NASA's Digital Information Platform to Accelerate the Transformation of the National Airspace System

Mohana Gurram[1], Pallavi Hegde[2], Swati Saxena[3]
NASA Ames Research Center, Moffett Field, CA, 94035, U.S.A.

In order to accelerate the digital transformation of airspace operations, a foundational framework and infrastructure for providing sustainable, data-driven, and cohesive decision-making digital services for both traditional and emergent air vehicles is being developed. The reference implementation of Digital Information Platform builds an ecosystem for the aviation community by providing access to a secure and trusted source of aviation data and services. Several key features and services have been implemented to enable secure data sharing, communication, and service registration on the Platform. The technical approach used to implement these features is presented here. NASA-developed integrated aviation data and machine learning based prediction services to optimize airspace operations are available on the Platform. These services are being evaluated in an operational environment by flight operators and the real-world benefits are being captured. The Platform fosters collaboration among industry and researchers to develop complex aviation services and the aim is to make it publicly accessible for consumption by the aviation community.

I. Introduction

U.S. National Airspace System (NAS) is going through the journey of digital transformation [1] and efforts are being made to modernize the technologies to accommodate new entrants such as autonomous uncrewed air vehicles along with traditional commercial passenger and cargo aircraft on a regular basis. There is also a need to develop and deploy digital aviation services to further optimize airport and airspace operations and contribute towards the aviation industry's goal of achieving net-zero greenhouse emissions by 2050 [2]. There are a number of existing challenges associated with the standardization of data and tools and access and management to support airspace operations. There is a need to address these challenges to accelerate the digital transformation of the NAS. Aviation data are often segmented, discontinuous and available from multiple sources in different formats. Relying on unfiltered data can lead to sub-optimal decisions. Access to secure and trusted aviation data is crucial for the development of new tools and technologies which optimize ground and enroute flight operations and contribute towards sustainable aviation goals. There is a need to make the information available in a secure manner for consumption by the service providers and the flight operators. Another challenge is related to data authentication and identity governance. The connection between the data source and the user must be established in such a way that the transfer of data is secure, and the data authenticity is maintained during the transfer. In addition to easy and secure access to aviation data, there is also a need for a foundational infrastructure to test, evaluate and improve digital aviation services for optimizing airspace operations.

The current work focuses on addressing these challenges by providing a foundational framework and infrastructure to register, develop and mature data integration and decision-making services. The Digital Information Platform (DIP) (referred as the Platform in the paper) research and technology development work focuses on designing, developing,

[1] Computer Engineer, Intelligent Systems Division, NASA Ames
[2] Software Engineer, Aviation Systems Division, NASA Ames
[3] Aerospace Engineer, Aviation Systems Division, NASA Ames, swati.saxena@nasa.gov, AIAA Senior Member

and demonstrating a cloud-based software portal to enable easy access to aviation data and services that will improve air traffic operations and enable sustainable aviation concepts. Potential aviation services rely on the availability of a variety of information such as integrated flight data, convective weather, turbulence, airspace restrictions. The vision of this platform is to enable easy, reliable, and secure access to aviation information. The Platform will address these data-specific challenges such as standardized data format, secure data storage and transfer between two entities. Through this capability, the operators and service providers will be able to gather and access real-time and historical flight data upon which data-driven services for NAS users are developed and operated. The services-based architecture will support the reusability of these solutions to advance technologies toward sustainable aviation operations. The cloud-based infrastructure will enable easier data hosting, accessibility and sharing among partners. It also enables adaptability, scalability, testing and validation of Machine Learning (ML) based digital aviation services across multiple metroplex areas.

The next section describes a few NASA-developed and non-NASA developed services that require high quality aviation data and can be used to build more complex services by making them accessible via the Platform. The details about other similar platforms and data sharing portals and their limitations are described. The challenges and the limitations addressed by the current work are presented. Details about the NASA developed data integration service currently available on the Platform are provided. The subsequent section presents the technical approach and design architecture used to build the reference implementation of the Platform. Finally, NASA-developed services currently available on the Platform are described along with conclusions and the plan for future work.

II. Background

There are several NASA-developed high value aviation services that need high quality aviation data for decision making. Details of these services will be given in section IV, In addition to the NASA-developed services there are several potential non-NASA services (both ground and flight deck) that can be used as building blocks towards more complex aviation services. As an example, non-NASA services that can potentially be used to build a contrail avoidance service include: human labeled contrail dataset [3], contrail radiative forcing estimation modeling [4] and contrail visualization model [5]. More mature services such as the ones provided by the Cavan Solutions [6] displaying current Traffic Management Initiatives (TMIs), weather, and airport configuration information or Boeing Jeppesen flight planning tool [7] are also viable candidates for the reference Platform.

There are several flight and weather data sources that are typically consumed by the aviation services. The challenge is that the present flight and air traffic data come from a variety of sources and facilities and are often segmented and discontinuous. The data coming from different sources can also be in different formats. It is therefor, a nontrivial effort to stitch different data streams together and fill gaps while maintaining the data authenticity and following standard data formats.

There are existing solutions that partially address the need for access to good quality aviation data and a software platform to build the services ecosystem. There are existing Application Programming Interface (API) Management Platforms to discover, connect and share APIs. RapidAPI [8] provides an API hub to find, test, and connect to thousands of APIs. Users can find the APIs for the project, embed the API into their app, and track usage of all APIs through a single dashboard. MuleSoft Anypoint Platform [9] provides software solutions to build and manage custom APIs. These platforms can be hosted on-premise or on-cloud depending on the use case. FAA API [10] portal hosted by FAA provides a list of APIs for developers to build their software applications and integrate FAA datasets. They manage two separate communities. Another example of a portal from FAA is FAA SWIM Industry-FAA Team (SWIFT) [11] portal which provides public access to flight and weather data via System Wide Information Management (SWIM) data products. CavanReports [12] from Cavan Solutions is a cloud-based analytics platform leveraging real-time flight data from SWIM and weather products to support trajectory-based operations value-added metrics. They provide web-based applications and APIs to support airspace operations. NASA’s Data and Reasoning Fabric (DRF) [13] project has built a research prototype to support the future air mobility needs for autonomous operations. Their research prototype provides a self-sustaining advanced air mobility data and reasoning marketplace for providers of weather services, mapping and surveying firms, air traffic management decision support, and others can offer their data and reasoning services to aircraft operators via digital agreements. DRF addresses the need for a system to operate and interact in a complex and dense airspace operational environment while our work focuses on addressing the needs of the traditional commercial aviation industry.

These existing solutions partially address the challenges related to data access and services. There is still a need for an open architecture framework which can be used by the aviation community to not only consume reliable aviation data from multiple sources but also leverage a microservices-based architecture to build more complex and scalable services for airspace operations. The Platform addresses this gap by developing the foundational framework to provide access to high quality aviation data, build connected, re-usable and scalable services to better meet the requirements of the traditional flight operators while paving the way for new entrants. The Platform also addresses the challenges with the current aviation data by hosting and providing access to the Fuser [14] data feed. The Fuser [14] service, developed in support of NASA's Airspace Technology Demonstration 2 (ATD-2) project [15], provides easy access to FAA SWIM [16] data collected from various sources. The service consolidates, filters, and reconciles data streams, making them available as an up-to-date, well-structured flight record. Through the partner evaluation activities, Fuser has further evolved to address high-level requirements that have been gathered from aviation industry stakeholders.

III. Platform Capabilities and Architecture

The Platform offers an ecosystem for partners to browse the services available in the Service Catalog or register their own services. Service consumers can browse the available services, sample the services, and if they wish, subscribe, and use the services. Service providers can use the platform to explore available data streams and Representational State Transfer (REST) API's and use them to create new services which can be provided to others for consumption.

The primary capabilities and features of the platform include a *Service Catalog*, *Reporting and Monitoring Dashboard*, *Data Archive Service*, *Data Streaming Service*, *Identity and Data Governance*, *Peer to Peer (P2P) Communication with the Platform Client* and *Distributed Architecture* (planned capability). The *Service Catalog* is a list of services that are available to consumers to use standalone or as a building block for more advanced services. The services can be classified under different categories and can be filtered accordingly. The *Reporting and Monitoring Dashboard* allows users to visualize the Platform and service performance metrics. There are both data archival and streaming services available on the Platform via *Data Archive Service* and *Data Streaming Service*. Near-real-time data is available as a WebSocket transport using Socket.IO [17] library that enables low-latency, bidirectional and event-based communication between a client and a server. The *Identity and Data Governance* provides the necessary APIs for entities to query, register, manage, and verify identities within the Platform ecosystem. The Platform enables peer-to-peer (P2P) interaction between the platform and partner infrastructure. The long-term vision for the Platform is to move towards a distributed architecture approach where the network traffic is distributed using P2P-enabled interactions.

The following sub-sections will present implementation details about these capabilities.

A. Platform Architecture

The platform's centralized architecture facilitates the direct integration of data and services and allows for standards-based communication and collaboration between contributors to and the users of the Platform. The secure data-driven cloud-based platform provides a broad range of features supporting the integration of data and services and providing widespread access. The Platform offers various features to enable users to achieve operational goals efficiently. The two major components include the Service Catalog and the API Gateway Proxy.

Figure 1 shows the high-level architecture of the Platform. The initial approach is to register and provide access to NASA-developed services in the Catalog and, after proof-of-concept, onboard services developed and provided by the industry. The Service Catalog provides information to access registered services and offers features to invoke a service API directly. The API Gateway proxies the requests to the appropriate service infrastructure internally.

The need for vendor-agnostic design and minimal Platform infrastructure maintenance cost are some of the areas that need further investigation. One of the methods to address some of these limitations is to transition the Platform to a distributed architecture from its current centralized architecture design.

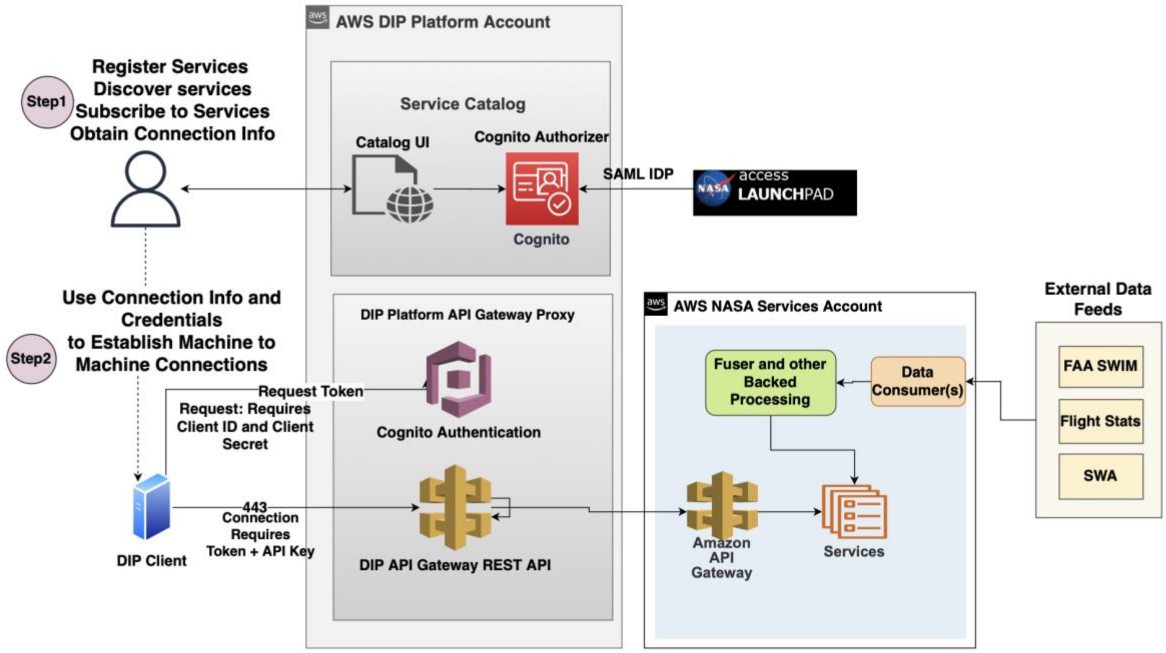

Fig. 1 Platform high-level architecture

B. Service Catalog

The Service Catalog is a repository of information about services provided by and available to the aviation community. The Service Catalog provides scalable, secure role-based access to service providers and service consumers to register, search for, try, and subscribe to services. Service providers can register their offerings with the catalog, providing service information and metadata to facilitate discovery and access to their services. Users can browse and search through the catalog, discover service APIs, and try out services. Users can subscribe to services and acquire the necessary connection information and required credentials for access to services. Figure 2 shows the typical catalog workflow on the Platform including the interactions between the Service Catalog components. The details of the current features available in the Service Catalog are provided below:

- **Authentication:** DIP uses Amazon Web Services (AWS) [18] Cognito authentication service and API Key tokens for identity verification. A valid client ID and secret-key are created generates and provided to the user via the Catalog; coarse-grained permissions are further defined and managed by OAuth custom scopes.
- **Administration (User Account & Company Account):** A user account is created on first sign-in using NASA Launchpad (Identity and Authentication Service) and prompts for Interconnection Security Agreement (ISA) Verified Token verification (a verification token provided by NASA to identify authorized users). NASA creates and maintains company accounts that lists company profile, credentials, and authentication policies.
- **Service Discovery (Search and browse):** The landing page of the web portal lists all available services offered in the Service Catalog. Users can browse and list all available services offered in the Service Catalog. Users can discover services in the search bar by entering specific categories, keywords, or both.
- **Service Registration:** Providers register services using a registration wizard and enter the required information such as service name, URL address, version, and access rights and authentication mechanism. They follow the screen prompts to enter the required fields describing the service. New services go through Service Registration approval process and the status of a service is indicated on the Service page.
- **Service Subscription:** Service consumers interact with a Service based on the subscription status of the subscription request. Users may browse all the available services on the Platform in the defined scope, regardless of subscription status, but the “Try It Out” feature will only return real results if a user’s company is subscribed to that service.
- **Platform Services:** The Platform provides NASA developed ML services, Data Access services, and Data Archive service in the Platform and users can subscribe to these services.

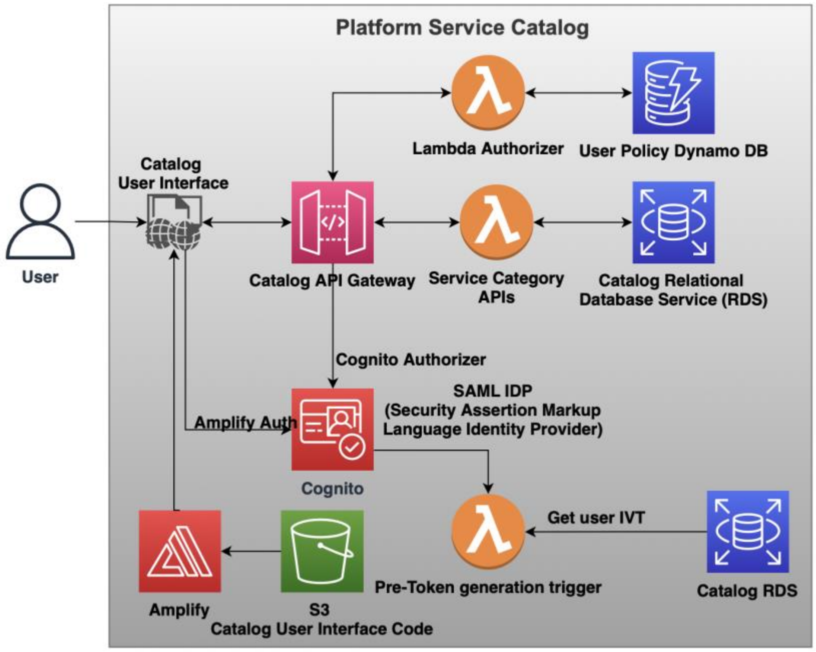

Fig. 2 Service Catalog architecture

C. Reporting and Monitoring Dashboard

Reporting and Monitoring dashboard provides valuable information about services to prospective and existing users. Quantifying the performance of the Platform and the registered services provides valuable insight needed to evaluate and ensure that the system is accessible, scalable, and performant. It includes Quality of Service (QoS) metrics such as availability, total response time, latency, and service processing time for registered services. The Platform calculates the static QoS metrics for all the services on the Platform. The Platform facilitates a dashboard to visualize metrics at various levels including dashboards, registration, subscription, and service information pages enabling transparency into performance data and building trust between the Platform and its entities.

Figure 3 indicates the Platform architecture for collecting QoS Metrics and Service Metrics. In addition to the QoS metrics, the Platform collects service provider metrics. Service providers can send performance metrics via API or send performance reports for registered services. The dynamic performance metrics are pushed to the Platform periodically to reflect current performance measures. A sample dashboard displaying QoS and service provider metrics is shown in Fig. 4.

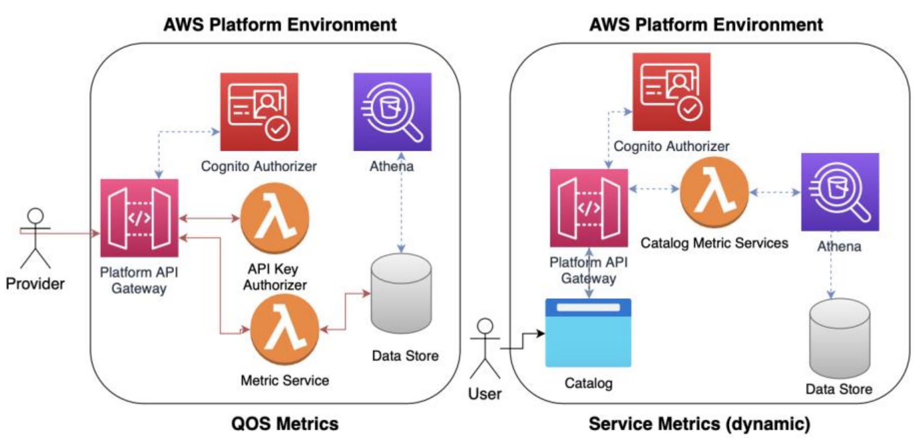

(a) QoS Metrics (b) Service Metrics (dynamic)

Fig. 3 Platform and Service Metrics

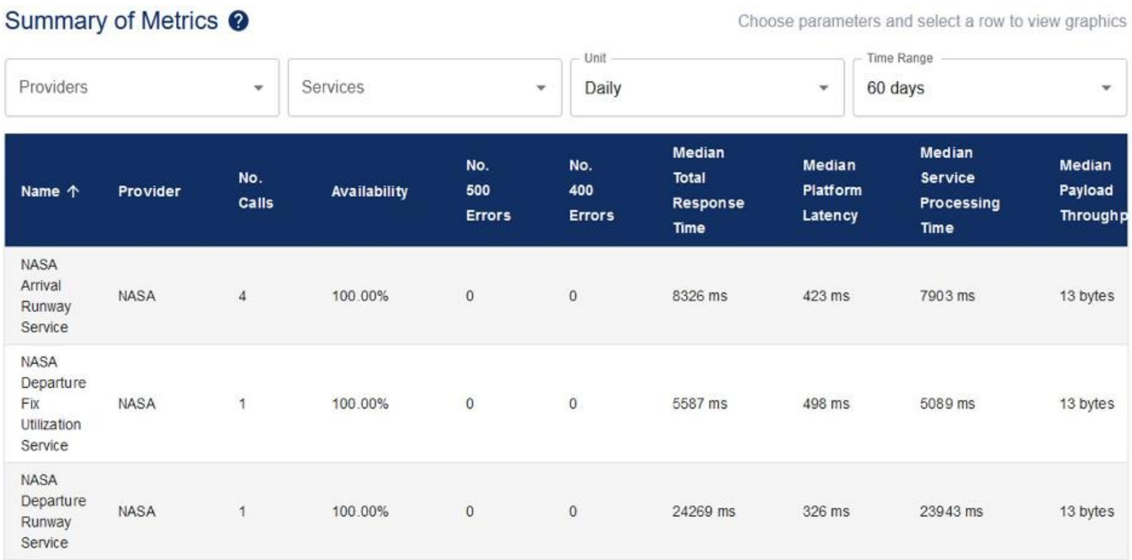

Name ↑	Provider	No. Calls	Availability	No. 500 Errors	No. 400 Errors	Median Total Response Time	Median Platform Latency	Median Service Processing Time	Median Payload Throughp
NASA Arrival Runway Service	NASA	4	100.00%	0	0	8326 ms	423 ms	7903 ms	13 bytes
NASA Departure Fix Utilization Service	NASA	1	100.00%	0	0	5587 ms	498 ms	5089 ms	13 bytes
NASA Departure Runway Service	NASA	1	100.00%	0	0	24269 ms	326 ms	23943 ms	13 bytes

Fig. 4 Platform Metrics Dashboard

D. Data Archive Service

The service providers providing data as a service must inevitably store large volumes of historical data on which consumers can perform analysis or train their own applications using the historical data. Data Archive Service provides an API with a generic data model for the service providers to describe the required information from the consumer based on their archived data environment. It is currently planned for Data Archive service to provide 6 months of data for large data transfer. Data is stored in AWS Simple Storage Service (S3) bucket and data can be transferred from an AWS S3 storage bucket to different data storage environments, cloud, or on-premises. S3 buckets are exposed to consumers over secure file transfer protocol (SFTP) using AWS Transfer family. The consumers can only access S3 buckets they are authorized to access as folders in their SFTP clients.

To monitor the data transfer over SFTP, the Platform enables Monitoring Data Transfer to prevent excessive data downloads. The communication between the Platform and the Partner storage infrastructure can be seen in Fig. 5. CloudTrail S3 Data Events are used to log specific read operations, and the logs are stored in an S3 bucket. An Athena table is used to query the log and roll up transfer totals by source/account. This allows the service providers to keep track of the data transfer and ensure that users are not downloading excessive amounts of data.

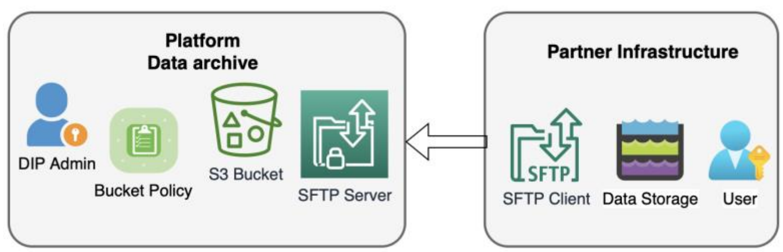

Fig. 5 SFTP Service

E. Data Streaming Service

Real-time data provision is an important aspect of many services, particularly those that require up-to-date information for analysis. The data streaming services on the Platform are designed to provide real-time data to the users. As mentioned in the Background section, Fuser live data stream is available on the Platform as a data integration service. In addition to Fuser data, there may be other data feeds that could benefit from real time provision in the future.

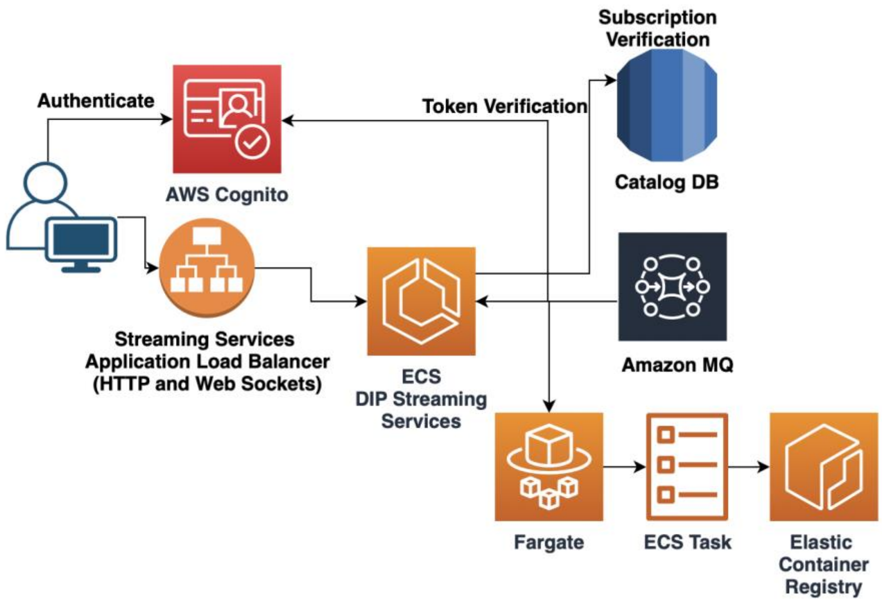

Fig. 6 Data Streaming Service (ECS: Elastic Container Service, MQ: Message Queue)

Data Streaming API is implemented using Socket.IO library that enables low-latency, bidirectional and event-based communication between a client and a server established with a WebSocket connection. Socket.IO is a software library which uses the WebSocket transport, but also adds features such as broadcasting, robust re-connection capability, and automatic message compression [17]. Figure 6 shows the Fuser I/O streaming API, which provides a near-real-time data stream. It consists of an Elastic Container Service Task, which runs a docker container with custom code that consumes data from the AWS MQ topic and emits messages to connected Socket.IO clients.

F. Distributed Architecture

Currently, the Platform employs a centralized architecture i.e., every interaction between the service provider and consumer is proxied through an API Gateway, relying on AWS API Gateway. Although this approach is robust, it comes with the limit on the network traffic that can be supported via the API Gateway. To alleviate this limitation, the plan is to move towards a distributed architecture approach where the network traffic is distributed to enable direct interactions between the service providers and the consumers. As a recommended best practice, the communication and workflow are tested on the centralized AWS API Gateway before testing it in a complex environment on the distributed architecture.

The Platform will provide, at the minimum, a reference implementation for the following capabilities: end-to-end encryption, network traffic distributions, and Platform data aspects like identities, governance, and service interactions metadata for auditing and reporting. Figure 7 provides a conceptual architecture as an extension of the architecture described in section A. Figure 8 provides a component view of Registered Identities, Governance, and Interactions Metadata capabilities that will serve as the building blocks to address the data security aspect of end-to-end encryptions and data governance aspects of identities and policies. Each component will provide APIs for the entities to interact with each other. The Platform will also offer an SDK reference implementation of the Platform client. The SDK will include plug-and-play components as sidecar microservices, enabling seamless integration with existing RESTful or Streaming data service providers. The SDK can also be used by Platform consumers that intend to build applications that interact with APIs provided by the registered services, both NASA and non-NASA. The following sections provide details on the implementation approaches that have been considered.

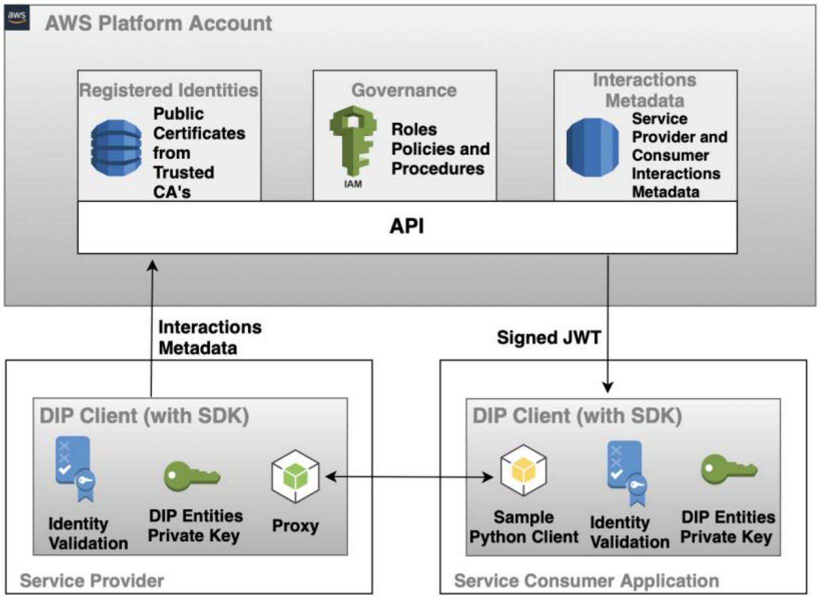

Fig. 7 DIP Distributed Architecture

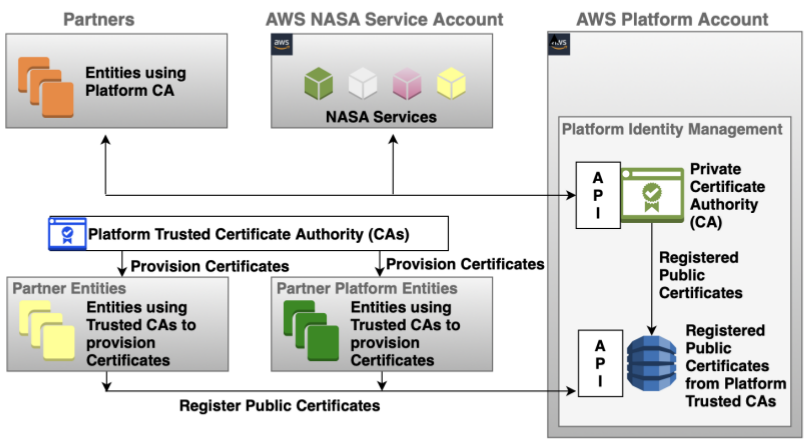

Fig. 8 DIP Entity Identity Management

G. Identity and Data Governance

As a platform, it's critical to govern the interactions within the Platform ecosystem. Therefore, the Platform should be able to uniquely identify every entity within the ecosystem and manage and retrieve the following with respect to each identity: Identity Description, Platform Access Activity, Associated Authorization Information (Roles and Policies) and Service Providers API interactions Metadata.

Identities within the ecosystem will not be limited to human users. They should extend to systems such as applications, external platforms (platforms with multiple services), service providers, services, and service instances, irrespective of the hosted environment, cloud, or on-premises.

The approach is to maintain a list of Trusted Certificate Authorities (CAs) that can issue certificates for the entities within the Platform ecosystem, making it convenient for entities procuring certificates from commercially trusted certificate authorities. The certificates can be self-validated for identity using appropriate public certificates provided by trusted CAs. For convenience to entities that do not want to procure certificates from commercial CAs, the Platform will use AWS [18] Private CA Service to issue PKI (Public-Key Infrastructure) credentials.

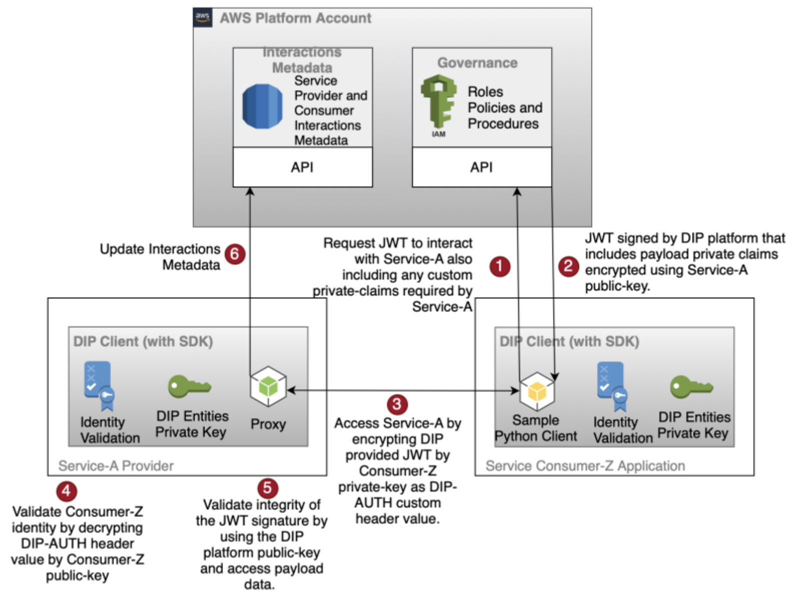

Fig. 9 DIP Identity Validations and Authorization Data

Figure 9 shows how the certificates can be managed and provisioned within the Platform ecosystem; every entity within the ecosystem can independently obtain digital certificates from the Platform trusted CA list. DIP maintains an active registry of Public Certificates associated with the entities within the Platform ecosystem. The Platform will expose an API endpoint for entities to register, manage and validate public certificates. For the Entity Identity, a standard public certificate including the public key associated with the certificate signed by the CA provisioning Public Key Infrastructure (PKI) Keys is used. The Platform Entities can be consumer of services, platform providers, service providers, and the service instances.

The Platform will also expose API endpoints to generate and sign JWT (JavaScript Object Notation (JSON) Web Token) payload data. The JWT payload will contain the following claims:

- **Registered claims:** As per the JWT standard specification.
- **Public claims:** As per the JWT name-collision prevention standard.
- **Private claims:** Following the custom platform data model, all claims in this section are encrypted by the service provider's public key.

The Platform generated and signed JWT token is passed as a custom Hypertext Transfer Protocol (HTTP) header, DIP-AUTH. It's important to note that none of the digital certificates discussed in the scope of this document are used to encrypt HTTPS network traffic; instead, they are used to validate identities and encrypt/decrypt private claims within the JWT payload private claims. DIP custom HTTP Header (DIP-AUTH) is a JWT Token with registered, public, and encrypted private claims. Figure 9 provides an overview of the steps required for entities to validate each other's identity implicitly. It is important to note that by enabling implicit identity validations, the architecture can now be scaled with distributed instead of centralized architecture principles. The Platform JWT HTTP Header (DIP-AUTH) will have all the necessary information to not only implicitly validate identity but also verify authorization to API requests. The user needs to take the following steps to validate their identity and establish connection with the platform:

1. A service consumer will request a JWT token from the Platform to interact with a service provider, and as part of the request, the consumer can include custom claims that are required by the service provider to access its API.
2. The Platform will generate and sign the JWT token, including a payload with private claims encrypted by the service provider's public key.
3. The consumer accesses a service provider's API by encrypting the JWT token using its private key and assigning it the DIP-AUTH custom HTTP header.
4. The service provider validates the identity of the consumer by decrypting the DIP-AUTH value using the consumer's public key.

5. The service provider also validates the integrity of the token by deserializing the token using the Platform's public key.
6. After the interaction is complete, the service provider's client updates the interaction metadata to the Platform.

H. P2P Communication with DIP Client

The Platform will provide a client Software Development Kit (SDK) with multiple components that can be enabled or disabled based on the use cases. The client will also log the interactions between the peers and sync the metadata to a centralized platform logging capability. Transactional metadata is logged to enable an auditing capability for disputed interactions. The Platform client will be designed to scale and easily plug in additional components in the future. Figure 10 captures the potential components to be developed as part of the Platform client SDK reference implementation.

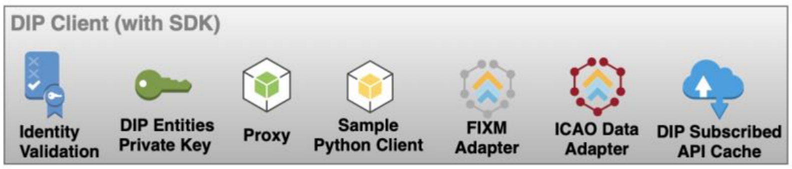

Fig. 10 Platform Client Components

1. **Identity Validation Component:** The component required by service providers implicitly validates the Platform entity identities and validates the integrity of the Platform signed JWT, as shown in Fig. 10. The component hides all the complexity of using PKI keys for encryption and decryption required in identity validation and will enable seamless integration with the existing Service Providers.
2. **Platform Entity Private Key Component:** Every entity will hold a private key associated with their public certificate registered with the Platform; the private key issued by the Platform trusted CA ensures the entity's authenticity because the private key is not shared with any other entity.
3. **Client Proxy:** The proxy component is required by the service provider to proxy the API access by consumers, as indicated the proxy functions are performed within the service providers environment enabling end-to-end encryption of interactions with the service consumer.
4. **Sample Python Client:** Service consumers require this client component to interact with service providers' DIP client proxy, the component also interacts with the platform API to request for JWT tokens to access service providers API as shown in Fig. 10.
5. **Aviation Data Adapters:** The Platform will support different aviation data exchange formats as supported by the community. For example, the Flight Information Exchange Model (FIXM) adapter will enable the service providers to map their data to FIXM standards, enabling interoperability among the Platform entities. International Civil Aviation Organization (ICAO) data adapter will enable service providers to map their data to ICAO standard data models, enabling interoperability among the Platform entities.
6. **Subscribed API Cache:** The component can be used by the service consumers to cache subscribed services' API specifications. This information can be downloaded from the Platform after subscribing to a particular service.

IV. NASA Digital Services on Cloud

One of the main objectives of developing the Platform is to support the aviation community and contribute towards aviation sustainability goals through a series of field demonstrations. To accomplish this, several high-value digital aviation services developed by NASA are currently available on the reference implementation of the Platform. They can be discovered via the Service Catalog. These services include both data integration (real time and archival) and prediction services to optimize pre-departure airport operations of aircraft. An example of the data integration service is Fuser data that provide near-real-time flight data from Fuser and access to its I/O streams. A cloud-based instance of these services can be accessed via the Platform by the users. This not only enables them to access the data in a secure and trustworthy manner but also get access to the most up-to-date version of the services registered on the Platform. To demonstrate these services in the live operational environment and capture the real-world benefits of the prediction services, the Platform is made available to a few selected US major commercial airlines who will use and evaluate the services via the Platform and provide their feedback. The performance of the Platform and the services will also be evaluated by the partners. Figure 11 shows the list of the NASA-developed services currently available on the Platform and the dependency and hierarchy among those services. Some of these services were developed during the NASA ATD-2 project and have now been deployed to the cloud and made available on the Platform.

The NASA Fuser service is an example of the data access (REST) service currently available on the platform that provides near-real-time data feed. The services return a collection of predictions and resources for a set of flights for a set period for a given airport. REST services allow the consumer to select the data elements they are interested into query. These services provide data from current time up to 30 days in the past. The API pages provide the ability to try out the service via data access REST APIs such as, Arrival/Departure Runway Utilization, Departure Fix Utilization, Arrival/Departure Runway Airport/Flight, ON/OFF Time Airport/Flight, Taxi IN Impeded/Unimpeded Airport/Flight and Taxi OUT Unimpeded Airport/Flight.

The ML models have been trained and tested successfully on historical flight data available at the North Texas Metroplex region and a few other select airports. The model performance and validation results are presented in the references [19] to [25]. The NAS Traffic Management Initiatives (TMI) REST Service returns a collection of current or historical TMIs for a given airport. The ML REST services can make various predictions at the airports the models are trained on. The models are currently trained for North Texas airspace but will be expanded to other airports at a later point in the project. They can be used in a what-if capacity to determine what the models would predict with given inputted parameters. The API pages provide the ability to try out the service and inspect its data schemas. These ML services include arrival configuration prediction [21], arrival/departure runway prediction [20], estimated ON time [22] and Taxi IN Impeded/Unimpeded AMA/Ramp Time prediction [19].

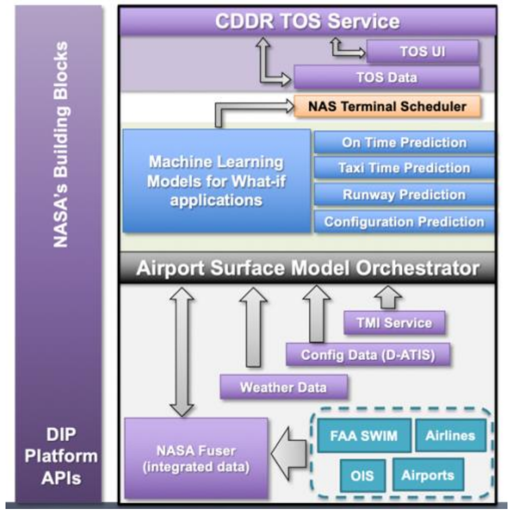

Fig. 11 Reference Implementation of NASA Developed Digital Services on the Platform (D-ATIS: Digital-Automatic Terminal Information Service, OIS: Operational Information System)

To support the first demonstration, the CDDR Trajectory Option Set (TOS) service has been tested and validated on-premise and is now deployed on the Cloud via the Platform. The CDDR solution is a ML based airport surface model to generate the unimpeded trajectory from gate to runway. This service is built upon the predictions of other ML based services and the dependency is outlined in Fig. 11. This is an example of how the individual services can be used by service providers as building blocks to enhance their own services.

V. Conclusions and Future Work

This paper presents an overview of the Digital Information Platform's (DIP) current architecture, capabilities, and the potential new capabilities to address the stakeholders' (traditional commercial flight operators and service providers) needs and requirements. The reference implementation of the Platform will contribute towards the digital transformation of the NAS by bringing together the aviation community and giving them access to a secure and trusted source of aviation data and services. The users will be able to compare, validate and identify services that provide them the maximum value. The service providers will be able to make their models available to the broader community for evaluation, testing and consumption. Academic institutions can use clean and integrated aviation data for research purposes and can potentially use the Platform as a sandbox to develop new prediction services. Easier access to the integrated and complete flight data will foster research in air traffic management.

The paper has described various platform capabilities such as Service Catalog, Reporting and Monitoring Dashboard, Archival Data Service and details of the NASA-developed services currently registered on the Platform. Features such as distributed architecture planned for future implementation and expansion of the Platform capabilities are also discussed. The Platform architecture will be further adapted to support payment gateway options which can be enabled for data monetization and business transactions once the reference implementation of the Platform is made publicly available and is adapted for commercial operations. There is also a plan to provide open and free access to the Platform and NASA-developed services to the community during the research phase of the project.

Acknowledgments

DIP research and technology development work is funded by the Air Traffic Management – eXploration (ATM-X) project supporting the NASA ARMD mission. The collaboration with the Federal Aviation Administration, Dallas/Fort Worth International Airport, and the current airline partners (American Airlines and Southwest Airlines) is greatly appreciated.

References

[1] "National Aeronautics Science & Technology Priorities" by Aeronautics Interagency Working Group of the National Science and Technology Council, March 2023. URL: https://www.whitehouse.gov/wp-content/uploads/2023/03/032023-National-Aeronautics-ST-Priorities.pdf
[2] United States 2021 Aviation Climate Action Plan. URL: https://www.faa.gov/sites/faa.gov/files/2021-11/Aviation_Climate_Action_Plan.pdf
[3] Ng, J. Y. H., McCloskey, K., Cui, J., Brand, E., Sarna, A., Goyal, N. & Geraedts, S., "OpenContrails: Benchmarking Contrail Detection on GOES-16", 2023, ABI. arXiv preprint arXiv:2304.02122
[4] Sanz-Morère, I., Eastham, S. D., Allroggen, F., Speth, R. L., and Barrett, S. R. H., "Impacts of multi-layer overlap on contrail radiative forcing", Atmos. Chem. Phys., 21, 1649–1681, https://doi.org/10.5194/acp-21-1649-2021, 2021.
[5] Avila, D., Sherry, L. and Thompson, T., "Design of Ice Super Saturated Region (ISSR) visualization tool for contrail planning," 2018 Integrated Communications, Navigation, Surveillance Conference (ICNS), Herndon, VA, USA, 2018, pp. 5B2-1-5B2-7, doi: 10.1109/ICNSURV.2018.8384904.
[6] Cavan Solutions tool for "Real-time Integrated System Evaluation (RISE)". URL:https://cavanreports.com/rise/
[7] Boeing Jeppesen Flight Planning tool "JetPlanner". URL:https://ww2.jeppesen.com/flight-planning-and-dispatch/jetplanner/
[8] RapidAPI. URL:https://rapidapi.com
[9] MuleSoft Anypoint Platform. URL: https://www.anypoint.mulesoft.com/login/
[10] FAA API Portal. URL: https://api.faa.gov/s/
[11] FAA SWIFT Portal. URL: https://portal.swim.faa.gov/
[12] Cavan Solutions CavanReports Portal. URL: https://cavanreports.com/
[13] Abdelbaky, M., Chen, J., Fedin, A., Freeman, K., Gurram, M., Ishihara, A., K., Joe-Wong, C., Knight, C., Krishnakumar, K., Reyes, I., Robinson, C., Shannon, P., Shetye, S. D., Tomljenovic, L., and Van Dalsem, W., "DRF: A Software Architecture for a Data Marketplace to Support Advanced Air Mobility", AIAA 2021-2287, Machine Learning and Big Data Applications to Air Transportation I, 2021 https://arc.aiaa.org/doi/10.2514/6.2021-2387
[14] Fuser Architecture Overview https://aviationsystems.arc.nasa.gov/atd2-industry-workshop/fuser/Fuser-Architecture-Overview_84377881.html
[15] Airspace Technology Demonstrations (ATD) Project Technical Interchange Meeting (TIM) https://nari.arc.nasa.gov/atd2021tim
[16] FAA System Wide Information Management (SWIM). URL: https://www.faa.gov/air_traffic/technology/swim/
[17] Socket IO Documentation. URL: https://socket.io/docs/v4/
[18] Amazon Web Services. URL: https://aws.amazon.com/
[19] Alexandre Amblard, Sarah Youlton and William J. Coupe, "Real-time Unimpeded Taxi Out Machine Learning Service", AIAA 2021-2401, AIAA AVIATION Forum, 2021, https://doi.org/10.2514/6.2021-2401
[20] Churchill, A., Coupe, W. J., and Jung, Y., C., "Predicting Arrival and Departure Runway Assignments with Machine Learning", AIAA 2021-2400, AIAA AVIATION Forum, 2021 https://doi.org/10.2514/6.2021-2400
[21] Rebollo, J., Khater, S. and Coupe, W. J., "A Recursive Multi-step Machine Learning Approach for Airport Configuration Prediction," AIAA 2021-2406. AIAA AVIATION 2021 FORUM. August 2021. https://doi.org/10.2514/6.2021-2406
[22] Wesely, D., Churchill, A., Slough, J., and Coupe, W. J., "A Machine Learning Approach to Predict Aircraft Landing Times using Mediated Predictions from Existing Systems," AIAA 2021-2402. AIAA AVIATION 2021 FORUM. August 2021. https://doi.org/10.2514/6.2021-2402
[23] Coupe, W. J., Jung, Y., Chen, L. and Robeson, I., "ATD-2 Phase 3 Scheduling in a Metroplex Environment Incorporating Trajectory Option Sets," 2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC*)*, 2020, pp. 1-10, doi: 10.1109/DASC50938.2020.9256509.
[24] Chevalley, E., Juro, G., Bakowski, D., Robeson, I., Chen, L., Coupe, W. J., Jung, Y., Capps, R. A., "NASA ATD-2 Trajectory Option Set Prototype Capability for Rerouting Departures in Metroplex Airspace," 2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC), 2020, pp. 1-10, doi: 10.1109/DASC50938.2020.9256640.
[25] Coupe, W. J., Bhadoria, D., Jung, Y., Chevalley, E., Juro, G., "ATD-2 Field Evaluation of Pre-Departure Trajectory Option Set Reroutes in North Texas Metroplex", 2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC), Portsmouth, VA, USA, 2022, pp. 1-10, doi: 10.1109/DASC55683.2022.9925876.