Ensuring Flexibility and Security in SDN-Based Spacecraft Communication Networks through Risk Assessment

Dylan Z. Baker
Department of Electrical and Computer Engineering
University of Massachusetts Dartmouth
North Dartmouth, MA, USA
dbaker1@umassd.edu

Dr. Hong Liu
Department of Electrical and Computer Engineering
University of Massachusetts Dartmouth
North Dartmouth, MA, USA
hliu@umassd.edu

Christopher Roberts
Exploration & Space Communications Projects Division
NASA Goddard Space Flight Center
Greenbelt, MD, USA
christopher.j.roberts@nasa.gov

***Abstract*— Software-defined networking (SDN) has enabled elastic networking and resource distribution in cloud computing. The centralization and separation of the Control Plane also offers a high degree of network configurability and management, which can be used to mitigate and manage threats to the network. Space communication networks have historically been restricted and circuit switching in these networks has been a manual process. This study evaluates the potential role of SDN in space communication networks from a networking security standpoint. The evaluation covers the networking security needs of spacecraft missions and their associated assets. The results from the evaluation lead to a risk assessment that identifies vulnerabilities in an SDN-based communications architecture. Security challenges introduced into the network from integrating SDN are also considered. A risk register summarizes the severity of the attack outcomes, as well as occurrence likelihood. The study identifies Denial-of-Service (DoS) attacks as a new threat (presently unmitigated by existing security controls) that would be prevalent in an SDN-based space communication environment. A Mininet-based emulation testbed is built to demonstrate the susceptibility of spacecraft flight software to a flooding DoS attack when on an interconnected SDN-managed network. This type of attack would be highly consequential to mission assets, and therefore SDN-based space communications would need to be resilient to such attacks. Future work will need to be performed to fully characterize DoS attack methods that can apply to the space communication scenario, as well as to devise a comprehensive DoS-resilient solution.**

Keywords—Software-Defined Networking (SDN), network security, space communications, OpenFlow, Ryu, Denial-of-Service attacks, spoofing attacks.

I. INTRODUCTION

Software-Defined Networking (SDN) has been a key driving technology behind cloud computing's growth in the past decade. SDN enables networks that are more scalable than traditional ones and facilitates the centralization of network management/configuration functions. It achieves this by decoupling the network's Data/Forwarding Plane (which forwards network traffic) from the Control Plane (which manages the network's configuration) [1].

The architectural advantages of Software Defined Networking could be applied to space communication networks. These networks have traditionally used pre-planned circuit switching for the transfer of spacecraft telemetry and commands [2]. This manual control of the network ensures reliability and security, although it hinders flexibility, scalability and on-demand routing. SDN has the potential to centralize spacecraft and ground support equipment (GSE) management. Doing so would give network/spacecraft operators sufficient control over their assets, while still supporting dynamic network configuration and routing.

Opening/consolidating networks for multiple spacecraft and enabling dynamic rerouting of data inevitably presents a greater attack surface. Various SDN protocols, as well as controller and network device implementations, provide security solutions. Despite this, there are both security challenges inherent to SDN and security challenges to implementing SDN with spacecraft systems.

II. BACKGROUND

A. Software Defined Networking

Networks can be described as containing different types of "planes". Two planes of particular interest are the Data Plane and the Control Plane. The Data Plane handles the forwarding of traffic. The Control Plane is used to define routing logic and other routing configurations. The Control and Data Planes are separated in Software Defined Networks in order to support the distribution of network configurations across devices. The third plane in SDN is the Application Plane, which communicates with the Control Plane for network configuration [3].

According to the Open Networking Foundation [3], an SDN includes the following high-level components and interfaces:

978-1-7281-5092-5/19/$31.00 ©2019 IEEE

- SDN Controller - This exists in the Control Plane, which is the middle layer between the Application and Data Planes. Logic is performed in the controller for network flow/route management between multiple network devices. A common interface is provided by the controller to SDN applications. The Control Plane manages network policy and monitors the network's performance.
- Network Devices – Physical networking devices or Network Elements (NEs) compose the Data Plane, which sits at the lowest logical level. The responsibility of the Data Plane is to hold network device configurations.
- SDN Applications – The Application Plane consists of a series of network applications which interface with the SDN controller. This top-level plane abstracts away the precise details of the network devices by the common (Northbound) interface.
- Northbound Interfaces - These are the interfaces between the Application and Control Planes.
- Southbound Interface/Control-Data-Plane Interface (CDPI) - This is the interface between the Data and Control Planes.

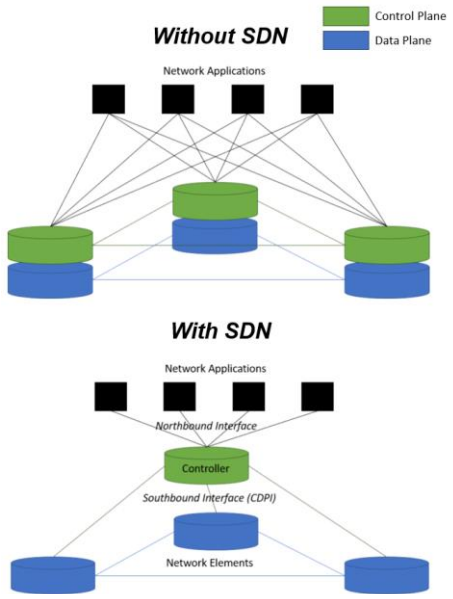

Fig. 1. Plane interface comparison between traditional networks and SDN.

B. *Space Networking Environment*

Space communication systems typically rely on Consultative Committee for Space Data Systems (CCSDS) protocols. CCSDS protocols are maintained by an international committee, with members from the world's major space agencies (including NASA, ESA, Roscosmos and JAXA). Commonly used CCSDS protocols include Space Packet Protocol [4] (which transmits data between logical data paths), AOS Space Data Link Protocol [5] (a link-layer protocol that manages connections between two nodes), and CCSDS File Delivery Protocol/CFDP [6] (which transmits files between nodes).

Spacecraft typically perform line-of-sight communication with ground stations and/or relay satellites. The ground stations exist as part of networks, such as the Deep Space Network (DSN), Near Earth Network (NEN) and Estrack. Tracking and Data Relay Satellites (TDRS) are an example of relay satellites, which forward data to the ground when line-of-sight isn't available. Missions typically schedule utilization of ground stations' time slots.

The ground segment often includes operation centers. A Mission Operations Center (MOC) is where a spacecraft is commanded, as well as where its health/status-related telemetry is tracked. A Science Operations Center (SOC) handles large mission-relevant science data and stores/distributes it. MOCs and SOCs communicate with spacecraft via terrestrial network connections and the aforementioned ground stations/relay satellites.

III. Problem Statement

Space communications has long centered on circuit-switched technology (with or without a bent-pipe transponder) to facilitate wireless communication between spacecraft and earth-based ground stations. A bent-pipe transponder simply redirects and amplifies a spacecraft or earth-based ground station's signal in situations with poor line-of-sight. Once on the ground, data is routed via TCP/IP terrestrial networks. This lack of true networking functionality in the space segment prevents numerous types of attacks associated with networking, although it restricts spacecraft communication to designated ground station antennas within allocated timespans.

NASA's Space Communication and Navigation (SCaN) program is seeking to combine three formerly separate space network service providers (Space Network, Near-Earth Network and Deep Space Network) into a single network in order to support spacecraft management via any of the networks and to bring down operational costs [7]. This integrated network will need to support a diverse set of nodes. These nodes, as indicated in Fig. 2, can include stratospheric balloons/High Altitude Platform Stations (HAPS), high and low-earth orbit platforms (including smallsats, communication satellites and space stations) and missions to both nearby and deep-space destinations.

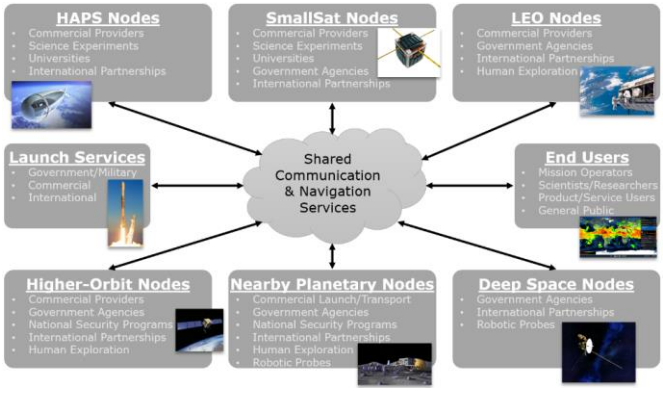

Fig. 2. Nodes/missions to access communication/navigation services, present and near-future.

Additionally, satellites, satellite constellations (clusters of multiple satellites), crewed vehicles, probes and other spacecraft are becoming more sophisticated and numerous (due to trends such as commercialization and miniaturization/smallsats). These advancements are both generating higher data requirements and requiring communication with more individual spacecraft. Therefore, the traditional direct/circuit-switched paradigm may no longer be feasible with some upcoming missions.

A standard TCP/IP network architecture is distributed in contrast to the centralized architecture of current space communication networks. Software-Defined Networking's separation of the control and data planes enables the type of centralized elasticity that will be required for spacecraft missions of the future to share common communication and navigation resources. This architecture has key advantages from a security standpoint (such as centralized control and flow sampling), although it may introduce security challenges.

IV. State of the Art

The OpenFlow SDN protocol has been widely-used in terrestrial networks across industry. Various software implementations exist, such as OpenDaylight and Ryu [14]. Alphabet's X Development, LLC has developed Temporospatial SDN (TS-SDN) technology for Project Loon (a project to provide internet access via stratospheric balloons). TS-SDN optimizes network topology and routing in cases where the physical networking infrastructure is moving in significant, but predictable ways (such as with atmospheric balloons and satellites/other spacecraft) [8]. Other proposed SDN architectures for space communications include SERvICE [9] and SDN-SAT [10].

Software-Defined Networking has been demonstrated to offer a number of security-focused advantages and capabilities [11]. Centralized network programmability (via the control plane) can aid administrators in deploying security policies and services across large networks. It also opens the opportunity for performing cyber forensics on a running network. The control plane can be configured to sample flows and alter data plane configurations in real time based on perceived threats. Research has also been performed on using fuzzy logic in SDN IPS software to detect anomalies to adapt network configurations.

SDN also introduces some security challenges. Due to its logically-centralized nature, the control plane can act as a single point-of-failure for an entire network. The controller itself can be a target for a Denial-of-Service (DoS) attack, which would render the controller unable to set flows correctly. AVANT-GUARD [12] was designed to throttle data being sent to the controller in order to prevent flooding. Another solution, CPRecovery, is designed to manage DoS attacks through controller failover [13].

The programmability of Software-Defined Networks can also be utilized by an attacker to adjust the flow of data if they are able to gain privileged access to the controller (potentially compromising both data availability and confidentiality, as well as endpoint integrity). Additionally, they can also redirect flows to flood devices with data. Proposed solutions are discussed in [11], including trust systems and role-based authentication [14].

Proper SDN controller and network configuration are critical, as a single misconfiguration via the Application Plane can result in an entire network being compromised. Contradicting security policies in large distributed networks can also potentially result in security holes. The handling of security policy conflicts and order-of-precedence/propagation across large networks can potentially result in weaker protections than certain devices need. Numerous solutions are discussed in [11] that are designed for detecting and handling configuration problems.

V. SDN-Based Space Communication Network

As part of this study, a small-scale space communication network was emulated with off-the-shelf SDN hardware/software (OpenFlow) for standard spacecraft flight and ground infrastructure. A terrestrial network (Fig. 3) includes three ground stations: Mission Operations Center (MOC), Science Operations Center (SOC) and Network Operations Center (NOC).

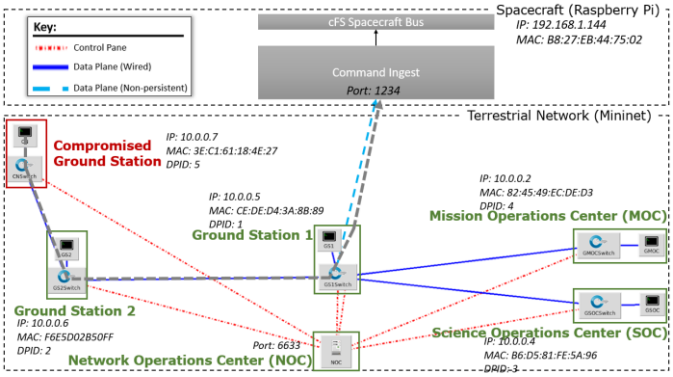

Fig. 3. Spacecraft network emulation using Mininet, OpenFlow and NASA cFS.

The NOC is an OpenFlow controller that sends control messages to each of the ground nodes and to the spacecraft segment. Each of the other nodes consist of a virtualized OpenFlow switch and a virtualized host (computer). The virtualized hosts all run an off-the-shelf Linux environment. The terrestrial network is tested in the Mininet SDN emulation environment. The spacecraft node runs on a Raspberry Pi and uses NASA's open-source Core Flight System (cFS) spacecraft software bus.

VI. Vulnerability Study

A risk analysis was conducted which identifies assets in an integrated SDN-based network that would be under threat, along with threat management and mitigation techniques. Fig. 4-7 contain attack trees which illustrate types of attacks that can be carried out on such a network. This study was performed as part of [15].

A. Confidentiality

Fig. 4 focuses on attacks to confidentiality. Data encryption should be practiced on any network to prevent interception/interpretation by an illegitimate recipient. There are two primary methods of compromising the network's data confidentiality: data redirection and unauthorized data access. A

spacecraft mission could involve confidential data being telemetered by a spacecraft. An attacker could redirect data by gaining Control Plane access (compromising endpoint integrity on the Control Plane) and using that access to modify data flow rules. They would be able to pose as the intended recipient (including during key exchange). Alternatively, an attacker can also compromise the endpoint integrity of one or more ground nodes and obtain trusted user credentials, which would give them unauthorized data access.

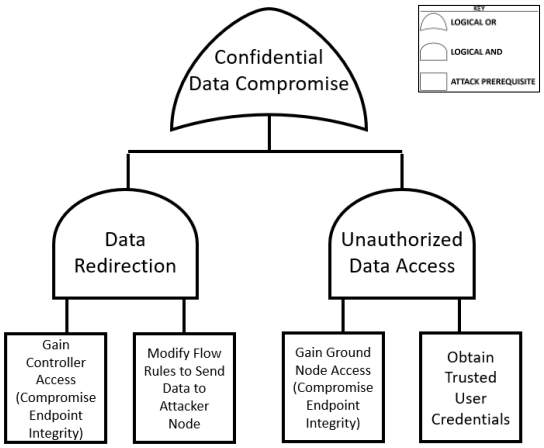

Fig. 4. Attack tree demonstrating the C.I.A security goal of confidentiality assurance.

B. Data Integrity

Compromises to data integrity are shown in Fig. 5. Attacks can be performed on mission telemetry, spacecraft commands and SDN nodes. All of these attacks would require a compromise to endpoint integrity (access to the network or controller).

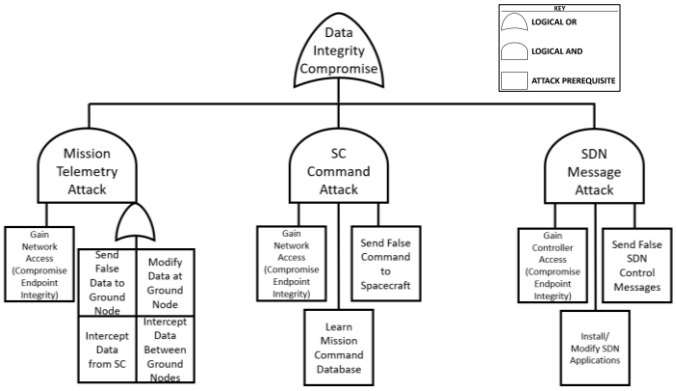

Fig. 5. Attack tree demonstrating the C.I.A security goal of data integrity assurance.

The first type of attack can be carried out on mission telemetry. With access to the network (an endpoint integrity compromise), the attacker could have the ability to either send false data to a ground node, modify data at a ground node, intercept data being sent from the spacecraft or intercept data being sent between multiple ground nodes.

Potentially even more damaging could be a spacecraft command attack. With unauthorized access to the network, an attacker could send a false command to the spacecraft if they had knowledge of that mission's command database. This could result in an attacker gaining control of the spacecraft.

The third type of data integrity compromise is an attack on the SDN itself. By compromising the endpoint integrity of the Control Plane, an attacker could send false SDN control messages if they have the correct SDN applications installed/configured. This could give the attacker the ability to route data incorrectly, or not route data at all.

C. Endpoint Integrity

An endpoint integrity compromise, as shown in Fig. 6, can impact a ground node, the spacecraft itself and an SDN controller. The three endpoint integrity compromise scenarios are a ground node compromise, spacecraft compromise and SDN controller compromise. In each of the scenarios, the attacks can be conducted through man-in-the middle attacks, session hijacking, modifying trust (through SDN flow rules) or spoofing of a particular endpoint (ground node, spacecraft or SDN controller). With this access, mission telemetry could be faked, modified and/or intercepted. If the attacker had knowledge of that spacecraft's command database, they could send bad commands to the spacecraft (which, in some cases, can cause the spacecraft to damage itself). Access to the controller would enable the attacker to send unwanted control messages to reconfigure the network.

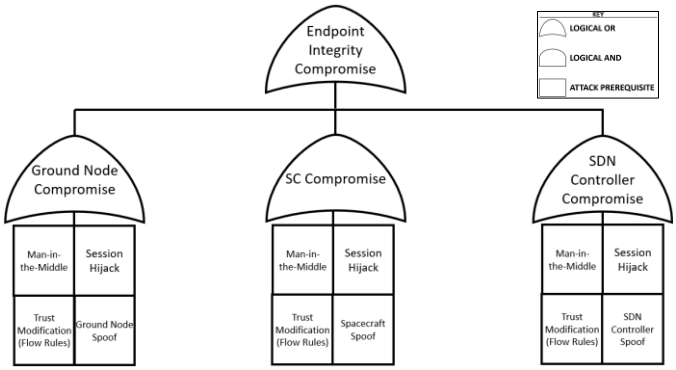

Fig. 6. Attack tree demonstrating the C.I.A security goal of endpoint integrity assurance.

D. Availability

An availability compromise can be carried out on a spacecraft, ground node and/or SDN controller (as shown in Fig. 7). The three availability disruption scenarios correspond to the three types of targets: spacecraft, ground node and SDN controller.

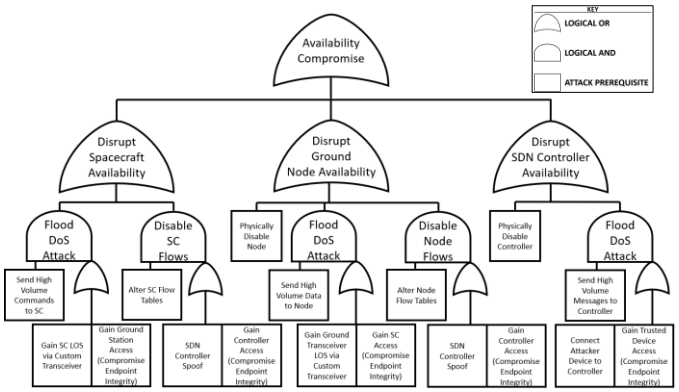

Fig. 7. Attack tree demonstrating the C.I.A security goal of availability assurance.

Spacecraft availability can be compromised by performing a flooding Denial-of-Service (DoS) attack or disabling network flows to/from the spacecraft. The DoS attack can be performed by either compromising ground station endpoint integrity or gaining physical line-of-sight with the spacecraft and sending a high volume of commands (overwhelming the spacecraft). Flows can be disrupted by either gaining control of or spoofing a controller to alter flow tables.

Ground node availability can be disrupted in similar ways to spacecraft availability: flooding DoS and flow disabling. Physical line-of-sight or compromising spacecraft integrity can be used to flood ground nodes with data. Flow tables can again be altered by controlling or spoofing controllers. In addition to these attacks, the ground node can also be physically disabled (i.e. by removing power or disconnecting cables).

The SDN controller/Control Plane is subject to similar types of availability attacks. With physical access, the controller can simply be disabled. Assuming no direct physical access, a flooding DoS attack can still be performed on the controller. If an attacker is able to connect a device to the network or control an existing device, they can send a high volume of messages to overwhelm the Control Plane.

VII. Results

The risk register in Table 1 was generated to classify the various types of risks to spacecraft and supporting assets in the SDN scenario. The risk register is based on the ISO 27000 Series of Standards on IT Security Techniques. The ISO 27000 Series defines Information Security Management Systems (ISMS) and provides recommendations for the identification, classification and management of information security risks [16]. The risks in the table are derived from the attack trees in Fig. 4-7. The table is sorted by descending risk priority.

TABLE I. Risk Register [15]

Asset	Threat/ Vulnerability	Existing Controls	Likelihood	Consequence	Level of Risk	Risk Priority
Reliability, availability, integrity of spacecraft	Attacks/ errors affecting spacecraft (i.e. DoS)	Space Data Link security; direct connection; command verification	Rare	Catastrophic/ Doomsday[a]	Extreme	1
Integrity and availability of ground nodes	Attacks/ errors affecting ground nodes	Space Data Link security	Unlikely	Moderate	Medium	2
Confidentiality of spacecraft telemetry/ commands	Interception of telemetry or commands	Data encryption	Unlikely	Moderate[b]	Medium	3
Integrity of spacecraft commands	Corruption or loss of command data	Error Detection & Correction codes	Possible[c]	Minor	Medium	4
Integrity of spacecraft telemetry	Corruption or loss of telemetry data	Error Detection & Correction Codes	Possible[c]	Minor	Medium	5
Integrity and availability of SDN controller	Attacks/ errors affecting controller; corruption/ loss of SDN control messages	Configurat-ion; controller authenticat-ion	Possible	Moderate	High	6

[a] Such an attack can result in mission failure, and in the case of crewed missions, potential loss of life.
[b] The severity is dependent on the mission type (a breach of classified data is significantly more severe).
[c] Corruption is likely to occur due to Single-Event Upsets (SEUs) experienced from ionizing particles in space, although EDAC & retransmission make data loss less likely.

Table 1 is broken down into 7 columns: Asset, Threat/Vulnerability, Existing Controls, Likelihood, Consequence, Level of Risk and Risk Priority. The "asset" defines the physical item, service or attribute that is being protected. "Threat/vulnerability" identifies a threat to that asset, while "Existing controls" identifies the controls (mechanisms) that are already in place to protect the asset from the identified threat. "Likelihood" characterizes how likely that compromise is to occur with existing security controls and "consequence" reflects the impact severity of the compromise occurring. "Level of risk" combines the "likelihood" and "consequence" fields to determine how much of a risk is posed by the threat to this asset. The "risk priority" column ranks the risks based upon the importance of addressing them.

The risks rated at a level of "Medium" are manageable with the current controls/best practices described under the "Existing Controls" column. Two risks (the risks to spacecraft reliability/availability and SDN controller integrity/availability) have risk levels beyond medium. Both risks are new in an integrated SDN-based space communications network.

A failure in spacecraft reliability/integrity, while rare, would have a catastrophic impact (end-of-mission/potential loss-of-life). Therefore, a security implementation plan for the spacecraft needs to center not only around lowering the likelihood of an attack, but also managing the occurrence of such an attack. The spacecraft and network need be able to identify and react to a DoS attack. SDN flow sampling can be used to identify DoS attempts in the network and the controller can react by limiting/removing flows along the attacker's route. Throttling can also be applied on the ground to prevent an excessive volume of data from reaching the spacecraft.

A loss of either the integrity or availability of the SDN controller could severely impact network operations. An integrity compromise can result in malicious Control Plane

978-1-7281-5092-5/19/$31.00 ©2019 IEEE

messages being sent to network nodes, while loss of availability effectively removes the control Plane.

VIII. Conclusion

The increased complexity of new spacecraft missions makes a true network architecture more necessary. While networks are inherently less secure than direct physical connections, SDN offers a flexible centralized solution that can be secure. The two greatest risks to be addressed are those to the reliability/availability of a spacecraft, and the integrity and availability of the SDN controller.

On a high-value asset, like the spacecraft (as well as the human lives that can potentially be dependent on it), threat mitigation alone is insufficient. While controls to reduce the likelihood of attack should be put in place to the maximum extent possible, the system (spacecraft and network) should be designed to gracefully handle an eventual attack. One single incident (no matter how rare) can have catastrophic consequences. Meanwhile, the threats to a space network's SDN controller can be mitigated with existing solutions. Due to the nature of implementing SDN into a space networking environment (most of the network's topology is unaffected), existing controls already sufficiently address threats in existing parts of the network.

Future work will need to be conducted to identify how, when and where to limit data flows in a space network in order to prevent entry into a denial-of-service state. SDN offers the potential to solve the DoS threat by both being able to centrally identify active attacks on the network (and react by isolating the attacker) and by being able to configure network throttling in a way that will prevent entry into a DoS state. Throttling would be necessary because flow sampling alone cannot guarantee that a spacecraft won't enter a DoS state, as heavy data volumes can disrupt the messages used to conduct flow sampling. Future work can also involve a formal vulnerability study using the STRIDE model with tools such as Microsoft's Security Development Lifecycle (SDL) Threat Modeling Tool.

References

[1] I. Ahmad, S. Namal, M. Ylianttila and A. Gurtov, "Security in Software Defined Networks: A Survey," IEEE Communication Surveys & Tutorials, vol. 17, no. 4, pp. 2317-2346, 2015.

[2] M. Sanchez, D. Selva, B. Cameron, E. Crawley, A. Seas and B. Seery, "Exploring the Architectural Trade Space of NASAs Space Communication and Navigation Program," in IEEE Aerospace Conference, Big Sky MT, 2013.

[3] Open Networking Foundation (ONF), "SDN Architecture Overview," Open Networking Foundation (ONF), 2013.

[4] Consultative Committee for Space Data Systems, "Space Packet Protocol Blue Book," September 2003. [Online]. Available: https://public.ccsds.org/Pubs/133x0b1c2.pdf. [Accessed 11 May 2019].

[5] Consultative Committee for Space Data Systems, "AOS Space Data Link Protocol Recommended Standard," September 2015. [Online]. Available: https://public.ccsds.org/Pubs/732x0b3e1.pdf. [Accessed 11 May 2019].

[6] Consultative Committee for Space Data Systems, "Recommendation for Space Data System Standards: CCSDS File Delivery Protocol (CFDP)," January 2007. [Online]. Available: https://public.ccsds.org/Pubs/727x0b4.pdf. [Accessed 11 May 2019].

[7] J. M. Reinert and P. Barnes, "Challenges of integrating NASAs space communication networks," in IEEE International Systems Conference (SysCon), Orlando, 2013.

[8] B. Barritt and V. Cerf, "Loon SDN: Applicability to NASA's Next-Generation Space Communications Architecture," in 2018 IEEE Aerospace Conference, Big Sky MT, 2018.

[9] T. Li, H. Zhou, H. Luo and S. Yu, "SERvICE: A Software Defined Framework for Integrated Space-Terrestrial Satellite Communication," IEEE Transactions on Mobile Computing, vol. 17, no. 3, pp. 703-716, 2018.

[10] S. Nazari, P. Du, M. Gerla, C. Hoffman, J. H. Kim and A. Capone, "Software Defined Naval Network for Satellite Communications (SDN-SAT)," in IEEE Military Communications Conference, Baltimore, 2016.

[11] S. Scott-Hayward, S. Natarajan and S. Sezer, "A Survey of Security in Software Defined Networks," IEEE Communication Surveys & Tutorials, vol. 18, no. 1, pp. 623-654, 2016.

[12] S. Shin, V. Yegneswaran, P. Porras and G. Gu, "AVANT-GUARD: Scalable and Vigilant Switch Flow Management in Software-Defined Networks," in ACM Computer and Communications Security Conference (CCS), Berlin, 2013.

[13] P. Fonseca, R. Bennesby, E. Mota, A. Passito and , "A replication component for resilient OpenFlow-based networking," in 2012 IEEE Network Operations and Management Symposium (NOMS): Mini-Conference, Maui, 2012.

[14] G. Yao, J. Bi and P. Xiao, "Source Address Validation Solution with OpenFlow/NOX Architecture," in 19th IEEE International Conference on Network Protocols, Vancouver, 2011.

[15] D. Z. Baker, H. Liu, C. Roberts and P. J. Fortier, Secure Software-Defined Integrated Space Communication Service Networks, University of Massachusetts Dartmouth, 2019.

[16] ISO/IEC, "ISO/IEC 27000:2018," February 2018. [Online]. Available: https://standards.iso.org/ittf/PubliclyAvailableStandards/c073906_ISO_IEC_27000_2018_E.zip.

978-1-7281-5092-5/19/$31.00 ©2019 IEEE