node can result in a postponement of transmission of one or more lower-priority bundles: the premises upon which a lower-priority bundle's best candidate route was selected might no longer hold. When such an

anomaly can be detected, new routes for the affected bundles should be determined by exercise of the route determination procedures defined in this Recommended Standard.

NOTE – Procedures for detecting and remediating such anomalies are an implementation matter.

3.4.2 CONTACT FAILURE

Route determination as defined in this Recommended Standard assumes that the information in the contact plan is reliable. However, actual periods of contact may not always conform precisely to the plan: when a planned contact begins later than planned, ends earlier than planned, or does not occur at all, or when the actual volume of a contact is less than the planned volume (e.g., because some contact time was consumed by unanticipated bundle retransmission), one or more bundles enqueued for transmission on the corresponding convergence-layer transmission may not be transmitted. In this event, new routes for the affected bundles should be determined by exercise of the route determination procedures defined in this Recommended Standard.

NOTES

- 1 Procedures for detecting and remediating such anomalies are an implementation matter.
- 2 The efficiency of SABR varies directly with the degree of agreement among the contact plans exposed to the nodes. Deficiencies in this agreement will introduce errors, resulting in suboptimal transmission of bundles to neighboring nodes.
- 3 In the general case it is impossible to determine whether forwarding along a given route will in fact result in bundle delivery at the destination node: storage resource availability and transmission opportunity commitments at ‘downstream’ nodes cannot reliably be predicted in a delay-afflicted network. Contingent forwarding and reforwarding procedures must therefore be in place at all nodes to minimize the likelihood of bundle expiration prior to delivery. Such procedures are beyond the scope of this specification. However, the developer is advised that extremely large bundles are in general more vulnerable to forwarding failures due to resource exhaustion than small bundles are. Proactive fragmentation into relatively small fragmentary bundles can improve bundles’ forwarding prospects at all points along the end-to-end path, by enabling brief contacts to be allocated to these smaller bundles.

3.4.3 CUSTODY REFUSAL

3.4.3.1 When a BP custody refusal signal is received, citing some previously forwarded bundle, a new route for the affected bundle may be determined by exercise of the route determination procedures defined in this Recommended Standard. In this event, backward propagation of the bundle shall be authorized as described in 3.2.5, above. Moreover, if the

node that sent the BP custody refusal signal is the neighboring node to which the bundle was previously forwarded, then that node shall be deemed an excluded neighbor as defined in 2.4.4, above.

3.4.3.2 Alternatively, a node's failure to forward a timely custody acceptance signal citing some previously forwarded bundle may similarly initiate determination of a new route for the affected bundle by exercise of the route determination procedures defined in this Recommended Standard.

NOTE – Procedures for detecting these conditions and initiating remedies are an implementation matter.

ANNEX A

PROTOCOL IMPLEMENTATION CONFORMANCE STATEMENT PROFORMA

(NORMATIVE)

A1 INTRODUCTION

A1.1 GENERAL

This annex provides the Protocol Implementation Conformance Statement (PICS) Requirements List (RL) for implementations of *Schedule-Aware Bundle Routing (SBAR)*, CCSDS 734.3-R-1. The PICS for an implementation is generated by completing the RL in accordance with the instructions below. An implementation shall satisfy the mandatory conformance requirements of the base standards referenced in the RL.

An implementation's completed RL is called the PICS. The PICS states which capabilities and options of the protocol have been implemented. The following can use the PICS:

- the protocol implementer, as a checklist to reduce the risk of failure to conform to the standard through oversight;
- the supplier and acquirer or potential acquirer of the implementation, as a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the standard PICS proforma;
- the user or potential user of the implementation, as a basis for initially checking the possibility of interoperating with another implementation.
- a protocol tester, as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

A1.2 NOTATION

The following are used in the RL to indicate the status of features:

A1.2.1 Status Symbols

M mandatory.

O optional.

A1.2.2 Support Column Symbols

The support of every item as claimed by the implementer is stated by entering the appropriate answer (Y, N, or N/A) in the support column:

- Y Yes, supported by the implementation.
- N No, not supported by the implementation.
- N/A Not applicable.

A1.3 INSTRUCTIONS FOR COMPLETING THE RL

An implementer shows the extent of compliance to the protocol by completing the RL; that is, the RL shows compliance to all mandatory requirements and lists all options that are not supported. The resulting completed RL is called a PICS. In the Support column, each response shall either be a response code selected from the indicated set of responses or else comprise one or more parameter values, as appropriate. If a conditional requirement is inapplicable, N/A should be used. If a mandatory requirement is not satisfied, exception information must be supplied by entering a reference Xi, where i is a unique identifier to an accompanying rationale for the noncompliance.

A1.4 REFERENCED BASE STANDARDS

The base standards referenced in the RL are:

- Schedule-Aware Bundle Routing (this document)

A1.5 IDENTIFICATION OF PICS

Ref	Question
1	Date of Statement (DD/MM/YYYY)
2	PICS serial number
3	System Conformance statement cross-reference

A2 IDENTIFICATION OF THE IMPLEMENTATION

Implementation name	
Implementation version	
Special Configuration	
Supplier	
Contact Point for Queries	
Implementation name(s) and Versions	
Other Information Necessary for full identification, e.g., Project developed for, name(s) and version(s) for machines and/or operating systems, system name(s).	

A3 IDENTIFICATION OF THE PROTOCOL

Protocol Name	SCHEDULE_AWARE BUNDLE ROUTING
Protocol Version	
Addenda Implemented	
Amendments Implemented	
Have any exceptions been required? (Note: A YES answer means that the implementation does not conform to the protocol. Non-supported mandatory capabilities are to be identified in the PICS, with an explanation of why the implementation is non-conforming.	Yes _____ No _____
Date of Statement	

A4 ICS PROFORMA TABLES

A4.1 CONTACT GRAPH ROUTING

A4.1.1 Contact Graph Routing Procedures

Item	Protocol Feature	Reference	Status	Support
SABR-CGR-01	Contact plan check	3.2.2	M	
SABR-CGR-02	Route pruning	3.2.3	M	
SABR-CGR-03	Route computation	3.2.4	M	
SABR-CGR-04	CGR preparation	3.2.5	M	
SABR-CGR-05	Populating the proximate nodes list	3.2.6	M	
SABR-CGR-06	Proximate nodes list check	3.2.7	M	
SABR-CGR-06	CGR forwarding	3.2.8 except 3.2.8.5	M	
SABR-CGR-07	Anticipatory fragmentation	3.2.8.5	O	

A4.1.2 Management Information Base Configuration Parameters

Item	Protocol Feature	Reference	Status	Parameter Value
LMIB-CGR-01	Contact plan	2.3.1	M	
LMIB-CGR-02	One-way light time margin	2.4.2	M	
LMIB-CGR-03	Excluded neighbors	2.4.4	M	

A4.2 ALTERNATIVE ROUTE DETERMINATION

A4.2.1 Alternative Route Determination Procedures

Item	Protocol Feature	Reference	Status	Support
SABR-ARD-01	Routing to neighbor	3.3.1	M	
SABR-ARD-02	Static routing	3.3.2	M	
SABR-ARD-03	Custody refusal	3.4.3, 2.4.4	M	

A4.2.2 Management Information Base Configuration Parameters

Item	Protocol Feature	Reference	Status	Parameter Value
LMIB-ARD-01	Neighbors	3.3.1	M	
LMIB-ARD-02	Static routes	3.3.2	M	

ANNEX B

SECURITY, SANA, AND PATENT CONSIDERATIONS

(INFORMATIVE)

B1 SECURITY CONSIDERATIONS

This document specifies how to determine the routes over which bundles should be forwarded, given information such as a set of static routes and a contact plan. How the contact plan(s) should be distributed to the various nodes in the network, and whether all nodes need exactly the same information or not is beyond the scope of this document. It should be noted, however, that the information used to generate / update the contact plan(s) should be secured. There are a number of instances of routing attacks in the terrestrial Internet (intentional and unintentional) to motivate this. One possible means of securing contact plan distribution is by using the CCSDS Streamlined Bundle Security Protocol (reference [C1]).

B2 SANA CONSIDERATIONS

This recommendation does not require any new registries to be instantiated, and does not use any registry information from SANA.

B3 PATENT CONSIDERATIONS

At the time of publication, CCSDS was not aware of any patents pertaining to the technology described in this document.

ANNEX C

INFORMATIVE REFERENCES

(INFORMATIVE)

- [C1] *CCSDS Streamlined Bundle Security Protocol Specification*. Draft Recommendation for Space Data System Standards. Forthcoming.

ANNEX D

ABBREVIATIONS

(INFORMATIVE)

<u>Term</u>	<u>Meaning</u>
CGR	Contact Graph Routing
DTN	Delay-Tolerant Networking
EID	endpoint ID
EVC	estimated volume consumption
OSI	Open Systems Interconnection
OWLT	one-way light time
RTT	round-trip time
SABR	Schedule-Aware Bundle Routing
TTL	time-to-live
URI	Uniform Record Identifier