18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025.

SpaceOps-2025, ID #478

Canadian Constellation Solutions for Extensible and Scalable Operations

Bryan Cooke*, **Jason Hill, Dale Swanston**, **Brent Clark**
Calian Advanced Technologies, Canadian Space Agency, 6767 route de l'Aéroport, Saint-Hubert, Québec, J3Y 8Y9, Canada. bryan.cooke@calian.com
* Corresponding Author

Abstract

Canada was one of the significant early contributors to satellite design, manufacturing and operations, becoming the third nation to successfully construct and operate an Earth-orbiting satellite when Alouette-1 was launched in 1962. The following decades saw both the Canadian private sector and government continue to expand Canadian space capabilities. Since the launch of the RADARSAT-1 satellite in 1995, Canadian Space Agency (CSA) missions have been operated from the CSA Mission Operations Centre in St-Hubert, Quebec. The ground control system for RADARSAT-1 was a bespoke system operated by its prime operations contractor, SED Systems (later Calian, Advanced Technologies). Calian Advanced Technologies has delivered to CSA their Generic Mission Operations Centre (GMOC) systems and operations services encompassing the SCISAT (launched in 2003), NEOSSat (launched in 2013), M3MSat (launched in 2016) and RADARSAT Constellation Mission (launched in 2019) missions.

Traditionally, satellites were constructed as single large-budget missions that required unique development to deliver them. Recently, the large growth in nano-satellite constellations has driven innovations in mass manufacturing and resulted in faster delivery of satellite-related products and services to the market. These constellations encompass larger-body and more complex satellite platforms in constellation deployments, necessitating further innovation in many disciplines, including automation, scalability, redundancy, artificial intelligence and machine learning. Among the areas requiring significant innovations are the Network Operations Centres (NOCs) through which these constellations provide their products or services to the millions of users worldwide. Leveraging decades of software engineering and mission operations expertise, a critical hub designed by Calian Advanced Technologies can provide a highly scalable and reliable platform for real-time satellite data processing, communication, orbit custody, and troubleshooting. It is equipped with autonomous technologies to support satellite telemetry, tracking, and control (TT&C), resource orchestration, flight dynamics, data processing and performance analysis, aiming to enhance operational reliability, mitigate risks, and optimize mission efficiency across large constellations.

The Advanced Technologies Operations Centre (ATOC) looks to target the growing need of organizations building constellations that require the robust and highly available connectivity demanded by customers while maintaining growth capabilities matching constellation evolutions and a flexible DevOps environment for custom and complex deployments. The ATOC provides multiple redundancy nodes of operations for highly available constellation performance, leveraging advanced software and hardware to monitor, control, identify and resolve issues promptly through autonomous means. The control centre will integrate with existing satellite infrastructure and provide seamless connectivity, emphasizing continually reducing the downtime margins and maximizing constellation availability for best-in-class service.

Keywords: (maximum 6 keywords) constellation, orchestration, efficiency, scalable, automated, high availability

Acronyms/Abbreviations

AI&T	Assembly, Integration & Testing
API	Application Programming Interface
ATOC	Advanced Technologies Operations Centre
CMP	Centralized Management Platform
COLA	Collision Avoidance
CSA	Canadian Space Agency
EMS	Element Management System

FD	Flight Dynamics
GMOC	Generic Mission Operations Centre
IAM	Identity and Access Management
IoT	Internet of Things
ITM	Intelligent Telemetry Monitoring
LEOP	Launch and Early Operations Phase
OC	Orchestration Controller
OD	Orbit Determination
ML	Machine Learning
MOC	Mission Operations Centre
Mon-A-Co	Monitoring And Control
PoC	Probability of Collision
RT3	Real-Time Telemetry, Tracking & Tele-Commanding
TT&C	Telemetry, Tracking, and Control
UI	User Interface

1. Introduction

For decades, the operations of orbiting satellites were highly manual undertakings. Operators monitored and controlled satellites through the LEOP phase to ensure appendage deployments and unit boot-up sequences were completed properly. Commissioning the satellite for routine operations required complex testing and verification sequences to ensure the mission could deliver the intended results to its customers. Routine operations followed with regular taskings based on users' needs necessitated manual monitoring of the ground-space-ground custody chain to ensure data was delivered promptly. As computing technologies advanced, more opportunities became available to streamline these manual processes, reduce manual errors and undertake missions with more complex designs. Automation has allowed satellite operators to improve ground operational efficiencies even further.

These ground systems are customarily referred to as Mission Operations Centres (MOCs), and contain all the ground infrastructure, hardware and applications necessary to operate satellite missions. Due to the rapidly growing numbers of satellites and constellations expected to be launched in the next half-decade [1], the traditional MOC concept needs to adapt to successfully service evolving customers' needs. Satellites are now able to be software-defined, allowing for edge innovations mid-mission. Constellation sizes are growing rapidly, and the deployment of inter-satellite optical links brings Internet of Things (IoT) connectivity to the world. Alongside ground stations and user terminals, satellites have become just another endpoint in a large network of distributed assets. As satellites and constellation networks become more advanced and more connected requiring minimal latency, the operations network controlling them must equally increase in capability.

The Advanced Technologies Operations Centre (ATOC) provides complete network operations functions using a much more modular, flexible, extendable and inter-operable architecture that constellations require. Calian's Centralized Management Platform (CMP), currently under delivery to Telesat-Lightspeed [2] as an Element Management System (EMS), provides the backbone architecture for the ATOC. It is a system designed to service network constellations of any size while providing high levels of redundancy, resiliency and minimal latency. Leveraging modern technologies while relying on proven heritage algorithms, the ATOC can be easily configured via common configurations and Kubernetes containers. Technologies include open-source tools such as Kafka for efficient product transfer and management and data visualization tools such as Grafana and Graylog for rapid evaluations of system health. Automated features alleviate pressure on operators as the constellation grows while sharing the same infrastructure. The generic and modular nature of this system substantially reduces the time required to standup an operations control centre for new missions and enables the easy integration of mission-specific components as necessary. Multiple ATOCs integrated across diverse geographic regions provide the operational redundancy and high service availability required for these constellations. The authors will describe ATOC and each of its subsystems in detail, with a focus on how the solutions can be deployed in various configurations depending on the needs of the client mission or constellation.

2. Discussion

2.1 Architecture

The ATOC will deploy using a multi-site architecture where each site will utilize Kubernetes as a container orchestration tool. Within individual sites, Kafka will be used as the main messaging system and will facilitate communications between each microservice in the overall system. The system will store its data within Kafka, InfluxDB, MongoDB, Prometheus, ElasticSearch, and a variety of other databases (Fig. 1).

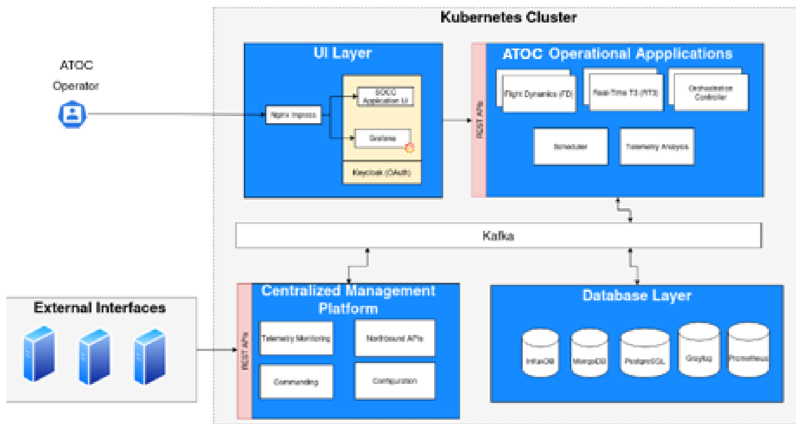

Fig. 1 Container Architecture

Connections for external interfaces (i.e. secondary control centres, payload/user terminals, ground station network, etc.) will be configured through the CMP. To support operations across multiple sites, critical software applications for the ATOC will operate in a Primary/Standby mode, where one site acts as a Primary instance while the other sites are in Standby mode. The primary ATOC site will handle incoming requests, process data, and operate the system, while Standby sites will be in a read-only state. Data integrity at each site is achieved using a variety of database replication as well as state replication mechanisms to ensure that the sites are in sync with one another.

2.1.1 Resiliency and Redundancy

The ATOC is designed to allow simultaneous operations across diverse geographic regions in a way that supports the catastrophic loss of an ATOC site without the loss of the overall system. This level of design can be supported via one of two deployment architectures – a stretch cluster or a replicated multi-cluster architecture. In either case, operations from a virtual site will opaquely interact with the ATOCs to look and behave like one system. In the case of a stretch cluster, the overall ATOC system looks and operates like one large ATOC sufficiently scaled to support the required load for all the satellites users, regardless of the user's location. Depending on the latency of the network between ATOC sites, certain services may be "pinned" to specific locations to keep service interactions performant. Additionally, data storage nodes will be present at multiple or all sites to prevent catastrophic data loss in the event of a site failure. Data is fully replicated across all data storage nodes at all applicable sites. Advantages of a stretch cluster over a multi-cluster deployment include a more simplified architecture and more flexible use of services. However, a stretch cluster is not feasible if network latency between sites is too high.

For the multi-cluster deployment, each ATOC site operates independently except for status reporting between clusters (for failure detection) and data replication. Advantages of a multi-cluster deployment over a stretch cluster is a cleaner and clearer separation between sites and less dependency on network latency. However, a multi-cluster deployment is more architecturally complicated and interactions between the users and ATOC sites is more involved and explicit.

Resiliency is the ability of a software system to continue operating effectively through power outages, hardware failures, and other disruptions within the system. A resilient system will maintain essential functionality and be able to quickly recover from any potential interruptions to ensure minimal impact on performance and availability. The system will employ several features for resilience including fault tolerance, redundancy, scalability and observability. Deployment sites will support fault tolerance by deploying any third-party applications (such as databases, messaging queues, etc.) in a clustered manner across an odd number of nodes (3, 5, 7, etc.). Utilizing the clustered nature of these services, the system can continue functioning even during the loss of one or more nodes. The number of failures that can be tolerated depends on the configuration of the service and how many members are part of the cluster. For applications that cannot be clustered and must operate individually, resilience will be provided through redundancy. Any standalone application that writes to disk will have its disk replicated to 3 (or more) nodes. This ensures that if an outage is experienced, the workload can be re-scheduled to a location where a replica of the workload's filesystem exists.

Resiliency will also be provided through the system's natural capability to scale. Because additional hardware can easily be added to the system at any point, new hardware can be added to the system to help mitigate potential failure zones if a more resilient system is desired. By adding new hardware to the cluster that exists in a different data centre or availability zone, the system's scalability can ensure that frequent outages in one zone are avoided by hardware in another zone. Finally, the key to resiliency is having good observability in the system. By monitoring the system's resource utilization, events, logs, traces and more, operators can be aware of potential upcoming problems and be able to mitigate them before they occur. For example, by observing CPU, memory, and disk utilization predictions can be made about disk utilization or potential hardware upgrades that might be required to mitigate upcoming problems.

2.1.2 Extensibility and Scalability

The ATOC system will support expansion in the quantity of managed satellites without impacting the system architecture nor interrupting ongoing operations. Adding additional constellations to the system will be as simple as configuring the details of the additional satellite constellations within the ATOC User Interface (UI). Once a constellation is added, it will be recorded in the databases, and the system responds to the additional constellation naturally. Being able to support this kind of dynamic growth is critical to supporting changes to a constellation. This is achieved by decoupling the message producers from workload processors (consumers). This decoupling allows the addition of additional replicas of a component to the system to handle an increased load without re-designing the system architecture. Furthermore, partitioning allows multiple consumer workloads to split up any incoming messages on the event bus. This design enables the ATOC system to handle all incoming messages in parallel when multiple consumers are listening to the same topic. As a result, the application load is distributed to handle multiple satellite constellations. Adding new services on the network to deliver new functions is and equally important aspect of ATOC.

In addition to the distribution of load across multiple software instances, the ATOC system can dynamically add more workloads (and, thus, add more consumers), based on cluster demand. By using Horizontal Pod Autoscalers, the system can ensure that if the software workloads need more computational power due to an added satellite constellation, the system can automatically spin up additional replicas to handle the workload. This allows the system to dynamically scale the ATOC applications to handle an increased load as additional constellation get added to the system. However, the considerations above do not consider hardware limitations. If a constellation is added but the ATOC is struggling to keep up, more hardware may be required. In these scenarios, the system can be horizontally scaled by adding additional servers to the cluster. Once the servers are added to the cluster, Kubernetes will manage the creation of additional workload containers on the new hardware which would allow the system to utilize the additional capacity to support the additional constellation. If using a cloud deployment, this behaviour could be automated to ensure new hardware instances are added as they are needed. As satellite designs and mission requirements evolve, CMP can also support the addition of new sub-systems or functional capabilities to maintain the innovation edge. New applications can be configured and integrated within the CMP architecture and tested in a group staging environment prior to operational deployment. This allows the inclusion of new functionality as a constellation evolves without re-designing the operations centre systems.

2.1.3 Centralized Management Platform (CMP)

The CMP is a feature-rich, sophisticated platform for providing core and enhanced services and features both of which are necessary and desirable for a satellite network management system. These core and enhanced services of the CMP, as well as Calian Mon-A-Co products, can be applied as the foundation for other types of systems within the satellite communications domain. This is the case for the ATOC. The CMP is utilized to provide core infrastructure and services to the rest of the ATOC applications, as well as providing interfacing software to integrate external interfaces between the ATOC and other (external) systems within a constellation ecosystem. The CMP architecture is highly scalable, performant, and extensible. It is a service-based architecture with flexibility in deployment, allowing various end-system architectures to be supported. The deployment architecture is based on Kubernetes, which allows it to support several operational ecosystems within private, public cloud, or hybrid network environments.

Fig. *2*, below, provides an illustration of the services provided by CMP to the ATOC and the relationship between these services and the other ATOC operational applications and external systems.

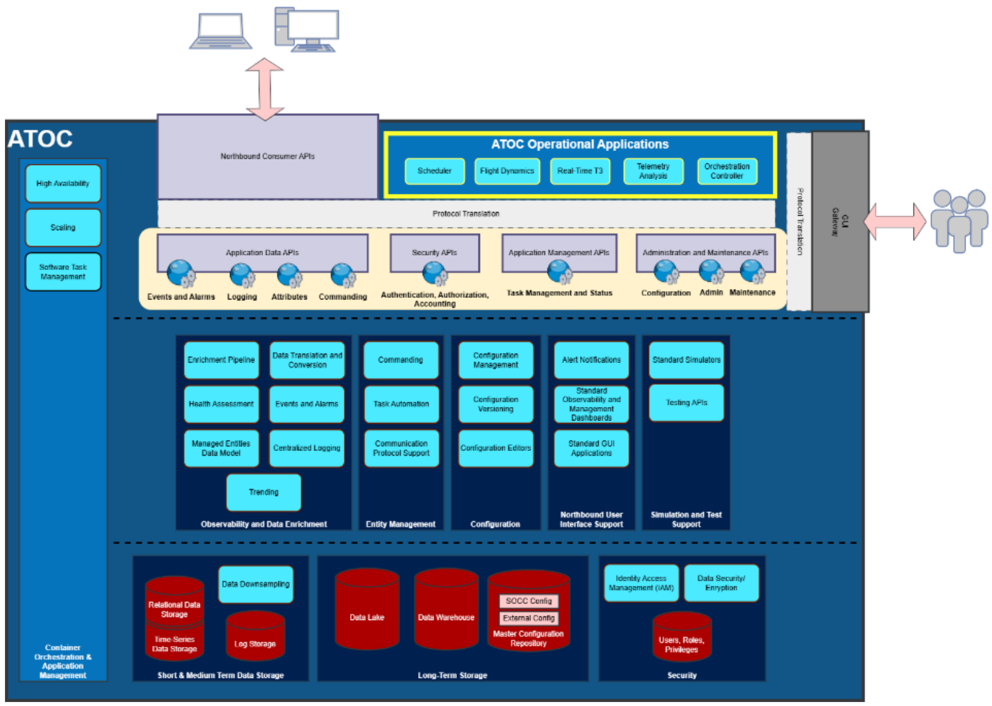

Fig. 2. ATOC Management Platform

2.1.4 Core Management

The core services provided by the CMP and are categorized into five categories: entity management, observability and data enrichment, configuration, Northbound user interface support, simulation and test support. Entity management services are focused on interfacing to external entities to be managed by the ATOC system. This includes gathering data, commanding and configuration, and lifecycle management. Additionally, this includes higher-level task automation features for creating and scheduling scripts to execute automation tasks of various levels of complexity. Services for observability and data enrichment take the data gathered through entity management and perform enrichment operations on it. This includes mapping it to the hierarchical overall data model, transforming it from raw form to a more user-friendly or specific-purpose format, assessing the health, and even creating new or derived data. Through the enrichment process, the system creates and manages events, alarms, and logs (Fig. 3).

Fig. 3. CMP Dashboards

The configuration services manage both the configuration of external entities and the internal configuration of the CMP. The CMP provides various editors to allow configurations to be viewed and modified. The full set of configuration data is stored in a version control system with change management features. Northbound services focus on upstream users of the system – both human and machine. The services provide API support for machine-machine interfaces. Additionally, the CMP provides a dashboarding system that supports user-created content, as well as provides sophisticated applications designed specifically for the platform. Alert notifications for alarms and other activities are also available. Simulation and test support services are core to the CMP which include standard simulator and infrastructure to help create additional more specialized simulators. These simulators can be used to support external interface development, system scaling, troubleshooting, and other activities. Additionally, the CMP provides an API to help support testing activities as they pertain to entities managed by the CMP.

2.1.5 Data Storage and Security

The CMP manages data storage for various types of data across different storage durations. The storage services of the CMP can be summarized into short/medium-term and long-term storage. Implementing short and medium-term storage depends on the type of data and use cases for the data as different data types will be stored in different types of databases. Most data are stored in time-series format that maintains the history of changes to that data. A down-sampling facility is available that allows data to be reduced in granularity and scope as it ages. Some data is better stored in a relational format. Finally, log data is stored in its own database that is specialized in storing and querying text-based data efficiently. As data ages, it is moved from short- and medium-term storage to a more appropriate longer-term storage data lake or data warehouse. Additionally, within this category is the configuration version management system used to manage both the configuration of monitored entities and the CMP, itself.

The CMP provides a high level of security for your data and operations through the use of high-grade third-party identity management software compliant with industry standards and best practices that include the use of TLS, OAuth2, managed certificates, encrypted account data, single sign-on, and more. Because the CMP uses industry-standard security practices, the integration of other Identity and Access Management (IAM) compliant systems is possible without extensive effort.

2.1.6 APIs and Protocol Translation

The CMP makes its core services, data storage, and security services available through consistent APIs and API gateways. The API gateways consolidate APIs from multiple into a single access point based on a theme or category. The API gateway categories are the Application Data APIs, Security APIs, Application Management APIs, and Administration and Maintenance APIs. The CMP also provides facilities for protocol translation and mapping. These facilities can be used to map API calls and/or data structures from one API to another and are often used to bridge external APIs to CMP APIs without code modification. This is especially useful for interfaces that cannot be modified or would be expensive/labour-intensive to modify. Additionally, the translation mappings are used to map data from external sources into CMP data models and vice versa. Finally, the translation facilities can be used to modify data from external sources into a more usable format within the CMP-based system.

2.1.7 Container Orchestration and Application Management

The Container Orchestration and Application Management service provide the primary facilities used to scale and expand the ATOC, as well as provide the fundamental high availability aspects of the system architecture. The Kubernetes container orchestrator is used to manage the ATOC clusters. The CMP supports multi-cluster and stretch-cluster deployments for site diversity. Each deployment architecture has pros and cons and the decision of which deployment to use is based on factors that include the operational use cases, required level of availability, network characteristics, operational facilities. The ATOC architecture provides containerized services loosely coupled through the core infrastructure provided by Kubernetes and Kafka. Applications of the ATOC are designed to either be stateless or require little time to rebuild their state from temporary storage. This allows the ATOC to support a horizontal scaling model (adding more application instances to handle a higher load), in addition to a more traditional vertical scaling model (add more hardware resources), to provide a robust and flexible method of scaling the ATOC from a few satellites to a large fleet or fleets of hundreds or thousands of satellites. In addition to the core services, the CMP provides additional task management, cluster monitoring, and troubleshooting capabilities. These additional capabilities are focused on providing operators with visibility into the system operation and tools to help alleviate problems reported to the operators, in addition to easing the maintenance of the system for administrators.

2.2 Operational Applications

The user-facing applications used to monitor and control the constellation work in an automated and synchronous fashion are shown in Fig. 4. These include Flight Dynamics (FD), Resource Scheduler, Intelligent Telemetry Monitoring (ITM), Real-Time Telemetry, Tracking & Tele-Commanding (RT3) and Orchestration Controller (OC). Many operations are event-based where sub-systems autonomously respond to messages received from the CMP message bus while some operations require time-specific executions.

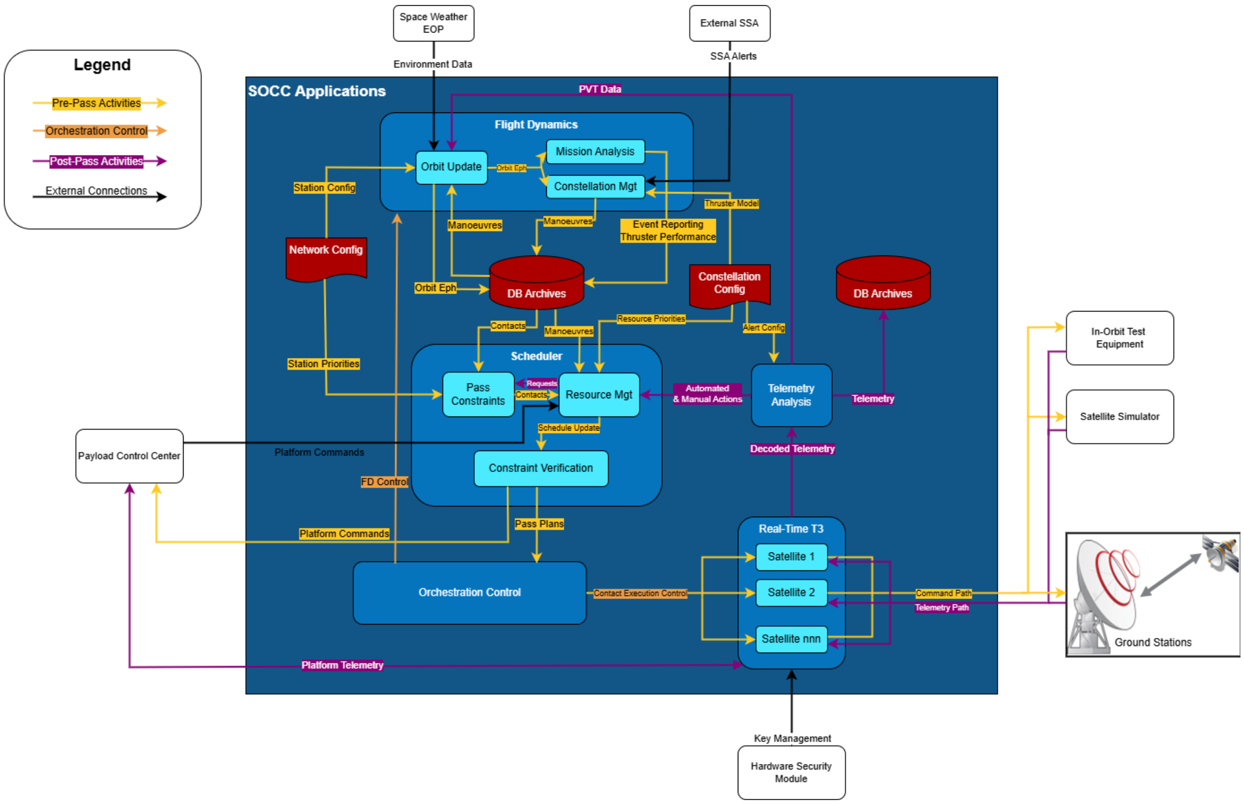

Fig. 4. ATOC Applications

2.2.1 Flight Dynamics (FD)

Orbital custody of the constellation is handled by the automated FD segment, which itself constitutes the Orbit Update, Mission Analysis, and Constellation Management services. Orbit Update performs orbit determinations (OD) for each satellite that has new tracking data available. Any executed manoeuvres are incorporated into the OD and their performance is evaluated. Definitive orbit ephemerides resulting from the OD solutions are generated. All definitive orbit ephemeris products and manoeuvre evaluation results are sent via the CMP messaging service to the archives for storage.

Constellation management executes at predetermined and/or event-driven times, generating all station-keeping and collision avoidance (COLA) manoeuvres. COLA manoeuvres are internally managed in cases of intra-constellation events or triggered by external requests for conjunctions with non-constellation objects. The latest orbit solutions are retrieved from the database and verified against the inter-satellite separation requirement (both in-plane and out-of-plane) within a predefined look-ahead period. If any separation criteria violations are detected within the look-ahead period, an avoidance manoeuvre is triggered to prevent the violation and added to the database. While the number of orbital plane intersections in a constellation are highly dependent on orbital inclination, it is reasonable to expect larger numbers of conjunction alerts at the Earth poles. Conjunctions against other satellite objects or debris would be processed ad-hoc against a defined level of probability of collision (PoC) risk.

Mission Analysis propagates the OD solution to provide an accurate prediction of the short and long-term constellation trajectory. Any manoeuvres defined by Constellation Management are incorporated into the propagated orbit to ensure accurate timing of events and ground station contact times. Mission Analysis contains a full representation of the constellation state where automated or manual reporting can be configured (such as eclipse predictions, yaw steering, attitude collinearities, etc).

2.2.2 Resource Scheduler

The Resource Scheduler is a mission planning system capable of autonomous operation using high-performance algorithms to support the operator and produce optimized mission plans. It is comprised of three services that enable a fully automated and optimized tasking plan to be generated for each satellite within the constellation. The Scheduler service is responsible for the automation of scheduling activities and the generation of pass plans. This takes into consideration the optimization of scheduling and a priority system to allocate as many activities as possible, given resource availability and constraints. The Configuration service is responsible for management and distribution of all shared configurations across the constellation or sub-sets of the constellation. This includes activity and pass constraints, priority schemes, and planning and scheduling periods. The Monitoring service is responsible for logging and generating alerts based on the Resource Scheduler's activity.

Given a set of inputs, the Scheduler creates a tentative schedule for optional approval by the operator. The operator can edit this schedule by adding, editing, or deleting activities before committing. Control over the optional schedule approval is controllable and can be managed through the Scheduler configuration. Regardless of the status of the approval step, the option to modify an existing schedule is always available via the user interface UI. Editing of pass plan products is supported for all items. Constraints are configurable and managed through the UI and can include ground station constraints (band, polarization, pre/post-pass activities), spacecraft constraints (state, power, storage), and activity constraints (parent/child relationships, parameter verifications). The Resource Scheduler supports a configurable priority scheme to guide the optimization process, allowing for spacecraft utilization to vary depending on the overall constellation demands and user needs.

2.2.3 Real-Time Telemetry, Tracking & Tele-Commanding (RT-3)

This system is responsible for the real-time automated contact operations between the ATOC and the respective satellite via the scheduled ground station. Individual pass plans generated by the Resource Scheduler are converted to spacecraft command procedures in preparation for upload, including timeline activities to be executed in the back-orbit. The constellation database defines the command and telemetry mnemonics along with any conversions and limits. A mission database management tool is used to modify database records within a simple UI. Live telemetry is displayed on user-defined pages which contain configuration parameters to colour-coded values based on status and/or limit violations specified in the database. The status of the ground station antenna and other equipment is received by RT3 and streamed into CMP. Using the CMP Observability and Data Enrichment facilities, this status is monitored for alerting conditions and displayed to the user (including the antenna path). New telemetry is extracted from the raw binary files upon reception of downlinked files. This telemetry extraction process can also convert raw telemetry counts into engineering values. The extracted telemetry is sent to Intelligent Telemetry Monitoring (ITM) for telemetry analysis and CMP for archiving.

2.2.4 Intelligent Telemetry Monitoring (ITM)

The ITM uses Calian-developed Machine Learning (ML) models [3] to autonomously detect anomaly signatures in satellite telemetry. Executed as soon as telemetry is downlinked and decrypted, the new telemetry is processed by the Time-Series Anomaly Detection (TSAD) service and compared against the expected nominal behaviour (Fig. 5). Any anomalous signatures are automatically identified and categorized (if possible) against known signatures. Automated notifications will trigger configured automated recovery actions within the Scheduler system. Dashboards are available for operator evaluation, signature confirmation or false positives and manual intervention when needed. Model re-training is available at any time to ensure satellite and constellation performance is properly monitored based on the expected behaviours. Examples of events that can trigger re-training include new anomalies that do not fit within already defined categories, spacecraft performance degradation, and flight software updates. These models can be adopted across the constellation as a whole, specific tranches of the constellation, or individual satellites depending on the specific needs. Newly learned signatures from a single satellite can be deployed across targeted satellites (such as the same generation) or even the whole constellation, providing shared intelligence across the fleet.

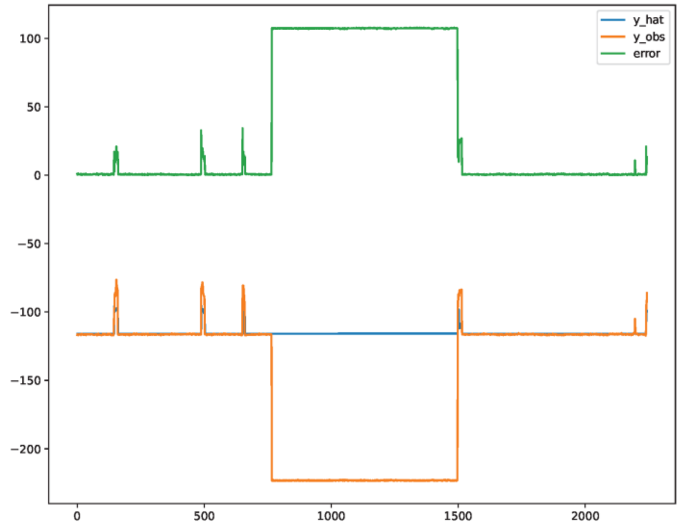

Fig. 5. Anomaly Detection

Analysis can also be performed on a configured and regular basis using the same machine learning models or using specialized models to monitor telemetry over longer periods of time (for example, long-term trending). The results of this analysis will permit engineering personnel to further investigate potential subsystem degradations and better assess the overall health of the spacecraft/constellation in a more efficient manner. Identifying these signs of degradation early is key for taking pre-emptive measures which can reduce any mission downtime and prolong spacecraft lifetime. Telemetry from Phase D AIT (Assembly, Integration & Testing) can also be used to train the ML models and identify any anomalies before launch. This can facilitate the development and testing new flight procedures, especially for new tranches or generations of satellites. By leveraging high-performance computing, telemetry evaluation can be executed rapidly upon reception of data ensuring timely recovery actions can be scheduled and implemented to support the extremely high levels of constellation availability required by the users.

2.2.5 Orchestration Controller (OC)

The OC is responsible for driving the subsystems of the ATOC that are not event-based. This predominately involves the RT3 sub-system but can include others for specific use cases. The OC initiates action in the subsystems as needed and provides monitoring and feedback control to ensure their successful completion. The Resource Scheduler sub-system generates many multiple pass plans for each satellite/ground contact pair, which must be delivered to the appropriate control path for execution. The OC ensures the distribution of the pass plans to the relevant RT3 station based on the resources being controlled by the RT3 and command authority.

The OC handles command authority between workstations. This includes the issuance of command authority based on role and satellite information, switching command authority between workstations, and the enforcement that each control shall only have a single workstation with command authority. The OC facilitates the switching of command authority between workstations by referencing role definitions and satellite identification tables to ensure the target workstations are issued the proper command authority. The OC also references the command authority assigned to all other workstations to ensure there is no overlap when the new command authority is issued. Conflicts are resolved by revoking the existing command authority from other workstations before issuing command authority to the target workstation.

2.2.6 User Interface (UI)

The ATOC UI provides rapid situational awareness of the constellation for everything from satellite health, orbit custody and network management. The UI is designed as a web-based experience for the various types of users who interact with it. The main everyday use cases are handled with a set of React-based screens, while Grafana supplies supplementary screens for things like logs and low-level charting in separate browser tabs. URL linking from button clicks on the React screens provides navigation to the Grafana screens, creating a feeling of a single pane of glass. UI functions include playback options, geographic mapping, telemetry and timeline charts, schedule, and constellation greenboard views (Fig. 6 and Fig. 7).

Fig. 6. Greenboard View

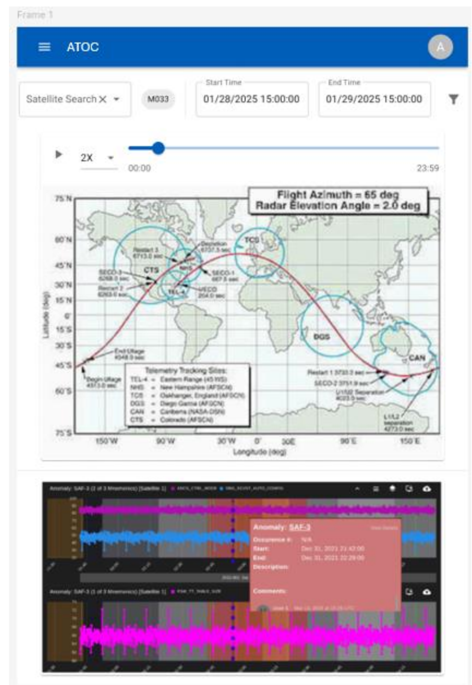

Fig. 7. Map View & Telemetry Charts

2.2.7 Configuration and Life Cycle Management

ATOC configuration management ensures the proper version control, change management and configuration baselines are maintained. Tracking all software, firmware, and hardware versions, including maintaining a record of updates, patches, and configurations within the ATOC, is necessary for proper system control and maintenance. Any changes made to the system (e.g., software updates andhardware upgrades) should be done systematically with a change management process that includes testing, approval, and documentation. Maintenance of an up-to-date configuration baseline prevents issues during upgrades and changes and ensures compatibility with existing systems. The life cycle maintenance encompasses the ongoing processes involved in ensuring the continuous operation, performance, and security of the centre's systems throughout their operational life. This includes maintenance strategies from deployment through to decommissioning and document and asset management to keep detailed records of all hardware, software, configurations, and versions.

3. Conclusions

By taking advantage of the capabilities now present in advanced computing technologies, Calian's ATOC provides key advances in efficient and secure satellite management. This infrastructure provides the high levels of redundancy, availability, scalability and flexibility that are essential for delivering "always on" mission availability and the rapid decision making that is required by the customers of today and tomorrow.

A highly available system is required to ensure redundant systems for communications, constellation custody, and data retention. Data handling systems must manage large volumes of telemetry and databases to rapidly provide the needed information to the relevant systems. Failover systems are essential to ensure continuous satellite operations, even during technical failures. Calian's CMP provides these backbone functions for the ATOC and delivering multiple distributed but connected centres can provide even more robust services. Automation across these various systems is crucial to ensure the high levels of availability demanded by users. Intuitive interfaces for operators to monitor satellite health, perform real-time adjustments, and view data visualizations are just as important as the underlying functions. Leveraging automation and AI for routine tasks reduces operator workload while maintaining manual override options. This allows operators to manage large constellations more effectively, leading to faster decision-making and, ultimately, mission availability for the users. Leveraging advanced software systems for mission planning, satellite health monitoring, constellation management, data analytics, and anomaly detection ensures the delivery of automated monitoring and response actions, rapid data processing and improved accuracy in the overall constellation situational awareness.

The ATOC is designed with the flexibility to integrate future satellite missions and new technologies without impacting performance. Scalability is crucial for handling changes in satellite constellation size, expanding data storage, or adapting to new mission requirements. And finally, these attributes of the ATOC architecture result in an operational cost reduction across the constellation, allowing for more focus on creating additional value with new products and services for the market. Calian's mix of technical experts (engineers, analysts) and operational staff who manage satellite health, troubleshoot issues, and handle mission operations are invaluable resources in designing and de-risking an advanced constellation operations centre. Combined with our experienced systems engineering, software development, data analytics and quality control experts, Calian is uniquely positioned to deliver a best-in-class constellation solution.

References

[1] GoldmanSachs, The global satellite market is forecast to become seven times bigger, 5 March 2025, https://www.goldmansachs.com/insights/articles/the-global-satellite-market-is-forecast-to-become-seven-times-bigger, (accessed 06.04.25).

[2] Telesat, Telesat Selects Calian to Develop and Deploy Element Management System for Telesat Lightspeed Network, 9 September 2024, https://www.telesat.com/press/press-releases/telesat-selects-calian-to-develop-and-deploy-element-management-system-for-telesat-lightspeed-network/, (accessed 07.04.25).

[3] W. Qian, J. DeJong, B. Cooke, Prompt Anomaly Detection for Small Satellites in Low-Earth Orbit Constellations: A Machine Learning Approach, 2024 Smallsat Conference, 2024