DSN Wide Area Network Architecture, Capacity and Performance

Timothy Pham, Roger Cortez, Jason Liao
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California, USA
e-mails: Timothy.Pham@jpl.nasa.gov, Roger.Cortez@jpl.nasa.gov, Jason.Liao@jpl.nasa.gov

***Abstract* - This paper discusses the architecture of the wide area network that connects key communications facilities within the National Aeronautic and Space Administration (NASA) Deep Space Network (DSN). Several considerations are given to the design of this wide area network to ensure a timely, reliable, and secure data delivery between the mission users and their spacecraft. The network star configuration simplifies data delivery to users and minimizes operational cost. The dual-path connections maximize the system reliability, with geographical diversity in data routing to avoid single point of failure. Data encryption enhances the protection of mission users' data. The system bandwidth is determined by balancing the needs to minimize the operating bandwidth cost and to have sufficient bandwidth to be able to deliver data to all users within the required. The DSN uses a class base weighted fair queuing (CBWFQ) method in its data delivery. This scheme guarantees a minimum bandwidth to each class of users and allows users to also access any unused bandwidth by other groups. The paper will also show the performance of system reliability and bandwidth margin.**

Keywords - DSN; network architecture; network design

I. INTRODUCTION

The Deep Space Network (DSN), an infrastructure of the National Aeronautics and Space Administration (NASA), is a global network that enables communications with far-flung spacecraft exploring the outer space. Its three main tracking facilities – Goldstone, Canberra and Madrid - are spread equidistance across the Earth, at both northern and southern hemisphere, to enable constant contact between Earth's mission controllers and their spacecraft. Antennas at each of these facilities enable communications with 35-plus spacecraft currently in operations. Data are exchanged in both directions - in the forward links where commands are sent to spacecraft and in the return links where spacecraft's housekeeping data and collected scientific data are sent back to Earth. In addition to spacecraft command and telemetry, radiometric measurements of spacecraft position and velocity, i.e., ranging and Doppler data also flow through this wide area network to mission navigation teams for orbit determination. Data collected at the tracking facilities also include observations/measurements for radio astronomy, planetary radar, and radio science. As shown in Figure 1, the three tracking facilities are connected to the Deep Space Operations Center (DSOC) at the Jet Propulsion Laboratory (JPL) in Pasadena, California via a wide area network, which is the focus of this paper. The DSOC facility in turn connects to mission operations centers (MOC) and science investigators.

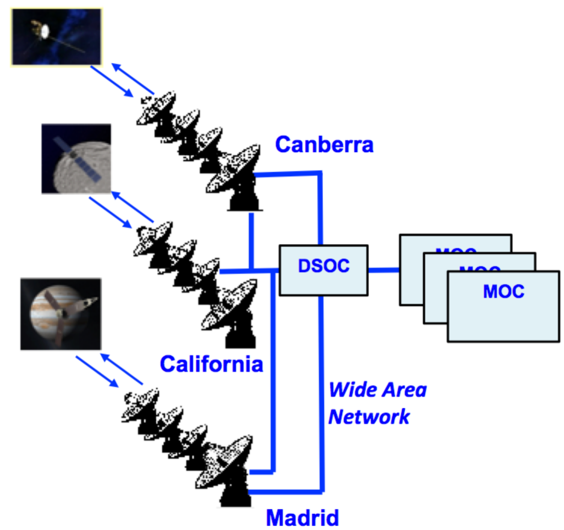

Figure 1. Topography of the DSN Wide area Network

Section II will discuss the architecture of the DSN wide area network. Several design considerations are important to maximize the system resiliency against possible outages and cyber security. The network capacity, driven by the needs of mission users, science users and DSN operations, are discussed in Section III. A scheme for fair allocation of bandwidth among users is described in Section IV. Section V provides a gleam of system performance in terms of bandwidth usage and margin, and connection reliability.

II. ARCHITECTURE

DSN engineers working with the NASA Communications Services Office (CSO) design the DSN wide area network architecture. The network connections are implemented, maintained and operated by a telecommunications company who has contract with the NASA Communications Service Office. The US-based prime contractor, in turn, works with domestic providers who have circuit access to Goldstone, and with international providers who have access to the other two overseas facilities at Madrid and Canberra. Past architecture relied on point-to-point dedicated circuits, based on Time Division Multiplex (TDM) technology. In the current architecture, the connections are done through the provider's IP VPN

 ISBN: 978-1-61208-545-6

(Internet-protocol virtual private network) solutions that run over their multi-protocol label switching (MPLS) backbone. Each site (DSOC, Canberra, Madrid and Goldstone) is connected to the carrier's MPLS backbone via dual access circuits. With all sites being connected to the MPLS backbone, virtual point-to-point circuits are created for connections between the three tracking facilities and DSOC.

This architecture offers a greater redundancy inherent in the network backbone, more flexibility to set virtual circuit for any-to-any node connection, no physical change required when there is a need to increase the capacity, and a lower cost for equivalent bandwidth compared to traditional TDM services. The drawbacks include the loss of end-to-end visibility and monitoring. The physical route across the MPLS backbone is not determinable, making it hard to verify the route diversity within the cloud. Note that although with the cloud, the realizable bandwidth is still limited to the physical capacity of the access circuits that connect the tracking stations to the cloud (i.e., last mile).

Star Configuration - The topology of the DSN wide area network resembles that of a star configuration where the three tracking facilities (Goldstone, Canberra, Madrid) are connected to a central DSOC node at the Jet Propulsion Laboratory, Pasadena, California. From DSOC, data are then delivered to external mission operations centers. This star configuration has both benefits and drawbacks. On the benefits, it provides a single-node connection to the mission users, regardless of where the tracking pass takes place. It simplifies the functions normally done at DSOC such as packet extraction and data interfacing with mission users. With a centralized configuration, the equipment only needs to be deployed once and operated by one team, reducing the hardware and operation costs. On the drawbacks, it makes the wide area network becomes a critical element to DSN operations and thus, requires funding to sustain it. Nowadays with Internet access being widely available and cheaply, a non-star configuration may offer a possible lower cost alternative. In such scenario, the DSN would deliver data to users directly from the tracking facilities. The users would connect to those tracking facilities. The cost of such connection would be born by the users, but not the DSN. The overall cost for NASA, however, may not change much since NASA pays for both mission operations cost and DSN operations cost. The second drawback, albeit minimal impact, is that data access and delivery are less direct. For example, a mission operation team for one of European missions, instead of being able to get to its mission data as soon as the data are received at Madrid, would need to wait for data to travel from Madrid to California before coming back to Europe. This delay impact is rather insignificant – an extra ~200 millisecond delay is quite small compare to the delay of several minutes/hours associated with signal travel in space before reaching Earth. System reliability is also a bit lower because there is more transferring nodes in the data delivery; however, in practice, the difference in reliability is quite small due to the built-in redundancy of the DSN wide area network, which is discussed next.

Redundancy - The tracking facilities are typically located in remote areas to minimize potential radio frequency interference (RFI) with terrestrial sources from nearby cities. As a result, the connections from the tracking site to the telecommunication provider's nearby point of presence are often of long distance, making it being both expensive and limited in choices of routing. In the past, even though the dual connections from the site to the nearby city were routed via two distinct fiber optic cables, both cables often ran through the same fiber bundle that extended over many tens of kilometers. There were instances where the roadwork or power crews accidentally cut both fiber optic links during construction and maintenance activities. Since these links are in remote area, recovery from outages often took many hours, and at times days. To avoid these problems, it is critical that the routing of signals from the tracking facility to DSOC is geographically distinct. The insulation between the two links starts at the tracking facilities where there are two equipment racks provided by vendor. The signals are routed via two nearby towns to two different metropolises, and in the case of Canberra and Madrid, traversing on two different undersea cables prior to reaching the United States. The VPN circuits arrive to JPL at two different entry points and terminate in two different buildings. All of these precautions help to ensure that even if there were a large-scale natural or man-made disruption affecting one of the cities or metropolis that one of the links traverses, the other link would remain operational. As shown in Figure 2, the Madrid tracking facility has one connection through the nearby village (Node 1) onto Madrid. The other connection is through a different village (Node 2) and onto Barcelona. Similarly, Canberra links are routed such that one going through Sydney and the other through Melbourne. The current dual paths from Goldstone have one common point of presence in a nearby town Victorville; however, with the effort spent over the past few years, another fully diverse route bypassing Victorville will become operational in 2017.

Between the two links, the prime path is normally the direct connection between Canberra/Madrid and DSOC. The backup path is routed via Goldstone. The primary driver for having the backup Canberra/Madrid connection via Goldstone is for the DSN Emergency Control Center. The ECC will be activated if DSOC is not available, and having direct links between Goldstone and Canberra/Madrid sites means that these two sites can maintain communication with the Emergency Control Center even if the DSOC node is down. As a result, the Goldstone-DSOC link is expected to carry Goldstone-own traffic and either Canberra or Madrid traffic in the event of failure of the prime connection at one of these two sites. As seen in Figure 2, the bandwidth capacity of Goldstone facility is greater than that of Canberra and Madrid. The capacity for Goldstone link is scoped to accommodate a failure in the prime link at either Canberra or Madrid, but not both. Such limitation is a tradeoff of bandwidth subscription cost versus the low probability of simultaneous failures at Canberra and Madrid.

Security - As indicated earlier, the information being sent to and received from spacecraft traverses through the provider's network backbone. It is important that the data

 ISBN: 978-1-61208-545-6

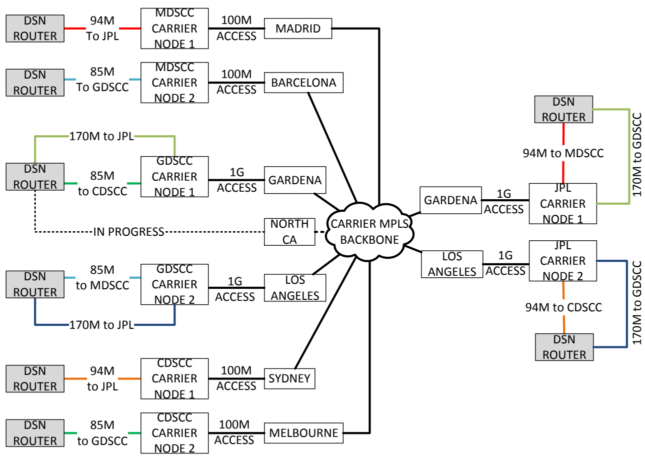

Figure 2. Dual path connections from tracking facilities to DSOC

confidentiality and integrity are not compromised. To further protect this information, beyond the vendor's IP VPN, encryption is applied to the data traverses over the DSN wide area network. IPSec tunnels are created between the sites, using the NSA Suite B Algorithms. In addition, the TCP maximum segment size (set to 1330 bytes) is being interjected by DSN routers to ensure that all frames for TCP connections stay under 1500 bytes to prevent fragmentation. This was required to accommodate the IPSec overhead. The encryption nominally adds about 6% bandwidth overhead.

III. Bandwidth Capacity

The consideration on bandwidth capacity is determined by two factors: sufficiently meet the need of users and incur minimum cost. Because of the remote location of the tracking facilities and the high availability requirement (99.95%) to meet mission critical need, the bandwidth cost for DSN wide area network is much higher than what is typically encountered in commercial business or residential rates. As such, it is important to correctly size the required bandwidth so that NASA would not have to pay the extra cost for idle bandwidth. Reference [2] provided a detailed description on the method of estimating the required bandwidth. In this paper, the estimated bandwidth reflects the greater demand from some recent missions with higher data rates.

There are many users of the DSN wide area network. In general, they can be divided into three groups – mission users, science users and DSN operations.

Mission Users – Telemetry data return from missions' spacecraft drives the required bandwidth because of the low latency requirement for proper mission operations. Command data to spacecraft, which flow in reverse direction from telemetry, and radiometric data needed for mission navigation consume much lower bandwidth. Most missions are tracked daily; thus, data that are downlinked on a given day need to be delivered within 24 hours so that the DSN would be cleared and ready for next day track. For some missions, the required latency is much shorter – in a few hours – so that mission operations team can determine the activity for the next day or for the follow-on tracking pass at the next tracking facility. Prior to 2017, the highest data rate among the current missions being supported by the DSN is 6 Mbps from the Mars Reconnaissance Orbiter. The aggregated downlink data rate for all missions tracked by various antennas within a tracking facility is typically 8 Mbps or less. In late 2017, the data rate will be significantly increased with the upcoming Terrestrial Exoplanet Survey Satellite (TESS) mission at 125 Mbps [3]. TESS tracking schedule is however infrequent, occurring roughly once every two weeks. The tracking passes for TESS is also rather short, ~2.5 hrs a day. The data collected over 2.5 hrs of TESS downlink pass can be delivered over a 24 hrs window at a much lower rate of 13 Mbps. In late 2018, another high rate mission - the James Webb Space Telescope (JWST) - will be launched. Its 28 Mbps downlink, although smaller compared to TESS, require a real-time data delivery [4]. The bandwidth driven by JWST is 28 Mbps. Looking forward into the future years in early or mid 2020's, there are likely missions with data rate upward of 150 Mbps, and even 300 Mbps. These missions will then drive a need for greater capacity in the DSN wide area network.

Science Users - Besides supporting standard telemetry, tracking and command functions in communications with scientific spacecraft, the DSN also serves as a science instrument platform for planetary radar science, radio astronomy, space geodesy via Very Long Baseline Interferometry (VLBI), and radio science. These observations typically require high-rate sampling of the wide-band received signal to extract information in the received signal spectrum. As a result, these science observations generate very high data volume. Fortunately, these observations are done less frequent (every few days or few weeks), the latency requirement is not as stringent as telemetry data. Currently, radio science and radar science recordings could be relayed via the network, but data from VLBI and radar science observations need to be delivered off-line via shipping of recorded disks. A better operation mode is to deliver data via the network. It would increase the timeliness of data access and reduce the burden of manual shipping of recorded data. In the current capacity, 15 Mbps is booked as reserved for science data transfer.

DSN Operations – Beside data delivery to missions and science users, there are other traffic flows necessary for DSN operations. Monitor data reflecting the performance and configuration of the ground equipment, and conditions of the received signals need to be delivered to the DSN Operations at DSOC and passed onto mission control centers.

Up to now, each tracking facility is responsible for the operations of equipment at its site. Toward the end of 2017 the DSN will be moving to a new operational concept called follow-the-sun operation in which one tracking facility (Goldstone, Canberra, or Madrid) would take control of the entire DSN operations on a rotating basis [5]. To support such operation, video streams of antennas that used to be flown internal within the facility will need to be routed through the wide area network to another facility in order to provide assurance of proper operations of the remote antennas thousands of kilometers away. This monitoring traffic is estimated at 20 Mbps.

In addition to monitor data and antenna video streams, a small portion of the bandwidth is needed for voice-over-IP traffic that is essential to maintain communications among operation teams at different facilities, computer network monitoring (via simple network management protocol), and occasional remote logins for system testing and diagnostic. This type of traffic is estimated about 5 Mbps.

In total, the near-term bandwidth need is estimated to be ~100 Mbps. This capacity has been recently implemented; enables the DSN to support upcoming TESS and JWST missions.

IV. Bandwidth Allocation and Access

Bandwidth allocation for DSN WAN traffic is done with a class base weighted fair queuing (CBWFQ) method [1]. Of the total bandwidth available, each class of users is given a minimum guaranteed allocation. The minimum allocation guarantees that the users within a specific class will have that bandwidth available to them, should they need it. If data flows within one class are below the allocation, the unused bandwidth then becomes available to users in other classes. Conversely, if the bandwidth demand within one class is more than the minimum allocation and there is unused bandwidth elsewhere, the users would be able to access it. Certain data traffics that are sensitive to delay (e.g., voice over IP) could be placed in a low-latency queue, which lets the data be sent out first.

Out of the current 100 Mbps capacity, telemetry allocation is guaranteed about 60 Mbps. Science users are given about 15 Mbps and DSN operations about 25 Mbps. Since science users typically desire to have access to greater bandwidth, an option to maximize the WAN resource is to flow science data via the available backup link that is normally idle.

V. Performance

To ensure that system achieves its performance as required, monitoring capability is built into the system to collect necessary metrics for assessing the connectivity and bandwidth utilization.

Figure 3 shows how network monitoring is done in the DSN. As telemetry and science data comes to JPL from Goldstone, Canberra and Madrid, a copy of data packet is routed to the Network Monitor server for traffic loading analysis. Based on the frequency and volume of incoming data packets for each class of traffic, the Network Monitor creates the network loading profile.

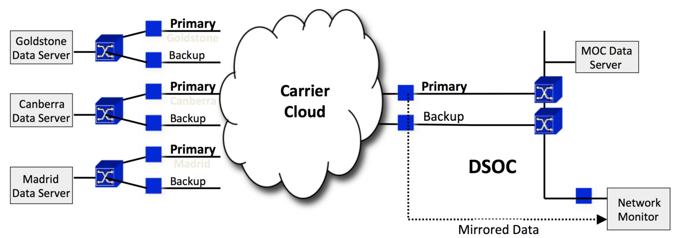

Figure 3. Peak flow of different data traffic

To monitor the status of link connection, the Network Monitor server pings the routers at Goldstone, Canberra and Madrid every 15 second. If there were a failure in the current active signal path, the router would automatically switch the data routing to the redundant path. Thus, the ping over the failed link would not be successful. Figure 4 shows a sample of IP connectivity over one month period for one of the Goldstone-JPL links. The blue line shows the actual availability, which is nearly 100% available. The red line shows the required 99.5% availability. Because of the dual-path redundancy, even if one of the paths is down, the overall service availability remains high at 100% level.

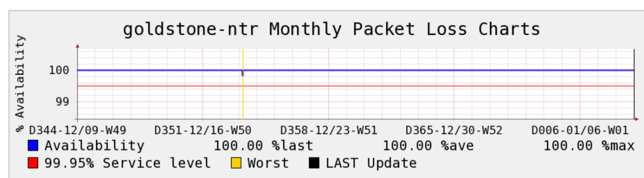

Figure 4. Service availability of a sample Goldstone link

On bandwidth utilization, Figure 5 shows the current average loading of ~15 Mbps (prior to the launch of the high-rate TESS and JWST missions). That represents a 15% usage or 85% margin of the link capacity. We expect however that the bandwidth usage will be significantly increased toward the end of 2017 when TESS mission is launched, followed by another increase with JWST launch in 2018.

One may also notice in Fig. 5 that different classes of traffic have different peak flow. Telemetry data return, shown in blue, is typically constrained by the spacecraft downlink data rates. Currently, the maximum downlink rate for a single mission is 6 Mbps, which is quite small compared to the link capacity. Toward the end of 2017, TESS will enter its mission operational phase and will downlink data at 125 Mbps. At that rate, telemetry data will flow at the maximum link capacity during the tracks and a few hours at the end of the pass. TESS tracking occurs infrequent, typically every 14 days. In between the two TESS passes, the maximum telemetry data flow - including JWST -would be at a lower peak, around 30 – 35 Mbps.

 ISBN: 978-1-61208-545-6

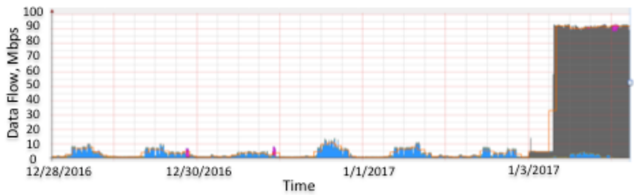

Figure 5. Peak flow of different data traffic

Science data, shown in dark gray curve, always reaches the maximum bandwidth of the wide area network. It is because science observations for planetary radar and radio astronomy studies typically involve the high-rate recording of wideband RF spectrum. The sampling data usually exceeds the capacity of the WAN.

VI. Conclusion

In summary, this paper discusses the characteristics of the wide area network of the Deep Space Network. Delivery of spacecraft telemetry to mission users is done via a single node at JPL. Features such as geographically distinct routing between the prime and backup path help to avoid a single point of failure and improve system reliability. Encryption helps to protect the information integrity, particularly relevant to the command data that are critical to spacecraft operations. The capacity of the WAN is set based on the balance between being able to deliver data to users within the required latency and minimizing the WAN cost. Bandwidth is allocated among users via class base weighted fair queuing scheme. The system has been recently upgraded to 100 Mbps capacity in preparation for upcoming support to a few high rate missions such as TESS and JWST. The current average loading at 15 Mbps shows that there is good margin of system capacity. The margin reflects DSN readiness to support upcoming missions in the next few years.

Acknowledgment

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] S. Floyd and V. Jacobson, “Link-sharing and Resource Management Models for Packet Networks,” IEEE/ACM Transactions on Networking, Vol. 3 No. 4, August 1995, pp. 365-386.

[2] T. Pham, “A Cost Effective Modelling of Near-Term Mission Bandwidth Demand”, 6th International Symposium on Reducing Cost of Space Ground Communications Systems Operations Symposium, Darmstadt, Germany, September 2005

[3] TESS – Transiting Exoplanet Survey Satellite, NASAfacts, National Aeronautics and Space Administration, NASAfacts, https://tess.gsfc.nasa.gov/documents/tess_factsheet_06-05-15.pdf, accessed March 15, 2017

[4] JWST Mission Operations Concept Document, Space Telescope Science Institute, January 13, 2006, http://spacese.spacegrant.org/JWST_Mission_Operations_Concept_Document.pdf , accessed March 15, 2017

[5] M. Johnston, M. Levesque, S. Malhotra, D. Tran, R. Verma, S. Zendejas, “DSN Automation Improvements in the Follow-the-Sun Era”, International Joint Conference on Artificial Intelligence Workshop on Artificial Intelligence in Space (AI Space, IJCAI 2015). Buenos Aires, Argentina. July 2015