On Endpoint ID schemas

Rick Taylor

IETF 113 March 2022

What does RFC 9171 tell us?

Bundle Protocol version 7 defines Endpoint IDs as a pair: `{schema,content}`

Where:

- `schema` is the numeric identifier of the schema in use.
- `content` is a CBOR encoded schema-specific value.

- Applications executing somewhere in the DTN are the receivers of bundles, and receivers can have multiple Endpoint IDs.
- Endpoint IDs may refer to zero or more receivers, i.e. they can be used for unicast (1), multicast (> 1) or as a 'black hole' (0)?
  - I shall define this value as the Endpoint ID's *multiplicity*.
- Node IDs are defined as a specialisation of Endpoint IDs that are explicitly unicast (*multiplicity* = 1), and identify individual Bundle Processing Agent (BPA) services.
- BPv7 defines 2 schemas, the *ipn* and *dtn* schema.
  - Neither defines how to determine the *multiplicity* an Endpoint ID.
  - The *ipn* scheme concatenates the name of the consuming service and an identifier of the BPA processing bundles on behalf of the service, as the pair `{id,service}`, introducing the concept of a *Service Number*.

Questions

I want to address these questions:

1. If the *dtn* schema is for universal use, how do I write a protocol handler for a tool such as cURL?
2. Given the *ipn* schema is restricted in its capability, what should this Working Group recommend for its usage?
3. If the IETF, or another SDO, wishes to standardise an additional Endpoint ID schema, how should this be done?

Q1: External use of the *dtn* schema

As the schema is registered in the IANA “Uniform Resource Identifier (URI) Schemes” registry, it implies that universal tools that support sending messages via different protocols, for example cURL, must be able to implement support for the *dtn* protocol.

Such support, at its most basic level, involves resolving a command such as:

```
curl --upload-file "file1" dtn://name/demux
```

Into a series of steps, including:

1. Resolve the URI `dtn://name/demux` into a transport protocol plus destination for that protocol.
2. Transmit the content of the source file using the resolved transport protocol.

But RFC 9171 does not provide sufficient information to perform step 1.

Proposal 1

When a *dtn* schema URI is used *externally*, by a so-called *User Agent*, to transmit to a BPv7 Endpoint ID, i.e. for ingress, the following rules are applied:

1. The URI is parsed, according to RFC 9171, into its `name` and `demux` components.
2. The `name` component is resolved using DNS to an IP (v4 or v6) address.
3. A TCPCLv4 (RFC 9174) session is established to the resolved address, and the bundle transmitted.

Example

The command:

```
curl --upload-file "file1" dtn://dtn.ori.co/some/service
```

Sends a BPv7 bundle, with the content of `file1` as payload, with destination EID `{1,//dtn.ori.co/some/service}`, via a TCPCLv4 session with the agent at 91.240.47.127

Obviously extra options are required for real-world use, including TLS certificate information and Primary Block fields.

Clarifications

This proposal does not prevent the following:

- BPA agents between the ingress and the destination:
  - If the user agent does not have access to a DNS resolver, then the bundle may be forwarded to some other DTN node(s) that are believed to have access to a DNS resolver.
- BPA agents after the node resolved via DNS:
  - The `demux` part of the Endpoint ID could resolve to the name of a DTN service 'deeper' in the network, e.g.: `dtn://dtn.ori.co/ipn:143.7`, with the `name` part simply identifying a gateway node.

Q2: Recommended use of the *ipn* schema

RFC 6020 does a good job of defining the *ipn* schema:

- `ipn:0` is the null EID, identical to `dtn:none`
- `ipn:`*`N`*`.0` is the bundle processing agent service itself.

ipn schema EIDs are concise and simple to process and manage, however using numeric identifiers for global addressability across all possible DTNs is unfeasible using this schema.

And it leaves the multiplicity of an *ipn* EID undefined.

Proposal 2

I believe the Working Group should recommend that the *ipn* schema should be used to identify endpoints within a single *Autonomous Naming Region*, as described by Scott Burleigh at IETF 112: https://datatracker.ietf.org/doc/slides-112-dtn-big-dtn/

By making the assignment of *ipn* EIDs, and the maintenance of the routing rules within an *ipn* region a local matter, deployment specific implementation differences can be enabled.

However, as Scott points out, some mechanism for ingress and egress to/from an autonomous naming region using the *ipn* schema is required.

Do we need to update RFC 6260?

Q3: Standardising additional EID schemas

The process for standardising additional DTN compliant Endpoint ID schemas is defined in RFC 9171.

To define another standard EID schema, one MUST:

- Document the schema, including addressing the multiplicity aspect.
- Make a formal request to IANA for a new schema value in the “Bundle Protocol URI Scheme Types” registry.