FAA BAA FINAL REPORT

Advance the UTM and its overall ecosystem to enable and advance BVLOS operations across the unmanned aircraft systems (UAS) Industry

SUBMITTED BY:

SWITCH LTD
SUPERNAP
7135 S DECATUR BLVD
LAS VEGAS NV 89118-4376

Contract Number: 692M15-20-C-00003
Solicitation Number: 692M15-19-R-00020

DEVELOPED FOR:

Federal Aviation Administration
William J. Hughes Technical Center
Atlantic City NJ 08405

EXECUTIVE SUMMARY

Switch and its partners, the Nevada Institute for Autonomous Systems (NIAS) and ANRA Technologies, constructed a hybrid cloud computing environment, built and deployed a standalone UTM system, and successfully subjected that UTM instance to 50,000 simulated UAS. These tests were conducted to better understand the performance, scalability, and cost factors associated with operating the UTM at scale. Switch is proud to present the findings of this important work to the FAA and the UAS industry as a whole to further the goal of UAS integration into the National Airspace System (NAS).

While conducting the largest FAA sanctioned UAS simulations to date, Switch saw an opportunity to conduct long range, BVLOS operations in dense UTM environments to better understand what UAS operators need to facilitate safe operations. This effort involved conducting UAS remote split operations (RSO) from Las Vegas in the Switch SIGHT command center and controlling an aircraft located at the Switch Citadel UAS Range in Reno, NV. These live de facto BVLOS missions were conducted in proximity to hundreds of simulated UAS connected to the UTM system. UAS operators generated valuable lessons learned relating to conducting RSO of an Unmanned Aircraft.

Switch recognizes that the UTM system is only as secure, reliable, scalable, and sustainable as the infrastructure it is built on. Our test results identify and highlight needed data infrastructure and technology conditions for FAA consideration highlighted below.

- 45 days of UTM data storage would require approximately 250 Petabytes
- Network capacity sufficient to handle high demand is required to deliver UTM services.
- Data storage will be a major cost center of the UTM system.
- Storing telemetry requires highly capable and high-capacity data storage solutions.
- UTM operating environments can function under and scale to meet extremely high loads.

These key technology infrastructure research findings will be useful as the FAA drives toward the goal of full UAS integration into the NAS. This effort was made possible by countless Switch team members as well as contributors from ANRA and NIAS. Switch utilized several product sets to enable this effort such as Switch VAULT for data storage, Switch CONNECT for carrier circuits, and Switch SIGHT the world’s only Tier 5® System + System network operations center (NOC), which is designed and equipped to provide state of the art command center technology for UAS operations.

The Switch Team took every measure to ensure that we provided the FAA with a detailed, meaningful and accurate Final Report that provides value to future FAA autonomous system operations and studies. Thank you for this opportunity to advance the FAA's research platform in Nevada and evolve this technology for national deployment. We look forward to working with the FAA in future projects.

Betsy Fretwell
Senior VP
Switch Smart Cities
Principal Investigator FAA BAA1

TABLE OF CONTENTS

1.0 Introduction

The Broad Agency Announcement (BAA), solicitation number 692M15-19-R-00020, was sponsored by the Federal Aviation Administration (FAA), Unmanned Aircraft Systems Integration Office (UASIO), UAS Program and Data Management Branch (AUS-410). This BAA included integration interests from the UASIO, which is located in Washington D.C. Under authority of the FAA Modernization and Reform Act (FMRA) of 2012 (P.L. 112-95) and the FAA Extension, Safety and Security Act (FESSA) of 2016 (P.L. 114-190), the FAA established seven UAS Test Sites (UASTS) to support the integration of unmanned aircraft systems into the national airspace system (NAS). Under the FAA Reauthorization Act of 2018 (P.L. 115-254), UASTS Pilot Program was extended to September 30, 2023. This BAA was issued under the FAA Acquisition Management System (AMS) under authority 49 USC 106(1) and (m), which provides for the selection of proposals submitted in response to this announcement. Under this authority, the FAA selected Switch's proposal to support their BAA1 efforts.

Switch is recognized as the independent world-leader in exascale data center ecosystems, edge data center design, industry-leading telecommunications solutions & next-generation technology innovation. Switch sustainably innovates the digital foundations of the connected world with a focus on enterprise class, emerging hybrid cloud technology solutions. Switch participated in the development of the UTM with NASA and the FAA on multiple UAS test initiatives such as NASA's TCL4 and the FAA's UPP. These areas of expertise and those testing experiences served as the bedrock for the BAA1 testing effort. They also informed our goal of advising the FAA regarding not only the UTM construct itself, but the underlying infrastructure that systems like UTM utilize and depend upon.

During our BAA contract period of performance, Switch and its partners, the Nevada Institute for Autonomous Systems (NIAS) and ANRA Technologies, constructed a hybrid cloud computing environment, built and deployed a standalone UTM system. This UTM system provided an opportunity to stress-test the UTM at scale with the added benefit of conducting live UAS RSO operations simultaneously.

2.0 Background

Prior to the execution of this contract, the UTM system had yet to be subjected to any realistic number of live or simulated aircraft. Once the UTM is implemented, this system must remain available, secure, cost effective, and scalable to meet demand. Switch recognized that prior to deployment of the UTM, that system needed to be tested with greater numbers of UAS. The demonstration efforts conducted under this contract allowed Switch to operate the UTM under production type loading to better understand the system's requirements. This effort allowed Switch to better gain insights into the future requirements of operating the UTM at scale. When deployed across the NAS, final UTM implementation will require constant availability and uptime, remain secure, be cost effective, and efficient. To achieve this, Switch built and deployed a hybrid cloud platform, where ANRA Technologies, a lead UAS Service Supplier (USS), deployed their USS and simulation tools. This enabled the ANRA USS to accept tens of thousands of simulated UAS to connect to its' UTM services and allowed an evaluation of the underlying infrastructure that will inevitably power the UTM system.

Another challenge to UAS integration into the NAS is the standardized approval for BVLOS Part 107 waivers without requiring visual observers. Supporting UAS operations in dense urban environments requires confidence in a UTM system that will provide for safe, complex operations in dynamic environments, without failures. This contract allowed the Switch team to study, develop, and deliver a Standardized Risk Management Template for Day/Night BVLOS Operations which included remote split operations (RSO), as well as a Command and Control (C2) Checklist to assist in future integration efforts. The Switch team successfully conducted day and night de facto BVLOS operations from the SwitchSIGHT Command Center, located over 325 miles south of the operations area.

3.0 FAA Broad Agency Announcement (BAA) Call #001

In June 2020, Switch, in partnership with NIAS and ANRA Technologies was awarded this FAA contract under Call #001 of the Broad Agency Announcement (herein referred to as BAA1). The Switch BAA1 contract was awarded to test the UTM system by deploying the ANRA USS to a hybrid cloud platform running inside Switch data centers and simulate tens of thousands of UAS connecting to that system. In addition to these simulations, Switch conducted remote split operations of de facto BVLOS aircraft located in Reno, NV and controlled from a command center in Las Vegas, NV. The objective of the Switch BAA1 was to test, not only the UTM ecosystem, but the performance of the underlying infrastructure. In addition, the BVLOS operation during BAA1 allowed Switch to understand how remote operators can safely operate and navigate in a dense UTM environment. The predicate for the BAA1 was based on Switch's experience in testing the UTM with the FAA and NASA in other past exercises. In May and June of 2019, Switch, in partnership with the FAA-designated Nevada UAS Test Site (NIAS), participated in NASA's UTM Technical Capability Level (TCL) 4 and the FAA's UTM Pilot Program (UPP) with a flight operations team and provided overall network architecture recommendations for each operation. The operational experience gained from these previous experiences served as a deep foundational knowledge base for the entire BAA1 team.

4.0 Switch BAA1 Team Partners

The Switch BAA1 team consisted of the following organizations:

Teammates	**Company Name and Description**	**Key Contributions**
switch	Switch: World-leading exascale technology solutions	SWITCH, LTD: Provided Principal Investigator, Primary program and project management. Provided Switch Citadel UAS Range and Switch SIGHT Command Center
ANRA TECHNOLOGIES	ANRA Technologies: Leading international provider of end-to-end drone operations and air traffic management solutions for unmanned system operators and airspace managers. ANRA offers intelligent and modular software capabilities as part of our SmartSkies™ family of solutions, enabling the creation of an unmanned ecosystem for compliant UTM and UAM operations.	ANRA Technologies: Provided USS, live aircraft and simulated aircraft integration services and built all simulated flights on this study
STATE OF NEVADA UNMANNED AIRCRAFT SYSTEMS TEST SITE	Nevada Institute for Autonomous Systems (NIAS): A state sanctioned 501(c)3 non-profit organization leading the FAA-designated UAS Test Site on behalf of the State of Nevada	NIAS: Provided Project Management, Airspace development, Waiver authorization development, Safety Case development, Test Plan and Test Card development, and range operations including command and control. Over saw all flight safety aspects and pilot training requirements
	Carbon Autonomous: Produces unmanned systems enhanced by seamless data analysis and visualization systems, and is guided by a team of experts with decades of combined experience in related fields.	Carbon Autonomous: provided engineering support and qualified flight crews. Carbon Autonomous was assigned as the primary flight crew

	Utah State University (USU): Through USU's Electrical and Computer Engineering Dept. comes the Aggie Air Program. Aggie Air have been long time teaming partners with NIAS on multiple major NASA and FAA contracts. See: http://aggieair.usu.edu/	USU: Provided engineering services, UAS aircraft and qualified flight crews. USU was assigned as the primary backup flight crew and aircraft
	Reno-Stead Airport continues to serve as the primary UAS test airfield for the State of Nevada having supported multiple NASA TCL events and FAA projects such as FAA UPP and FAA BAA1	Reno-Stead Airport: Provided office and airfield space on demand for pilot currency training, software updates and meetings at no cost

5.0 Personnel Roles and Responsibilities

The following chart depicts the key POC's contributing to this effort and email contact information.

Name	Company	Role and Responsibility	Email
Betsy Fretwell	Switch, LTD	Principle Investigator/ Team Leadership	betsy@switch.com
Wes Dye	Switch, LTD	Program Manager /Prime Contract Lead	wdye@switch.com
Brent Klavon	ANRA Technologies	Senior POC for all ANRA efforts	bklavon@fkyanra.com
David Murphy	ANRA Technologies	Primary Technical POC	dmurphy@flyanra.com
Blair Smith	NIAS	Senior POC for NIAS efforts	blair.smith@nias-uas.com
Chris "Oco" O'Connor	NIAS	Primary NIAS PM	Chris.oconnor@nias-uas.com
Kevin "Lucky" Fallico	NIAS	Project Manager; Primary Remote Pilot	kevin.fallico@nias-uas.com
Dan Cassidy	Carbon Autonomous	Primary Pilot in Command	dcassidy7@gmail.com
John Hammond	Carbon Autonomous	Pilot in Command	dragonflight@charter.net

6.0 Operational Range Area

UAS operations for Demonstration 3 were conducted at the Switch Citadel UAS Range, located approximately 15 miles East of the Reno-Tahoe International Airport. The Switch Citadel UAS Range combines the most essential elements required to conduct safe, productive, and valuable UAS testing. The Switch Citadel campus consists of approximately 2000 acres of Switch controlled property adjacent to tens of thousands of acres of unoccupied land, offering a safe testing location with minimal manned aviators operating in the area. Given that Switch's Citadel Campus contains the world's most advanced data centers, the Switch Citadel UAS Range offers unparalleled network, power, data storage, and infrastructure capabilities. All of this combines to make the Switch Citadel UAS Range one of the most advanced and flexible UAS test ranges in the country.

Figure 1: Live Flight Operational Area located at the Switch Citadel UAS Range in Reno, NV

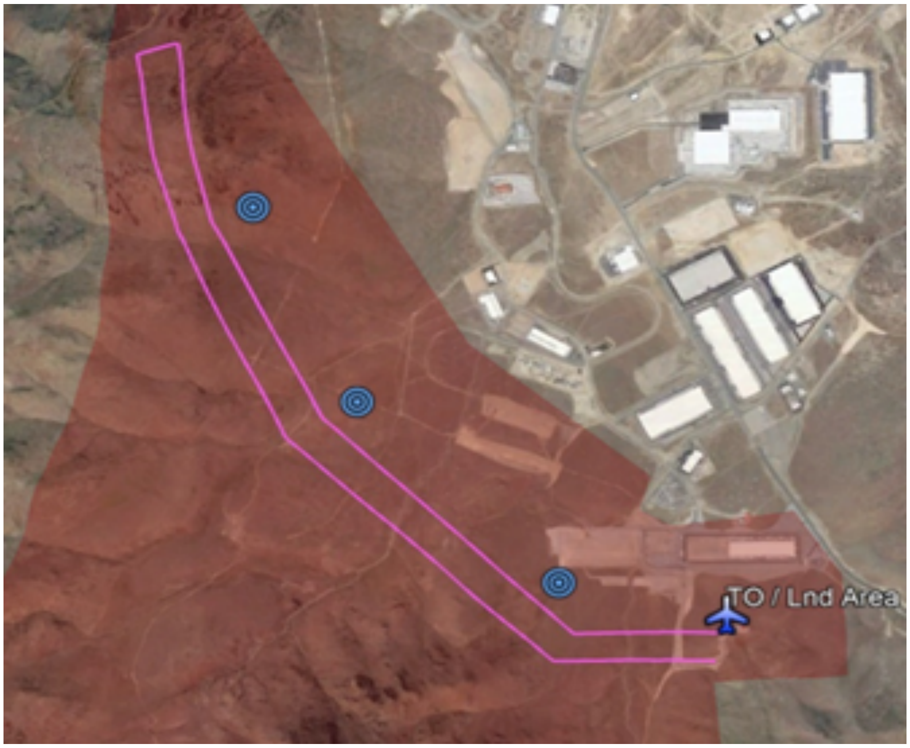

7.0 Switch SIGHT Command Center

Switch SIGHT is the world's only Tier 5® System + System network operations center (NOC), which is designed and equipped to provide state of the art command center technology for hands on mission-critical operations monitoring, telecom operations monitoring, ground drone piloting services, aerial drone piloting services, edge data center monitoring, and on-prem data center monitoring. The human-in-loop monitoring is provided via two-way communications for operations support in the field, nationwide.

The Switch SIGHT Command Center is located at the Switch LAS VEGAS 9 data center and consists of up to 10 UAS operator stations, 10 Gbps connectivity, and un-interruptible power. These elements make for a highly secure, highly connected, and extremely resilient location for UAS operations. By implementing these redundancies in the Switch SIGHT Command Center, Switch reduces risk and increases resiliency during long-range BVLOS missions.

Figure 2: Switch SIGHT Command Center

Figure 3: Connection to UA in Reno, NV to Command Center in Las Vegas

8.0 BAA1 Test Plan

Demonstrations 1 and 2 Test Card Overview

The Demonstration 1 test card was designed to exercise a broad range of UTM functions and scenarios with the expectation that results would highlight functional flows and/or architectural components that would merit further study in Demonstration 2. Figure 4 provides an overview of the Demonstration 1 and 2 test cards. Section 13 provides an analysis of Demonstration 1 results that informed the changes to the Demonstration 2 test card. Test card parameter permutation applications were split into test environment setup and test execution.

Figure 4: Demonstration 1 Test Card Overview

Environment Setup		
	Demonstration 1	**Demonstration 2**
# of Simultaneous Operations	250	10,000 +
Operation Density	2.7 ops/mi^2, 4 ops/mi^2	2.7 ops/mi^2
Aircraft Characteristics	Baseline Quad-Copter, Altered Quad-Copter	Baseline Quad-Copter
Test Execution		
	Demonstration 1	**Demonstration 2**
Plan Operations	Plan 250 simultaneous operations	Plan 10,000+ simultaneous operations
Initiate Telemetry	Initiate telemetry flow for all simulated operations	Initiate telemetry flow for all simulated operations
Initiate Mid-Flight Replans	Replan 1-2 of the active operations	*(Removed from test card)*
Initiate Off-Nominal Operations	Cause 1-16 of the activated operations to enter an off-nominal state	Cause 1,000 of the activated operations to enter an off-nominal state
Plan Conflicting Constraint	Plan a constraint that conflicts with 3%-16% of planned operations	*(Removed from test card)*

The left-hand column of Figure 4 defines the various test permutations of the test cards. These are described in greater detail below:

Environment Setup

- **Number of Simultaneous Operations:** The number of simulated simultaneous operations flying throughout the duration of the test.
- **Operation Density:** The number of operations flying within a square mile. This density is applied to the entirety of the operations for each test.
- **Aircraft Characteristics:** The BAA test objectives sought to explore the effects of different aircraft characteristics on UTM performance. For Demonstration 1, most flights were conducted with a "Baseline Quad-Copter" which embodied the baseline flight characteristics implemented in the physics engine of the ANRA Batch Operation Simulator. For Demonstration 1, Test 3, aircraft airspeeds and navigational performance were altered to generate the "Altered Quad-Copter" characteristics to be explored in this effort's analysis.

Test Execution

- **Planning Operations:** Planning operations consists of identifying the waypoints and altitudes that represent the desired flight path, generating associated operational intent volumes, recalling relevant operator information, generating a drone model to execute the operation, and submitting the operation to the UTM system.
- **Initiating Telemetry:** Sending telemetry for an operation signifies the commencement of flight within the UTM system. It also initiates conformance monitoring services that track a given simulation's telemetry against the planned operational intent volumes.
- **Initiating Mid-Flight Replans:** When an operation is initially planned, the UTM software conducts a strategic deconfliction assessment to determine if the requested operation is in conflict laterally, vertically, and temporally (4D) to any other operation or constraint. If a 4D conflict is detected, then the operation cannot be planned. This same logic applies to the replanning of operations that are mid-flight. As such, this permutation explores the performance of these interchanges within different operational environments.
- **Initiating Off-Nominal Operations:** An operation is considered off nominal if it exits its operational intent volume. When this occurs, the UTM software identifies the operation as "non-conforming." Should the operation remain in the "non-conforming" state beyond a set period of time, the operation is identified as "contingent."
- **Planning Conflicting Constraints:** Similar to operations, constraints are defined in terms of their 4D parameters. Should a constraint overlap with an operation in both space and time, then that operation is determined to be in conflict with the constraint. This test parameter defines the percentage of planned operations that are in conflict with the generated constraint.

Demonstration 3 Test Card Overview

Demonstration 3 preparations involved an iterative, collaborative process among Switch, ANRA Technologies, and NIAS. Demonstration 2 illustrated that the USS software and simulator could adequately manage 10,000 concurrent operations.

Demonstration 3 operational scales reverted back to 250 simulated UAS and 1 live operation. The testing focus was placed on remote C2 and operator display factors as opposed to USS software and IT infrastructure performance.

To support Demonstration 3, Switch provided NIAS access to the Citadel UAS Range in Reno, NV and the Switch SIGHT Command Center located in Las Vegas, NV. Live operations were conducted with the remote pilot-in-command located on-site with the UA in Reno, NV and the remote operator located in the Switch SIGHT Command Center in Las Vegas, NV. The connection between The Citadel UAS Range and the Switch SIGHT Command Center was enabled by the Switch SUPERLOOP, consisting of diverse fiber pathways with less than 7ms of round-trip latency. This latency was imperceptible to UAS Pilots during the operations.

Demonstrations 4 Test Card Overview

Demonstration 4 was designed to explore UTM failure. Early preparations for this demonstration sought to establish a means through which UTM failure could be quantified in a test environment. Typically, a software systems' sub-components do not exhibit a binary "functional" or "non-functional" state. Rather, sub-components might enter state of degraded performance that would still constitute a failure of the UTM to provide its' services. The challenge then becomes identifying performance values wherein degraded function could be classified as failure.

Given this context, Demonstration 2 results would serve as a baseline for defining Demonstration 4 failure modes. As is discussed in Section 12, the Demonstration 2 metrics served as a tool for monitoring system performance and identifying USS software and IT infrastructure components that were underperforming or functioning in a degraded state. By repeating Demonstration 2 execution methods at higher loads and lower IT infrastructure resource allocations, resulting Demonstration 4 metrics could be compared to Demonstration 2 results to better quantify "failure." This in turn would highlight components that needed to be optimized and retested.

9.0 Switch Hybrid Cloud Platform Development and Deployment

Network Infrastructure Deployed

In order to support this contract, Switch deployed an enterprise class network that allowed for highly available, high throughput connectivity between UTM services and out to the internet. For firewalls, Switch utilized a two Palo Alto PA-850 next gen appliances, these devices are capable of up to 1.2 gigabits/second (Gbps) of throughput with intrusion prevention enabled, creating a highly secure test environment. Downstream from the firewalls, Switch deployed a pair of Cisco 10 Gbps appliances to serve as firewall switches to distribute traffic to downstream devices. The hybrid cloud environment was then supported by a pair of Dell top-of-rack (ToR) switches capable of up to 100 Gbps. These Dell switches connected physical nodes of the Dell VxRail cloud environment discussed in the next section. Other points of presence such as the Switch SIGHT Command Center or the Switch Citadel UAS Range were supported by 48-port Cisco switches capable of 1Gbps per port with redundant 10 Gbps uplinks to the firewall.

Compute Environment

During this effort, Switch deployed a Dell VxRail to support UTM compute requirements. The Dell VxRail platform, being a hyper converged infrastructure (HCI) solution, consists of compute and storage on the same physical hardware. In support of this contract, this came in the form of four Dell R740 servers with a combined capacity of approximately 225 physical cores, 1.5 terabytes (TB) of memory, and 20TB disk space. These servers also were equipped with four 25 Gbps ports each to connect to the Dell ToR switches. The Dell VxRail environment utilizes VMware's vCenter application as a hypervisor to virtualize the hardware environment and supports the deployment of virtual machines (VMs). Using these VMs, Switch deployed a highly available instance of software from Rancher Labs to create and manage Kubernetes clusters. Kubernetes provided support for containerized applications and Rancher allowed for the deployment of test workloads into the Switch environment or into the public cloud, providing a truly hybrid cloud approach.

Monitoring

For environmental monitoring, Switch deployed the Prometheus application to scrape available endpoints for data and the Grafana application to visualize that data. These solutions were deployed using a Rancher provided Helm chart.

10.0 Simulation Software Development

SmartSkies™ CTR

Each test was executed utilizing ANRA's SmartSkies™ CTR platform that provided UTM services and was deployed on the Switch hybrid cloud platform. UAS and their corresponding telemetry were simulated using ANRA's batch operation simulator deployed on a VM. In addition, specific instrumentation on the SmartSkies™ CTR platform and underlying IT infrastructure was implemented to capture metrics data in an automated manner during test execution. While the primary focus of the test cards and corresponding metrics was the back-end performance of UTM functions and supporting IT infrastructure, the following figures provide front-end insights to better contextualize the test execution.

Figure 5 displays the SmartSkies™ CTR operator dashboard that provides an aggregated view of the various operations under the purview of the operator. This particular screen shot was taken over the course of Demonstration 1. As is highlighted by the red box, at the time of this screen capture, 398 simultaneous operations were active.

Figure 5: SmartSkiesTM CTR Dashboard

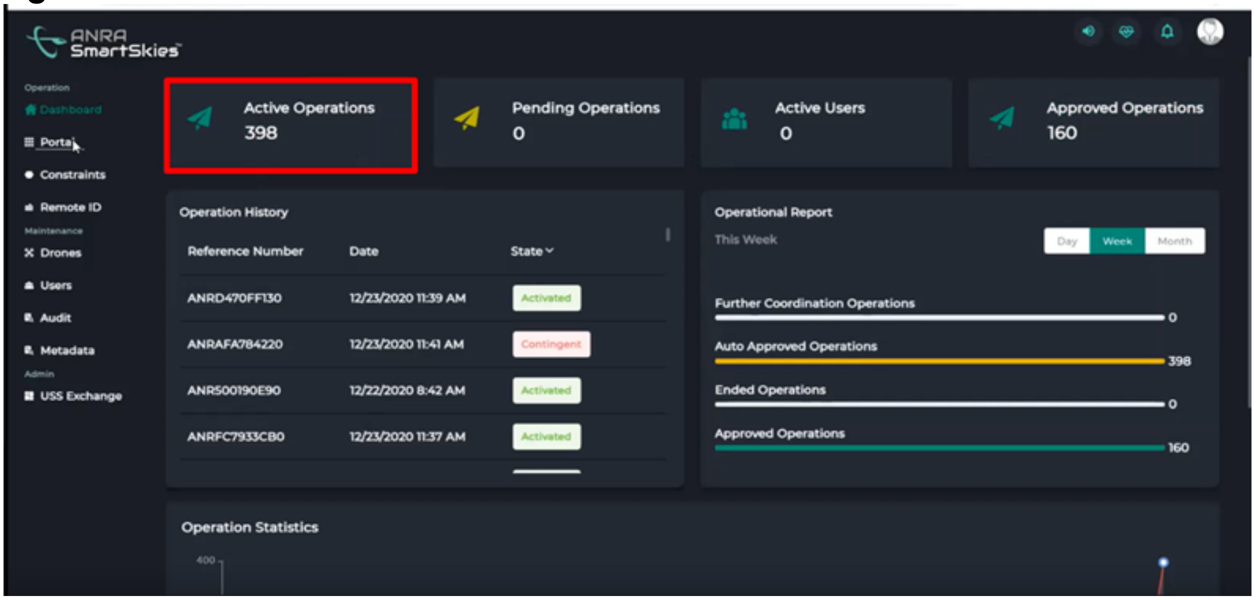

Figure 6 provides a user's view of notifications when a UA leaves its operational volume. Should a UA leave its operational volume, the SmartSkiesTM CTR software identifies the off-nominal situation and automatically updates the mission status to "non-conforming." Should the UA remain outside its operational volume for longer than a set period of time, the mission is automatically transitioned to the "contingent" state. These state changes are communicated to the user through status awareness in the left-hand operation pane and through push notifications, as depicted on the right.

Figure 6: SmartSkiesTM CTR Off-Nominal Operation Notification

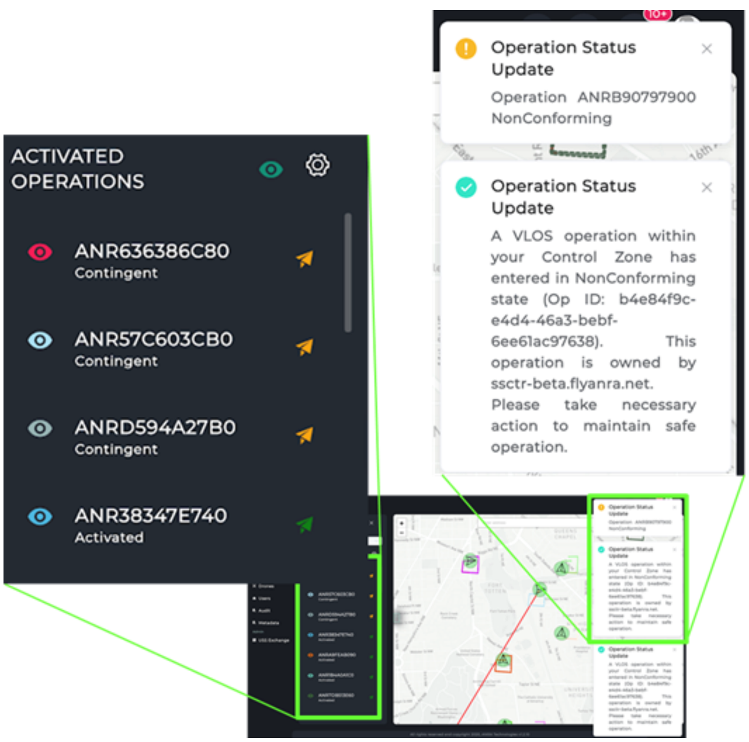

Figure 7 depicts the user's view of activated operations. In order to manage large numbers of operations, this portal view is paginated to display a subset of the total operations at any given time. This view also displays constraints that are planned in the area. The red polygon in Figure 7 depicts such a constraint.

Figure 7: SmartSkies™ CTR Operations Portal

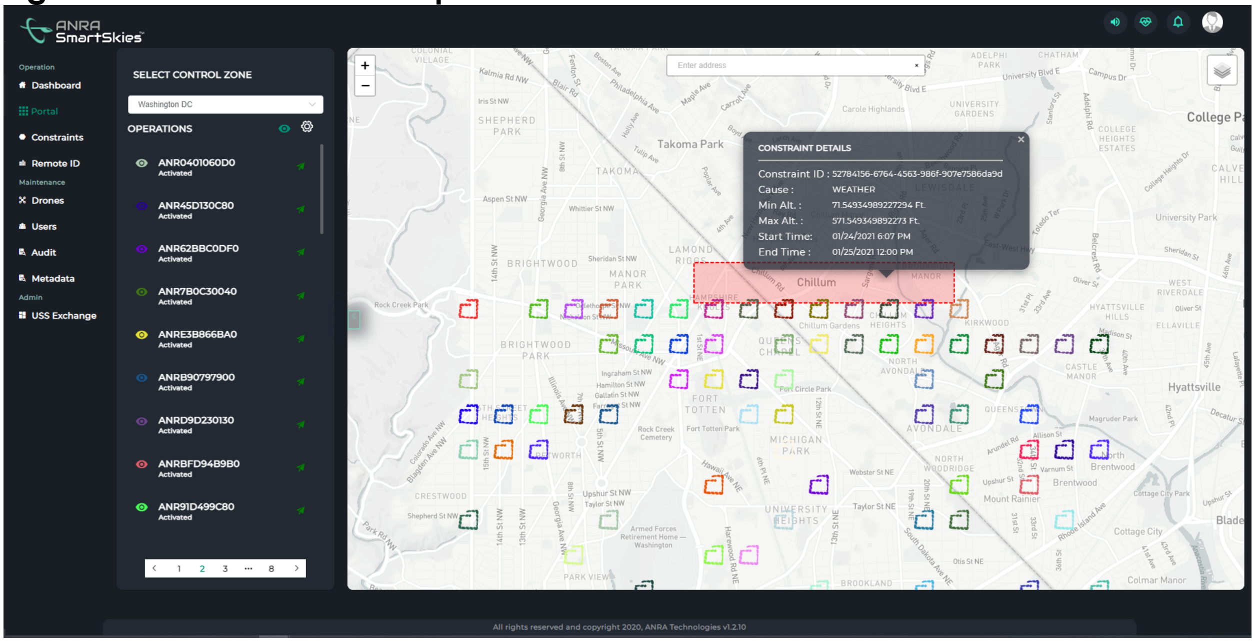

SmartSkies™ SIM

ANRA's SmartSkies™ SIM was utilized to generate all operational details and telemetry necessary for SmartSkies™ CTR operation planning and conformance monitoring functions. Figure 8 highlights the core functions of SmartSkies™ SIM that allows for simulating high-volume operations. The operation planner functions enabled dynamic flight planning which could be applied individually or as a batch on numerous missions. The network was architected to optimize for high volume data exchanges and feeds over the course of the demonstrations. Underlying all telemetry was the core physics engine which ensured realistic flight paths of each individual operation. Finally, the time clock capabilities enabled synchronization of telemetry feeds for the simultaneous operations.

Figure 8: SmartSkies™ SIM Functional Overview

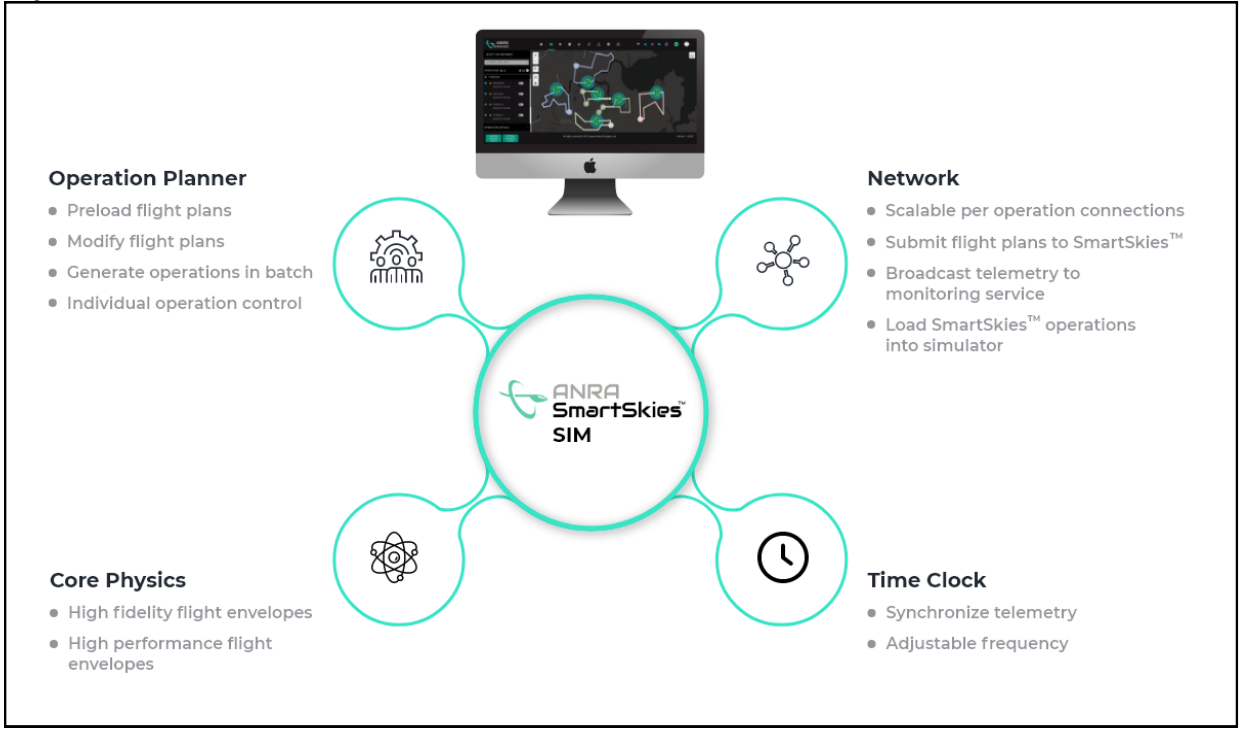

11.0 Comprehensive Safety Plan

Safety Plan and Risk Assessment Template

The Comprehensive Safety Plan was submitted and accepted as a separate deliverable to the FAA in August 2020. As a direct result of our upstream safety efforts there were no downstream accidents, violations, or injuries during the execution of this contract. In compliance with the Comprehensive Safety Plan, several flight days were postponed due to weather. Per the contract statement of work, a Risk Assessment Template for Day/Night BVLOS Operations with associated instructions is included as Appendix 5 in this report.

Remote C2 SOP Operations

Although not specifically required under this effort, Switch provides to the FAA the standard operating procedure (SOP) checklist developed and used by NIAS while conducting remote C2 and RSO missions as Appendix 6 in this report.

12.0 Data Collection Metrics

Data Capture Instrumentation

Data collection was a core aspect of demonstration execution. Data capture tools were implemented within SmartSkies™ CTR, SmartSkies™ SIM, and on the hybrid cloud platform. This instrumentation was adapted for the specific needs of each demonstration. Metrics for each demonstration are outlined below.

Demonstration 1 Metrics

In conjunction with the development of the Demonstration 1 test cards, specific metrics were defined that would allow the team to measure the relative performance of the UTM software and served as the baseline metrics for follow-on demonstrations. Demonstration 1 consisted of 250 simulated UAS connecting to the UTM.

Success of Mid-Flight Replan
- Units: % (dimensionless)
- Test Execution Focus: Initiate Mid-Flight Replans
- Description: If a mid-flight replan is determined to be strategically deconflicted from other operations and constraints, then it is considered a "success." This metric is captured in order to understand the % of successful replans that are not rejected due to 4D conflict detection.

Number of Impacted Operations
- Units: Number of Impacted Operations
- Test Execution Focus: Initiate Off-Nominal Operations
- Description: This metric captures the total number of operations that are impacted by the off-nominal operations outlined in the test card.

Time to Notify for Constraint Conflict
- Units: Seconds
- Test Execution Focus: Plan Conflicting Constraint
- Description: This metric captures the time between the planning of a constraint and the detection of any operational conflict.

Receive Bandwidth
- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure received data throughout testing
- Description: This metric captures the amount of data received during the test execution.

Transmit Bandwidth
- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure transmitted data throughout testing
- Description: This metric captures the amount of data transmitted during the test execution.

CPU Usage
- Units: vCPUs (Virtual Centralized Processing Units)
- Test Execution Focus: Measure vCPUs used throughout testing
- Description: This metric captures the number of vCPUs utilized during the test execution.

Memory Usage
- Units: Gigabytes (GB)
- Test Execution Focus: Measure memory used throughout testing
- Description: This metric captures the amount of system memory utilized during the test execution.

Disk Usage
- Units: Terabytes (TB)

- Test Execution Focus: Measure disk space required for telemetry retention
- Description: This metric captures the amount of disk space required to retain a 7-day running telemetry history with a given number of operations.

Demonstration 2 Metrics

Demonstration 1 highlighted several opportunities for how metrics could be adjusted in order to better describe UTM performance during Demonstration 2. As part of this metrics analysis, the team referenced the draft ASTM UTM standard for meaningful performance measures and approaches within the UTM domain. Demonstration 2 consisted of 10,000 simulated UAS connecting to the UTM. The metrics below were derived from this analysis and captured during Demonstration 2.

Time to Create and Submit Planned Operations

- Units: Number of Operations/Time
- Test Execution Focus: Plan Operations
- Description: This metric explores the time required to create an operation. This aggregated metric is comprised of the time to generate a simulated UA and its metadata, generate the operational intent volume, and submit the operation.

Time to Initiate Planned Operations

- Units: Number of Operations/Time
- Test Execution Focus: Plan Operations
- Description: This metric explores the time required to establish connectivity between the simulated UA and the UTM software for each planned operation across all operations running in the simulation.

Time to Transition Operation State

- Units: Number of Operations/Time
- Test Execution Focus: Initiate Off Nominal Operations
- Description: This metric captures the time required for the simulator to transition an operation from the "activated" state to a "non-conforming" state across all operations running in the simulation.

Receive Bandwidth

- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure received data throughout testing
- Description: This metric captures the amount of data received during the test execution.

Transmit Bandwidth

- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure transmitted data throughout testing

- Description: This metric captures the amount of data transmitted during the test execution.

CPU Usage

- Units: vCPUs (Virtual Centralized Processing Units)
- Test Execution Focus: Measure vCPUs used throughout testing
- Description: This metric captures the number of vCPUs utilized during the test execution.

Memory Usage

- Units: Gigabytes (GB)
- Test Execution Focus: Measure memory used throughout testing
- Description: This metric captures the amount of system memory utilized during the test execution.

Disk Usage

- Units: Terabytes (TB)
- Test Execution Focus: Measure disk space required for telemetry retention
- Description: This metric captures the amount of disk space required to retain a 7-day running telemetry history with a given number of operations.

Demonstration 3 Metrics

Demonstration 3 metrics were the same as Demonstration 2 but scaled back to 250 simulations with 1 live operation. Testing focused on remote C2 and operator display factors rather than USS software and IT infrastructure performance.

Demonstration 4 Metrics

Demonstration 4 focused on UTM services and IT infrastructure performance with greater numbers of simulated UAs connecting to the USS. With the lessons learned implemented from Demonstrations 1-3, Demonstration 4 consisted of 50,000 simulated UAS connecting to the UTM. The metrics detailed below were captured during this demonstration.

Receive Bandwidth

- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure received data throughout testing
- Description: This metric captures the amount of data received during the test execution.

Transmit Bandwidth

- Units: Megabits per Second (Mbps)
- Test Execution Focus: Measure transmitted data throughout testing
- Description: This metric captures the amount of data transmitted during the test execution.

CPU Usage

- Units: vCPUs (Virtual Centralized Processing Units)
- Test Execution Focus: Measure vCPUs used throughout testing

- Description: This metric captures the number of vCPUs utilized during the test execution.

Memory Usage

- Units: Gigabytes (GB)
- Test Execution Focus: Measure memory used throughout testing
- Description: This metric captures the amount of system memory utilized during the test execution.

Disk Usage

- Units: Terabytes (TB)
- Test Execution Focus: Measure disk space required for telemetry retention
- Description: This metric captures the amount of disk space required to retain a 7-day running telemetry history with a given number of operations.

13.0 Analysis

Demonstration 1:

Demonstration 1 was designed to explore numerous UTM operational permutations to better understand USS performance at higher operational volumes and to refine testing for subsequent demonstrations. The test permutations, associated metrics, and testing results are summarized in Figure 9. The primary finding of Demonstration 1 was the identification of deterministic test parameters. More specifically, it was observed that Operation Density, Constraint Conflicts, and Mid-Flight Replans generated results that were solely affected by the test parameters themselves. This is most clearly evidenced by the Success of Mid-Flight Replan metric, which was 100% for each test. This was the case because each operation replan could be pre-determined to not conflict with another operation. Even deciding to purposefully replan an operation in conflict with another or applying randomized replans would generate predetermined corresponding results. For this reason, the mid-flight replan was removed from subsequent test cards.

Similar behavior was observed regarding constraint conflicts. Constraints could be pre-determined to conflict with a certain percentage of active operations based on its size and orientation to the operational intent volumes. The test permutation of Operation Density also affected the number of operations impacted by a given constraint but, again, this impact was observed to be largely deterministic: The number of impacted operations was solely a function of where a constraint was placed in regard to the outlay of operational intent volumes. For these reasons, Operation Density and the application of constraints were removed from subsequent test cards.

Aircraft characteristics were adjusted on test iteration #3. As expected, this permutation had a minor impact on the flight trajectory of a given operation through the operational intent volume but had no other impact on system performance or other observed metrics. As such, it was determined that a single baseline quadcopter physics model would be used for subsequent demonstrations.

Fluctuations in the Time to Notify of Constraint Conflict metric were observed over the course of the Demonstration 1 tests. As can be seeing in Figure 9 fluctuations were relatively minor and mins/maxes could not be easily attributed to any test permutation. It was the team's conclusion that these fluctuations were most likely a result of minor variations in IT infrastructure handling of the data exchanges between the simulator and USS software. This metric was used as a model to inform the metrics of Demonstration 2.

Figure 9: Demonstration 1 Metrics Overview

Environment Setup								
Test #	1	2	3	4	5	6	7	8
Operation Density	2.7 ops/ mi2	4 ops/ mi2	2.7 ops/ mi2	2.7 ops/ mi2	4 ops/ mi2	4 ops/ mi2	2.7 ops/ mi2	2.7 ops/ mi2
Aircraft Characteristics	Baseline Quad-Copter	Baseline Quad-Copter	Altered Quad-Copter	Baseline Quad-Copter	Baseline Quad-Copter	Baseline Quad-Copter	Baseline Quad-Copter	
Test Execution								
Constraint Conflict %	3.04%	2.88%	3.40%	16.67%	3.43%	5.34%	4.89%	3.22%
# of Mid-Flight Replans	2	2	2	1	1	1	1	1
# of Off-Nominal Operations	2	2	2	1	3	3	9	16

Metrics								
Success of Mid-Flight Replan	100%	100%	100%	100%	100%	100%	100%	100%
Impacted Operations from Off-Nominal Operation	21	27	4	19	34	32	30	65
Time to Notify for Constraint Conflict	0.794 s	0.578 s	0.572 s	0.578 s	0.647 s	0.317 s	0.369 s	0.346 s
Receive Bandwidth (Mbps)	107.9822 Mbps							
Transmit Bandwidth (Mbps)	105.0164 Mbps							
CPU Utilization (vCPUs)	25.3625 vCPUs							
Memory Utilization (GB)	50.4334 GB							
Disk Use for Storing Telemetry (7days of 1Hz updates) (TB)	2.520 TB							

Demonstration 2:
Lessons Learned from Demonstration 1, as documented in Appendix 1, identified several of the test steps and permutations requiring adjustment to offer more

valuable insights to UTM performance at high operational volumes. As such, for Demonstration 2 the test sequence was updated to focus on computationally expensive UTM functions that are foundational to the UTM concept. At a high level, these functions included operation planning, telemetry streaming, and conformance monitoring. Figure 10 details IT infrastructure performance during Demonstration 2

Figure 10: Demonstration 2 Metrics Overview

Demonstration 2 Metrics (10,000)	
Receive Bandwidth (Mbps)	1081.0856 Mbps
Transmit Bandwidth (Mbps)	1095.4344 Mbps
CPU Utilization (vCPUs)	40.6531 vCPUs
Peak Memory Utilization (GB)	75.7543 GB
Disk Use for Storing Telemetry (7days of 1Hz updates) (TB)	100.8 TB

Figure 11 shows system performance as the simulator initiates operations within the USS software. As illustrated, the system experienced an initial surge of operation connections, then the completion rate declined for the remaining operation connections. It is believed that this behavior indicates the system reaching a maximum processing capability then plateauing to address remaining operation connections as others are completed. Despite this behavior, the fact that all 10,000 operation connections were established over the course of approximately 10 seconds indicates a very high performing system. During testing, average rates of operation connections established routinely exceeded 1000 operations per second.

Figure 11: Demonstration 2 Time to Initiate Planned Operations

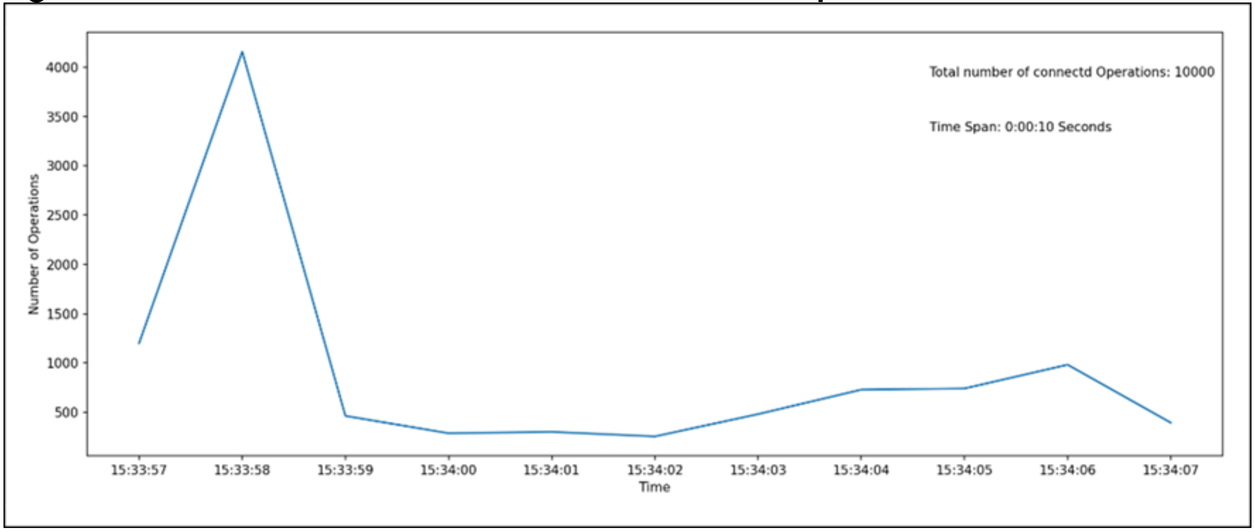

Figure 12 depicts the time lapse of transitioning 1,000 of the 10,000 operations into an "off-nominal" state. As can be seen, the observed behavior represents a steady-state processing of these transactions between the simulator and the USS software. Additionally, all states were transitioned over the course of 0.03 seconds, indicating a very high-performance translation.

Figure 12: Demonstration 2 Time to Transition Operation State

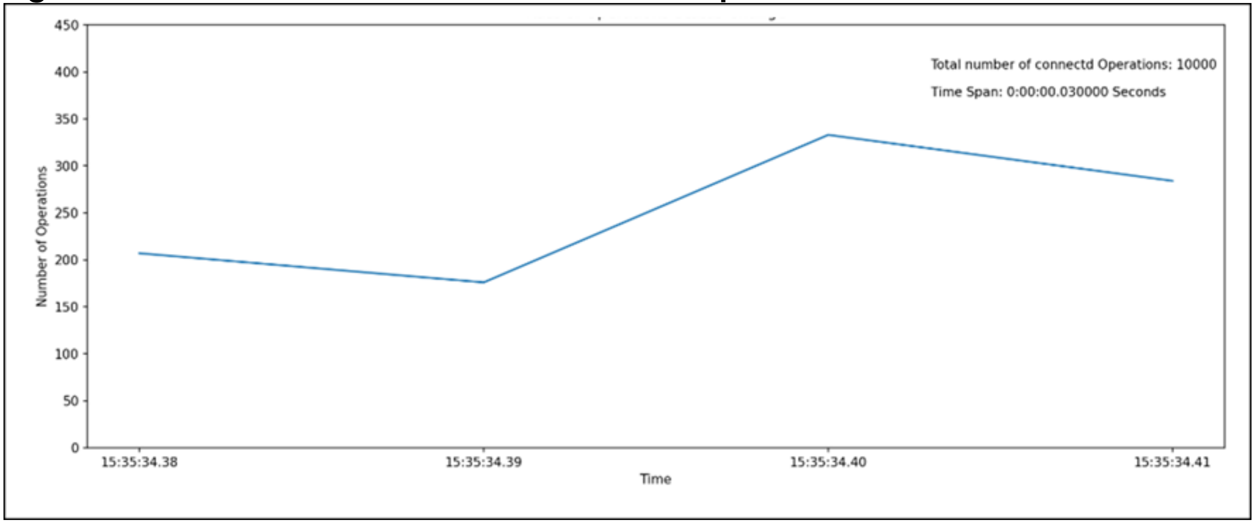

Demonstration 3:

Demonstration 3 consisted of 250 simulated UAS with 1 live UA connecting to the UTM. This demonstration focused on developing the Risk Assessment Template for Day/Night BVLOS Operations, included in this report as Appendix 5. From an IT perspective, Demonstration 3 was conducted the same as Demonstrations 1 and 2 with the addition of 2 physical locations. These locations included:

Reno, NV - Switch Citadel UAS Range: The aircraft, Pilot in Command (PIC), visual observer, and NIAS Range Control Officer were stationed at the Switch Citadel UAS Range. The PIC maintained visual line of sight of the aircraft and was prepared to take over with manual control while the aircraft was operated by the remote pilot in the SwitchSIGHT Command Center, located 325 miles away.

Las Vegas, NV - SwitchSIGHT Command Center: The NIAS Mission Director and NIAS Remote Pilot were stationed in the SwitchSIGHT Command Center. These de facto BVLOS/RSO missions were directed and flown from this command center in Las Vegas, NV while the aircraft and PIC remained at the Switch Citadel UAS Range.

All sites were connected utilizing Zoom teleconferencing software, Slack messaging platform, and UgCS software for UA C2.

Figure 13: Demonstration 3 Flight Schedule

Date	Aircraft	Take-off Time	Land Time	Total Time	Mission	PIC	Remote Pilot	Notes
12-Apr	PARROT1	2:45 PM	2:52:57	7min 57 s	Day local	D.Cassidy	K. Fallico	
12-Apr	PARROT1	2:54 PM	2:53:25	1 min 25s	Day local	D.Cassidy	K.Fallico	
12-Apr	PARROT2	3:12 PM	3:12:16	16 s	Day local	D. Cassidy	K. Fallico	
12-Apr	PARROT2	3:13 PM	3:14:38	1 min 38 s	Day local	D. Cassidy	K. Fallico	
12-Apr	PARROT2	3:15 PM	3:19:15	4 min 15 s	Day local	D. Cassidy	K. Fallico	
12-Apr	PARROT2	3:19 PM	3:24:27	5 min 27 s	Day local	D. Cassidy	K. Fallico	
12-Apr	PARROT2	3:25 PM	3:27:24	2 min 24 s	Day local	D. Cassidy	K. Fallico	
14-Apr	PARROT1	3:06 PM	3:08:28	2 min 28s	Day local	D.Cassidy	K.Fallico	
14-Apr	PARROT1	3:12 PM	3:14:27	2min 27s	Day local	D.Cassidy	K.Fallico	
16-Apr	PARROT1	7:00 AM	7:04:01	4 min 1 s	BVLOS/RSO	D.Cassidy	K.Fallico	
16-Apr	PARROT1	7:09 AM	7:13:02	4min 2 s	BVLOS/RSO	D.Cassidy	K.Fallico	
19-Apr	PARROT1	5:51 AM	5:54	3min 59 s	BVLOS/RSO	D.Cassidy	K.Fallico	Night BVLOS/RSO
19-Apr	PARROT1	6:02 AM	6:05:59	3min 59 s	BVLOS/RSO	D.Cassidy	K.Fallico	Night BVLOS/RSO
19-Apr	PARROT1	6:27 AM	6:30:41	3 min 41 s	BVLOS/RSO	D.Cassidy	K.Fallico	
20-Apr	PARROT1	6:34 AM	6:35:08	1 min 8 s	BVLOS/RSO	D.Cassidy	K.Fallico	
20-Apr	PARROT1	6:35 AM	6:35:53	53 s	BVLOS/RSO	D.Cassidy	K.Fallico	
20-Apr	PARROT1	6:55 AM	6:56:05	1 min 5 s	BVLOS/RSO	D.Cassidy	K.Fallico	
20-Apr	PARROT2	7:06 AM	7:07:39	1 min 39 s	BVLOS/RSO	D. Cassidy	K. Fallico	
20-Apr	PARROT2	7:11 AM	7:12:39	1 min 39 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT1	6:08 AM	6:12:46	4 min 46 s	BVLOS/RSO	D. Cassidy	K.Fallico	Night BVLOS/RSO

Apr-21	PARROT2	7:08 AM	7:11:41	3 min 41 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT2	7:13 AM	7:16:40	3 min 40 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT2	7:19 AM	7:22:51	3 min 51 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT2	7:25 AM	7:28:41	3 min 41 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT2	7:32 AM	7:35:40	3 min 40 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-21	PARROT2	7:37 AM	7:41:10	4 min 10 s	BVLOS/RSO	D. Cassidy	K. Fallico	
Apr-22	PARROT2	5:15 AM	5:19:20	4 min 20 s	BVLOS/RSO	D. Cassidy	K. Fallico	Night BVLOS/RSO
Apr-22	PARROT2	5:22 AM	5:23:02	1 min 02 s	BVLOS/RSO	D. Cassidy	K. Fallico	Night BVLOS/RSO
Apr-22	PARROT2	5:26 AM	5:30:12	4 min 12 s	BVLOS/RSO	D. Cassidy	K. Fallico	Night BVLOS/RSO

Demonstration 4:

Demonstration 4 explored various software and hardware limitations and their impact on overall system performance. Through early demonstration testing, it was determined that the operation creation and submission function of the USS software was one of the more computationally expensive functions. As such, operation creation was a primary focus for Demonstration 4 load testing. The operation creation and submission function includes generating a simulated drone and associated metadata, generating the operational intent volumes, then submitting the operation request to the USS software. The following series of figures breaks this complex function down to its subsidiary components and explores the simulation and USS software performance for each component.

Figure 14 details IT infrastructure performance during Demonstration 4.

Figure 14: Demonstration 4 Metrics Overview

Demonstration 4 Metrics (10,000 – 50,000 UAS)					
# of Simulated UAS	10,000	20,000	30,000	40,000	50,000
Receive Bandwidth (Mbps)	773.3995 Mbps	1517.929 Mbps	1731.7071 Mbps	2379.0460 Mbps	3773.6463 Mbps

Transmit Bandwidth (Mbps)	772.2311 Mbps	1515.1678 Mbps	1605.7002 Mbps	2379.3614 Mbps	3625.4800 Mbps
CPU Utilization (vCPUs)	40.6531 vCPUs	46.2828 vCPUs	42.4023 vCPUs	36.6748 vCPUs	54.7153 vCPUs
Memory Utilization (GB)	79.3529 GB	87.0912 GB	95.6100 GB	109.3711 GB	132.4825 GB
Disk Use for Storing Telemetry (7days of 1Hz updates) (TB)	100.8 TB	201.6 TB	302.4 TB	403.2 TB	506.4 TB

Figure 15 depicts a histogram of the time required to generate an individual drone while processing a batch of 10,000 drones in Demonstration 4. As can be seen, most drones were created in less than 0.25 seconds each. A subset of the total were created in approximately 1.00 second each with an anomalous minority of drones required greater than 1.25 seconds. Through analysis of these data, it is believed that the subset of drones requiring greater than 1.00 second to be generated is an indicator of a USS software service node approaching its capacity limits. Detailed metrics such as these offer unique insights to the performance of each component within a system of systems.

Figure 15: Demonstration 4 Drone Creation Duration

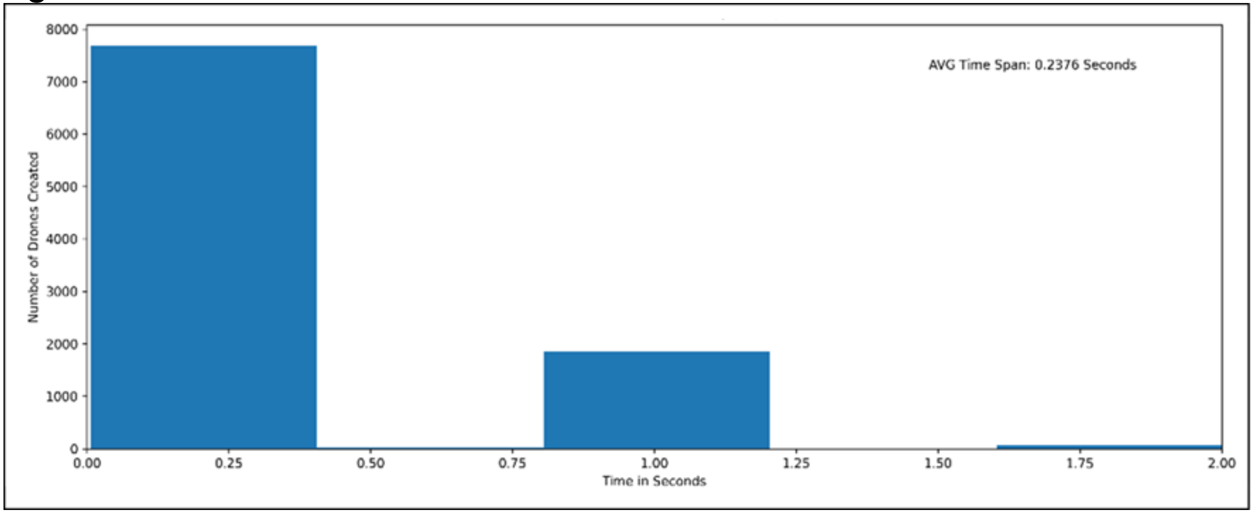

Figure 16 depicts a histogram of the time required to generate the volumes associated with each operation. This function is considerably more complex, and therefore more computationally expensive, than the process of generating drones. This is evidenced by the higher average time to create the volumes across all operations, as compared to the drone creation duration.

Figure 16: Demonstration 4 Operational Volume Creation Duration

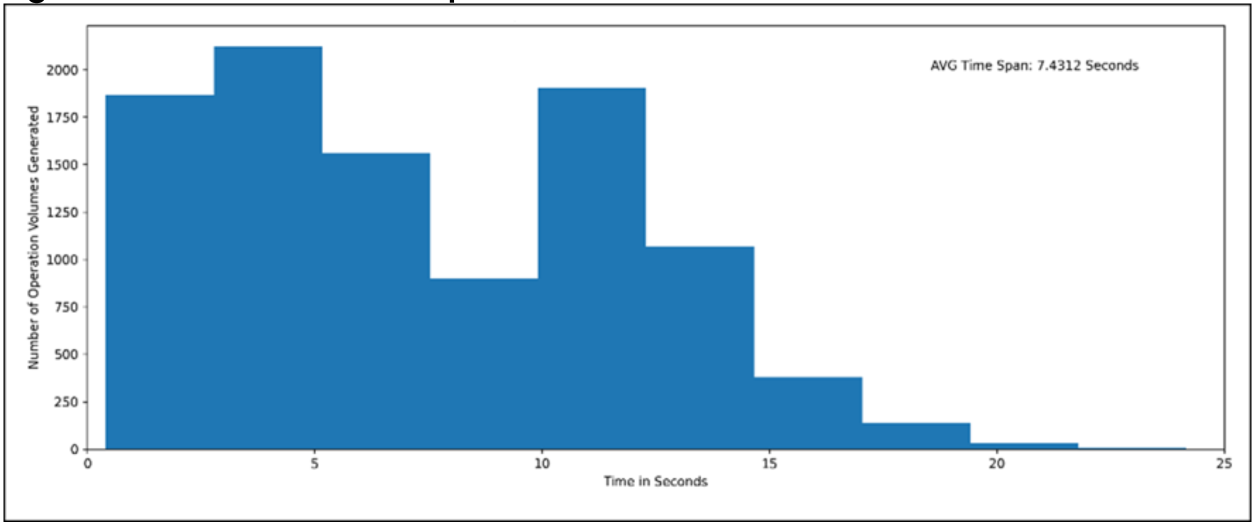

Figure 17 depicts a histogram of the time required for the simulation software to submit an operation to the USS software. The performance of this function models the performance of the drone creation function. As can be seen, a subset of the operations equal to what was observed for the drone creation data required additional time to process. This repeated degradation of performance across two independent functions supports the hypothesis of a node on the USS software services suite approaching its capacity limits.

Figure 17: Demonstration 4 Operation Submission Duration

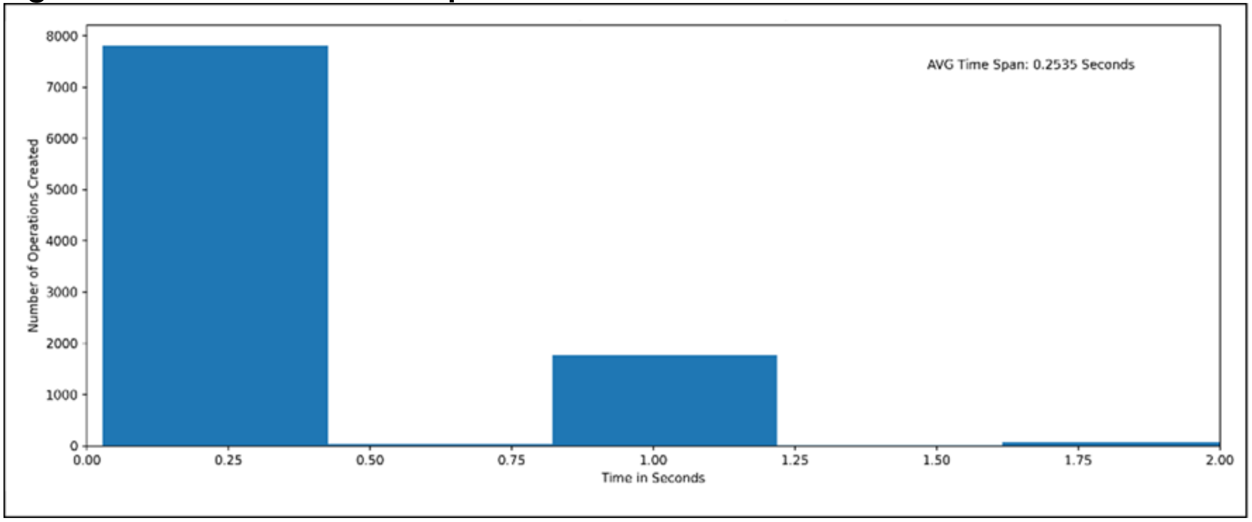

Figure BD depicts a histogram of the total time required to create and submit an operation to the USS software. In other words, this is a cumulative total of the histogram data presented in Figures 15-17. As the figure depicts, the average time to submit an operation was approximately 15 seconds with a distribution of duration performance spanning 0.5 - 35 seconds.

Figure 18: Demonstration 4 Total Time Duration for Operation Creation and Submission

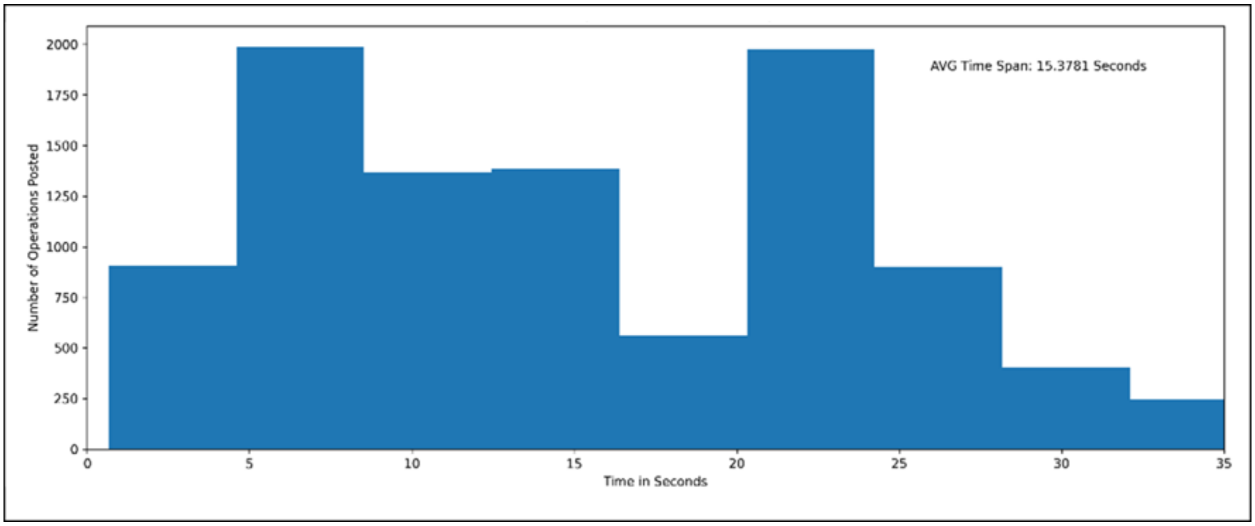

The operation creation and submission durations observed in Figure 18 posed certain challenges in Demonstration 4 load testing due to the total time required to generate new operations for each test iteration. As such, focus was placed on optimizing the operation creation and submission flow, centering on the most computationally expensive step of volume creation. Figure 19 depicts a histogram of the time required to create and submit 10,000 operations after the implementation of the optimized function. Comparing the average time spans in Figures 18 and 19, it can be seen that the optimized function decreased overall operation creation and submission time by approximately 12.5 seconds per operation (~80% improvement). This significant reduction in operation creation and submission time greatly facilitated the ability of rapidly testing up to 50,000 simultaneous operations over the course of Demonstration 4.

Figure 19: Demonstration 4 Improved Algorithm for Operation Creation and Submission

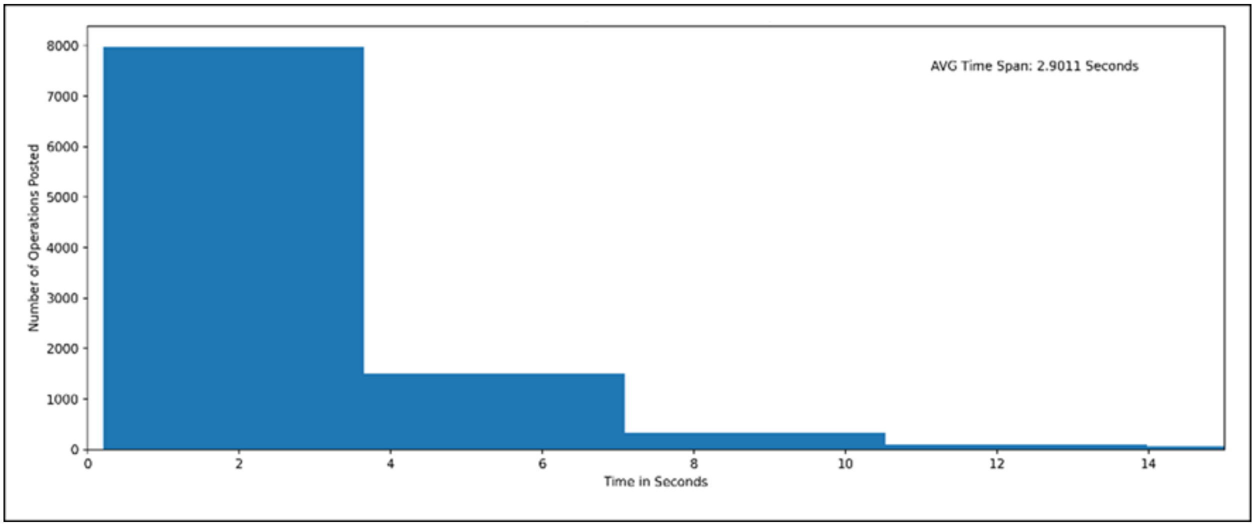

14.0 Lessons Learned and Recommendations

Lessons Learned:

Lessons learned throughout each demonstration were captured and previously submitted to the FAA as separate deliverables. The lessons learned highlighted below are detailed in Appendices 1-4.

- IT Lessons Learned
  - Cloud native architecture is key to building scalable and portable applications like UTM.
  - Network capacity sufficient to handle high demand is required to deliver UTM services.
  - Deployment automation allows for rapid scalability of UTM into different environments.
  - Packets per second rates can be a limiting factor during periods of high demand.
  - Overall network design should be considered when deploying the UTM.
  - USS services issues can be caused by connectivity disruptions.
  - Data storage will be a major cost center of the UTM system.
- USS Software and Simulator Lessons Learned
  - Aircraft characteristics had no appreciable impact on UTM performance.
  - Simulation automation is crucial for rapid iteration and scaling of testing.
  - Live flights verses simulation had no impact on USS software functions or performance.
  - Storing telemetry requires highly capable and high-capacity data storage solutions.

    - UTM operating environments can function under and scale to meet extremely high loads.
- RSO/ BVLOS Lessons Learned
    - Standardized C2 procedures and checklist development resulted in streamlined operations.
    - UTM training provided crews an opportunity to enhance efficiency prior to testing.
    - Adequate field operations equipment is vital to a successful test environment.
    - Contingency plans and backup options were critical to successful RSO/BVLOS operations.

Recommendations:

Recommendations moving forward with this work include:

- Continue to evaluate a diverse array of telecommunications, infrastructure, and cloud provider solutions to provide the UTM as a service, meeting and exceeding forecast requirements.
    - o Performance and cost of UTM will be dependent on deployment methods and technology solutions utilized. The FAA should continue to assess the deployment architecture for UTM such as public, private, and hybrid cloud resources.
- Continue to study the command, control, and communications (C3) requirements for RSO operations.
    - o Due to the complex yet highly valuable nature of BVLOS and remote split operations, the FAA should continue to study methods to lower barriers to entry for industry to conduct these types of operations. FAA approval of part 107 waivers proved to be challenging during the execution of this contract, resulting in a necessity to conduct "de facto" BVLOS rather than true beyond visual line of sight operations. Industry and FAA standardization of risk mitigation methods, aircraft certifications, and standard operating procedures for RSO and BVLOS operations may assist in streamlined processing of waiver applications while maintaining safety and compliance in the NAS.
- Establish data storage requirements for UAS operations and telemetry in the NAS.
    - o Currently, USS systems like ANRA's SmartSkies are not required to store telemetry data per the UTM construct. The ability for the FAA to retrieve UTM data will be critical for conducting accident investigation, trend analsys, and maintaining historical databases as well as maintaining the safety of the NAS
- Study UTM deployment methods that physically and logically segregate inter-USS, USS to FIMS, and UAS operators to their own UTM data planes.
    - o By isolating different types of data traffic, each layer of the UTM system can be deployed to infrastructure tailored to its intended purpose. Where high

throughput and security are required, USS to FIMS and inter-USS communications can be connected on dedicated circuits. Where public access is required, such as for UAS operators, traditional ingress methods can be used. Evaluating these options results in a more secure and effective UTM system.

15.0 Final Thoughts and Future Implications

This BAA effort conducted a detailed exploration and analysis of high-volume UTM operations and the associated impacts to USS software and IT infrastructure. The optimizations applied and lessons learned produced have resulted in a USS deployment and IT configuration that can scale to 50,000+ simultaneous operations without significant user or system challenges.

Future research efforts should further the developments of BAA1 by focusing on the performance of UTM software and infrastructure in a UTM operational environment with multiple USS deployments. This inclusion would introduce numerous, coordinated data exchanges between the USSs. Exploring these exchanges at high operational volumes through new test flows and metrics would offer valuable insights, as it is currently envisioned, to the federated UTM architecture at scale. This type of test exercise would provide additional actionable data to the FAA regarding the performance and scalability of the UTM.

Further tests would demonstrate the impact of inter-USS communications on system performance. This could come in the form of a dedicated USS to USS network, a USS to FIMS network, and the public internet to support individual UAS connections to the UTM. By removing inter-USS and USS to FIMS communications from the internet, performance and security are greatly increased while simultaneously decreasing cost. As for the future implications brought to light during BAA1, data storage stands out among the rest.

During testing, Switch and ANRA Technologies identified that storage of UAS telemetry will be a major cost to the UTM system. Data generated by each UA would need to be stored such that the FAA could assess UA involvement in an incident or accident as well as for statistical analysis. To store this data in one second intervals for a 7-day history of 1,000,000 UAS, the FAA would need to procure approximately 40 petabytes (PB) of data storage (40,000 terabytes, or 40,000,000 gigabytes). This 40 PB consists of multiple mirrors, and additional capacity for high demand times. When using public cloud offerings, this would alone represent a > $800,000/ month expense to support the UTM whereas utilizing hybrid cloud solutions may result in a lower cost of ownership. If expanded to support a greater retention requirement such as 45 days, this would grow to approximately 250PB of storage. To support this amount of highly accessible data, enterprises regularly utilize on-premises solutions to meet their requirements and the FAA should consider the same.

Exploration and implementation of alternative solutions for data storage, IP services, and UTM compute will ensure the scalable, reliable, and sustainable growth of the UTM and allow for full integration of UAS into the NAS.

16.0 Appendices

Appendix 1: Demonstration 1 Lessons Learned

1. Lessons Learned Regarding IT Infrastructure

Deploying the Switch hybrid cloud platform with support for Kubernetes clusters provided an excellent learning experience that will carry over to follow on demonstrations. In the preparation phase for Demonstration 1, Switch deployed an enterprise network, a hybrid cloud using Dell VxRail, and built Kubernetes clusters using Rancher. Following that preparation, Switch assisted ANRA Technologies to deploy the SmartSkies CTRTM platform, and supported ANRA during simulations. This list below identifies key lessons learned regarding the underlying infrastructure that must be considered and/or acted upon as the effort transitions to Demonstration 2.

- **Cloud Native Architecture:** Software designed to be cloud native and built to run in Kubernetes clusters allowed for applications to be deployed rapidly, modified, and scaled to meet client-side demand.
- **Network Capacity:** During Demonstration 1, network throughput appeared to have the most significant increases during periods of high demand with increases of up to 700% above baseline values. In other environments network performance could prove to be a limiting factor to provide services to large numbers of UAS. In follow-on demonstrations, network throughput will be closely monitored and more data will be collected to identify what types of services see increased network demand during testing. Current peak values recorded were approximately 100Mb/s. Switch's network for this project is capable of up to 25Gb/s of throughput and will be able to support much more traffic than is required under this contract.
- **Deployment Automation:** During preparation for Demonstration 1, Switch rebuilt and redeployed the SmartSkies CTRTM platform no less than 3 times in order to take advantage of new features available from services like Rancher. This process can be time consuming if done manually. Switch has begun to build automation scripts that allow the deployment of a single deployment file that will build the entire SmartSkies CTRTM in a designated Kubernetes cluster. In the event that clusters need to be rebuilt and redeployed, this capability will give greater flexibility and greatly reduce deployment turnaround time.
- **Kubernetes Persistent Storage:** When first deployed, the original Kubernetes cluster did not have the capability to provision persistent storage volumes in vCenter. Given the ephemeral nature of Kubernetes pods, data does not persist when individual pods are restarted or redeployed. During the initial testing phase, it became apparent that persistent storage would be critical to ensure data retention during testing. Following that finding, Switch deployed a fresh instance of Rancher and built new Kubernetes clusters to support the ANRA SmartSkies CTRTM platform that had the ability to claim persistent volumes from a cloud provider, in this case, our local instance of vCenter. Any future clusters built under this effort will have this capability.

- **Automated Metrics Parsing:** During Demonstration 1, Switch collected data regarding CPU, Memory, and network throughput from our Kubernetes cluster. These data points come in the form of a comma separated values file that can be opened in standard spreadsheet software like Microsoft Excel. The first data collected was manually parsed in Excel to extract key data points. During this exploration of the data, it became evident that this approach would not be scalable as future demonstrations unfolded. Switch created parsing software that allowed these data to be captured and parsed in batches for every test card. Automation metrics parsing dramatically decreased time required to capture key metrics from the active Kubernetes clusters. This software will be continually updated to ensure it can be used for metrics parsing during all future demos.

2. Lessons Learned Regarding USS Software and Simulator

Demonstration 1 preparations involved an iterative, collaborative process between ANRA Technologies and SWITCH. These preparations highlighted several challenges involving the deployment of the USS software, SmartSkies CTRTM, on the IT Infrastructure. As outlined above, these challenges were resolved or remediated in order to accomplish the objectives of Demonstration 1. The list below identifies the key lessons learned regarding the USS software, simulation capabilities, and test methodologies that will be acted upon as this effort transitions to Demonstration 2.

- **Human in the Loop Test Execution:** The Demonstration 1 test card was built with the expectation that a human would trigger the major test events. This was a conscious decision as having a human in the loop avoids constructing a deterministic test environment and allows for more flexibility exploring various test parameters and permutations. Having a human in the loop does introduce a certain level of test variability. To address this, metrics were derived to adequately capture trigger event timestamps and other input parameters and thus contextualize the results. From preliminary Demonstration 1 results, it seems that these test data capture methods produce sufficient data to generate meaningful results. As such, this human in the loop approach will be maintained for Demonstration 2.
- **Operation Update Success Rate:** Demonstration 1 captured a metric that measured the success rate of operation updates. Through testing and further exploration of the UTM topic, the team believes that this metric should be reconsidered or perhaps removed. Because this is a closed UTM operational environment we are aware and in control of all operations taking place. As such, the operation update success rate can ultimately be predetermined based on how and where the updates are taking place. Demonstration 2 tests should consider refocusing this metric to the number of impacted operations or perhaps the time required to notify impacted operators. This would provide greater insight into the overall impact of the operation update and the performance of the underlying architecture.

- **Aircraft Characteristics:** One of the major independent variables explored in Demonstration 1 was the impact of different aircraft characteristics. Demonstration 1 data revealed that varying the aircraft characteristics had no appreciable impact as observed through reported metrics. This aligns with the team's expectations on this topic: Aircraft performance might alter the flight path a drone follows to execute an operation, but it will not alter the packet rate or size of telemetry that must be generated by the simulator, ingested by the USS, and handled by the IT infrastructure. Demonstration 2 tests should consider removing this test variable to allow for deeper concentration on variables that offer more insights to UTM performance.
- **USS User Interface:** For Demonstration 1, testers benefitted from utilizing the SmartSkies CTRTM user interface to observe the progress of the tests. Throughout the tests, automated scripts captured the necessary data to enable reporting of metrics, so use of the user interface was purely an aid for overseeing the tests. While the user interface was a helpful aid for 250 operations, it will not be as useful for 10,000 operations. Visual tools such as operation lists, telemetry maps, and push notifications do not scale to these levels. As such, the background automation will be relied upon to monitor and measure the Demonstration 2 test execution. Alternatively, certain user interface components might need to be altered in order to account for the high volume of operations.
- **Operation Creation Performance:** For Demonstration 1, it would take a couple seconds to create a drone, generate an operational volume, strategically deconflict the operation, and finally submit the operation to the Discovery and Synchronization Service (DSS). Generally speaking, this performance aligns with expectations considering the numerous steps and communications among the various UTM entities. Regardless, this operation creation performance must be drastically improved in order to effectively manage the test iterations outlined for Demonstration 2. Performance gains will be pursued through IT Infrastructure scaling and through improved batch processing in the simulator.
- **Post-Hoc Metrics Analysis:** As mentioned above, specific instrumentation was implemented to automatically capture the raw data used to feed the identified metrics. Through the analysis of the data captured in Demonstration 1, several additional metrics of interest were identified that required post-hoc processing of the core data and primary metrics. This analysis was conducted manually to affirm the value of the resulting information. For this information flow to persist into Demonstration 2, this processing will have to be automated.

Appendix 2: Demonstration 2 Lessons Learned

1. Summary

Demonstration 2 preparations involved an iterative, collaborative process between Switch and ANRA Technologies. These preparations highlighted several architectural challenges involved in scaling of the USS software, SmartSkies™ CTR, on the Switch hybrid cloud platform. As outlined below, these challenges were resolved or remediated in order to accomplish the objectives of Demonstration 2. The list below identifies the key lessons learned regarding the Switch's hybrid cloud platform and IT infrastructure as well as ANRA's USS software, simulation capabilities, and test methodologies.

In order to support the large increase in scale from Demonstration 1 to Demonstration 2 (250 operations to 10,000 operations), Switch re-architected numerous services within the Kubernetes cluster to support the ANRA USS. These changes allowed for increased scale and scalability, paving a pathway for high-capacity stress testing in Demonstration 4 and provide insightful lessons learned about the nature of deploying applications like the UTM in a production environment. The lessons learned for Demonstration 2 are listed below and involve simulating over 10k concurrent operations connecting to the USS platform. Switch and ANRA also successfully simulated up to 20k simultaneous operations during further testing.

These lessons learned along with past and future lessons learned from this effort will be used to influence the final report. Switch and subcontractors continue to evaluate the UTM system during this effort to ensure that the final report provides meaningful guidance regarding the scaling of the UTM system to support the NAS.

2. Lessons Learned Regarding IT Infrastructure

- **Kubernetes ingress services:** During post simulation analysis of Kubernetes cluster telemetry, Switch and ANRA discovered network loading inconsistencies across individual services running the ANRA USS. Upon further investigation, it was determined that service ingress into the cluster must be configurable and routable to the individual Kubernetes pod level to ensure consistent load across the application services. These changes were implemented and resulted in even distribution of load across the cluster allowing for more efficient operation of the USS given the same amount of compute resources.
- **Load balancing:** In addition to the Kubernetes ingress changes made, Switch and ANRA implemented further load balancing across the application. These load balancing changes complemented the ingress changes discussed in lesson learned 2-1. This enables inbound data to be routed to individual Kubernetes pods and evenly shared across sets of pods. Switch also implemented a high availability configuration on all load balancers. This practice ensures constant application availability during high load and in the event of a service interruption, client devices would experience no downtime.

High availability load balancer deployment consists of a pair of redundant load balancers that share the exact same configuration. This pair of load balancers share one virtual IP address using virtual router redundancy protocol (VRRP), making either load balancer available at one IP address in the event of a failure.

- **Packets per second rates:** Data transfer rates in terms of packets per second (pps) were of particular note during Demonstration 2. For example, across the entire cluster running the ANRA USS, Switch observed over 1.6 million pps during simulations of 10,000 UAS. Switch observed rates of over 40,000 pps during Demonstration 2 on individual components such as telemetry ingress.

  If extrapolated to support 1 million UAS connecting to a given USS, Switch expects to see pps measurements of approximately 4 million pps for telemetry ingress alone. These high throughput events could appear as a type of denial-of-service attack without the proper network design and components. During Demonstration 2 Switch and ANRA experienced a network denial of service event during simulations of 10,000+ UAS due to the sheer number of incoming telemetry streams. This was mitigated by allowing known network traffic, adjustments in firewall configuration, and rearchitecting how data flowed from simulation virtual machines to the USS.
- **Overall Network Design Considerations:** Due to high throughput required on the network, it is necessary to utilize highly capable hardware able to handle the load of hundreds of thousands of concurrent operations sending telemetry to the USS as well as components able to handle the additional traffic between USS services. In order to scale the UTM to a nationally available service, it would be advisable to implement solutions such as geo-load balancing or a separate inter-USS network that does not traverse the public internet. Geo-load balancing allows for distributed ingress based on the traffic source location and the inter-USS network could allow for secure, very high throughput, and highly reliable connection between USS providers and the FAA. For example, a UTM ingress point could be deployed in major metropolitan areas where any traffic originating from the surrounding areas would enter the UTM network with as little latency as possible. This also allows for UTM services to be hosted at the network edge, providing for a redundant and high performance UTM system. These solutions combine to create a dynamic, scalable, and secure network that can support the UTM across the Nation.

3. Lessons Learned Regarding USS Software and Simulator

- **Adjustment of Test Sequence:** One of the lessons learned from Demonstration 1 identified that several of the test steps and permutations needed to be adjusted to offer more valuable insights to UTM performance at high operational volumes. As such, for Demonstration 2 the test sequence was updated to focus on computationally expensive UTM functions that are foundational to the UTM concept. At a high level, these functions included operation planning, streaming of telemetry, and conformance monitoring.

- **Adjustment of Test Metrics:** Similar to the lesson learned above, Demonstration 1 highlighted several opportunities for how test metrics could be adjusted in order to better describe UTM performance during Demonstration 2. As part of this metrics analysis, the team turned to the draft ASTM UTM standard as a reference for meaningful performance measures and approaches within the UTM domain.

  This metrics update also serves to set the stage for Demonstration 4. One of the challenges identified as part of Demonstration 2/4 planning was the definition of "UTM Failure." As was experienced throughout Demonstration 2 preparations and testing, failure rarely took the form of a downed service or blocked function. Rather, it took the form of degraded performance within a certain service or function. Utilizing performance requirements-based metrics will enable the team to establish performance values measured during Demonstration 2 which can then be used to inform failure thresholds in Demonstration 4. This approach provides meaningful metrics for each test event, maintains continuity between the two test events, and provides a means to derive quantitative definitions of "UTM Failure."
- **USS User Interface:** One of the Demonstration 1 lessons learned identified the challenges associated with scaling user interface (UI) software and functionality to support the large number of operations within Demonstration 2. Seeking to explore some of these challenges, the team made adjustments to the user interface that enabled UI use throughout the entirety of Demonstration 2. These adjustments embodied several human factor design iterations that focused on how to properly segment and present large amounts of information to a human consumer. Through these updates, the team achieved UI functionality and user comprehension at high operational volume. Additionally, similar to Demonstration 1, the UI proved to be extremely useful in back-end design iterations and testing.
- **Simulation Automation:** The simulator functionality suite was augmented for Demonstration 2 to enable automated orchestration of the various steps of the test sequence. This automation was necessary as manual test sequence triggers utilized in Demonstration 1 would not scale for Demonstration 2. In developing these orchestration functions, focus was placed on the operation creation process. Operation creation is a multi-step process that includes drone creation, volume generation, and operation creation. To execute this workflow for the loads tested in Demonstration 2, the simulator was finely tuned to execute parallel processing pathways while maintaining proper workflow sequencing. Through these efforts, all operations within each test iteration could be rapidly generated. During Demonstration 2, the team was able to generate, submit, and start 10,000 operations in less than 60 seconds, making rapid iteration of testing possible.
- **System of Systems Optimization:** The Demonstration 2 software test bed embodies a highly complex system of systems. As mentioned above in the IT Infrastructure lessons learned, there are numerous platforms working in unison to achieve the processing power necessary to support Demonstration 2 loads. In addition, the USS Software itself consists of numerous microservices working

in conjunction to plan, fly, and monitor each operation. The stress testing of this architecture through Demonstration 2 has yielded several discoveries of how to more optimally configure the environment and services to achieve higher performance. Some key examples include:

- Automated Pod Scaling: As mentioned above, the Demonstration 2 test sequence was adjusted to focus on computationally expensive UTM functions. However, these functions are not consistently loaded over the course of the test. For example, flight planning functions and associated services are only utilized at the onset of the test. Automated pod scaling enabled the services infrastructure to dynamically scale in response to the immediate processing needs. This approach supports required processing levels without the need to maintain underutilized pods.

- Database Optimization: For Demonstration 2, the database connections were optimized to work with load balancers and multiple pods of the USS software services. Database pooling was tuned to achieve high performance and reduced data insertion failure rate.

Appendix 3: Demonstration 3 Lessons Learned

1. Summary

Demonstration 3 preparations involved an iterative, collaborative process among Switch, ANRA Technologies, and NIAS operators. Demonstration 2 illustrated that the USS Software and Simulator could adequately manage 10,000 concurrent operations. With Demonstration 3 operational scales reverting back to 250 simulated operations and 1 live operation, the testing focus was placed on remote command and control and operator display factors as opposed to USS Software and IT Infrastructure performance. The list below identifies the key lessons learned regarding the USS software, simulation capabilities, live flight operations, and test methodologies.

To support Demonstration 3, Switch provided to NIAS access the Citadel UAS Test Range in Reno, NV and the SwitchSIGHT Command Center located in Las Vegas, NV. Live operations were conducted with the remote pilot-in-command located on-site with the UA in Reno, NV and the Remote operator located in the Switch Sight Command Center. The connection between the Citadel Test Range and the SwitchSIGHT Command Center was enable by the Switch SUPERLOOP, consisting of diverse fiber pathways with less than 7 ms of round-trip latency. This additional latency was imperceptible to UAS Pilots during the operations. During live operations, ANRA also simulated 250 UAS in the surrounding airspace.

These lessons learned along with past and future lessons learned from this effort will be used to influence the final report. Switch and subcontractors continue to evaluate the UTM system during this effort to ensure that the final report provides meaningful guidance regarding the scaling of the UTM system to support the NAS.

2. **Lessons Learned Regarding Live Flight Operations**

- **Standardization of Command and Control (C2) Procedures:** Throughout the shakeout period of Demonstration 3 NIAS had the opportunity to study the communications between the remote pilot stationed in the command control room over 350 miles away (Las Vegas) from the on-sight pilot in command and aircraft (Switch Citadel Campus Reno, NV). Many communications issues were observed and debriefed. Each day of the shakeout the team continued to improve on the below checklist to eliminate or control the potential errors. The multiple lessons learned in remote split operations (RSO) are captured in the standardized RSO checklist for the command-and-control mission director and remote pilot.
- **UgCS and USS Familiarity Training:** Key to the success of Demonstration 3 was the shakeout week of pilot training dedicated to pilot and mission director familiarity with UgCS and USS software training given by ANRA. Even with the week dedicated to training on these two software applications, given the opportunity to conduct similar operations in the future we would recommend an increase in the time allotted to ten days. The NIAS and Utah State University flight crews were highly experienced and were veterans of NASA's TCL 3, 4, and the FAA's UPP, and the timeline was tight despite the experience level.
- **Multiple Applications Require Multiple Displays**: Initially the on-site flight crew was only provided with one GCS display monitor for at least five applications. With only one GCS display monitor the on-site, the flight crew became overwhelmed and overtasked leading to frustration. The on -site flight crew team added a second and then a third monitor which reduced the workload switching between applications. As an initial recommendation, the team recommends a maximum of 2 apps per one monitor for human factors considerations.
- **Backup Plans Critical to RSO/BVLOS Operation:** Throughout Demonstration 3 the need for backup plans became readily apparent. The most pressing need for backup plans were in the areas of aircraft, flight crews, and location. NIAS tested out the Drone America built NAVX aircraft for Demonstration 3 and those tests were initially successful. Over time and software upgrades, the NAVX aircraft became increasingly unreliable triggering a team decision to move the Parrott aircraft as primary. Backup qualified flight crews and locations were added as well to ensure program success. NIAS selected longtime teammate Utah State University (USU) and their aircraft the Great Blue as the primary alternate. Switch, LTD. added FLIR as a secondary backup. The alternate flight crews and aircraft were developed, trained, and treated as though they would become the primary flight crew. Having qualified and ready flight crews in different locations (i.e., Reno, NV. and Logan, UT) allowed for mission flexibility especially as it related to weather considerations.
- **Route Hygiene:** On two different occasions, the team learned the need to clear out old routes from the UgCS software such that new routes could be recognized. The lesson learned is that old routes can back up the UgCS system with corrupted files and lead to delays. Proper hygiene includes clearing out old routes every day from the UgCS and uploading new routes every day.

- **Connectivity Checks Are Most Time Consuming:** Future shakeout efforts regarding RSO/BVLOS operations should allow for several weeks of connectivity testing. With proper upstream connectivity testing downstream results are better achieved. As thought and planned, connectivity proved to be the most time consuming and critical component of RSO/ BVLOS operations.

3. Lessons Learned Regarding USS Software and Simulator

- **Operator Functions and Associated Training:** Through the hotwash tests leading up to Demonstration 3 flight operations, several operator checklist items were identified and incorporated into the USS software workflow. This included a cross check of UgCS connectivity to the USS Software prior to flight as well as incorporating proper hygiene regarding the termination of operations. Once incorporated into the checklist, these functions were cleanly accomplished by the operators, resulting in uneventful USS functionality for the duration of each operation. These lessons learned have been incorporated into ANRA training materials for future deployments.
- **Remote Pilot User Interface:** ANRA's USS software offers a specific user role for remote pilots that enables users to focus on their respective operations. This role was utilized by the Demonstration 3 remote pilots for planning and monitoring their live operations. From remote pilot accounts, this role and its specialized user interface functions proved effective in maintaining situational awareness throughout the live flights even while 250 simulated flights were concurrently operating in the same area.
- **Elevated Access User Interface:** In addition to the remote pilot user role, ANRA's USS software offers an elevated access user role and associated viewing capabilities. Through this role, all operations within a specified geographic area can be viewed, regardless of the pilot conducting the operation. Throughout Demonstration 3, this user role was utilized to understand the challenges associated with overseeing and monitoring 251 simultaneous operations within a proximate geographic region. The following offers specific findings from these tests: This role proved helpful overall accessing operational details and in maintaining general situational awareness of all operations within the operating environment.
    - Status update push notifications were too numerous for monitoring status changes across 251 operations. These findings have informed more aggregated methods of identifying status updates at higher operational volumes.
    - Methods of more effectively filtering and/or searching for operations of interest have been highlighted through this experience.

- **UAS Integration with USS Software:** For Demonstration 3, ANRA provided its custom-built client for integrating the UgCS ground control system with the USS Software. This client proved to be an effective means of establishing and maintaining UAS telemetry feeds throughout the live flights. This integration workflow requires a human in the loop to establish the telemetry feed. While this method was effective for one live flight, this approach may not effectively scale to higher numbers of live flights. As live flight operations scale, cloud-based system-to-system interfaces will have to be established to automate the telemetry connection process.

- **Live Flight Impacts:** Inclusion of live flights in Demonstration 3 incorporated a unique characteristic as compared to the other demonstrations within this effort. It was observed that this factor had no impact on the USS Software functions or performance. The USS software does not distinguish between live and simulated operations or telemetry. As such, the inclusion of live flights did not create or require any difference in processing methods. Human factors associated with the live flights have been outlined in the previous Lessons Learned.

Appendix 4: Demonstration 4 Lessons Learned

1. Summary

Demonstration 4 preparations involved an iterative, collaborative process between Switch and ANRA Technologies. As observed in Demonstration 2, Demo 4 introduced several new architectural challenges involved in scaling of the USS software, SmartSkies™ CTR, on the Switch hybrid cloud platform. As outlined below, these challenges led to a greater understanding of some of the nuances of running the UTM at a national scale. Demonstration 4 focused on stressing the IT Infrastructure, USS Software, and Simulator beyond the thresholds tested in Demonstrations 2 or 3. The lessons learned below highlight the primary findings regarding UTM failures and challenges in a high load operating environment.

During the execution of Demo 4, Switch and ANRA conducted an iterative process of adding simulations in increments of 10,000 operation until a system failure was identified. Then the issue was addressed, documented, and the system was reset. This process allowed for the capture of data in the above stated increments from 10,000 to 50,000 operation giving a clear picture of system requirements with increasing load. This process also produced insight into what components of the UTM system would be impacted by high demand.

These lessons learned along with past lessons learned from this effort will be used to influence the final report. Switch and subcontractors continue to evaluate the UTM system during this effort to ensure that the final report provides meaningful guidance regarding the scaling of the UTM system to support the NAS.

2. Lessons Learned Regarding IT Infrastructure

- **UTM Service Issues Caused by Connectivity Disruptions:** During Demonstration 4 after multiple iterations of submitting tens of thousands of operations to the USS and resetting for follow on tests, Switch and ANRA observed telemetry handling issues when interrupting connections to said services. This observation, while intentionally and knowingly caused while resetting tests, highlighted the need for high capacity/availability connectivity to the internet. The UTM endpoints that individual operations send to must remain available and be immune to service interruptions caused by upstream carrier networks. This can be accomplished by building services with multiple connections from numerous internet service providers. Additionally, services should be deployed within infrastructure built for fault tolerance and high performance, allowing for consistent delivery of services.

- **Data Storage Costs:** As also discussed in Demonstration 4 lesson learned 3.1, data storage of historical telemetry data will be a large cost center in supporting the UTM. For example, 10,000 operations generate about 10 Gigabytes (GB) of data per minute. If a 7-day running history were required for the UTM running at scale, storage requirements grow rapidly. Supporting a 7-day history for one million operations would require approximately 10 Petabytes (PB) of storage. With mirrors for redundancy and additional capacity for periods of high demand, this could easily grow up to 40 PB or 40,000,000 GB. Costs from data storage must be carefully considered when evaluating solutions for meeting this need.
- **Ingress Throughput:** Telemetry streams inbound to the UTM system of 60 Megabits per second (Mbps) were observed per 10,000 operations. The relationship between the number of operations and the inbound data rate were linear as expected. These data rates represent the inbound throughput required to accept telemetry from a given number of operations. In this implementation, the Switch hybrid cloud environment is capable of up to 25 Gigabits per second (Gbps) since the simulated telemetry and the UTM cluster occupied the same environment. In a production implementation however, networking infrastructure would be required to handle a given number of expected operations feeding telemetry into the UTM. Given the rates observed, to support one million concurrent operations, the UTM network would need to be capable of at least 5Gbps.
- **Cluster Throughput**: Throughput performance internal to the Kubernetes clusters that run UTM services will be crucial when building scalable UTM applications. During simulations of up to 50,000 operations connecting to the UTM, Switch observed combined throughput rates of up to 7.4 Gbps inside the Kubernetes clusters across all services. The relationship between the number of operations and the cluster throughput appeared to be linear. High performance clusters will be essential to building a production UTM. Other methods to reduce required cluster throughput and increase scalability and redundancy would be a combination of geo-load balancing and automated deployment of new clusters as the demand increases.

3. Lessons Learned Regarding USS Software and Simulator

- **Storing Telemetry**: Ongoing UTM CONOPS and associated standards discussions often refer to the need to store operational data for auditing purposes. Demonstration 4 highlighted the challenge USSs and/or IT Infrastructures will face in storing telemetry data for high volumes of operations. As a rough metric, 10,000 operations operating for one minute generated approximately 10 GB of data during Demonstration 4. Extrapolating this out to higher numbers of operations and longer operational durations, one can see the eventual challenges in storing and processing telemetry data. Data retention policies will also be a major contributor to these challenges. It is intended that this lesson learned offers some context while considering policies and practices for operational data logging and auditing.

- **Operation Planning**: Operation planning algorithms and associated data exchanges have been the focus of optimization efforts throughout the Demonstration series as they represent some of the more computationally expensive UTM functions. This effort yielded several updates in operational planning data models and processing sequences that achieve high planning rates while maintaining USS interoperability compliance.
- **Operation Volume Scaling**: IT Infrastructure, USS Software, and Simulator improvements made over the course of Demonstrations 1-4 achieved a point where the number of simultaneous operations supported by the USS Software scaled linearly with the devoted IT Infrastructure hardware. This software environment characteristic was validated up through 50,000 concurrent operations. This finding demonstrates that with optimized USS Software and adequate supporting IT infrastructure, UTM operating environments can be configured to effectively scale to extremely high loads (as compared to today's operating levels).
- **Minimum Configurations for the USS**: Conversely to the linear IT infrastructure scaling that led to Demonstration 4 lesson learned 3.4, several Demonstration 4 tests were devoted to exploring minimum IT infrastructure resources necessary to maintain stable USS function. Through this testing, the messaging brokers were identified to be the software bottlenecks while exploring UTM failures. While parameterizing Demonstration 4 failure cases at 10,000 concurrent operations, it was determined that a minimum of three user datagram protocol (UDP) nodes were necessary for managing the telemetry feed from the simulator to the message broker and one pod was necessary for managing the interservice message broker exchanges.

Appendix 5: Risk Assessment Template for Day/Night BVLOS Operations
Risk Assessment Template for Day/Night BVLOS Operations

UAS Lost Link / RTH

	RTH & Autoland / Smart RTH	RTH	'kill' switch	Orbit	Manual Only
Flying Under:					
COA	1	2	2	3	4
Section 44807	1	2	2	3	4
Part 107	1	2	2	3	4
None	3	4	4	5	NOGO

PIC Experience

	Total Time	>100	50-100	20-50	<20
Time On Type					
>10		1	1	2	3
5-10		1	2	3	4
2-5		2	3	4	5
<2		3	4	5	6

Characteristics of UAV

Weight-lbs	<5	5-25	25-55	55+
Mission Speed (mph)				
<5	1	2	3	4
>5-20	2	3	4	5
>20-50	3	4	5	5
>50-100	4	5	6	6
>100	5	6	6	6

Operator Experience

	Total UAV/ Sensor	>20	10-20	2-10	<2
Sensor Type					
>20		1			
10-20		1	2		
2-10		2	3	4	
<2		2	4	5	6
1 man crew		3	4	5	6

Detect & Avoid / See-Be Seen

	Operating w/in 500'	> 500'	> 2500'	W/in 5000'	BVLOS
Other	1	TBD	TBD	TBD	TBD
Sense-&-Avoid	1	1	1	1	2
Anti-Collision Lights	1	1	1	2	4
Hi Viz Markings	1	1	2	3	5
None	1	1	3	4	6

Safety/Visual Observer

	Total	>20	10-20	2-10	<2
UAS					
>20		1			
10-20		1	2		
2-10		2	3	4	
<2		3	3	5	6
None		3	4	5	6

Highway Crossings

	UAV <2.5lbs	2.5 – 5 lbs	5 – 20 lbs	20 – 55 lbs	>55 lbs
0	-	-	-	-	-
1 - 2	1	1	2	3	4
3 - 6	1	2	3	4	5
7 - 10	2	3	4	4	6
>10	2	3	5	5	6

Remote Split Operation

	Native / Manufacturer C2	Integrated C2	Public Infrastructure
None	-	-	-
Local TO/Land	2	3	4
Local PIC w/ C2 Override	2	3	5
Local PIC w/o C2 Override	3	4	5
RPIC w/ grnd crew only	4	5	6

Weather

	Forecast Visibility	>3	>1	1	<1
Ceiling					
>1,500		1	2	3	NOGO
1,001-1,500		1	2	4	NOGO
1,000		2	4	5	NOGO

Complexity Factors

1st time crew (<5 hours)	1	
Ops Max Performance	1	1 pt / characteristic (Weight, speed, maneuvers, etc.)
Multiple Aircraft in Area	1/ac	(Must have multiple crews)
>400'AGL	3	
Payload delivery	3	
Operator Training	1	E.g., Currency, Mission
Unqualified Crew	3	/ Unqual crew member
Dusk / Night	2 / 3	
BVLOS – Flight planned route	3	
BVLOS – Non-flight plan	6	

Circadian Shift

	+/- Hours Outside Duty	0 Hours	1-4 Hours	5-8 Hours	9-12 Hours

Expected Crew Duty Day

	<8 Hrs	8-10 hrs	11-12 hrs	12-16 hrs	>16

Normal Duty Hours		1			
Outside normal hrs			2	4	5

Show time thru End of Day	1	2	4	5	Ops Dir Req.

Time of Day / Area Familiarization			
	Familiar Ops Area	New Ops Area	
Day	1	3	
Night	3	5	

Ops Over Infrastructure				
	Level of Coord	Coordinated & Sanitized	Coordinated Only	Not Coordinated
No		-	-	-
Yes		1	3	5

Total Points:

Mitigation Worksheet

Total Unmitigated Risk Points from Page 1: ___________________________

Unmitigated Risk Level:

LOW<15	16 AVERAGE 20	21 MODERATE 25	HIGH >26

Average or Low risk is acceptable without mitigation

MITIGATION MENU

<100' AGL	-2
Rural – No People, Structures, Vehicles	-3
Additional Observer (per observer)	-1
Observer has Pilot Cert.	-3
PIC has Instrument or CFI Cert	-3
Crew Time > 25 flight hours	-2
>250' from hazards	-3
Geo-fence software	-6
Flight planned route	-3
Perpendicular HWY crossings only	-1
HWY crossing <2 sec (generally 30mph+)	-1
Visually clear HWY prior to crossing (VO or sensor scan)	-2
Ops over infrastructure is continuously monitored	-1
Hwy is traffic controlled (subtract up to)	-3

Total Mitigation Points: ______________________

Unmitigated Risk Points ____________ - Mitigation Points ________________ = Residual Risk __________________

LOW<10	10 AVERAGE 17	18 MODERATE 25	HIGH>25

LOW: No further approval necessary.

AVERAGE: No further approval necessary. Preflight briefing must emphasize possible hazards identified.

MODERATE risk must be evaluated and approved by the Operations Officer/Chief Pilot. Mitigating options must be thoroughly briefed to the crew and client.

HIGH residual risk must be evaluated and approved by Director prior to operations

Instructions

Risk assessment and mitigation can only be effectively accomplished by providing objective answers to various mission risks. RPICs shall answer each section as they relate to the planned mission. The RPIC should also reassess for dynamic changes in the plan, environment, or situation.

Answer Matrix

UAS Lost Link / RTH

Match the mission flight authorization to the UAS Lost Link / Return to Home capabilities. The more stringent the authorization and better LL/RTH capabilities result in a lower score.

Characteristics of UAV

Match the maximum aircraft speed to the aircraft takeoff weight. Higher speed and higher weight aircraft will result in a higher score. Note: if the maximum speed and the planned mission speed differ, the RPIC may mitigate by substituting the maximum planned speed.

Detect & Avoid / See and be Seen

Match the on-board detection system capabilities to the maximum planned range of the mission. A more capable detection system will result in a lower score. The RPIC may further lower the score by reducing the planned maximum range of the mission.

Highway Crossings

Match the number of highway crossings by the weight of the aircraft. Fewer crossing and lower weight aircraft will result in lower score. The RPIC may further lower the score by minimizing the number of highway crossings or using the lightest possible UA to accomplish mission objectives.

Weather

Match the ceiling to the visibility. Greater visibility and higher ceilings will result in a lower score. RPICs must use the weather conditions at the operating area and not rely on forecast conditions or even macro weather conditions.

Circadian Shift

Match the time of mission, as it relates to normal duty day, to the number of hours of the duty day shift in time. For example: if normal duty day is 0800 – 1800, and the operation begins at 2300, then there is a 5-hour circadian shift. A 5-hour shift results in a score of 4. The RPIC may reduce the score by allowing the crew to adjust the duty day prior to the planned mission.

PIC Experience

Match the RPIC's hours on the type of UA to the PIC's total UA hours. More time on type and total time will result in a lower score.

Operator Experience

Match the Sensor Op or other crewmember's hours with the sensor to the SO's total UA hours. More time on sensor and total time will result in a lower score. Note: a multi-person crew will result in less risk and therefore a lower score.

Safety / Visual Observer

Match the Safety / Visual Observer hours on the type of UA to the VO's total UA hours. More time on type and total time will result in a lower score. Note: a multi-person crew will result in less risk and therefore a lower score.

Remote Split Operations

Match the planned type of remote split operation to the type of command and control in use. For example: the mission is planned to have a local RPIC and that RPIC has capability to override the off-site PIC. The mechanism to accomplish the C2 at both locations is conducted via a commercial, third party communications network (cellular, fiber, satellite, etc.). The resulting score is 5. If the same communications architecture is in use but the local crew is not 107 certified, the resulting score is 5. The score will be reduced by using local crews for takeoff and landing functions and using native applications vs public infrastructure.

Complexity Factors

UAS operations seldom are routine. The use of the additional complexity factors section allows for additional variables to be accounted for and therefore more accurately reflect actual mission risk. The RPIC should select all complexity factors that are applicable and tally them for the total score. The more routine the mission and the fewer additional factors will result in a lower score.

Expected Crew Duty Day

Match the expected length of day from crew show time to the end of flying to get the section score. Shorter duty days will result in a lower score.

Operations over Infrastructure

If the planned operations are over infrastructure, match the level of planned safety and coordination to get the score. For example: if the planned mission is over/near power infrastructure for the purposes of a survey for the owner and the area is free of non-participants, the resulting score is 1, but if the area has a likelihood of incursion by non-participants, the score is 3. Sanitizing and maintaining awareness of the infrastructure and surrounding area will result in a lower score.

Tally Total Score

This results in an unmitigated total risk score. A score that falls within the Low or Average ranges is acceptable. The RPIC may proceed without any further mitigation.

If the unmitigated total risk score is in the moderate or high range, the RPIC must attempt to mitigate risk by determining if the planned operation includes any items from the Mitigation Menu.

Mitigation Menu

There are situations, environments, and techniques that RPICs may employ to mitigate risk. The RPIC should select all Mitigation Menu items for the planned mission and crew and use available mitigation techniques to try to bring the residual risk to the low or moderate range.

If the residual risk, after mitigation, is low, the RPIC may continue. If the residual risk remains in the average range, the RPIC may continue, but must emphasize the specifics mission risks in the crew or pre-mission briefs.

If the residual risk remains moderate or high, the RPIC must gain authorization at the appropriate level in order to conduct the mission. Even with higher level approval, the RPIC must emphasize the specifics mission risks in the crew or pre-mission briefs.

For moderate level residual risk, the Operations Officer /Chief Pilot must sign and acknowledge the level of risk.

For high level of residual risk, a director level official must sign and acknowledge the level of risk.

Note: organization have differing structures and may substitute an equivalent level of authorization that matches their organizational structure.

Appendix 6: Command and Control (C2) Remote Split Operations (RSO) Checklist

C2 Start Up Checklist for RSO Operations

1. Power on all applicable computers and monitors
2. Log into all applicable software applications (Mission Planner (UgCS), USS, Cameras, Chat, Video Conf, etc.)
3. Verify UgCS software version
4. Verify remote cameras are powered and set for greatest situational awareness
5. Verify connectivity with remote site
6. Verify USS and any simulations are prepared
7. Ensure NOTAM filed, as required
8. Launch any flight simulations

C2 Mission Director Brief

Confirm simulations are airborne

1. Verify remote site GCS checks are complete
2. Verify common UgCS software versions at all locations
3. Verify remote site flight crew brief has been completed
4. Verify manual controllers are powered and with the on-site PIC
5. Verify remote site performed RF spectrum analysis
6. Verify flight operation area is free of hazards (cars, people, wild animals)
8. Ensure all required data recording systems are recording (Video Conf, Chat windows, Remote cameras, payload sensors. e.g., Zoom, Site cameras, aircraft camera)
9. Conduct C2 Pilot Brief
   a. **Verify with remote site correct route being uploaded to UgCS**
      i. Weather Brief
      ii. EPs
      iii. Communications
      iv. Aircraft Handoff Procedures
      v. Verify NOTAM Filed
      vi. Any reason we should not fly today?
      vii. Crew Task Management
      viii. Alternate plans
      ix. Aircraft performance
      x. Weight and balance
      xi. Unique risks
      xii. Sterile cockpit procedures

C2 Pilot Take Off Checklist

1. Verify Battery Life
2. Verify GPS
3. Verify RF Strength
4. Check Winds
5. Takeoff at C2 pilot's discretion

C2 Pilot After Takeoff Checklist

1. Verify aircraft is climbing and on course with remote site
2. Verify aircraft has reached enroute altitude with remote site
3. Verify each point on route with remote site
4. Perform safety check on each leg, at minimum (Battery state, GPS State, Connectivity)

C2 Pilot Descent Checklist

1. Verify on site PIC ready to take over controls
2. Verify battery health, GPS and Connectivity
3. Verify aircraft descending
4. Monitor and call out altitudes on descent.
5. Remote site PIC call aircraft landed

C2 Pilot Post Flight Checklist

1. Verify aircraft motors are shutdown
2. Conduct C2 Pilot Debrief
3. Log aircraft discrepancies
4. Log aircraft flight times
5. Prepare for next sortie, if applicable