ESA ITT 8809

Delay Tolerant Networks for flexible communication with Earth Observation satellites

ESR - Executive Summary Report

GMV-INS-DTN_EO-REP-1007
13 December 2018

Signatures

Jeremy Pierce Mayer
C=DE, O=GMV INSYEN AG, CN=Jeremy Pierce Mayer, E=jpmayer@gmv-insyen.com

12-12-2018 14:45:03

Prepared: ______________________ Date: ______________________

Jeremy Pierce Mayer

Project System Engineer

Digitally signed by Michael Staub
DN: cn=Michael Staub, o=GMV Insyen AG, ou=AEROSPACE, email=mstaub@gmv-insyen.com, c=DE
Date: 2018.12.12 14:12:57 +01'00'

Approved: ______________________ Date: ______________________

Michael Staub

Project Manager

Digitally signed by Sebastian Martin
Date: 2019.03.19 14:13:06 +01'00'

Accepted: ______________________ Date: ______________________

Sebastian Martin

ESA Point of Contact

Distribution List

Name	Organization	Copies
Sebastian Martin	ESA	1
Document Management System	GMV INSYEN	1

Contents

Figures

Tables

1. INTRODUCTION

This document summarizes the findings of the DTN-EO project, including a brief overview of results and challenges. It is not intended to be an exhaustive list of the results; an interested reader should read the various project deliverables outlined in section 1.2.

1.1. PURPOSE AND SCOPE

This document may be used as a reference for mission developers of Earth Observation missions, in order to determine if the work performed by the DTN-EO consortium is applicable to their mission.

1.2. APPLICABLE DOCUMENTS

[1] DTN-EO Consortium, "D2 - Scenario Report - GMV-INS-DTN_EO-REP-1003," GMV INSYEN AG, 2018.

[2] DTN-EO Consortium, "D3 - Scenario Analysis - GMV-INS-DTN_EO-REP-1004," GMV INSYEN AG, 2018.

[3] DTN-EO Consortium, "D4 - Roll Out Strategy - GMV-INS-DTN_EO-REP-1005," GMV INSYEN AG, 2018.

1.3. LIST OF ACRONYMS

AOS	Acquisition Of Signal
BOT	Beginning Of Transmission
BP	Bundle Protocol
BPSec	Bundle Protocol Security
CCSDS	Consultative Committee for Space Data Systems
CFDP	CCSDS File Distribution Protocol
CGR	Contact Graph Routing
CLA	Convergence Layer Adapter
DupACK	Duplicated ACKnowledgement
DTN	Disruption/Delay Tolerant Network
ECOS	Extended Class Of Service
EDRS	European Data Relay Satellite
EID	Endpoint IDentifier
EO	Earth Observation
EOT	End Of Transmission
ESTRACK	European Space Tracking (network)
ETO	Earliest Transmission Opportunity
FSO	Free-Space Optical (communication)
GB	Giga Byte
GEO	Geostationary Earth Orbit
GS	Ground Station
HOSC	Huntsville Operational Support Center
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ION	Interplanetary Overlay Network
IP	Internet Protocol
ISL	Inter Satellite Link
LADEE	Lunar Atmosphere and Dust Environment Explorer
LEO	Low Earth Orbit
LLCP	Lunar Laser Communications Demonstrator
LOS	Loss Of Signal
LOP-G	Lunar Orbiting Platform – Gateway
LTP	Licklider Transmission Protocol
MB	MegaByte
MCC	Mission Control Centre
METERON	Multi-Purpose End-To-End Robotic Operation Network
NASA	National Aeronautics and Space Administration
OCP	Optical Communication Payload
OWLT	One Way Light Time
PID	Process IDentifier
PDU	Protocol Data Unit
RF	Radio Frequency
RT	Real-Time
SAR	Synthetic Aperture Radar

SAT	Satellite
TC	TeleCommand
TCP	Transmission Control Protocol
TM	TeleMetry
UDP	User Datagram Protocol

2. STUDY OVERVIEW

2.1. STUDY OUTLINE & GOALS

The DTN-EO study was initiated in order to research the performance of DTN in future missions, particularly Earth Observation and constellation-type missions. In the original proposal, ESA stipulated 3 scenarios which were to be tested: Earth Observation, Constellation, and Exploration missions. These scenarios built from each other, starting with an Earth-Observation centric mission such as the Sentinel missions, before moving to a large constellation. Ultimately, the DTN network from the first two scenarios was expanded to represent a deep-space exploration mission. These three scenarios provided a cross-section of future mission goals and requirements, and were analysed to reveal best practices for missions using DTN.

2.2. DTN BACKGROUND

In order to determine the relevance of DTN within future missions, it is imperative that mission designers have a high-level understanding of the Bundle Protocol, in order to accurately define benefits and potential pitfalls that may arise during the design, integration, and operational phases of the mission. The remainder of this section serves to provide an overview of some important components of a delay/disruption network implemented via the Bundle Protocol.

DTN has been utilized in a variety of experiments throughout Europe, and the United States. For example, the ESA METERON missions have showcased an advanced and robust DTN network, intended for the evaluation of deep-space missions. In the satellite domain, the first satellite within the United Kingdom Disaster Management Constellation demonstrated the bundle protocol on a GEO satellite. The NASA Lunar Atmosphere and Dust Environment Explorer (LADEE) mission carried the Lunar Laser Communications Demonstrator (LLCP), which demonstrated the feasibility of deep-space communication via an optical link. Within the LLCP mission, DTN was evaluated, and was found to provide reliable communication.

2.2.1. CONVERGENCE LAYERS

The Bundle Protocol is envisioned as an overlay network, riding on top of existing network protocols and topology. Therefore, a key component is the utilization of Convergence Layers, which provide a mechanism to encapsulate BP data over various links, both IP & non-IP based (such as SpaceWire). The most important of these of these CLA's is the Licklider Transmission Protocol (LTP), which provides a stateful transport layer protocol designed for long-distance and lossy links, such as space-to-ground links. It was designed based upon lessons learned during the development of CFDP, and is the CCSDS suggested CLA for space-ground connectivity.

LTP functions by segmenting data into into Protocol Data Units (PDUs) of a pre-defined size. These PDUs are appended with a header, including a serial number before being transmitted. LTP contains two transmission pathways: red and green. The red pathway provides reliable transmission by allowing the receiver to respond with specially-formatted report segments containing information about the segments which were successfully received. This allows the transmitter to limit the scope of retransmission, and removes the overhead required by the DupACK -based retransmission as is used in TCP. Segments are held by the transmitter until a report arrives which indicates their successful reception. The green pathway is unreliable, and simply transmits segments to the receiver before deleting them from transmitter memory. The selection between the red and green pathways can be made by a user or application developer.

2.2.2. CUSTODY TRANSFER

In order to provide additional robustness towards packet loss, the bundle protocol supports a delay-tolerant mechanism for hop-by-hop and end-to-end verification of transmission. In a typical DTN network, bundles are transferred from node to node. Provided that a given bundle is utilizing custody transfer, each intermediate node becomes the "bundle custodian" for that bundle, and is responsible for sending a custody acceptance report to the previous custodian. If (for reasons such as storage constraints, etc.) a node is unable to accept a bundle, it must notify the current custodian. This mechanism is designed to function over any convergence layer adapter, as long as the prior custodian node can be reached via some path.

Custody transfer also allows for transmission of a custody signal to the original sender, allowing for that sender to trigger a retransmit if the bundle gets lost in transit.

2.2.3. SECURITY

End-to-end security is paramount for satellite operations and must be integrated throughout the DTN system. The requirements for a security system within a wide-ranging DTN network are fundamentally distributed into several classes of requirements:

- Payload data security – For proprietary payload data and/or commanding
- Operational Authorization – Methods by which the payload user may be authorized to transmit to their payload.
- Network Traversal – International partnerships may result in European satellite data traversing non-trusted networks.

In order to accomplish these goals, wide-ranging usage of the Bundle Protocol Security (BPSec) must be evaluated.

3. SCENARIO RESULTS

3.1. SCENARIO A: EARTH OBSERVATION MISSIONS

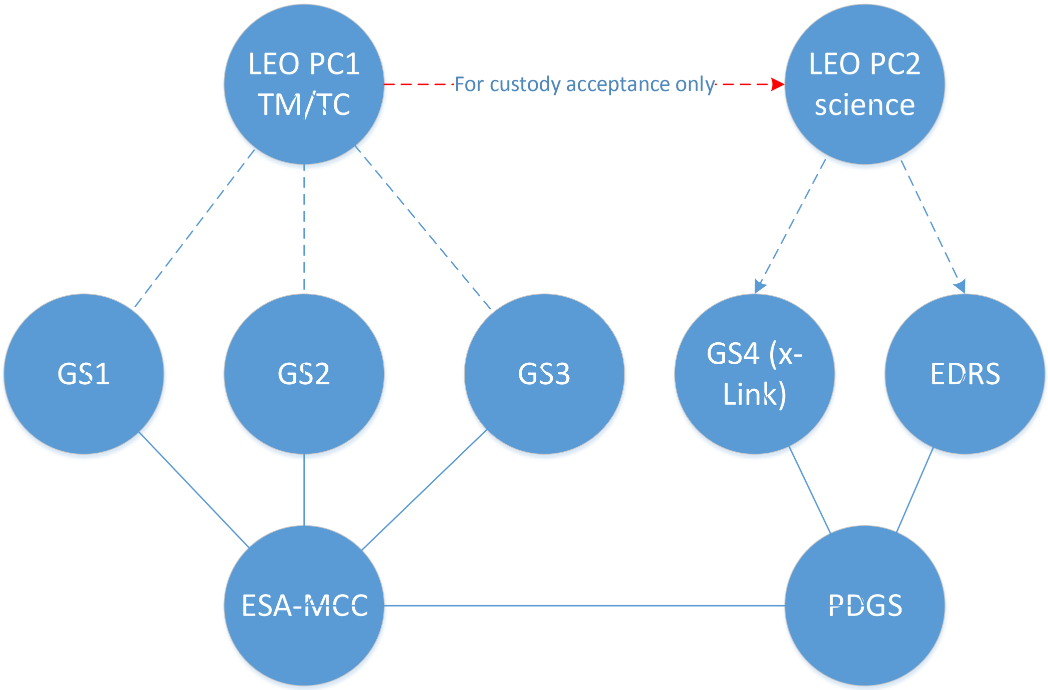

Figure 3.1 Scenario ABis logical layout (from [1])

3.1.1.RESULT SUMMARY

Scenario A (bis) provided a high-fidelity emulation of the Sentinel-1 Earth Observation mission, including all connectivity options. The entire duration of the space-to-ground contacts were effectively utilized in order to transmit SAR data to earth. Meanwhile, the store-and-forward behavior of the bundle protocol were seen in both the X-Band and EDRS ground segments, as they forwarded data to the PDGS at the maximum possible rate. At the Mid-Term Review, it was decided to further enhance this scenario to provide a more realistic emulation of an Earth Observation mission; the resulting scenario was deemed Scenario A Bis.

The most significant finding within scenario A was that DTN provides a noticeable increase in total (per-day) throughput, once the safety margins required by the Sentinel satellite & program are taken into account, as outlined in [1]. Although no loss was added in the scenario simulations, the design of Scenario A was intended to run the underlying network at (or near) capacity. The priority-based forwarding mechanisms available within the Bundle Protocol functioned as expected. High-priority traffic was forwarded in advance of lower priority data. This mechanism facilitates the transmission of high-priority TM/TC via DTN networks. This also theoretically validates the critical priority level within the bundle protocol. A bundle using this priority level will be transmitted via all available paths, therefore increasing the likelihood of reception at the node to its theoretical maximum; if there is any possible method to communicate with the spacecraft, it will be attempted. While this creates the potential for a denial of service attack, this functionality has been deemed essential for spacecraft applications.

Unlike the other scenarios, a full throughput and comparative analysis was performed, in order to verify the value of DTN in Earth Observation missions. Due to the nature of DTN-based networks, particularly concerning custody transfer and LTP retransmission due to segment loss, timeouts, etc., determining throughput is challenging. Therefore, the authors opted to analyze the delivery time at various nodes, as a metric to determine the total throughput. This approach can be extended in order to perform a high-level analysis of the total throughput, which is granular enough to see the total arrival time at various endpoints.

Per the scenario report [1], telemetry flows are constantly generated, although they are only transmitted during S-Band contacts. Table 1 and Table 2 show the science data throughput for scenario A. For science data, the throughput bottleneck existed between the ground segments and the PDGS, as outlined in the scenario report. Analysis revealed that DTN efficiently utilized the available bandwidth. The total through traffic from the on-board SAR source to the ground segment remains largely consistent, and uses the majority of the available connection throughput.

Table 1 Min/Max/Average Throughputs (bits/sec)- Telemetry

	Min	Max	Average
GS-1	163,840	2,457,600	377,924.3
GS-2	163,840	163,840	163,840
GS-3	163,840	163,840	163,840
ESA MCC	163,840	3,932,160	574,985.7

Table 2 Scenario A - Min/Max/Average Throughputs (MBit/second)

	Min	Max	Average
ipn:204.0	400.00	400.00	400.00
ipn:205.0	400.00	400.00	400.00
ipn:232.0	400.00	800.00	404.94

Another method which was used in order to analyze the performance of a DTN network is a comparison of available (per-contact/per-day) throughput. Performing such a calculation must take the safety margins of each link type into account, relative to the total number of bundles which were sent during the emulation. Table 3 provides a summary of this analysis. In this table, the “Theoretical Max” and “Safe Max” are calculated based upon the Sentinel contact schedule, while “DTN Max” is derived from the actual data transferred.

Table 3 Scenario A - Daily Link Capacities

Downlink	Theoretical Max (MB)	Safe Max. (MB)	DTN Max. (MB)	Delta data volume (between safe and DTN)
EDRS	552,960	475,520	497,512.5	21,992.5
X-Band	868,352	817,152	848,559.375	31,407.375
Total	1,421,312	1,292,672	1,346,072	53,400

As can be seen from Table 3, DTN provides an increase of approximately 53,400 MB/day of data transfer. This equals approximately 94% of the maximum utilization of the two high-rate links. The remaining 6% of link capacity can easily be explained by the 1 second granularity of DTN/ION timing, as well as the protocol overhead for individual bundles. First, 1 the 1 second granularity prevents sub-second AOS transitions. Once the AOS transition opens, it may take up to a second for that route to be considered by ION. Secondly, in order to provide the highly-configurable prioritization utilized within this scenario, the Extended Class of Service block was added to all bundles. Therefore, the overhead for bundles & links could be considered to be the sum of UDP, LTP, BP, and ECOS. Since some of these headers are added at different points (bundle starts, fragments, etc) as others, it is difficult to calculate a per-bundle overhead. However, the previous analysis should showcase that, although there is additional overhead, DTN achieves a high degree of link utilization and efficient end-to-end data transmission.

3.1.2. ANALYSIS

The performance analysis conducted for this scenario showcased the great potential that DTN implementations can offer with respect to traditional solutions based on IP or other CCSDS-specific standard solutions but not providing any store-forward capability. The major points of this advantage are shown in [2].

The advantages "theoretically" (i.e. based on the protocol specification) provided by the DTN architecture have been validated with the current status of the ION implementation, while showing that some improvements must be introduced in order to solve some bugs or ineffective behaviour of the DTN implementation. The close cooperation with NASA (either through personal contacts or during CCSDS meetings) helped solve some of these issues and new solutions are expected to be incorporated in future releases of ION, in order for the new implementation to respond to agency or mission requirements. For instance the management of green/red parts in LTP blocks is expected to be changed in the near future, making the corresponding implementation lighter and the overall mission configuration more efficient.

The separation between CCSDS and IETF standardisation bodies with respect to DTN development path does not seem to be of critical importance. The future versions of ION will still be based on the standardisation documents developed within CCSDS and the compliancy with IETF new documents (e.g., RFC 5050-bis) will be observed only when the new specification will be consolidated and well supported by manufacturers. Still on this point, the fact that the custodial transfer option is no longer part of the draft RFC 5050-bis is not a critical point, since it has been simply moved to the Bundle-in-Bundle Encapsulation document, hence preserving his functionalities and more importantly its value within the DTN architecture.

3.2. SCENARIO B: CONSTELLATION-CLASS MISSIONS

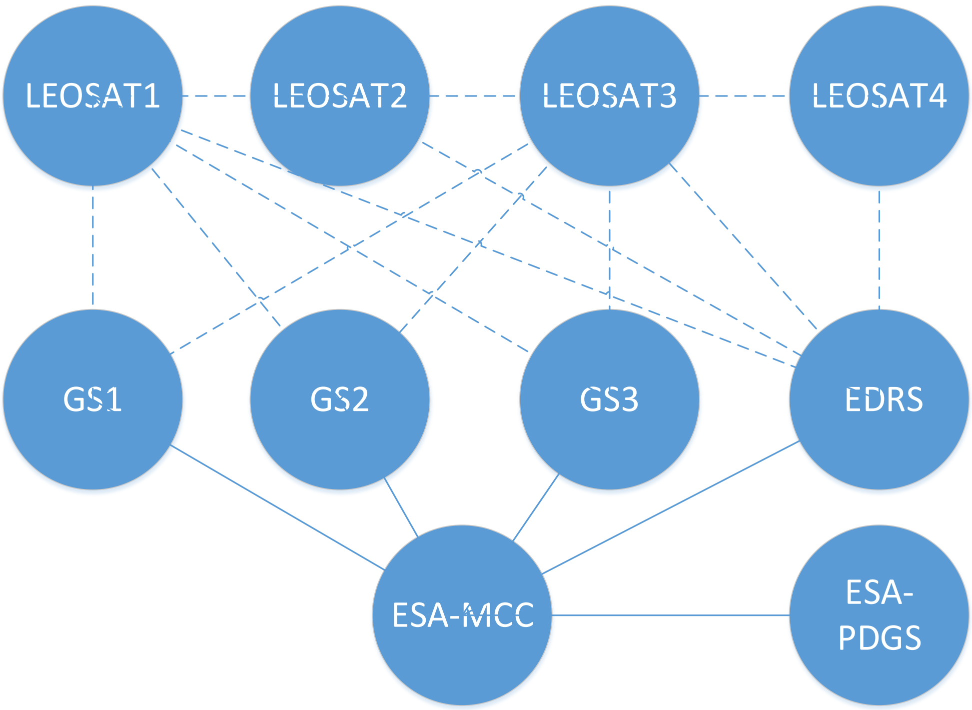

Figure 3.2 Scenario B logical layout (from [1])

3.2.1. RESULT SUMMARY

Unlike scenario A, scenario B was exclusively emulated as a DTN-based mission. Therefore, no comparative analysis was made with regards to individual node throughput or total network utilization. It was decided that the calculation of such an analysis would be extremely difficult and inherently incorrect, as there are a multitude of ways which a convoy mission could be implemented, such as file-based forwarding between nodes or naïve bent-pipe connections.

However, even without that analysis, some conclusions can be drawn from the data collected; most importantly, and with few exceptions, CGR made the correct routing decision, allowing data to be consistently received by Earth, within the constraints of physics.

3.2.2. ANALYSIS

Initial scenario design was complicated by the Inter-Satellite links; the computational complexity of Contact Graph Routing increases exponentially with the number of available nodes, due to the requirement that all proximate routes are considered. Attempts have been made to minimize this complexity by only calculating a single route, although that was found to be sub-optimal and has been revised in recent versions of ION and the CGR specification.

In the traditional non-convoy implementation of a constellation, each node should be able to connect to a large subset of fellow satellites. In turn, those satellites may have multiple space-to-ground connection options. It is difficult to calculate the exact number of required routes, as they are highly dependent upon the geometry of the constellation and connection mechanisms. However, it is safe to assume that they will grow exponentially.

If the assumption is made that all (or a significant subset) of available routes must be traversed in order to find a near-optimal final solution, this can take a significant amount of time. Furthermore, CGR routes are cached after calculation, in order to speed later bundle forwarding, which results in a sizable amount of memory utilization. These conclusions have led the authors of the DTN-EO study to hypothesize a set of potential requirements for a future DTN implementation, which are outlined in section 5.1.

3.3. SCENARIO C: DEEP-SPACE MISSIONS

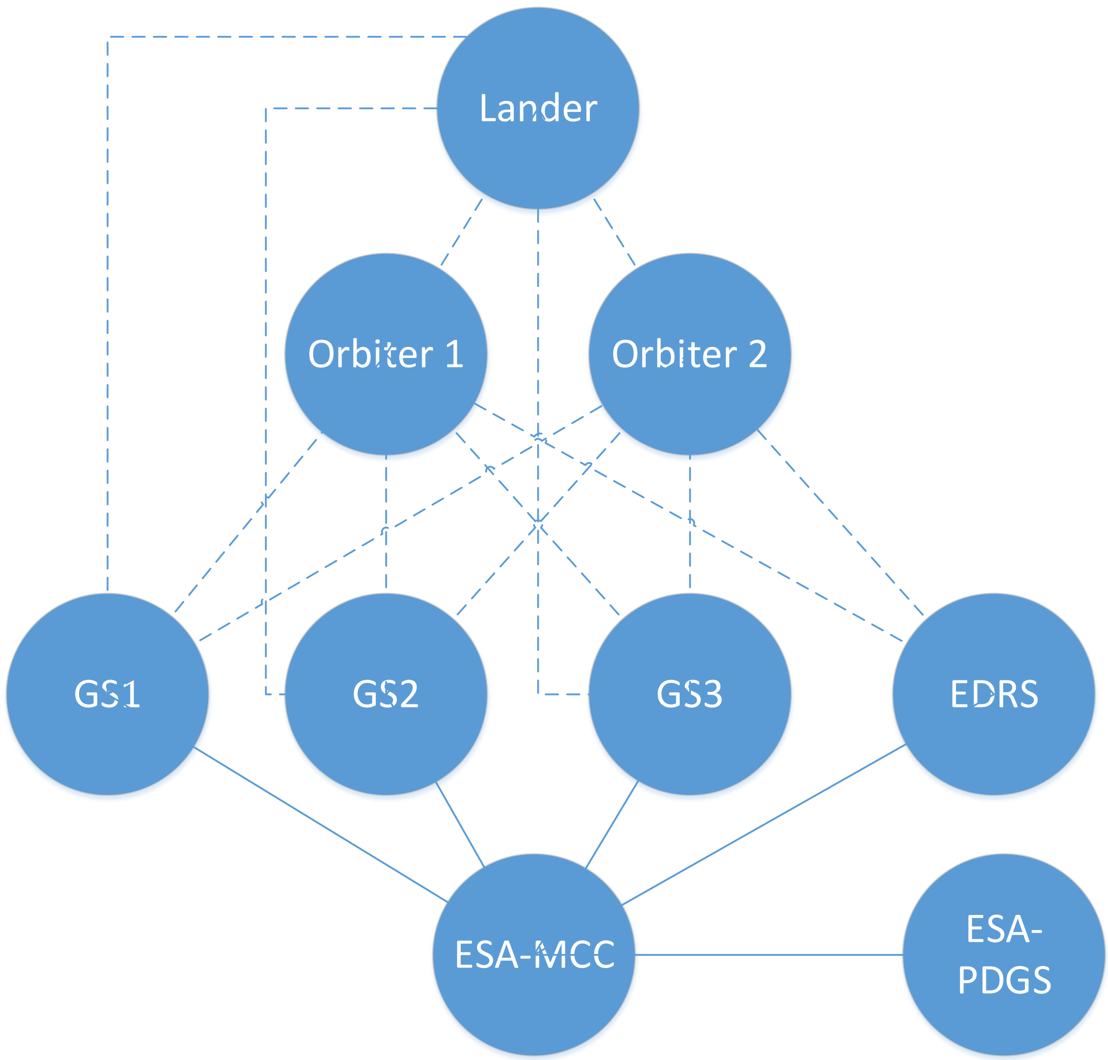

Figure 3.3 Scenario C logical layout (from [1])

The deep-space scenario (scenario C) considers a Martian scenario, where a rover is located on Mars. Two relay satellites orbit Mars, while one relay satellite is located in Earth's geostationary orbit. Assumptions applicable to constellation missions (from Scenario B) are generally re-used in this scenario. In order to simplify scenario analysis, although the LeoSat cluster from Scenario B has been removed and the rover has been replaced with a stationary lander. It is envisioned that this lander will forward data from an un-emulated rover.

Scenario C was particularly complex. It involved 9 DTN nodes, 13 concurrent data flows, a contact plan with 68 contacts on intermittent space links. The primary focus was on CGR/LTP ability to cope with multiple flows of very different characteristics (priority, bundle dimension, data rate), multiple paths and traffic very close to the maximum network capacity.

This scenario stressed the current state-of-the-art in DTN implementations, producing initially mediocre results. Careful and methodical analysis allowed the team to discover a set of underlying problems, including a few bugs in ION, which, when taken together, caused the issues seen in early testing.

Once the issues discovered in this scenario were mitigated, by modifying the initial specifications or and by adding new fixes to the ION code, the complex network from a Martian mission could be simulated successfully, with only few issues remaining. This showcases the fact that, as DTN protocols (including CGR) and implementations are evolving, accurate system-level reliability tests, as those done in our analysis, are in order before actual deployment.

4. ROLL-OUT

As discussed in the previous sections, DTN’s dynamic store-and-forward routing offers possibilities beyond the current point-to-point communication which is used in current spacecraft communication mechanisms.

At the same time, such new concepts and approaches can only be introduced gradually until they have proven themselves and shown some benefit for missions. Based upon this insight, the following roadmap can be considered for the integration of DTN into missions, both within the flight hardware and their respective ground segments.

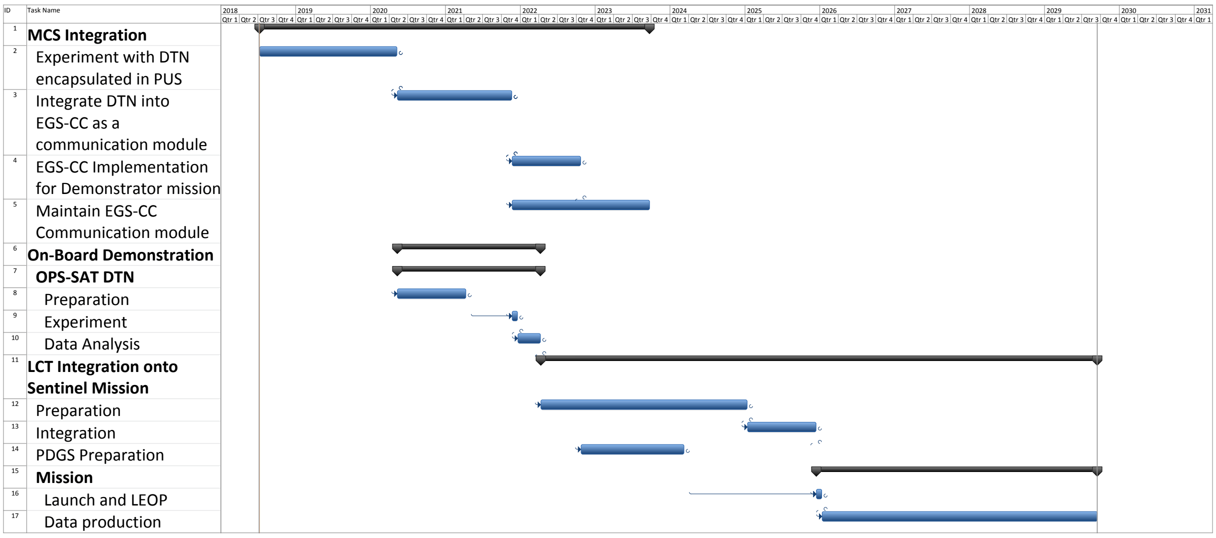

Figure 4.1 Overview of Roll-Out Process

4.1. ON-BOARD ASSETS

The design of communication and avionics infrastructures are dependent upon several factors: the long-time delay, which is a prime benefit of DTN implementation and the high bit error rate due to environmental conditions, electromagnetic interference, and electronic degradation/failure. The importance of these factors within infrastructural design is highly dependent upon the intended use of the avionics system: interplanetary, inter-satellite, or satellite-ground communications. Other variables, such as the required throughput, anticipated data generation capacity, and number of potential peer nodes (in a convoy or constellation) will affect other design choices, such as the mass memory unit devoted to communication storage, which may be omitted in some missions. The objective of this trade-off is an evaluation of avionics systems that may be suitable for current and future satellites, space probes or other spacecraft working under extreme conditions. Therefore, radiation-tolerant devices or even radiation-hardened devices will be considered.

The roll-out plan for ground-based infrastructure is closely tied to the choice of missions which use DTN. Therefore, in WP4000, the consortium analyzed multiple possible ground segment infrastructures [3]. Particular focus was paid onto the integration of DTN within a modern Monitoring and Control system, in order to easily support multiple future missions.

Within the DTN-EO study, two different requirement analyses were performed: the first, in WP4000, aimed to provide a high-level overview for future mission developers. The WP4000 work was performed prior to the completion of the scenario analysis, so further refinement of these requirements was performed in WP3000 [2], where the consortium outlined potential requirements for differing mission classes. For brevity, these requirements have been omitted in this document. For more details on these requirements, see [2].

4.2. RISKS

In order to advance DTN into a mission-ready state, many individual components must be advanced and integrated. As a result, risks and difficulties are inevitable. Particularly during the initial DTN-ready missions, it is foreseen that multiple integration steps may create difficulties and schedule slip. These issues exist on the avionics, ground segment, and mission planning.

Within the avionics environment, early implementations will have to cope with a new set of interfaces between DTN and the payloads, as well as those which connect DTN to low-level communication hardware. These bus-and-mission specific interfaces must be well-defined, either via existing standards or Memorandums of understanding between all relevant stakeholders.

The risks faced by the ground segment are slightly different, but equally significant. Integration of DTN into a MCS system is a non-insignificant task, and one which must be developed using methodical approaches while looking towards the future utilization of the MCS on new missions. Therefore, once the DTN communication module for a given MCS system is developed and tested, it is imperative that it is provided as part of the standard release for that system, and maintained and upgraded as-required. As a result, new missions shall be able to use that module without issue.

Finally, the partially-opportunistic nature of DTN creates some difficulties in deterministic scheduling of data product arrival, etc. Therefore, network behavior shall be simulated in the early phases of the mission, in order to provide insight and education to planning and scientific teams. As DTN networks become more complex, some scheduling guidelines may have to change to follow the opportunism of the network, although the details and impacts of such changes cannot be evaluated at this time.

5. CONCLUSION

The scope of the DTN-EO mission was extensive, and required excellent coordination and technical skills. That being said, all study goals were accomplished, and DTN was shown to be an excellent future mission enabler.

DTN was shown to provide efficient data transmission in all scenarios, although some design decisions were discovered in the ION DTN implementation of various DTN components, including LTP and Contact Graph Routing (CGR) which caused sub-optimal performance during scenario testing. These issues were mitigated by the development of patches to ION, either by the DTN-EO team or via NASA contacts, who were extremely supporting of the study.

Furthermore, an extensive rollout study was conducted, providing multiple potential pathways for the integration of DTN on a future mission.

5.1. FUTURE WORK

It is the belief of the consortium that, in order to minimize the long-term increase of risks relating to DTN implementation, particularly in high-performance applications, an agency shall create an agency-wide common DTN implementation for avionics systems. As a baseline and following the ION requirements, the implementation must provide the following capabilities:

- Written in a flight software approved language, using agency toolkits and programming practices.
- Modular design, targeting multiple operating systems.
- A well-defined and stable API in the same language as the implementation. In the vent of breaking changes, adaption layers (or new API calls) shall be provided.
- Internal hooks shall be made for the:
    - Storage system, to interface with mass memories.
    - The routing system, to interface with FPGA’s.
    - Convergence Layer Adapters and derivative components, to allow for FPGA offloading of CLA functionality.
- A (separate or combined) management interface API, to allow for management.
- CGR
- Security
- Long-term support contract for commercial users

Optimally, the implementation would be released as an open-source solution, allowing for student groups, academic users, and Small & Medium Enterprise (SME) users to learn about DTN. It is understood that this adds additional steps on the road to DTN integration, but believes that, in the long term, the advantages are formidable. The manpower requirements for this implementation are intentionally omitted, as they are highly dependent upon the release schedule and milestones for individual releases.

The development of a modular DTN implementation becomes significant with the results of Scenario B: It was found that the amount of CPU & memory required to perform routing calculations in a larger network may create a necessity for constellation-based systems (both on-board and within the ground segment) to be designed around parallelization, especially for routing calculations.