

Use of Open Networks and Delay-Tolerant Protocol to Decrease WAN Latency of EOS Near Real-Time Data

Bruce McLemore¹

KBRwyle / GSFC, Greenbelt, MD 20771, USA

Terri Wood²

NASA GSFC, Greenbelt, MD 20771, USA

Kevin Kranacs³

NASA GSFC, Greenbelt, MD 20771, USA

Since 1999, NASA's Earth Observing System Data Operations System (EDOS) project at Goddard Space Flight Center (GSFC) has provided high-rate data capture, level zero processing, and product distribution services for a majority of Earth Observing System (EOS) high-rate missions, including Terra, Aqua, Aura, ICESat, EO-1, SMAP, and OCO-2. EDOS high-rate science and engineering (150-300 Mbps) data-driven capture systems are deployed at 7 worldwide ground stations which are connected via both private (closed) and public (open) wide area networks (WANs) to the centralized EDOS Level Zero Processing Facility (LZPF) located at GSFC, where the data is processed and Level 0 products are distributed to users worldwide. All data transferred over the open networks to GSFC traverse an IPSec tunnel, providing the same level of security as a VPN connection. EDOS produces both time-based and near real-time products (session-based). Near real-time data products are produced from a single ground station contact; time-based products are produced from multiple ground station contacts. EDOS is the primary supplier of EOS Level 0 data to the NASA near real-time user community known as the Land, Atmosphere Near real-time Capability for EOS (LANCE). For the past few years, EDOS has streamlined its systems to reduce WAN latency for near real-time data delivery, including implementing Quality of Service (QoS), expanding closed network bandwidth, adding open network connections with more bandwidth, and implementing a delay-tolerant protocol to mitigate long round-trip times to remote ground stations.

I. Introduction

NASA's high-rate data capture and distribution system EDOS has supported Terra science downlinks at White Sands Complex since December, 1999. Previous SpaceOps papers have described the evolution of this data-driven, multi-mission high-rate system to include additions of Aqua, Aura, ICESat, EO-1, OCO and ALOS to the mission set (see references 1, 2, 3). Even after the loss of OCO, ALOS, and EO-1, EDOS has continued to add new missions: SMAP, OCO-2, and more recently SNPP and JPSS-1 (NOAA-20). ICESat-2 has also been added to the mission set awaiting launch scheduled for later this year. This will bring the EDOS operational mission set up to 8 concurrent missions with modular capacity to add even more missions in a cost-effective manner. Supporting a large number of missions from a diverse set of ground stations and network interfaces places increasing demands on EDOS to continue to streamline systems and architectures to efficiently use network resources and improve near real-time product latency.

II. Background

EDOS, under the management of the Earth Science Mission Operations (ESMO) Project, GSFC Code 428, is responsible for capture and initial processing of science and engineering data for selected high-rate EOS

¹ GSFC-2 EDOS Project Lead, GSFC/ Code 428, Greenbelt, MD 20771, bruce.d.mclamore@nasa.gov

² EDOS Project Manager, GSFC/ Code 428, Greenbelt, MD 20771, terri.wood-1@nasa.gov

³ Networks Manager, GSFC/ Codes 423/428, Greenbelt, MD 20771, kevin.m.kranacs@nasa.gov

spacecraft. EDOS provides capabilities for spacecraft data that adhere to recommendations established by the Consultative Committee for Space Data Systems. EDOS elements are distributed over various ground stations (NASA and commercial stations) and 2 LZPF facilities at GSFC. EDOS ground station elements are located at the White Sands Complex near Las Cruces, New Mexico; the Alaska Satellite Facility at the University of Alaska, Fairbanks, Alaska; the Fairbanks Command & Data Acquisition Station at Gilmore Creek, near Fairbanks, Alaska; and the Wallops Flight Facility in Wallops Island, Virginia. EDOS operates in a data-driven mode (without schedules) and autonomously captures science data at the ground station and performs front-end processing. The system front-ends support spacecraft downlink rates from 150 Mbps to 300 Mbps per channel, with 2 channels per system. The EDOS front-end processors support both ECL (serial clock-data) baseband and analogue intermediate frequency interfaces. Each channel is configured to constantly listen to a single multi-mission antenna feed in data-driven mode. The system front-ends automatically identify the mission by the spacecraft ID and frame length detected from the telemetry, which indicates the mission-specific processing configuration to be used. Front-end processing includes demodulation, de-randomization, frame synchronization, and decoding. After data capture begins, the front-end automatically initiates transfer of science data to GSFC over the WAN.



Figure 1. EDOS High-Level Multi-Mission Data-Driven Architecture

Several key objectives over the years have been to reduce mission operations costs, increase productivity by enabling automation for all nominal operations from data capture to product distribution, and improve product latency to end-users. At the LZPF EDOS mission-specific processors ingest data received from the ground stations and perform Level 0 processing per mission specific interface requirements. EDOS produces various expedited near real-time session-based products from a single spacecraft contact session as well as time-based products from multiple contact sessions which contain an observation time interval of time-ordered packets. The Level 0 data products are distributed to the end-users with a variety of protocols and formats and archived at the LZPF on physical media based on mission requirements. The LZPF also provides a short-term 30-day on-line data storage capability for reprocessing of data. EDOS delivers roughly 1 Terabyte of Level 0 products worldwide (more than 20 external customers) on a daily basis. EDOS also maintains a complete backup operational LZPF facility at an alternate location for disaster recovery which also has network connections to all ground stations and end-users.

Recently EDOS began receiving SNPP and JPSS-1 data via the NWAVE network peered with the EOS network. This relayed data is received by EDOS on an intermediate system identical to the relay systems used by JPSS, and is forwarded to the LZPF so that it appears to EDOS that it came from a ground station. Although this interface is working well, EDOS has no control of the latency from the ground stations used by JPSS, so the methods described herein to reduce latency over the WAN are not directly applicable to SNPP and JPSS-1 on EDOS.

III. Quality of Service

One of the first mechanisms EDOS implemented to improve latency from the ground stations was Quality of Service (QoS). QoS is a set of rules to control bandwidth utilization on a network. At ground stations at which EDOS supports multiple antennas in a data-driven mode, multiple missions (e.g., Aqua, Aura) may overlap in contact times at the same site and compete for the same network resources. Without a method to prioritize transfers over the WAN, all missions would just share the available bandwidth equally (see Figure 2). Since not all missions have the same latency requirements, EDOS deployed QoS at ground stations with multiple EDOS-supported

antennas to enable prioritizing of one session over another. In addition to near-real time traffic, replay traffic may be given the highest priority so that additional latency is kept to a minimum. EDOS QoS configurations are extremely flexible providing assignment of an arbitrary number of priority levels, with each priority receiving a configurable amount of bandwidth. High-rate traffic at any priority will be given the full bandwidth if there is only one session in progress.

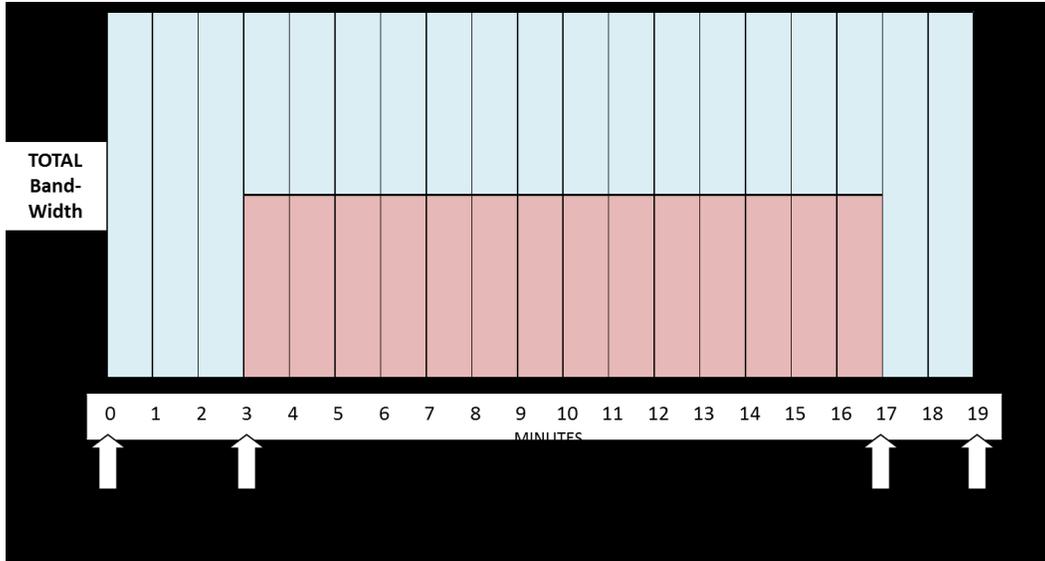


Figure 2. WAN bandwidth utilization between ground station and central site without QoS

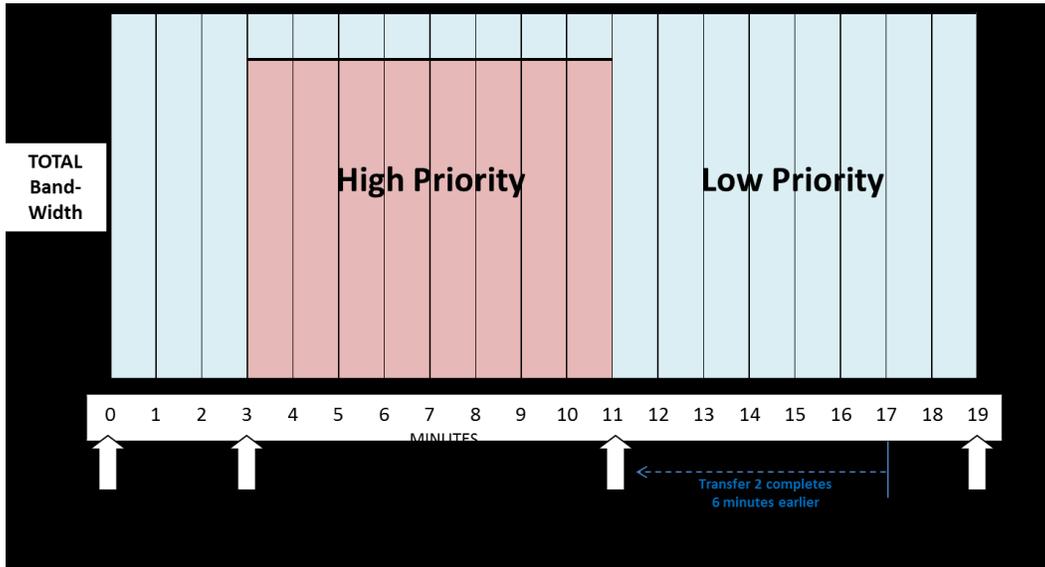


Figure 3. WAN bandwidth utilization between ground station and central site with QoS

The EDOS QoS appliances are small servers running only the SUSE Linux operating system (SUSE) with no additional software required. Other traffic management devices were considered; including commercial appliances such as *Packeteer* or *NetEqualizer*, but the open source Linux kernel version satisfied the requirement and was extremely cost effective. The QoS Box has only two network interfaces. It functions as a router in forwarding packets from one interface to another. It also classifies and shapes the packets based on the TCP destination port. Traffic shaping occurs on the data egress interface, unidirectional. The QoS queuing mechanism is known as Hierarchical Token Bucket (HTB) classful queuing, built right into the standard Linux kernel. EDOS QoS is configured in Linux using Traffic Control (tc) commands, included with the Linux kernel. The Hierarchical Token

Bucket (HTB) queuing discipline was chosen among the available Linux queuing disciplines because it was best suited to meet EDOS high-rate WAN needs:

- (1) Able to rate limit
- (2) Allows sharing of bandwidth
- (3) Able to prioritize packet flow

In the EDOS QoS implementation, there is a single "root" class that holds all of the bandwidth. This root class holds multiple "leaf" classes that hold a subset of bandwidth. The leaf classes are configured so they can borrow bandwidth from the root class. Each leaf class is associated with a different TCP destination port (or ports) which are configured for each mission.

For redundancy to eliminate a single point of failure, each ground station network interface requires 2 QoS boxes: one primary and one backup unit. All EDOS front-end systems have 2 high rate interfaces: high-rate A and high-rate B (see Figure 4). In the event of a failure of a QoS box (or its associated EOS firewall), the entire set of EDOS front-end systems is switched to the other high-rate interface to connect to the other QoS. For QoS to function properly all outbound traffic must be routed through only one QoS box; traffic cannot be split between them for the outbound network flow.

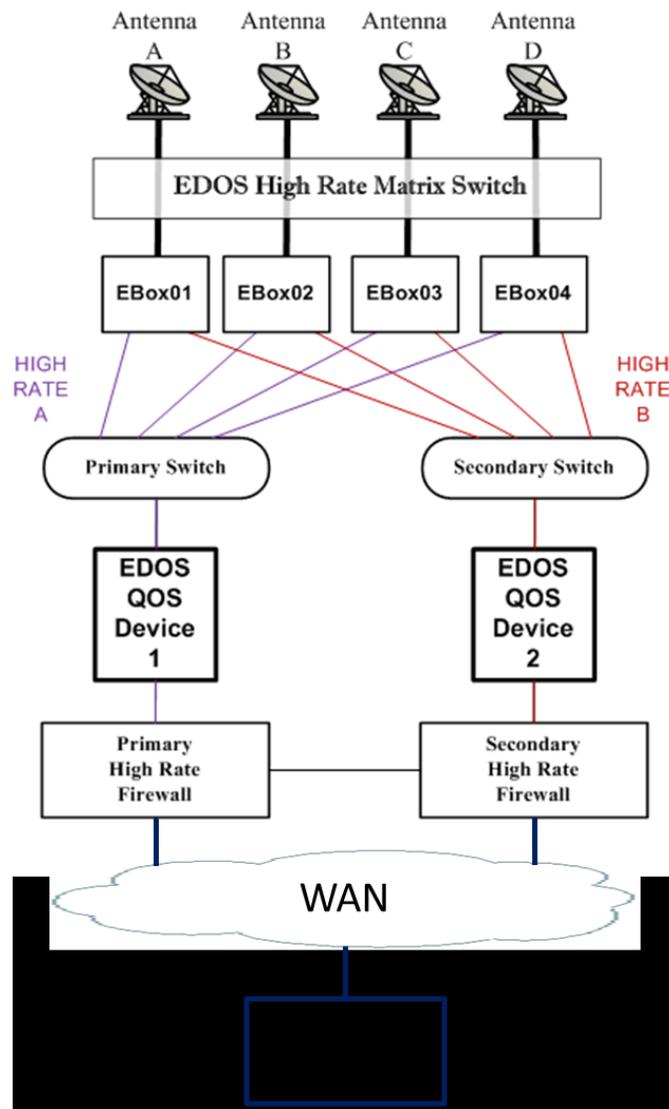


Figure 4. Typical EDOS QoS Configuration at Remote Site

EDOS QoS implementation at the remote sites provides a flexible means to prioritize dataflow to the central site and lower the latency of higher priority data. Network traffic is prioritized based on TCP port numbers which map to established priority classes. Each inbound packet is classified according to the destination TCP port; port ranges are allocated for each mission. Class priorities are defined for each mission. Streams of same class will share available bandwidth, and prioritization only occurs if there are concurrent high-rate data flows from the same station. Missions with high priority near real-time users are given high QoS priority.

IV. EDOS Ground Capability for EOS (Earth Observing System) (LANCE)

EDOS Ground Capability for EOS (Earth Observing System) (LANCE) provides global data and imagery from selected EOS satellites in less than 3 hours from satellite observation to meet the needs of the near real-time (NRT) applications community who needs data much sooner than what routine science processing offers. LANCE was developed in response to the need for timely satellite observations by applications users, operational agencies and researchers (see Reference 4). Especially prominent NRT users are the emergency support and rapid response personnel, such as firefighters and first responders. LANCE defines latency as the time from satellite observation to product delivery. To deliver satellite data products with sufficient latencies to meet the 3 hour needs of the near real-time (NRT) user communities, aspects of many EOS capabilities were modified: from geo-location (attitude and ephemeris) data to ground systems (including EDOS) and in some cases to science algorithms. Figure 5 below shows a simplified overview of the LANCE dataflow from satellite to users, including the role of EDOS in supplying the Level 0 data to the science data processing elements:

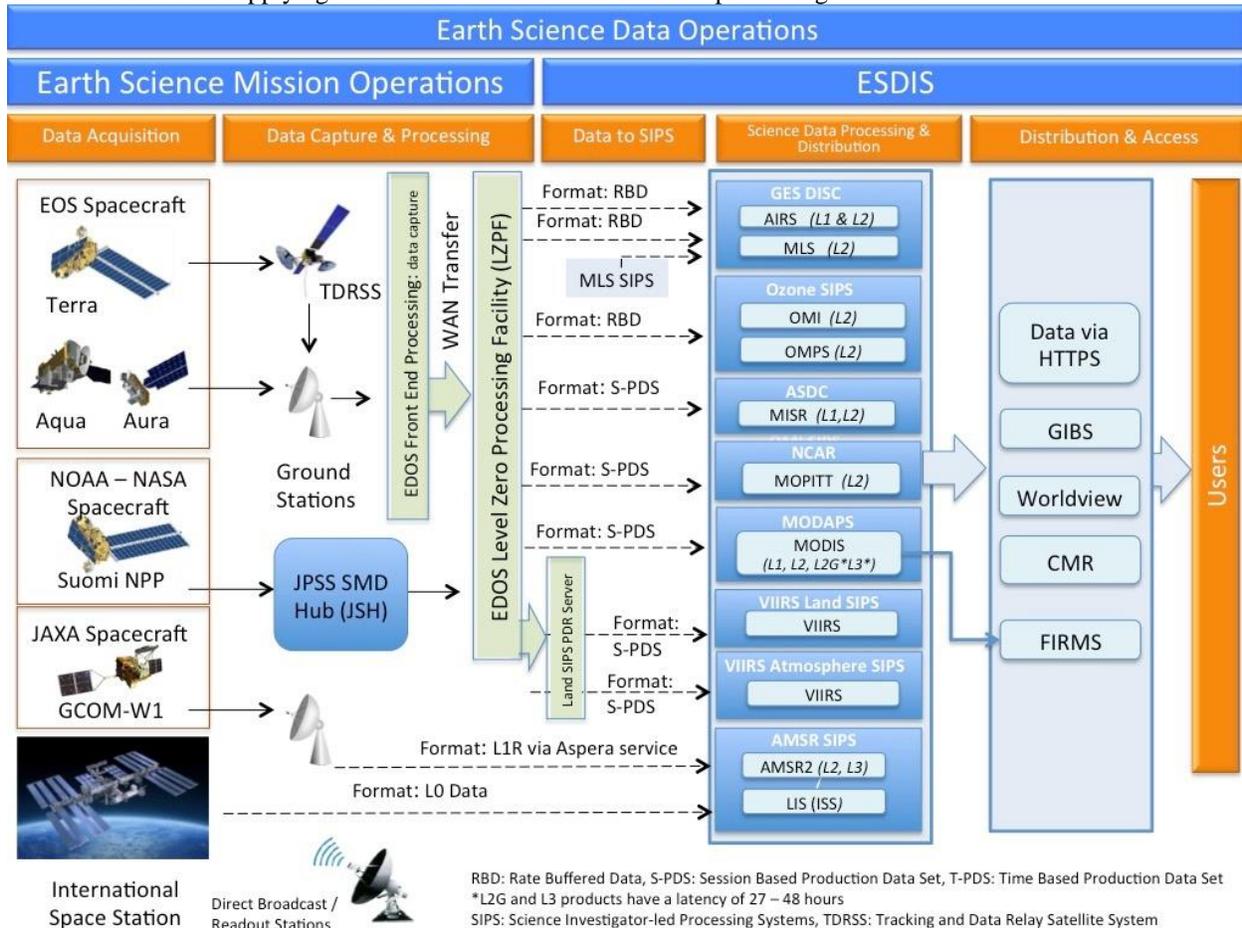


Figure 5. EDOS Role in LANCE

EDOS is responsible for capture and initial processing of science and engineering data from the EOS spacecraft. Once downlinked and captured at the ground station, EOS satellite data is delivered over WAN high-rate lines to the central site for processing and distribution to users.

contact session is sorted, processed and Level 0 data products are delivered in an expedited manner to designated Science Investigator-led Processing Systems (SIPS). EDOS produces session-based data sets especially for LANCE NRT use; this data is sent to dual LANCE destinations as part of the standard redundancy requirement for LANCE elements. EDOS-provided Level 0 data is processed into higher-level products at LANCE science processing facilities located at designated SIPS, where products are then made available to the NRT user community.

Figure 6 shows the approximate breakdown of latency times of the various EDOS ground system elements from data acquisition to distribution of Level 0 science products to LANCE elements for Aqua, Aura, and Terra. Note that the WAN transfer is not post pass, but is initiated and continues during data capture.

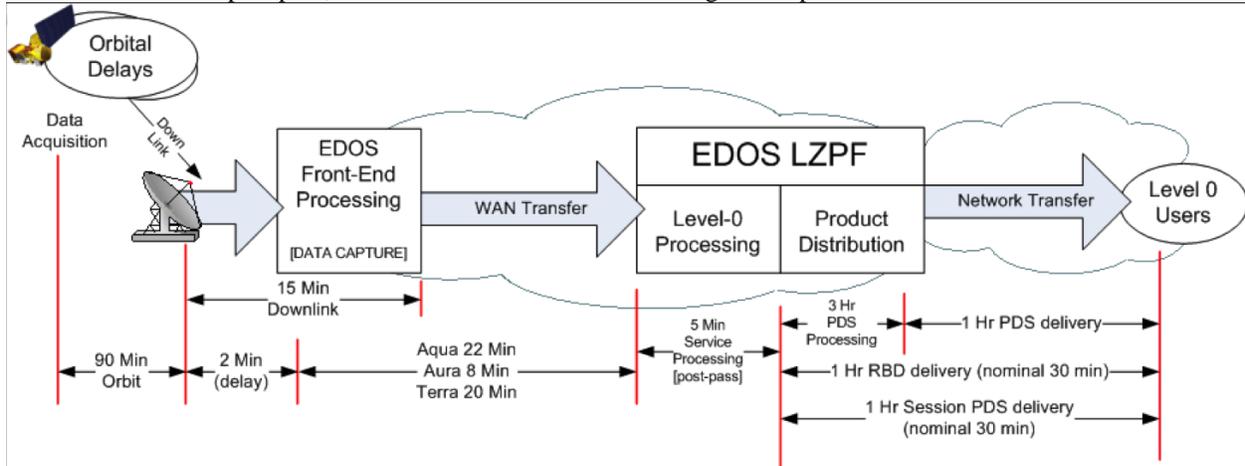


Figure 6. Breakdown of Latency from EDOS Data Capture to Level 0 Distribution

The primary EDOS contributor to latency is the transfer of the data across the high-rate WAN. Any decrease in WAN transfer time directly decreases the overall LANCE latency. In 2011, EDOS implemented two major latency enhancements focused on decreasing the time taken to transfer data to the LZPF:

- ◁ Removal of Reed-Solomon decoding bits from the transfer frame (128 bytes/frame)
 - Ø *Benefit:* 12% average decrease in WAN transfer time
- ◁ Implementation of lossless compression/decompression
 - Ø *Benefit:* 20% average decrease in WAN transfer time

The amount of compression varies by mission, depending on instruments and science payload content. For Aqua, the average WAN transfer time decreased from 23:08 to 14:24 minutes, a 37.8% overall improvement. These enhancements were achieved when the available bandwidth was limited on the private EOS network to about 50-60 Mbps, compared with the downlink rate of the primary EOS missions (Aqua, Aura, and Terra) of 150 Mbps.

EDOS continues to reduce latency times by using expanded network bandwidth, as it becomes available, with a goal of obtaining enough bandwidth to transfer the data from the ground station to LZPF in near real-time. EDOS has been successful in supporting the NRT requirements. The typical latency for most Level 0 products is under 2 hours, as shown in Figure 7:

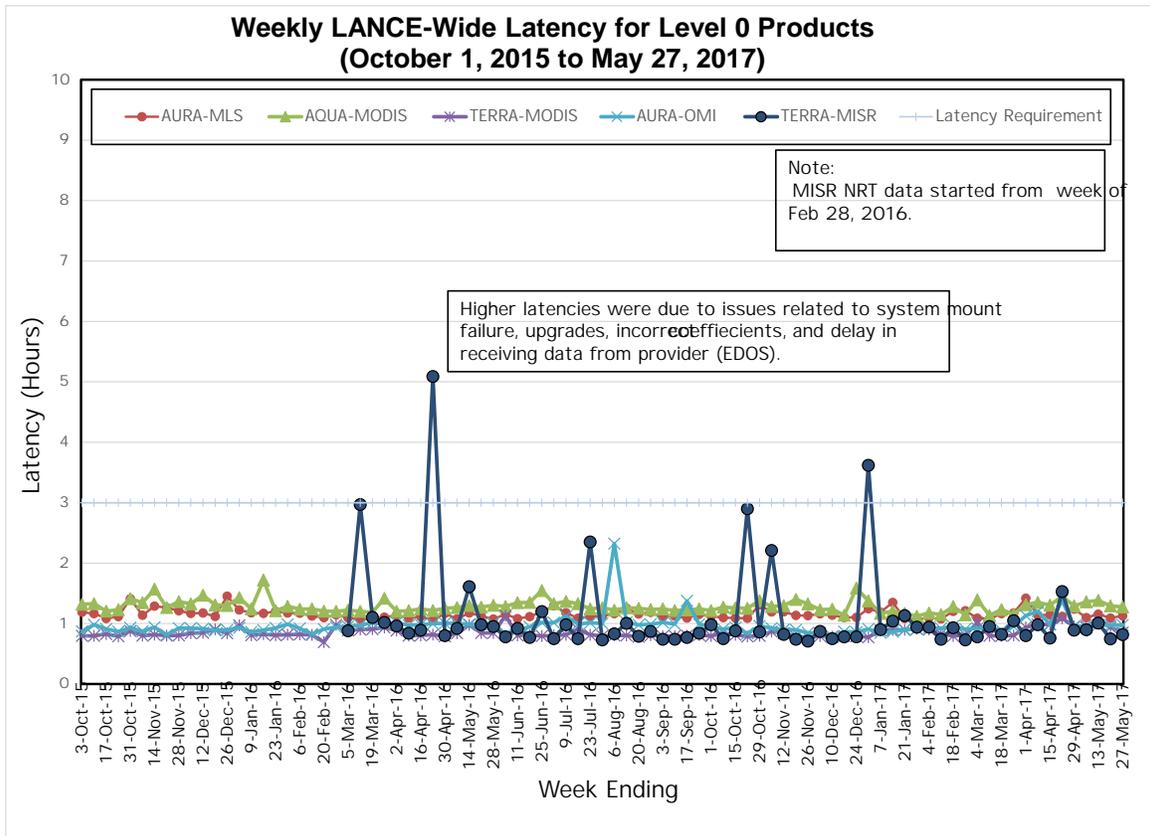


Figure 7. Weekly LANCE-Wide Latency for Level 0 Products

V. Open Networks and EDOS Hybrid Architecture

Expanding the bandwidth of the closed networks is a slow and costly process. But more bandwidth implies stations, where feasible. The addition of open (public) networks was only initially available at certain ground stations, namely Alaska Satellite Facility via Internet2, White Sands and Wallops via the NASA Corporate Network, and TrollSat via a satellite to Internet link (see Figure 8). In 2015 the Alaska fiber triangle was deployed, interconnecting all 3 Alaska ground stations supported by EDOS and extending the open and closed networks to all 3 ground stations. The open networks provided increased bandwidth with lower latency and at minimal increased cost.

Significant architectural changes were needed for EDOS to make use of the available open networks in addition to the existing closed network connections. Both ground station front-end processors (EBoxes) and the central site Level 0 processors (RBoxes) were located on the closed network. Some EBoxes at each ground station with an open side network connection needed to be moved to the open network since physical separation between open and closed networks is required. Also, all of the RBoxes in the central site needed to be moved from the closed network to the local open network to accept connections from both closed and open EBoxes. This new EDOS architecture connects to RBoxes now located on the open network. Figure 10 below shows the typical hybrid configuration of EDOS at the ground station.

Figure 10. Typical EDOS Ground Station High-Rate Hybrid Architecture

The hybrid LZPF architecture supports both the closed and open network ground station interfaces with no changes to Level 0 processing of the mission data received at the RBox. The open network connections do not replace the closed network connections; they augment the network configuration. Where both networks are available, the 2 network interfaces serve as backup for each other, and can even be used in parallel to augment WAN bandwidth for concurrent mission downlinks. Additionally, by using the high-rate matrix switch the same data can be routed to both closed and open EBox front-end systems (see Figure 10 above) and data can be routed in parallel to RBoxes on both the operational and backup LZPF systems. This is extremely useful for testing a new release on the backup system with live data, transitions from the operational system to the backup system, testing EBoxes at the ground station, and testing new antennas at a ground station.

An important additional benefit of the hybrid architecture development is the realization that EDOS can be deployed at new ground stations worldwide where high-rate open connections are available. The goal of the hybrid architecture was to increase bandwidth to transfer data to the LZPF at the same rate the data is captured at the ground station, but with an increased number of higher rate new missions the need for additional bandwidth continues to grow.

VI. Delay-Tolerant Protocol

The new open side network connections provided additional bandwidth and reduced latency for White Sands, Wallops and Alaska (see Figure 8). The round-trip delay and low BER on these networks was still small enough to continue to use TCP from these ground stations. From White Sands and Alaska, EDOS is able to keep up with single EOS missions at 150 Mbps downlink (132 Mbps without Reed-Solomon coding). However, the new open

side connection to TrollSat, Antarctica had a much longer delay due to a satellite link from Antarctica to Norway before being relayed to Goddard. This large round-trip time of 660 milliseconds between Goddard and TrollSat posed a big latency problem using the normal TCP protocol. Even with tuning the TCP window size, the rate never exceeded 10 Mbps out of the 50 Mbps allocation on the circuit.

EDOS researched available simple delay-tolerant protocols that overcame the delays inherent in TCP acknowledgements and chose UDT (UDP-based Data Transfer Protocol) as an optional TCP replacement suited for ground stations with long round trip times. UDT was selected because it is a stable, open source protocol that satisfied the primary need of being independent of network delay (see reference 5). UDT is available from SourceForge at udt.sourceforge.net. Even though it uses UDP for data transport, UDT functions similarly to TCP by being a connection-oriented, reliable, duplex, unicast data streaming protocol. Although it supports configurable congestion control, the default congestion control algorithm was chosen since congestion was not an issue from TrollSat. Reliability control is provided by sequencing and acknowledgment between the sender and receiver, in order to provide duplex data transfer. The receiver sends back acknowledgments and loss reports according to packet arrival and lost packets are retransmitted. UDT uses messaging between the sender and receiver to manage the UDT connection. Any EDOS front-end processor (EBox) can be configured to use TCP/IP or UDT whether it is on the open or closed network.

During initial testing of UDT (version 4.10), defaults were used for all configuration parameters except UDT packet size, the UDT window size (similar to a TCP window), and the Maximum rate the protocol will attempt to transfer (MaxBW). Test results over the open networks indicated that, in general, larger UDT packet size yielded better performance. The UDT window size as well as the Ethernet framing limit of the IPsec tunnel appeared to have little or no effect. The key to tuning UDT appeared to be selection of the correct UDT packet size (UDT_MSS), and setting the Max bandwidth (UDT_MAXBW) parameter to be somewhat below the expected link limit. Achieving peak performance must be done empirically since it is difficult to predict the exact parameter values where throughput will peak, on a particular link, in advance.

Without significant congestion and low BER, EDOS has been able to achieve throughput between 80 to 95% of the available bandwidth on the open networks with UDT through an IPsec tunnel. EDOS uses UDT exclusively from TrollSat on a 50 Mbps link (with no congestion) and routinely runs at about 48 Mbps (96% bandwidth utilization). A comparison of actual UDT and TCP data flows from TrollSat to EDOS LZPF at Goddard is shown below in Figures 11 and 12. Note the solid 48 Mbps output rate for the UDT data flow.

Figure 11. TrollSat WAN Performance with TCP