

DP Enhancements to the EECS network, an extension of SolarNet, 2025

The management has discovered that there are three additional points to make with respect to the specification. The first two are optional additions that you may utilize; the third is a constraint (not optional).

1. **Optional modifications to primary block information** in a bundle:
 - a. A forwarding node can optionally extend the lifetime of a bundle.
 - b. A forwarding node can optionally downgrade the priority of a bundle.
 - c. Note that these actions should **NOT** be done arbitrarily. The sending application has an expectation of delivery time and priority, so only under somewhat exceptional circumstances should you consider these options. If you choose to use them, be sure to explain and justify their use.
2. **Added optional queue:** The management has realized that there may be need for a fourth queue, in addition to the three priority queues. This fourth queue is a *control-plane queue*. It is intended for management of the network and its resources, in contrast with the original queues which are *data-plane queues*. There are several constraints and issues with respect to the control-plane queues.
 - a. Bundles are only sent directly between nodes. They may never have a different destination address than the node to which they are being sent directly.
 - b. They are small and relatively infrequent. This queue is absolutely not to be used to move bundles carrying the data (telemetry, software, email, etc.) between nodes. This queue is only for management of the communication between neighboring nodes.
 - c. The management has not figured out exactly how the processing of the control-plane queues will be handled in relation to the priority queues you are already working with. They are looking for recommendations from you.

Note that you are not required to utilize this queue at all. If you choose not to, you should still justify that decision.

3. **Storage constraint:** The storage on the LEOComs and GEOComs is provided by two slightly different devices, each providing half the storage. One of the two fails intermittently by overheating. It takes 10 minutes to recover and completely resets itself (losing anything on it) as it recovers. As a result, at any given 10% of the nodes have their storage capacity reduced by 50% to .25TB.