Exploring New Frontiers in Space Communications: Enhancing Delay Tolerant Networking through Cloud and Containerization

Blake LaFuente*, Rachel Dudukovich†, Thaddeus Kollar‡, Nadia Kortas†,
Daniel Raible†, Brian Tomko†, Alan Hylton§

*National Aeronautics and Space Administration, NASA Kennedy Space Center
†National Aeronautics and Space Administration, NASA Glenn Research Center
‡Banner Quality Management Inc, Cleveland, OH
§National Aeronautics and Space Administration, NASA Goddard Space Flight Center

***Abstract*—The High-rate Delay Tolerant Networking (HDTN) project at NASA Glenn Research Center has developed software that enables more flexible, reliable, and efficient space internetworking by using modern computing techniques such as cloud services, microservices, network function virtualization, software defined networking, and a distributed architecture [1]. HDTN is built upon the Bundle Protocol and related convergence layers which have been developed to mitigate the challenges of the space networking environment including long delays, asymmetric data rates, and intermittent connectivity. The HDTN implementation employs asynchronous message processing tasks which allow for non-blocking operations as well as deployment in both centralized and distributed architectures.**

This paper investigates deploying HDTN in a containerized approach on the NASA Goddard's Mission Cloud Platform using Amazon Web Services Elastic Compute Cloud (EC2). Commercial cloud computing will lower operating costs, provide flexible resource allocation, and allow for interconnectivity between multiple NASA centers as well as external partners. Containerization using Docker will enable greater portability and scalability for HDTN to be deployed into a variety of environments. We discuss possible NASA missions and use-cases such as the Laser Communications Relay Demonstration (LCRD) where the services provided by HDTN (reliable transport, high-rate message processing, and store-and-forward capabilities) will be enhanced through cloud computing and containerization. In addition, we describe the HDTN architecture and possible microservice-based networking approaches that can be obtained via HDTN's configuration capabilities. Finally, we detail the EC2 specifications needed to achieve data rates greater than 1 Gbps to support optical communication missions such as LCRD.

***Index Terms*—Delay tolerant networking, cloud computing, containerization, network virtualization, optical communication**

I. INTRODUCTION

The High-rate Delay Tolerant Networking (HDTN) [1] project at NASA Glenn Research Center has been evaluating the latest techniques for performance optimized space networks for many years, including approaches based on distributed architectures, micro-services, and virtualization [2], [3], [4]. The HDTN architecture [5] was designed to support high-rate optical communication demonstrations such as the Integrated LCRD Low Earth Orbit User Modem and Amplifier Terminal (ILLUMA-T) technology demonstration [6] and the Laser Communications Relay Demonstration [7]. This paper discusses a series of performance benchmarks of the HDTN software in a cloud-based containerized environment. Compute resources, network data rates, and cost modeling are discussed to evaluate the suitability of a cloud-based ground infrastructure for high-rate Delay Tolerant Network (DTN).

The Laser Communications Relay Demonstration (LCRD) is an optical communications satellite payload with the purpose of advancing optical communications technology toward infusion into deep-space and near-Earth operational systems. LCRD experiments will demonstrate that optical communications are a superior solution when compared to radio frequency because of reduced size, weight, and power (SWaP) requirements while also achieving higher bandwidths. In addition, LCRD's architecture will enable it to serve as a developmental testbed for advanced communication techniques including adaptive optics, symbol coding, link layer protocols, and network layer protocols. LCRD was launched on December 7, 2021, and the payload was powered on January 12, 2022. It will demonstrate optical communication in an operational environment for a minimum of two years.

A concept for a cloud-based ground infrastructure is shown in Fig. 1. Data can be relayed from any ground station to another through LCRD. LCRD will support both optical and RF links and ground stations can be connected to users via the cloud and/or terrestrial internet. A commercial cloud is centrally accessible to all ground stations as well as the user mission operations center. Cloud services such as Amazon Web Services support inter-region data flows that will connect ground stations to users throughout the globe. One motivating idea to keep in mind is that while LCRD can support greater than 1Gbps links, the optical ground stations are remote - HDTN then offers a robust and standardized way to offer *store, carry, and forward*, meaning its buffering approach can overcome any possible link asymmetries.

The future of space communication and networks has been evolving towards a partnership between government and

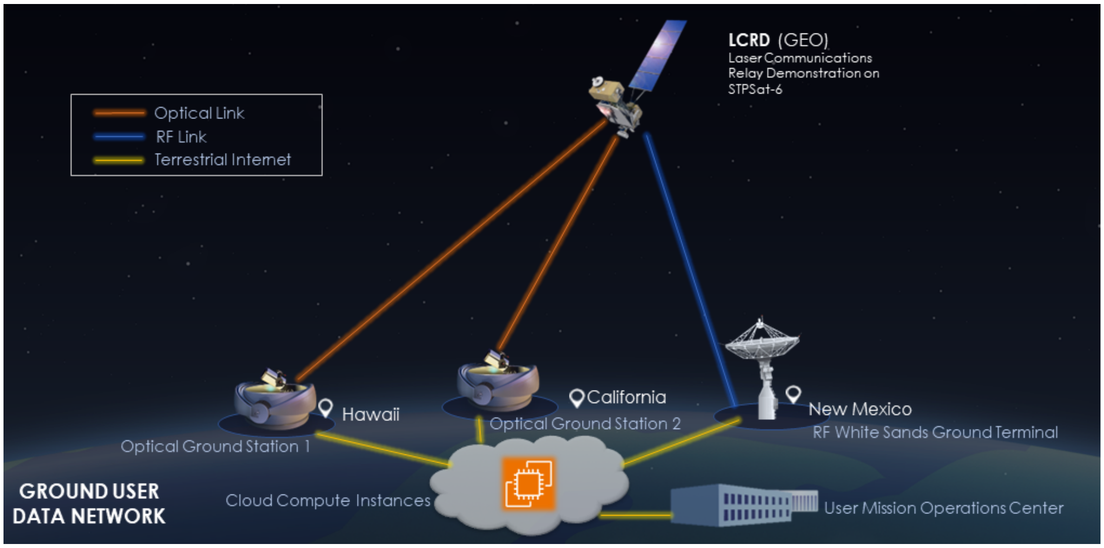

Fig. 1. LCRD Cloud Concept

commercial services. Commercialization efforts build upon previous research done by government space agencies such as NASA, but will improve system availability, accessibility, and cost efficiency. NASA has engaged with industry to expand near-Earth direct to Earth (DTE) as well as relay services, and determined a variety of synergies for future technology maturation and infusion opportunities. Delay tolerant networking (DTN), integration of optical communications ground terminals and integration of cloud computing assets into the near space network (NSN) architecture have been evaluated by the Space Communications and Navigation (SCaN) Commercial Innovation and Synergies (CIS) Office [8]. Cloud computing in particular establishes a simplified interface between a mission and service provider, reduces infrastructure complexity, and improves reliability based on asset distribution and redundancy [9].

II. BACKGROUND

A. HDTN Architecture

The HDTN software implementation has been developed to meet the primary requirements of high-rate networking performance and system reliability. The primary unit of data in a DTN is the *bundle*, which may be of any size. Moreover, DTNs operate as an overlay network on top of *any* underlying network/protocols. Hence the interpretation of "high-rate" is that bundles of any size (including less than 1k) can be sent at line-rate, which may be in excess of 1Gbps. The goal, then, is to create a lightweight application that takes advantage of modern hardware platforms to substantially reduce latency and improve throughput compared to today's DTN operations. HDTN's architecture also supports hooks to replace various processing pipeline elements with specialized hardware accelerators. This offers improved Size, Weight, and Power (SWaP) characteristics while reducing development complexity and cost. The HDTN source code is publicly available on the NASA open-source GitHub [10].

HDTN was designed with a modular architecture and consists of the following main modules: Egress, Ingress, Router, Storage and Telemetry Command Interface. Fig. 2 shows the main HDTN components, internal messaging, and associated data flows.

The Ingress module intakes bundles received by the HDTN node and then decodes the header fields to determine the source and destination of the bundles. If a link is available, Ingress will forward the bundles in a cut-through mode straight to Egress. Otherwise the bundles are sent to the Storage module. Bundles are also sent to Storage if custody transfer is enabled.

The Router determines if a given bundle should be forwarded immediately to Egress or sent to Storage. It does so by reading a contact plan which is a JavaScript Object Notation (JSON) file that defines all the connections between all the nodes in the network. The Router also calculates the optimal route using one of the routing algorithms in the routing library and sends a `RouteUpdate` ZeroMQ message to Egress to update its next hop. Unexpected link changes and contact plan reloads are also handled by the router which sends link changes events and updates the route dynamically.

Storage is a multi-threaded implementation distributed across multiple disks. It receives messages from the Router to determine when stored bundles can be released and forwarded to Egress. It is also responsible for custody transfer.

The Egress module is responsible for forwarding bundles received from Storage or Ingress to the correct outduct and the next hop based on the most optimal route for the bundle. If a connection is lost unexpectedly, Egress sends a status message to the Router requesting a new optimal route.

The Telemetry Command Interface module exchanges messages with the Graphical User Interface (GUI) which displays graphs of the data and statistics for network troubleshooting. One of the pages displays a live system view of HDTN and the different inducts and outducts linked to the current node. Another displays the current configuration settings.

B. Cloud Computing on the NASA Near Space Network

NASA's Space Communications and Navigation (SCaN) program and the Near Space Network (NSN) are in the process of creating a cloud environment suitable for the entire network architecture of the SCaN program [9]. The goal is to build a system that is cost-effective, scalable, robust, and resilient. Other benefits include improving data processing time and ease of access. Future missions planning to use cloud for the data delivery architecture include NASA-ISRO Synthetic Aperture Radar (NISAR), Plankton, Aerosol, Cloud, ocean Ecosystem (PACE), and Roman Space Telescope (RST).

The DTN Engineering Network (DEN) consists of labs at NASA Centers and occasionally external partners, connected through VPNs managed at NASA Glenn Research Center. Work is underway to add a gateway to an AWS Virtual Private Cloud so that instances in EC2 can communicate with hosts participating in the DEN. This connection between the DEN and the AWS cloud will allow for access between the NASA labs and AWS ground stations, NASA operated ground stations, and may serve as a common test environment. Fig. 3 shows the DEN and planned AWS integration via the Goddard hosted Mission Cloud Platform.

C. Containerization

As a requirement for working with LCRD, containerization functionality was implemented with HDTN using Docker. Containers can now be easily spun up with software images that build HDTN along with all of its necessary dependencies. These Dockerfiles are configured to run using Ubuntu or Oracle Linux. Docker Compose files exist so users can easily create both sender and receiver nodes. For tests and experiments that require a larger number of HDTN nodes, users can utilize Kubernetes for orchestration.

III. EXPERIMENT

A. Amazon EC2

Amazon Elastic Compute Cloud (EC2) is a web service that enables the use of virtual compute environments, called instances, using Amazon Web Services (AWS) Cloud. These instances are configured with a range of different hardware capabilities: CPU Cores, RAM Size, and Storage Size (all of which are Solid State Drives). The network performance capabilities also differed between instances. The HDTN team chose to use the t3 instances because they offered the widest array of hardware specifications and were both cheaper and faster than the previous generation t2 instances. The specifications for those used are listed in Table I as well as the cost per hour for users to run them.

Amazon EC2 was hosted using NASA Goddard's Mission Cloud Platform (MCP). MCP offers scalable and consolidated mission cloud services within NASA [11]. The MCP model lowers barriers to cloud entry by implementing compliance guardrails and security tools while supplying fully customizable services to support mission-specific needs.

TABLE I
EC2 NODE HARDWARE CAPABILITIES

Node Type	CPU Cores	RAM (GB)	Storage (GB)
t3.nano	2	0.5	16
t3.micro	2	1	16
t3.small	2	2	16
t3.medium	4	2	16
t3.large	8	2	16
t3.xlarge	16	4	16
t3.2xlarge	32	8	16

B. Test Setup

For each type of EC2 configurations used, HDTN nodes were set up on two instances with equivalent hardware capabilities; one to send a one gigabit data file and the other to receive it. The data transfer rate was measured for sending the data using the Licklider Transmission Protocol over User Datagram Protocol (LTP over UDP) convergence layer [12]. This process was then repeated using the next pair of instances to measure the effect of changing the hardware.

Fig. 4 shows the experiment topology. HDTN provides the file transfer applications BPSendFile and BPReceiveFile that convert files on disk into bundles. The Simple TCP (STCP) convergence layer is used to transfer bundles to and from the file transfer applications to the local HDTN bundle agent. This topology is similar to what is used on the ISS DTN network and what is anticipated to be used for LCRD [13]. STCP or TCP convergence layers are used for onboard local area networks or terrestrial networks. LTP is used for space to ground links and is specifically designed for long delays, asymmetric links, and link disruptions. This same test setup was used to run checkout tests for HDTN to be used during a flight experiment with LCRD.

IV. RESULTS

A. Rates based on Instance Type

TABLE II
FILE TRANSFER RESULTS

Node Type	Transfer Rate (Mbit/sec)
t3.nano	302.87
t3.micro	336.41
t3.small	334.76
t3.medium	250.46
t3.large	573.03
t3.xlarge	230.60
t3.2xlarge	604.70

The rates of data transfer for the different EC2 Nodes is listed in Table II. The rate was determined by taking measurements of the transfer rate and averaging all but the first and last values. HDTN was successfully able to transfer the data file between nodes regardless of the hardware capabilities.

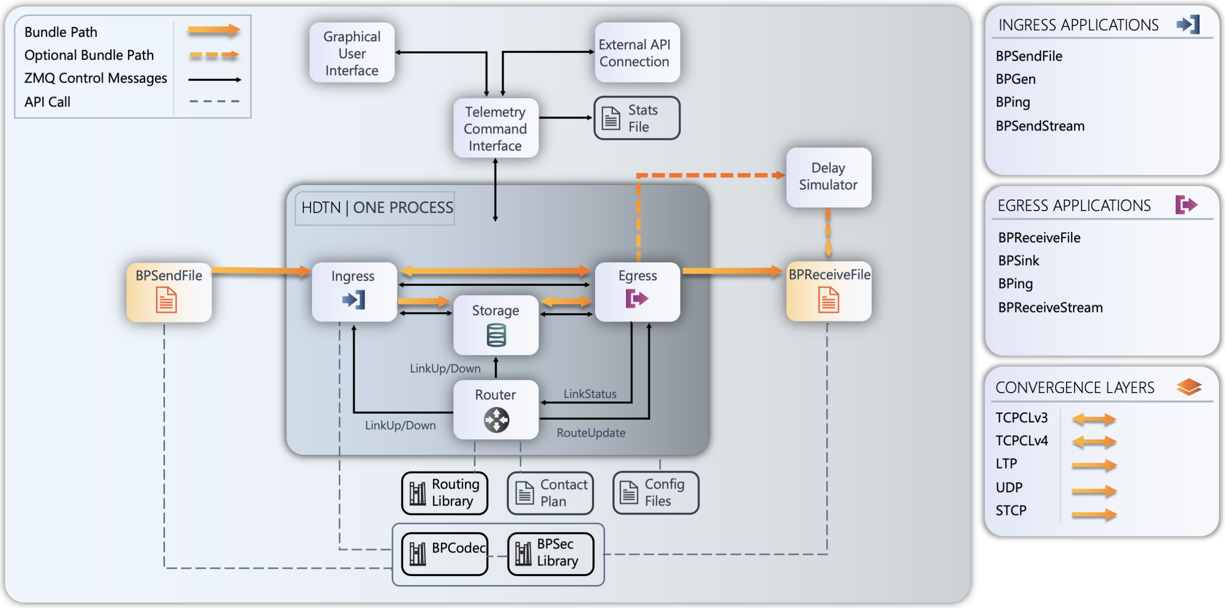

Fig. 2. HDTN software architecture and modules interactions.

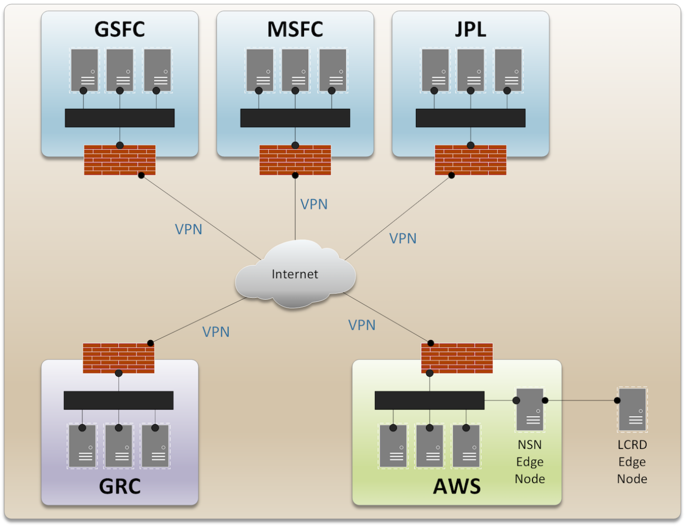

Fig. 3. Planned EC2 integration with the DTN Engineering Network.

The use of checksums confirmed that the data received was complete and without errors.

As the results show, the transfer rate did not increase with improvement in the hardware capabilities. This is apparent with the t3.medium and t3.xlarge nodes which experienced the slowest transfer rates. Each test was repeated to ensure that results were consistent. While Network Burst Bandwidth for all t3 nodes is listed at five gigabits per second, it appears that baseline performance is quite variable and unlisted by AWS.

B. Cost

In order for adoption for Delay Tolerant Networking to be accepted as a potential commercial option, cost must be optimized as well as performance. A major benefit of using AWS EC2 is the ability to test with a range of cloud-based hardware options with the price of running them given up-front. The cost of running each of the nodes per hour is listed in Table III.

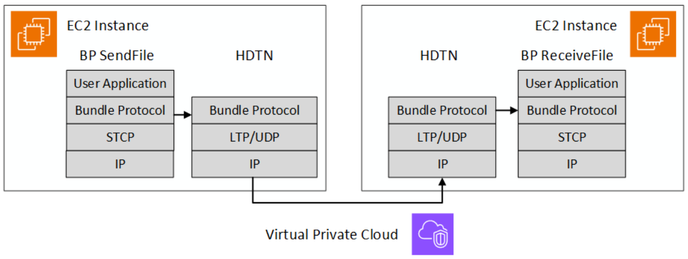

Fig. 4. Experiment Topology

Given the price of running two nodes and the transfer rate of sending the 1Gb file, one can easily calculate the true cost of the file transfer as shown in Table III. Based on these metrics, in order to gain the capabilities of doubling the transfer speed, as seen between the t3.nano and the t3.2xlarge in Table II, the costs increases by a factor of 32.

C. LCRD Checkout Testing

During checkout tests for LCRD, two HDTN containers were created using Docker with a 32Gb data file to send between them using the test network for LCRD. HDTN was able to transfer the data at 2.4 gigabit per second using the same convergence layer as these cloud-based test (LTP over UDP). Data transfer was successfully demonstrated between the two HDTN nodes via the HDLC (High-level Data Link Control) Encapsulated Interface for Data Interchange (HEIDI)

TABLE III
FILE TRANSFER COST

Node Type	**Node Cost/Hour ($)**	**Cost of 1Gb File Transfer ($)**
t3.nano	0.0052	0.0000095
t3.micro	0.0104	0.0000172
t3.small	0.0208	0.0000345
t3.medium	0.0416	0.0000922
t3.large	0.0832	0.0000807
t3.xlarge	0.1664	0.0004009
t3.2xlarge	0.3328	0.0003058

for the following four scenarios: Bundle Protocol version 6 with and without custody transfer, Bundle Protocol version 7 and with and without Bundle Protocol Security (BPSec).

FUTURE WORK

HDTN has undergone a series of tests in emulated space environments as well as this current work focused on support for LCRD and developing the future NSN ground infrastructure. Containerization has helped to create a standardized method of deployment in both cases. HDTN is planned to demonstrate high-rate space inter-networking capabilities on the ISS and via LCRD, both of which use containerized or virtualized environments. In addition to containerization, cloud infrastructure allows for a low cost, scalable, and distributed network. HDTN has shown to be at the forefront of this future commercialization trend.

The use of cloud-based ground infrastructure will allow for standardized test environments that are accessible to government agencies, industry, and academia. Our future work includes plans to further develop the DEN and extend its accessibility to a variety of collaborators, with cloud solutions offering a simple and secure interface to the network.

ACKNOWLEDGMENT

The authors would like to thank the NASA Space Communications and Navigation (SCaN) program, and in particular Philip Baldwin (NASA/HQ) for supporting this work. They would also like to thank John Nowakowski for his incredible leadership of the HDTN team these past few years.

REFERENCES

[1] NASA Glenn Research Center. *High-Rate Delay Tolerant Networking (HDTN) Project*. URL: https://www1.grc.nasa.gov/space/scan/acs/tech-studies/dtn/.

[2] Nathan Drzadinski et al. "Space Networking Implementation for Lunar Operations". In: *37th Annual Small Satellite Conference*. SSC23-IX-08. 2023.

[3] Rachel Dudukovich et al. "A distributed approach to high-rate delay tolerant networking within a virtualized environment". In: *2021 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)*. IEEE. 2021, pp. 1–5.

[4] Rachel Dudukovich et al. "Microservice architecture for cognitive networks". In: *2020 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2020, pp. 39–44.

[5] Alan Hylton et al. "Rising above the cloud: Toward high-rate delay-tolerant networking in low earth orbit". In: *Advances in Communications Satellite Systems. Proceedings of the 37th International Communications Satellite Systems Conference (ICSSC-2019)*. IET. 2019, pp. 1–17.

[6] Farzana I. Khatri et al. "System level TVAC functional testing for the Integrated LCRD Low-Earth Orbit User Modem and Amplifier Terminal (ILLUMA-T) payload destined for the International Space Station". In: *Free-Space Laser Communications XXXV*. Ed. by Hamid Hemmati and Bryan S. Robinson. Vol. 12413. International Society for Optics and Photonics. SPIE, 2023, 124130E. DOI: 10.1117/12.2649326. URL: https://doi.org/10.1117/12.2649326.

[7] D. Israel et al. *Early Results from NASA's Laser Communications Relay Demonstration (LCRD) Experiment Program*. URL: https://ntrs.nasa.gov/api/citations/20230000923/downloads/LCRD_SPIE%20West%202023_final%20corrected.pdf (visited on 01/19/2023).

[8] Gregory W Heckler et al. "NASA's Efforts to Pursue Commercial Communications Services for Missions in Near Space". In: *28th Ka and Broadband Communications Conference (Ka) and the 40th International Communications Satellite Systems Conference (ICSSC)*. 2023.

[9] P. Baldwin et al. *NASA's Implementation of Cloud Services for Human Space Flight*. URL: https://ntrs.nasa.gov/api/citations/20230001336/downloads/NASA_Cloud_SpaceOps23_R4%20-%20No%20Logos%20for%20NTRS.pdf (visited on 11/16/2023).

[10] NASA Glenn Research Center. *High-Rate Delay Tolerant Network*. URL: https://github.com/nasa/HDTN.

[11] C. Tresslar. *Mission Cloud Platform (MCP)— Transforming the Mission Environment at NASA*. URL: https://www.nasa.gov/wp-content/uploads/2016/06/490184_-_july-sep_2021_it_talk_design.pdf.

[12] M. Ramadas, S. Burleigh, and S. Farrell. "RFC 5326, Licklider Transmission Protocol - Specification ". In: *IETF Network Working Group* (2008). URL: https://datatracker.ietf.org/doc/html/rfc5326.

[13] D. Raible et al. *Developing High Performance Space Networking Capabilities for the International Space Station and Beyond*. Technical Memorandum. NASA, 2022. URL: https://ntrs.nasa.gov/api/citations/20220011407/downloads/TM-20220011407.pdf.