Data & Reasoning Fabric

Minimum Viable Product

William R. Van Dalsem,[1] Kenneth Freeman,[2] Sandeep D. Shetye,[3]
Aditya Das,[4] Kalmanje S. Krishnakumar,[5]
NASA Ames Research Center, Moffett Field, California, 94035

Aaron J. Swank,[6]
NASA Glenn Research Center, Cleveland, Ohio, 44135

Peter Shannon,[7] Luka Tomljenovic[8]
Radius Capital, San Francisco, California

Sandy Lozito,[9] Helen L. Euler,[10] Supreet Kaur,[11] Channon Wong,[12] &
Joseph J. O'Brien,[13]
NASA Ames Research Center, Moffett Field, California, 94035

A Data & Reasoning Fabric (DRF) is envisioned to enable the full potential of air mobility by providing all data and reasoning where they are needed. The DRF provides a marketplace based on an open foundational ecosystem of data and reasoning exchange between the many systems that must seamlessly interplay to manage the complex and dense airspace operations required to achieve advanced air mobility goals. The DRF marketplace is decentralized and will not be owned by any single party. To deliver reference designs and development environments that catalyze broad private and public sector buy-in and self-sustaining development of DRF and associated standards, DRF activities will collaboratively test these technologies, open standards and architectures, and an integrated framework with end-users. The DRF activity is developing a Minimum Viable Product (MVP) called the DRF Accelerator which will enable the rapid deployment of candidate data and reasoning services to test and evaluate the DRF core system functionality. The DRF Project is engaging with government and industry end-users and stakeholders, to: (1) assess the technical feasibility of the DRF Accelerator; (2) assess the likelihood of adoption, through shared test and evaluation; and (3) identify what data and services can be shared across DRF.

---

[1] Chief Strategy Officer, Aeronautics Directorate
[2] Aerospace Engineer, Aeronautics Projects Office
[3] Computer Scientist, Intelligent Systems Division
[4] Research Scientist, Aviation Systems Division, AIAA Member
[5] Associate Division Chief, Intelligent Systems Division, AIAA Associate Fellow
[6] Research Engineer, Information and Signal Processing Branch
[7] Founder, Managing Director
[8] Partner
[9] Chief, Aviation Systems Division
[10] System Engineer, Aeronautics Projects Office
[11] System Engineer, Aeronautics Projects Office
[12] Project Specialist, Aeronautics Projects Office
[13] Associate Chief, Aeronautics Projects Office

I. Nomenclature

AAM	=	*Advanced Air Mobility*
AI	=	Artificial Intelligence
CAS	=	Convergent Aeronautics Solutions Project
COTS	=	Commercial Off-The-Shelf
DRF	=	Data & Reasoning Fabric
FAA	=	Federal Aviation Administration
IoT	=	Internet of Things
IT	=	Information Technology
JSON	=	JavaScript Object Notation
ML	=	Machine Learning
MVP	=	Minimum Viable Product
NOTAM	=	Notice to Airmen
TACP	=	Transformative Aeronautics Concepts Program
UAM	=	Urban Air Mobility
UAV	=	Unmanned Aerial Vehicle

II. Introduction

Currently, when data is needed to support airspace operations, the required data sources are typically identified by human experts and custom software systems are developed to access or aggregate the needed data into a single software system. Decision support functions are then built into this software system for the specific airspace application. Each airspace application is addressed with the same approach, resulting in a series of very capable, but isolated, vertical integrations or silos. As we move to more data sources and more airspace applications this approach becomes more and more brittle.

As more data sources are provided by an increasingly diverse range of industries, relying on human experts to remain aware of and make safe decisions from all the relevant and often dynamic data sources will increasingly challenge human capabilities and may heavily constrain air mobility scalability as well as the ability to continuously extend airspace system capabilities to support technology advances and new mission types. The early internet evolved from a reliance on a human expert's knowledge of available websites, through human-maintained website lists/books, to self-sustained website catalogs and search engines. Access to airspace data sources needs to similarly evolve to be self-sustaining and scalable. It will also be important to provide data customers (for example, a company that wants to provide air mobility functions) the capability to solicit and vet new data sources, as well as to support reliable and consistent exchange and interpretation of data between software systems.

A similar scaling issue is arising with reasoning services. Historically, reasoning algorithms were built into each airspace application. However, reasoning services, ranging from speech recognition services to advanced machine-learning models, are emerging as standalone services. As with data sources, airspace applications need a scalable way to be aware of and integrate this new class of reasoning services. Additional concerns arise related to the current trajectory in that many systemic challenges are not addressed, such as the overall cost of developing many isolated vertical integrations/silos, systematic data quality tracking, and the overall security of the resulting landscape of many airspace application silos.

A Data & Reasoning Fabric (DRF) is envisioned to enable the full potential of advanced air mobility by providing all data and reasoning where they are needed [1]. The DRF (Fig. 1) provides a marketplace based on an open foundational ecosystem of data and reasoning exchange between the many systems that must seamlessly interplay to manage the complex and dense airspace operations required to achieve advanced air mobility goals. The DRF marketplace is decentralized and will not be owned by any single party. DRF activities will identify, test and - as needed - research and develop critical core technologies. To deliver reference designs and development environments that catalyze broad private and public sector buy-in and self-sustaining development of DRF and associated standards, DRF activities will collaboratively test these technologies, open standards and architectures, and an integrated framework with end-users.

Fig. 1 Data & Reasoning Fabric (DRF) for Advanced Air Mobility.

In this paper we present the minimum viable DRF product and its features in detail. The DRF Project is developing a MVP called the DRF Accelerator which will enable the rapid deployment of candidate data and reasoning services to test and evaluate – the DRF core system functionality. The DRF Project is also engaging with government and industry end-users and stakeholders, to: (1) assess the technical feasibility of the DRF Accelerator; (2) assess the likelihood of adoption, through shared test and evaluation; and (3) identify what data and services can be shared across DRF.

The objectives of the DRF Project are to:

- Deliver a MVP implementation of the reference DRF software, data, and reasoning architectures that will be presented in this paper. This MVP, the DRF Accelerator, includes prototype instantiations of key DRF features, including data and reasoning services, a DRF marketplace software architecture and other required functions sufficient to engage end-user representatives of key stakeholder communities in representative use-cases.
- Exercise the DRF Accelerator with representatives of key service providers and consumer stakeholder communities to demonstrate that the DRF Accelerator provides key data and reasoning functions, enables the future traditional and advanced airspace objectives, and informs a business plan for a self-sustaining advanced air mobility data and reasoning marketplace

The DRF Accelerator as shown in Fig. 2 is a prototype implementation of the DRF Framework, which enables the rapid deployment of candidate data and reasoning services to test and evaluate the DRF core system functionality - enabling the delivery of data and reasoning services to DRF consumers across the cloud-to-edge continuum. The DRF Accelerator showcases the core capabilities and objectively demonstrates the value of DRF across identified use-cases. It enables co-simulation of various DRF components including an initial edge simulation platform (e.g., simulated edge nodes), real and simulated sensors, and simulated vehicles. Further, the DRF Accelerator enables the ability to rapidly plug in and replace components. Finally, the DRF Accelerator provides the necessary DRF functionality to engage with representative customers and stakeholders.

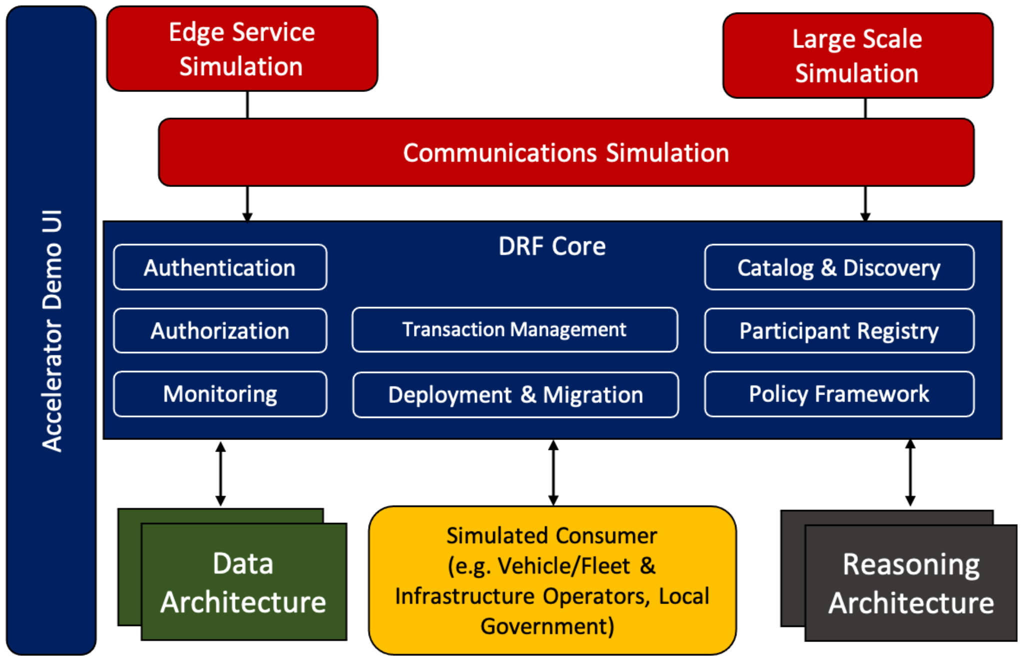

Fig. 2 DRF Accelerator MVP.

III. DRF Core

Fig. 3 shows the architectural overview of DRF Core, the central part of the DRF Framework and DRF Accelerator. Various functional requirements including registration, service discovery, authentication and authorization, transaction management, service monitoring, and policy framework are offered as core capabilities. Fig. 3 also shows the interaction of the fundamental stakeholders with DRF Core.

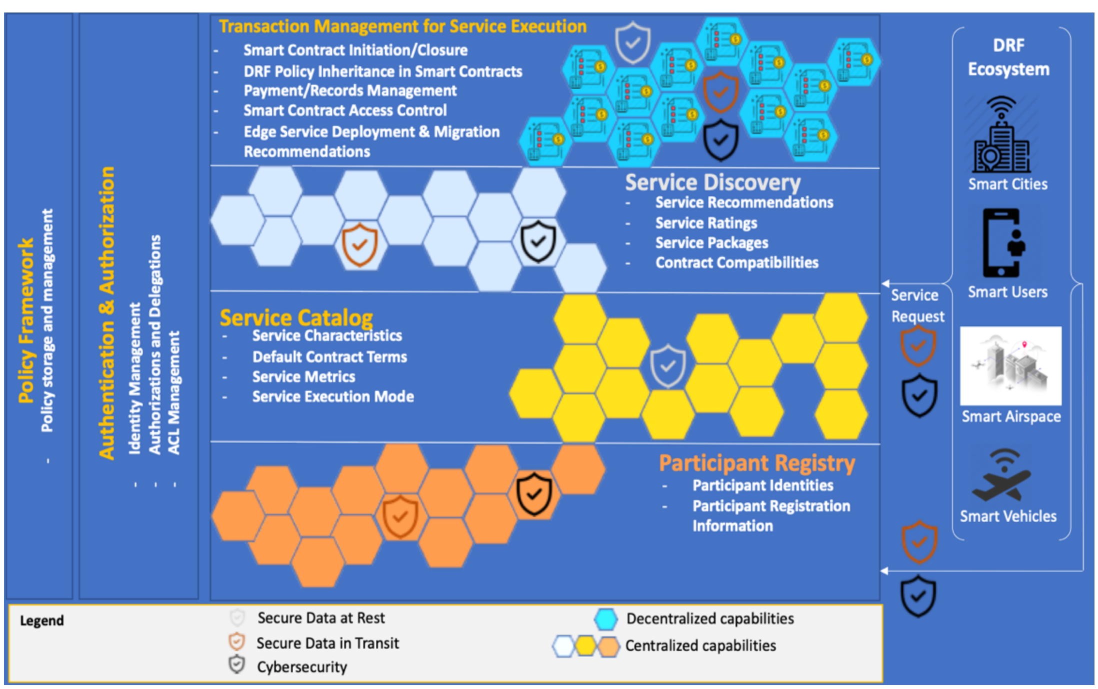

Fig. 3 Architecture Overview of DRF Core.

In particular, the DRF Core:

- Manages participant interactions with DRF (e.g., registration, authentication, authorization).
- Provides a service catalog for service registration and discovery.
- Enables distributed and decentralized service execution.
- Provides transaction management capabilities for capturing business logic and service agreements.
- Provides a policy management framework to ensure all participants operate according to specified policies.
- Monitors services and their quality of service (QoS) and service level agreements (SLA).

DRF Core exposes all its functionality to participants via APIs. The core APIs can be used by providers and consumers using standard HTTP libraries. To make it easier for the participants, we also plan to provide a reference implementation as well as client libraries which can be easily integrated with existing software. Below we highlight the key components of DRF Core. Further details of the software architecture are presented in a separate paper [2].

A. Participant Management

A core function for DRF is to validate DRF participants and maintain secure access to services during service deployment and exchange. To achieve these guarantees, the DRF Core includes a participant registry and an authorization and authentication framework. Every participant in the DRF marketplace will have a unique identity within the marketplace. To ensure the trust in the marketplace, only vetted and registered participants can be part of the marketplace. DRF Core also maintains and enforces access control for participants according to policies and predefined agreements.

B. Service Publishing and Discovery

DRF stakeholders require the ability to post, discover, sell, or buy data, reasoning, or edge infrastructure easily. All data, reasoning, and edge infrastructure providers register their available services by providing metadata about the services as well as service characteristics. This information is stored in a service catalog, which is then used by consumers to efficiently discover services based on service characteristics and consumers' access control. The service catalog acts as an app store for data, reasoning, and edge infrastructure buyers and sellers.

C. Transaction Management

Another important capability of DRF Core is transaction management. All transactions and agreements for service exchange are recorded in DRF using immutable verifiable logs. DRF utilizes smart contracts to programmatically enforce and guarantee that service exchanges are being carried out according to predetermined agreements. The role of the DRF Core is to facilitate interactions between consumers and providers. Fig. 4 shows an example sequence of interactions between these entities.

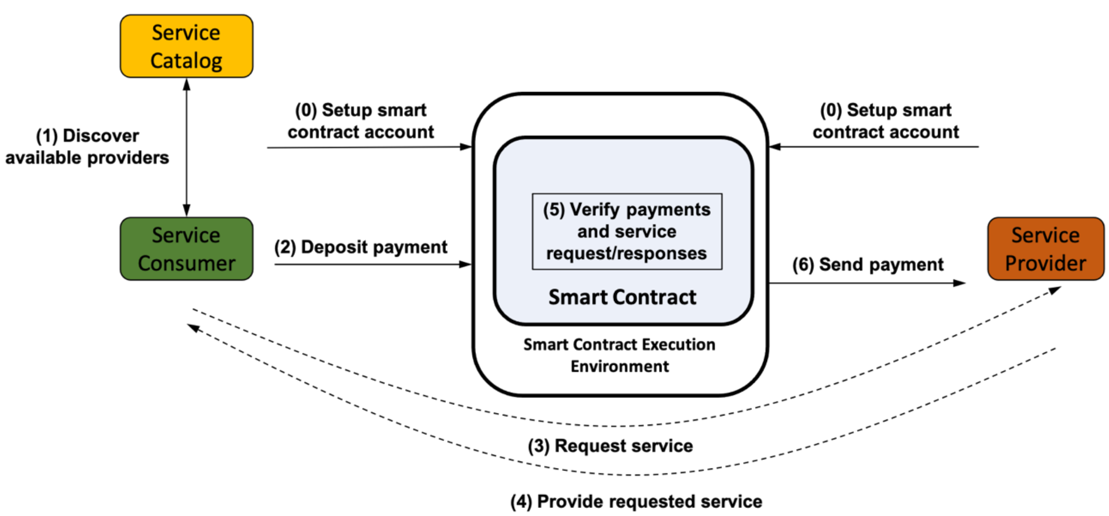

Fig. 4 An example of a smart contract negotiation sequence.

D. Policy Framework

DRF Core policy framework allows governing entities to regulate access and disseminate policy updates to participants. The policies will be stored and managed within DRF Core so that DRF marketplace participants can search and retrieve applicable enforced policies and associated documents. The policy framework is an important closed feedback loop that DRF supports with multiple end benefits.

E. Data Privacy and Security

We recognize that in some cases, data and reasoning that is exchanged in the marketplace need to remain confidential. DRF team is working on various techniques to preserve the privacy of data and the goal is to gradually add functionality into DRF Core. Initially we plan to introduce two different mechanisms that will enable data and reasoning owners to maintain the privacy of their data. One of the mechanisms will be an order preserving encryption mechanism which will be exposed as part of the DRF API library. DRF Core will also support compute-to-data which will allow data owners to share data without transferring it to the consumer.

IV. DRF Data Services

The goal of the DRF data architecture is to ensures consistency, availability, usability, safety-relevant quality, and democratization of data for future air mobility operations. The data produced, traded, and consumed by the diverse future air mobility ecosystem is going to be complex and the data functionalization model has multiple facets to it. Therefore, the data architecture needs to take into consideration the utilization modes and derived benefits from the perspective of different stakeholders to make it widely acceptable and usable for all future air mobility stakeholders.

The current effort towards a minimum viable DRF system includes a set of reference data services demonstrating functional interoperability among data and reasoning service demands. A successful demonstration of the service life cycle using commercially available data integration and access layer resources in the beginning, has been identified as a strategically suitable pathway towards substantial enhancement in the appeal, effectiveness, and accessibility of the DRF services to the wider user pool.

The DRF data services are essentially datasets commercialized on an open marketplace and offered by pre-approved service providers. The services could comprise of real-time data or historical data, delivered via streaming or spot/batch upload methods. Typical examples include weather data, traffic data, vertiport status, notice to airmen (NOTAM), sensor data onboard vehicles and other infrastructures, network data and so on. In the envisioned open marketplace setting, such data is generated/gathered and packaged for commercialization by service providers. Prior to onboarding a specific data service to the DRF marketplace, the service metadata undergoes an automated review process to identify the service domain, usage details, contractual requirements, and provider information. Metadata supports machine-to-machine exchanges of data, including information addressing data formatting, how it can be used, and data quality metrics. The metadata section is further augmented with market statistics by DRF to offer better selectivity for prospective service users.

With the advent of new DRF services as well as the usage methods for them in the future, the above list of metadata is expected to evolve to deliver a holistic insight about the services. It is noteworthy that most, if not all, service discovery, contracting, fulfillment, and feedback collection on DRF will be executed autonomously by smart/unmanned/AI agents. Therefore, the structuring of the metadata package requires special attention to make it machine interpretable and platform/software independent. Furthermore, modularization of the metadata package is highly desirable given that different users will use it for different purposes.

With the above requirement in mind, to describe the DRF data services, a json (JavaScript Object Notation) schema approach is used to store and process the service metadata. A json schema is simply a declarative format that is used to describe the structure of the actual metadata. Through this schema it is easy to describe the surface structure of the DRF service metadata and automate the validation of the service metadata. While the services themselves can be different in their constructs, input/output models, and other functional and commercial attributes, through a set of taxonomic models, it is hypothesized that a reduced set of metadata descriptors can be used to identify and describe the services sufficiently and effectively. These metadata descriptors provide the foundation for service interoperability among users, applications, and platforms.

As in case of the lifecycle of any other data modeling exercise, the DRF data modeling process also goes through three main stages, i.e., conceptual data modeling, logical data modeling, and physical data modeling. In the conceptual data modeling stage, the data requirements for specific DRF services are gathered through review of the current state of the art in comparable technologies and discussion with subject matter experts and stakeholders. Next, the conceptual data model is normalized and further refined into a logical data model that identifies the entity relationships and attributes. Finally, in the physical data modeling stage, the data model schema is designed considering the technical constraints and performance requirements for target applications. For DRF data services specifically, a NoSQL database architecture is used that provides flexible schemas for the storage and retrieval of semi-structured and unstructured metadata. The NoSQL database architecture is used to enable specific DRF core features such as: massive scaling, multifarious service offerings, minimization of infrastructure cost in cloud implementation.

As mentioned earlier, the DRF service models use json schemas to store and validate metadata packages. While the service implementation models, such as the input/output parameters and prescribed usage information, are unique to the service, several other metadata sections, pertaining to contracting, discovery/selection, auditing etc. can be broadly generalized. Following this idea, the overall metadata package for DRF services is divided into two parts. The first part, termed as marketplace metadata, holds a set of generalized metadata that are used for discovery, selection, requisition, auditing, and other similar governance functionalities. The second part, termed as service metadata, on the other hand, houses functionalization metadata that are used for the actual technical implementation of the service. Figure 5 shows a block diagram of the overall strategy to build such a modularized metadata schema. To allow further interoperability, metadata sub-schemas are used for describing input/output models for logically similar services. In the current effort, research is being conducted to build semantic definition for operational metadata elements of DRF services, based on taxonomic models for service domain, criticality, and use cases. In addition to exploring interoperability methodologies at the schema level, the team is investigating the methods to achieve interoperability at the record and repository levels with minimum loss or distortion of data. While not a lot of standards and guidelines for AAM specific metadata are available at present, it is anticipated that in the near future, with the introduction of more and more unmanned aerial commercial services into the national airspace, AAM data/metadata standards bodies will emerge. At that point, the current efforts towards incorporating DRF service metadata modularity will be instrumental in quickly and securely onboarding diverse set of AAM services.

As indicated above during the discussion on metadata, a key aspect of the DRF services is that for ubiquitous UAM operations with system-wide automation optimization, alternatively identified as UAM maturity level 6, where tens of thousands of simultaneous operations can take place, it is expected that such operations will be handled mostly by unmanned/smart/AI agents/machines. For reasoning services especially, as discussed in the following section, selection, delivery, and utilization of services involve not only the review of corresponding metadata at the time-of-service consumption but also tracking of the same over time to ensure the efficacy of the reasoning models. For this purpose, a custom set of performance metadata is used in conjunction with the descriptive, structural, and administrative metadata elements. For example: a machine learning based service is attached with a performance metadata sub-schema to indicate the model accuracy, recall rate, f-score, r-squared value, meta-architecture etc.

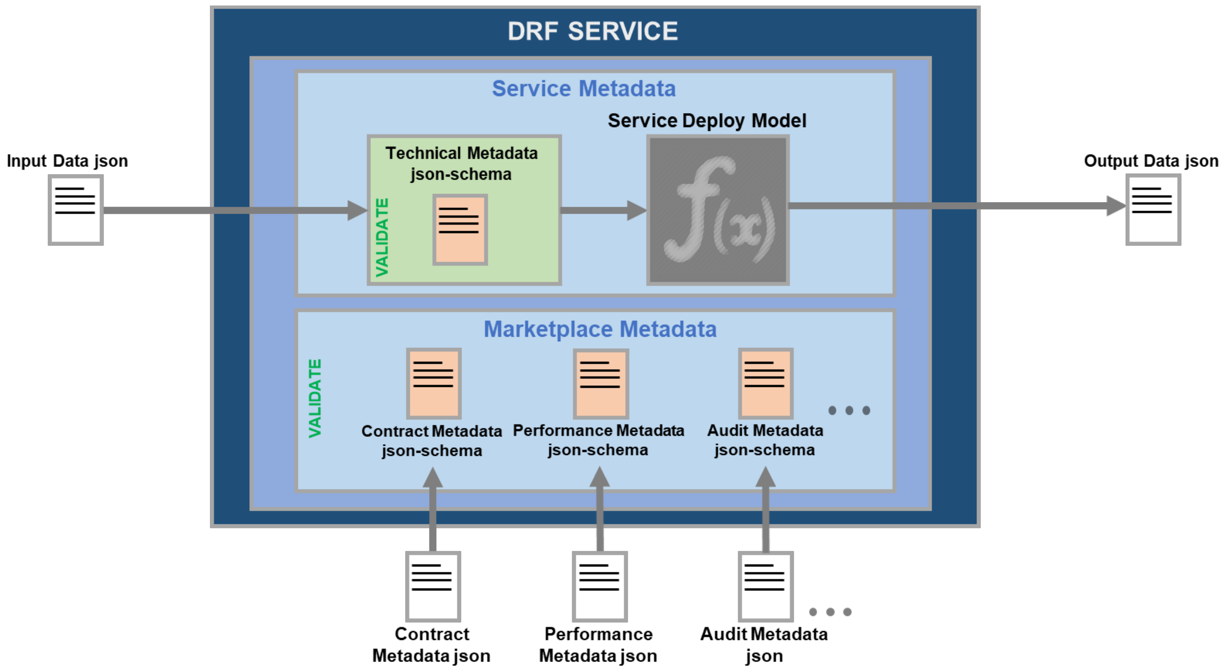

Fig. 5 Modular metadata json schema design strategy.

V. DRF Reasoning Services

To establish a DRF architecture that enables future air mobility operations, we consider the reasoning needs of highly automated aircraft operations where the human operator might be supervising many aircraft. The notion of 1:N operations is therefore a central driving factor in arriving at a reasoning architecture. Fig. 6 shows a prototypical reasoning architecture in which the services, the users, and the DRF core software architecture interact in real-time for efficient operations.

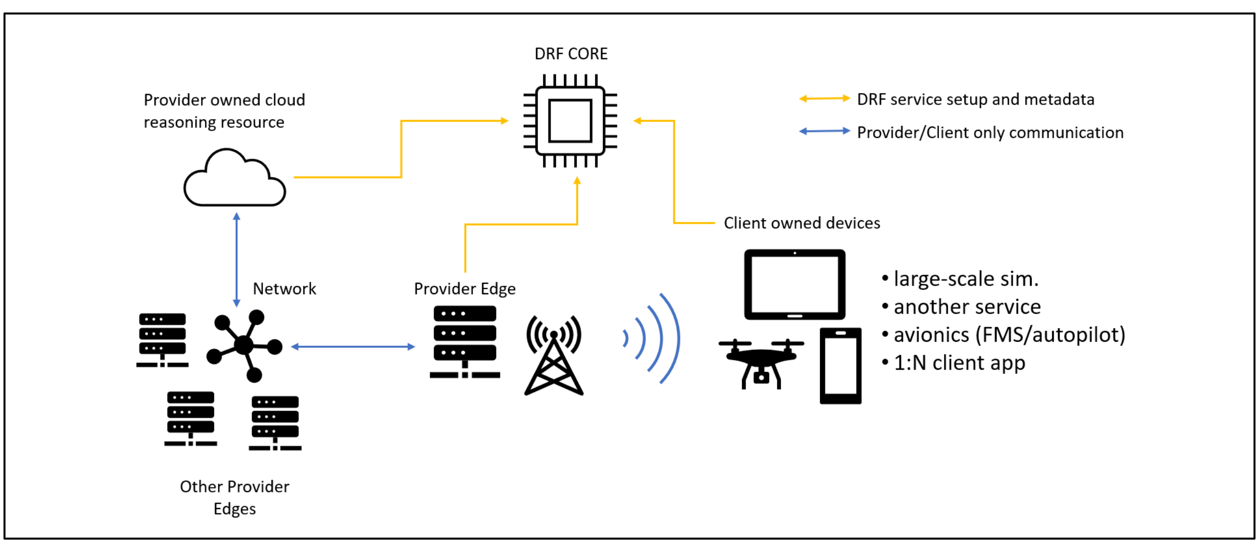

Fig. 6 A Reasoning Service Architecture.

The reasoning services which will be typically hosted on a cloud platform will have to be optimally located/moved to nearby edges to serve the user needs. In addition to quality and security of the services, the communication latencies between the services and the user will have to be managed in a dynamic fashion to maintain the safety needs of the user.

To better understand reasoning services requirements, we have arrived at four different reasoning services categories which will be required at a minimum to fill the needs of an operator of many automated vehicles:

(1) Mission Planning (MP) -- Reasoning for multi-vehicle trajectory generation, with airspace and operational constraints, and distributed contingency management.
(2) Health Monitoring (HM) -- Diagnostics and prognostics (constraints and predictions) of systems and components based on data, models, and simulations.
(3) Information Management (IM) -- Autonomy of model learning from data and physics, with active cross-checking, updating, and validation of uncertainty and performance.
(4) Situation Awareness (SA) -- Environmental (terrain, obstacles, no-fly zone, weather) and airspace (vehicle state, planned trajectory, destination) maps, risk assessment

Definition of reasoning elements and the metadata needed for service discovery are broadly classified based on latencies, computational complexity, and Edge/Cloud/Physical Layer representation. Details of a sample reasoning service are further discussed below.

Trajectory Reasoning Services

As an example of the type of integration that will be required for a successful DRF MVP, Fig. 7 and 8 detail the interaction between the user and the various reasoning services.

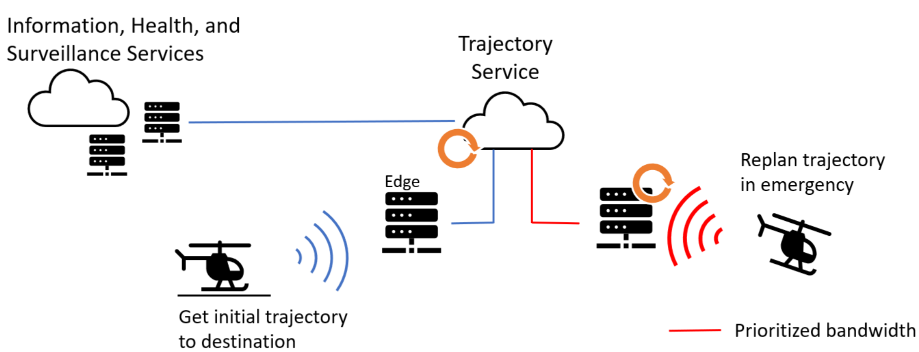

Fig. 7 Trajectory Services.

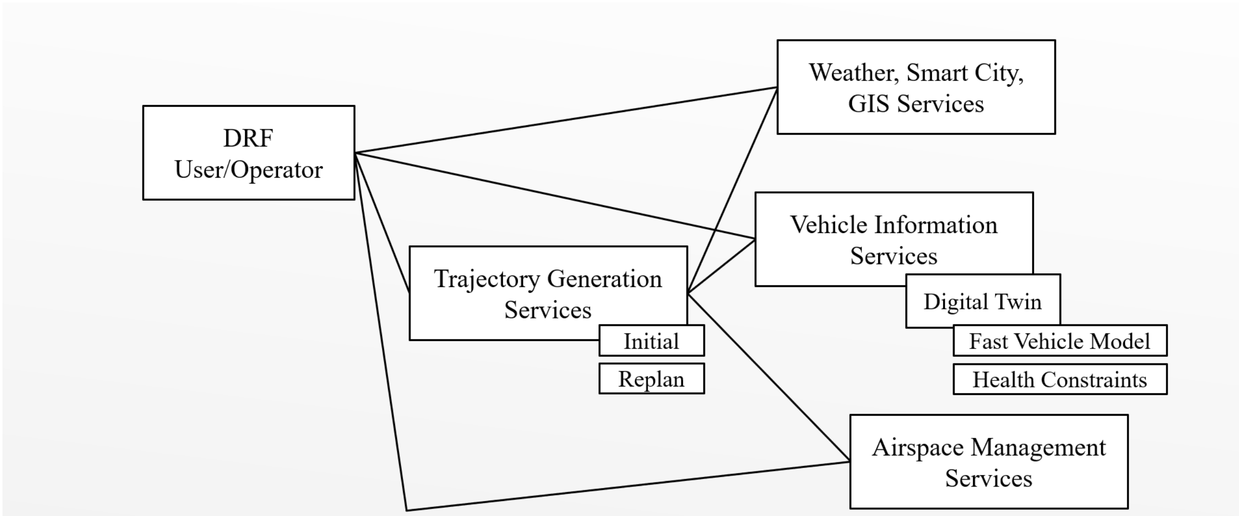

Fig. 8 A subnet needed to initiate a trajectory generation service.

The table below provides the characteristics needed to realize a vibrant trajectory reasoning service.

Reasoning Service	Service Characteristics	Service Metrics
Generate initial ATC approved vehicle trajectory for a mission	• User configurable cost function (prioritize time vs. energy, route aggressiveness). • ATC approved trajectory, with scheduled start time window. • Supports trajectory reservations. • List of vehicle types supported, etc. • Dependencies: Weather service, terrain SA service, smart city service, vehicle model and health service, airspace management service. • Outputs local optimal solution, with guaranteed constraint compliance. May return infeasible.	• Service latency in < 1sec average, 10 seconds max. • Low bandwidth required <1Mbps
Provide fast trajectory replan	• Monitors vehicles' progress every 2 seconds. • Cloud computes highest risk-factors for non-compliance, updating every 2 seconds. • Risk factors pushed to 1:N operator stations • Cloud continuously computes trajectory contingencies from current position and tentatively approves them with ATC. • Contingency trajectories are pushed to edge nodes nearest current vehicle location. • Guaranteed availability of at least one feasible reroute trajectory (possibly to emergency landing zone).	• Trajectory re-route pushed to vehicle <30ms. • Low bandwidth required <1Mbps • Edge node always within 1 mile of vehicle position.

Table 1 Trajectory Reasoning Service characteristics.

In addition to the characteristics, metadata needed for service discovery need to be identified. As an example, we outline the metadata potentially needed for two categories of trajectory services.

1. Generate initial approved vehicle trajectory:
    a. Required Data: Vehicle ID, and state; Intended start time; Vehicle mission objectives.
    b. Service Output: Local optimal trajectory solution, with guaranteed constraint compliance, and start time window. Note: May return infeasible.
    c. Dependencies: Airspace management service (Required); Weather service; Terrain, obstacle Situational Awareness (SA) service; Smart city service; Vehicle information service.
    d. User Options: Configurable cost function to trade-off time, energy, ride-quality; Vehicle options; Shadow mode use for planning.
    e. Metrics: Service latency < 1 sec average, Max 10 seconds; Low bandwidth required <1Mbps.
    f. Other Features: Output approved with airspace management service; Supports in-advance trajectory reservations; Vehicle types supported by the service.
2. Monitor vehicle progress and readies contingency trajectories:
    a. Required Data: Vehicle ID, and state; Vehicle objectives.
    b. Service Output: Trajectory reroutes for most likely contingencies; Service does not return infeasible.
    c. Dependencies: Airspace management service (Required); Weather service; Terrain, obstacle SA service; Smart city service; Vehicle information service.
    d. User Options: Configurable cost function to trade-off time, energy, ride-quality.
    e. Metrics: Cloud computes highest risk-factors for non-compliance, updating every 2 sec; Candidate trajectory re-routes pushed to edges nearest vehicle every 2 sec; Contingency trajectory downloadable to vehicle within (<30 ms); Low bandwidth required for data transfer (<1Mbps); Edge node always within 1 mile of vehicle position; Risk factors pushed to subscriber operator stations (10 second refresh).
    f. Other Features: Trajectory contingencies are tentatively approved with airspace management; Contingency trajectories are pushed to edge nodes nearest to current location; Guaranteed availability of at least one feasible trajectory reroute option (possibly to emergency landing or ditch zone).

<u>Technical challenges to be addressed:</u>

For successful DRF reasoning service architecture and integration, there are key technical challenges to be addressed:

- Access to edge infrastructure and other technology stacks and capabilities for service migration is a challenge that needs to be studied and understood. Automatically learning where best to execute each service (cloud/edge/physical) is important to understand and implement.
- Easy to understand service implementation via design patterns as starting points for building-up full services with industry partners.
- Service demonstration meets initial requirements for operation on the DRF, including measurement of data bandwidth usage and latency.
- Service migration technologies that are implementable for initial DRF needs and identification of gaps for new technologies for future expansion.

VI. DRF Simulations:

The DRF Accelerator simulates some aspects of the DRF Framework. As shown in Fig 2, the simulations will be leveraged to validate the viability of the DRF Accelerator. A large-scale vehicle simulation and visualization capability will be developed.

A. Large-Scale Vehicle Simulation

Implementation of a large-scale vehicle simulation capability using a mix of high-fidelity models and simpler kinematic models is important for validation of the DRF architecture. Fig. 9 shows the interactions between the operator, the DRF, the vehicles, and the service providers.

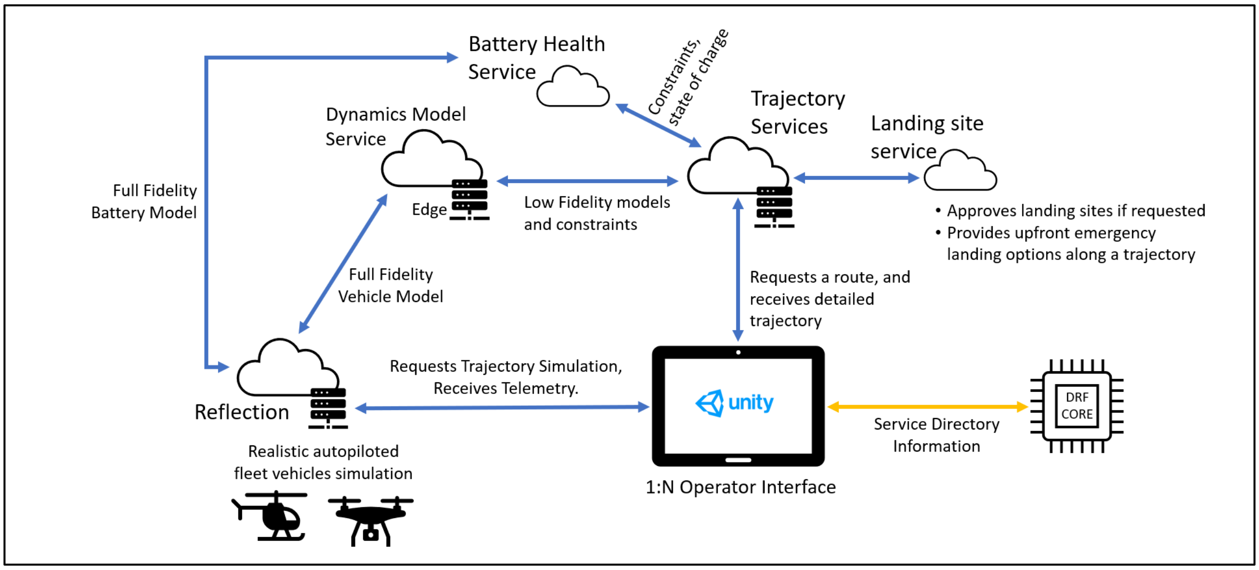

Fig 9. Operator – DRF - Simulated Vehicles - Reasoning Services interaction.

The large-scale vehicle simulation development focuses on the demonstration of many (~100) point-mass kinematic models executing reasoning services along with at least two full-dynamic UAM model-based vehicles that have on-board automated flight execution options served by DRF via the operator's wireless communications with the vehicles. Visualization includes aircraft executing trajectories devised by DRF services that are communicated via the operator. The DRF reasoning services include trajectory services with pre-approved landing options; real-time battery health and high-fidelity battery model information; and varying fidelity vehicle model service for designing the automated flight execution options such as dynamic autopilots and decision-making.

B. Edge Simulation

Edge computing provides infrastructure that is closer than cloud resources and may be used for navigation of aerial vehicles, which has similar safety requirements to navigation of terrestrial vehicles. There are two types of edge resources in DRF, the first type is provided as a service by large scale external providers (e.g., Amazon or Google) and the other type is more focused on smaller edge providers with finite capacity (e.g., roof top units, roadside units, in vehicle nodes). The DRF team is going to evaluate both types as part of DRF Accelerator. However, due to the lack of available large scale edge computing testbeds and due to cost concerns, the DRF team will use simulations to validate the DRF Edge architecture at scale. For the DRF Accelerator, we have identified qualities and related metrics to simulate edge nodes and edge services. The qualities considered for the simulation are according to ISO/IEC standards and include:

- Time behavior - response time, turnaround time, and throughput
- Resource utilization - memory utilization and bandwidth utilization
- Capacity - transaction processing capacity

C. Communication Simulation

Simulation of the network connectivity between the various components, including the edge nodes and mobile vehicles is a necessary component for evaluating the DRF Accelerator. A communications simulation will be leveraged to validate the capabilities of the DRF core, in conjunction with simulations for end-user interactions, edge computing and large-scale vehicle simulations.

The DRF may consist of a number of diverse communications elements, each with unique capabilities and requirements. For example, terrestrial, non-mobile, or backhaul connectivity can generally move large amounts of data without disruption. The performance capability is built into the ground network infrastructure. Terrestrial, or local area network connectivity to a mobile vehicle, for instance at an airport/vertiport, may experience communication disruption due to multi-path effects, vehicle dynamics and changes in environment. For ground operations, the required communications performance and data requirements are different from in-flight operations. Ground-based connectivity can be designed separately from an air to ground link to provide the required communications resilience.

Wireless connectivity to an aerial vehicle is the most bandwidth constricted communications element and may be supported through ground-based wireless communication and/or via satellite communication paths. Satellite communication paths typically exhibit more latency as compared to ground to air links, but can operate at locations in absence of ground infrastructure. Disruptions in connectivity can occur at different phases of flight due to changes in signal strength resulting from vehicle dynamics, airframe blockage and antenna patterns. Satellite communication links, for example, operate with very low receive signal levels and are therefore highly prone to obstructions and reflections.

In a network topology containing mobile vehicles, connectivity to the end user is certain to include disruptions and fade-outs. Communication quality of service (QoS) characteristics at the Internet protocol (IP) layer such as data throughput rates, packet loss and latency may be simulated with network emulators. Modeling QoS characteristics at the IP layer is sufficient for validating the capabilities of the DRF core and modeling/testing at the radio frequency (RF) level is out of scope for the initial efforts. There are a number of existing software tools to support the modeling, including both open source and commercial solutions. A summary of the various simulation options with varying level of capability and fidelity is summarized in Fig. 10.

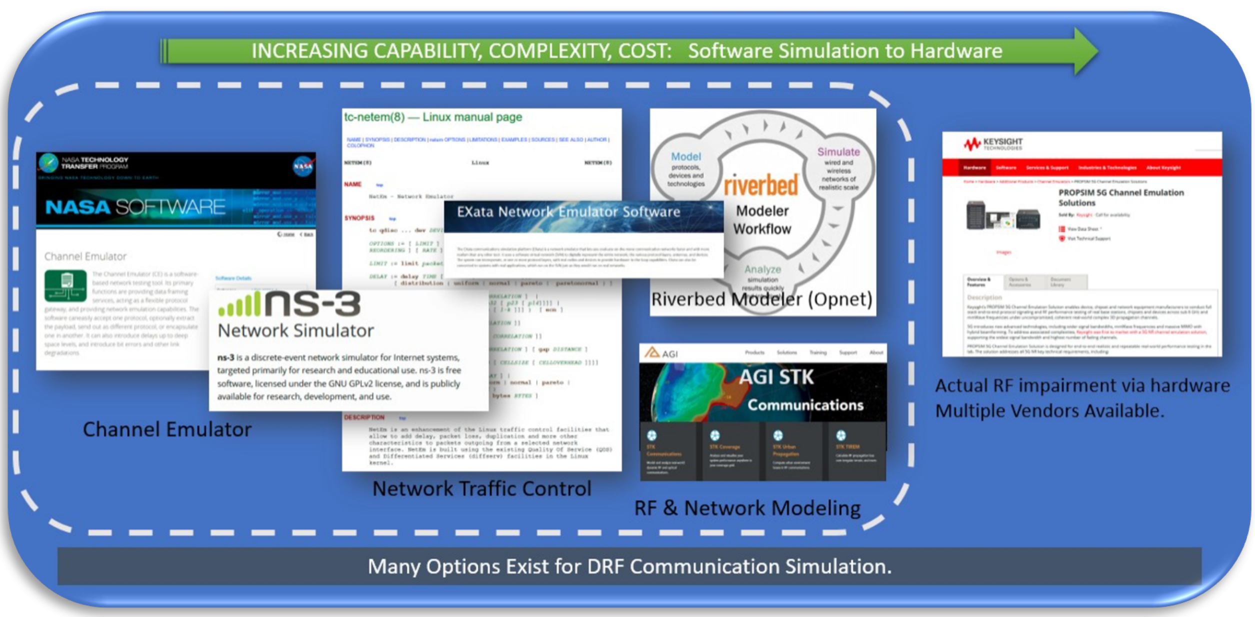

Fig. 10 DRF Communication Simulation Options

For initial network emulation, the open-source ns-3 network simulator [3] software will be used in combination with a Docker container to shape network traffic. The end goal is to simulate the various communications characteristics of each communications element, with an initial focus on QoS and data timeliness. The desired simulation characteristics include:

1. **Quality of Service & Data Timeliness:**
   - throughput, data rate, packet loss, latency, and timeliness

2. **Performance Diversity:**
    - link throughput/capacity and latency for different phases of flight and/or situational conditions
    - variable data rates in support of different vehicle classes, with varied communications capabilities
3. **Transitions:**
    - transitions between communication elements including for example satcom, wireless and ground networks as a vehicle moves through the DRF network
4. **Prioritization:**
    - allocation and prioritization of communications resources
    - prioritization of data streams

VII. DRF MVP Engagement

The DRF Project, which is funded by the NASA Aeronautics Mission Directorate, will engage with internal NASA partners and partners external to NASA to support the completion of its objectives. The DRF Project team will leverage partnerships with four types of organizations: (1) technology providers, (2) data service providers, (3) reasoning services providers and (4) end user organizations. These engagements will help the DRF Project assess the technical feasibility of the DRF Accelerator and assess the likelihood of adoption.

The goal of DRF is to enable a self-sustaining marketplace where data and reasoning service providers can "advertise" their services which the end user can discover, vet, purchase, aggregate, and use to achieve their air mobility missions. This will be accomplished by providing a secure software infrastructure and the tools, protocols, governance, and policies to implement, administer, manage, and operate data sharing and reasoning services across all air mobility needs. As the DRF Core and its associated simulations are being developed, the DRF team will partner with technology providers for implementation and to test objectives. For instance, to support large-scale vehicle models, the DRF Project will leverage common reference models (CRM) provided by the NASA Revolutionary Vertical Lift Technology (RVLT) Project. The NASA RVLT Project has developed several concept vehicles models for UAM technology development [4]. Additionally, the DRF team will seek collaboration with external technology, and strategy organizations in the industry to host and secure a unique, agnostic, decentralized solution to manage and connect sets of data services, operations, applications, and devices - giving enterprises the freedom to manage operations and communications, deploy applications and expand operations as air mobility needs will evolve.

The DRF team will wrap services provided by external data and reasoning providers to validate the capabilities of the DRF Core. User engagements will identify data and reasoning service providers who are willing to collaborate with the DRF team to utilize and test the execution of the DRF accelerator for deployment in support of end-user use cases. Data and reasoning service providers will be identified after end-user use cases are solidified to assure that service providers are providing data relevant to one of the end-users use cases. Given the agnostic and decentralized nature of the DRF Core, it will not be hosted in a centralized repository nor managed by a single entity, enabling privacy-preserving fair data exchange for air mobility needs. Use cases and data and service providers will help validate the DRF Core functionality and value during planned demonstrations over two years.

DRF is designed to support high-scale and high-density operations, as well as a diversity of mission types with potential benefits to society in areas of emergency survey and response, integration of advanced technologies of local and federal airspace and healthcare transportation in the air mobility environment. The DRF team seeks collaborations for future and near-term engagement with anticipated stakeholders who will play an active role in the establishment and sustainability of a healthy, productive, and growing marketplace for advanced air mobility data and reasoning services. The engagements will include, identifying organizations willing to collaborate with NASA on the end user application of DRF and demonstrating DRF capabilities in dense urban environments. The end-user use cases may be in a "shadow-mode" where a DRF-based application shadows an actual end-user application, from which the relevant benefits and disadvantages of the DRF-based application can be derived. Or, a DRF-based application could be providing decision-support to an actual application, in which case the DRF-based application is an auxiliary data feed to an actual operation. Though not ruled out, it is not required that a DRF-based application is the enabler of an end-user use case.

VIII. Conclusion

The DRF marketplace will be enabled by an open foundational ecosystem of data and reasoning exchange between the many systems that must seamlessly interact to manage the expected highly complex and dense airspace operations. The DRF activity is developing a MVP called the DRF Accelerator which will enable the rapid deployment of candidate data and reasoning services to test and evaluate the DRF core system functionality. The DRF Project is engaging with government and industry end-users and stakeholders, to: (1) assess the technical feasibility of the DRF Accelerator; (2) assess the likelihood of adoption, through shared test and evaluation; and (3) identify what data and services can be shared across DRF.

Acknowledgments

The DRF activity is funded by the NASA Aeronautics Research Mission Directorate (ARMD) Transformative Aeronautics Concepts Program (TACP) Convergent Aeronautics Solutions (CAS) Project.

The DRF activity includes many more contributors than the authors listed on this paper and the authors would like to acknowledge all of their contributions which will more fully described in subsequent more in-depth DRF publications. For this paper, the authors would like to specifically acknowledge contributions made by Moustafa Abdelbaky and Abraham Ishihara towards the DRF Core and Edge Simulation, Eric Chen and Pallavi Hegde towards Data Services, Stefan Schuet towards the Reasoning Services and Large-Scale Simulations, Michael Cauley for his contributions and expertise in aeronautical communication and Dennis Iannicca and Casey Bakula for their insight in communication and network simulation.

References

[1] Van Dalsem, W., Shetye, S., Das, A., Krishnakumar, K., Lozito, S., Freeman, K., Swank, A., Shannon, P., Tomljenovic, L., "A Data & Reasoning Fabric to Enable Advanced Air Mobility," AIAA SciTech Forum, January 2020

[2] Abdelbaky, M., Chen, J., Fedin, A., Freeman, K., Gurram, M., Ishihara, A., Joe-Wong, C., Knight, C., Krishnakumar, K., Reyes, I., Robinson, C., Shannon, P., Shetye, S., Tomljenovic, L., and, Van Dalsem, W., "DRF: A Software Architecture for a Data Marketplace to Support Advanced Air Mobility," AIAA Aviation 2021 Forum, 2021, to appear.

[3] NS-3 Network Simulator, https://www.nsnam.org

[4] Silva, C., Johnson, W. R., Solis, E., Patterson, M. D., and Antcliff, K. R., "VTOL Urban Air Mobility Concept Vehicles for Technology Development," 2018 Aviation Technology, Integration, and Operations Conference, American Institute of Aeronautics and Astronautics, 2018. doi:10.2514/6.2018-3847.