RA231 PRODUCT

SEXTANT

AI-Assisted Mission Planning & Model Calibration

Where prediction meets reality.

Chris Bomhardt
CEO, Right Ascension 231
chris.bomhardt@ra231.com

Jonathan Morton
CTO, Right Ascension 231
jonathan.morton@ra231.com

STScI Digital Engineering & MBSE Symposium

May 20–21, 2026
Baltimore, MD


Model Authoring


Plan Evaluation


Calibration

Who we Are

NASA GSFC Space Science Mission Operations

We fly 8 spacecraft for Helio, Planetary, & Astrophysics in all orbital regimes

We build and manage ground systems for those missions using vMMOC

Collaborating with GSFC, APL and AMMOS

THE PROBLEM

Mission Planning Is Constrained By Human Bandwidth, Not Computational Capability.

RA231

Models Take Months

Subsystem experts build activity models independently — power, thermal, data volume — then reconcile across teams. Months of work before a plan can be simulated.

Calibration Is Quarterly

Telemetry streams continuously. But comparing predictions against flight data is manual and tedious. Calibration happens every 3–6 months. Models drift between reviews.

Conflicts Caught Late

Planning errors surface during review or simulation runs — sometimes overnight. As missions scale to constellations, this bottleneck compounds.

The cycle — model, plan, simulate, fly, compare, calibrate — is too slow. Telemetry carries insight that sits unused between quarterly reviews.

ISP Architecture: Three Main Layers

Mission Applications

Mission Operations Capabilities

Telemetry, commanding, mission planning, flight dynamics, and data analytics

Mission Services

Unified Data Architecture

Data Lakehouse, event streams, object store, distributed queries, and analytics

Mission Platform

Infrastructure Automation

GitOps workflows, policy enforcement, continuous delivery, and self-healing

Three architectural areas working together to provide comprehensive ground system capabilities.

FOUNDATION

Built on Intelligent Space Platform

TELARIS DATA PIPELINE (The Sextant Enabler)

Sextant does not build a new pipeline — it places an AI intelligence layer on top of the Telaris architecture already in production across NASA missions.

PlanDev Integration: SSMO/RA231 has been collaborating with the NASA AMMOS team for the last 2 years to bring the entire AMMOS catalog running on top of this platform.

CAPABILITY 1 OF 3

Intelligent Model Authoring

Natural language → validated PlanDev activity models.

Describe spacecraft behavior in plain English. Sextant drafts syntactically valid PlanDev activity types — with resource effects, temporal constraints, and parameter bounds — then validates them against the PlanDev schema before delivery. Engineers review and approve. The AI drafts; the engineer validates.

TODAY

- Multiple subsystem teams build models independently
- Reconciliation rounds add weeks of review
- Model development: 2–6 months typical
- Requires deep PlanDev expertise to author

WITH SEXTANT

- Describe behavior in plain English — Sextant drafts
- Schema-validated before delivery — no syntax errors
- Model development: hours, not months
- Learns your mission's patterns over time

Human role: Review and approve every generated model. The AI drafts; the engineer validates and deploys.

CAPABILITY 2 OF 3

Automated Plan Evaluation

RA231

Pre-uplink constraint checking via real PlanDev simulation. Instant feedback.

EXAMPLE OPERATOR INTERACTION

Operator: "Check tomorrow's plan for the 14:00 UTC pass."

Sextant: "I found 3 issues:

① Battery margin drops to 38% at 14:32 (min 40%). Suggest: delay imaging 12 min.

② Antenna conflict: X-band overlaps calibration slew by 45s. Suggest: extend downlink +1 min.

③ SSR 2.1 GB over capacity. Suggest: reduce targets from 4 to 3.

Want me to simulate the plan with these adjustments?"

Real simulation, not inference
Every answer is grounded in an actual PlanDev simulation run — not AI guesswork.

Catches complex multi-constraint conflicts
Cross-subsystem issues that manual review misses — especially in multi-instrument plans.

Operator decides
Sextant identifies and explains. The operator chooses which suggestions to accept. No plan is modified without explicit approval.

Human role: The operator decides which suggestions to accept. The AI never modifies a plan without explicit approval.

CAPABILITY 3 OF 3

Continuous Model Calibration from Flight Data

RA231

Telemetry vs. PlanDev predictions, continuously. Calibration shifts from quarterly to daily.

EXAMPLE CALIBRATION RECOMMENDATION

Sextant: "Across 12 reaction wheel desaturation events:

Predicted power draw: 4.5W avg → Observed: 5.1W avg (+13%)
Predicted duration: 120s → Observed: 135s avg

Recommend updating model:
power_draw: 4.5 → 5.1 W duration_estimate: 120 → 135 s

Confidence: HIGH (consistent trend, no outliers). Improves battery margin predictions ~2% for plans with 3+ desats/orbit."

The Compounding Impact

1. Plans built on accurate models have fewer conflicts
2. Fewer conflicts → less rework → faster time-to-uplink
3. Better predictions → tighter margins → more science per orbit
4. Models trained on one mission accelerate the next

Human role: Approve every parameter change. The AI identifies what needs updating and by how much; the engineer confirms it makes physical sense.

STATUS & DEPLOYMENT

Where It Runs — And Where We Are

Cloud-Hosted

AWS Bedrock / GCP Vertex AI
Pay-per-use. No GPU infra.

On-Premise

Air-gapped, classified missions.
Self-hosted, no internet.

Hybrid

Train in cloud.
Deploy air-gapped via secure transfer.

WHAT WE'RE BUILDING NOW

- AI agents with tool-calling access to the Telaris data mesh (PlanDev, YAMCS, Trino MCP servers)
- Cloud-hosted and local model serving with pluggable backends (Bedrock, Vertex, on-premise LLMs)
- Agent runtime connecting AI models to operations tools via Model Context Protocol (MCP)
- Fine-tuning pipeline for mission-specific AI specialists
- Full PlanDev and YAMCS tool integration — model authoring, plan evaluation, continuous calibration

Backup Information

ISP Architecture: Three Main Areas

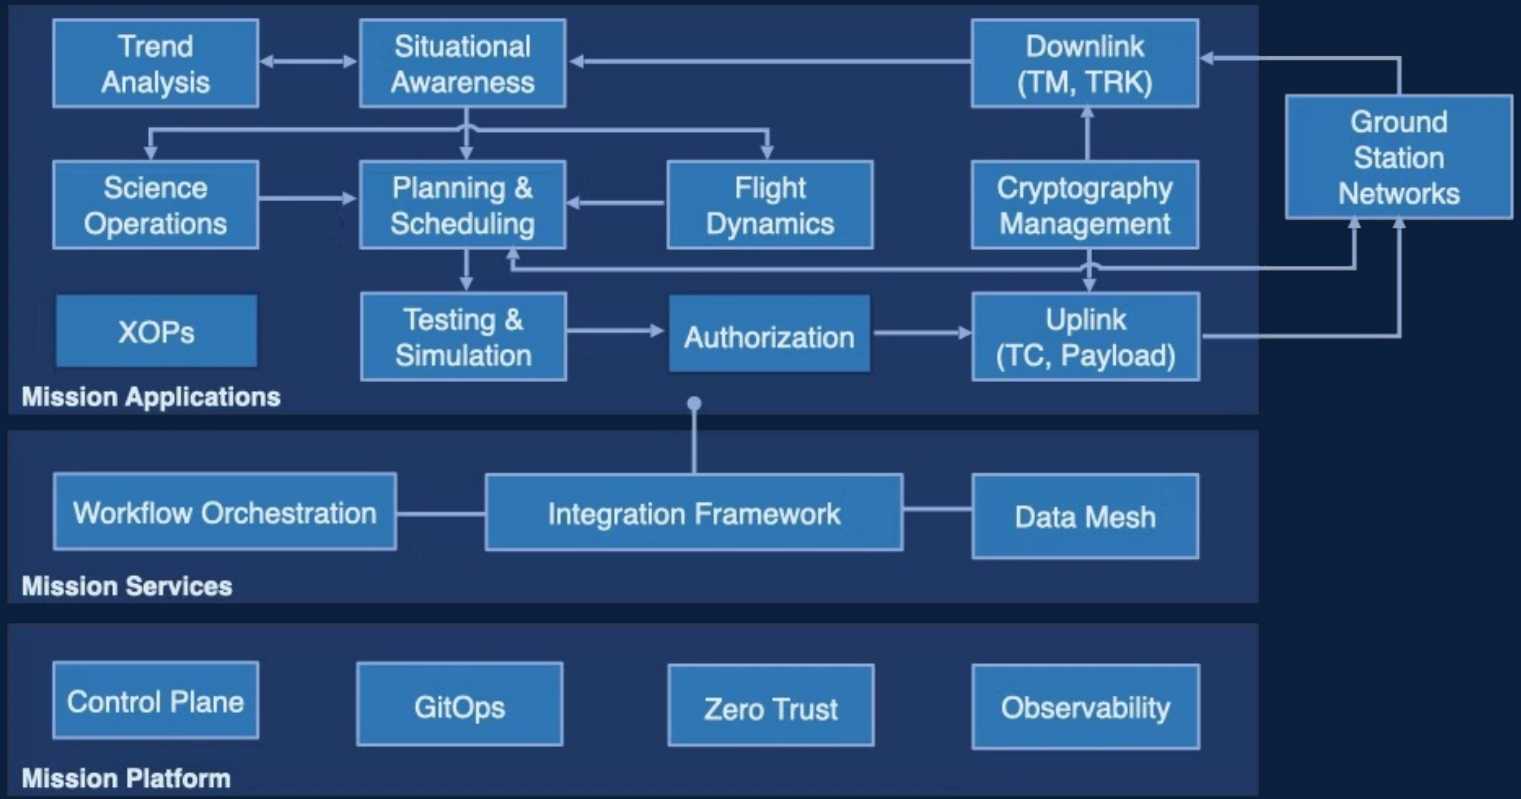

Application Layer: Mission Operations

Capabilities: Complete mission operations from telemetry to commanding

Mission Automation

- Observe (Tracking, Conjunction, Telemetry)
- Orient (Recognition, Conjunction Analysis, Trending)
- Decide (Planning & Scheduling, Simulation, Verification)
- Act (Commanding, Authorization, Procedures)
- Manage (Policies, Identity, Incident Reporting/Alerting)

Standards: OpenAPI (sync), AsyncAPI (async), JSON Schema, CCSDS Mission Ops ontology

Application layer provides complete OODA loop for mission operations with standards-based integration.

Services: Unified Data Architecture

RA231

Capabilities: Real-time streams + Analytics = unified architecture

Integration Framework

• **Serverless Functions (Knative)**	Scale to Zero adapters
• **Persistent Streams (Kafka)**	Event backbone
• **Workflow (Synapse)**	Orchestration Engine
• **Stream Processing (Flink)**	Real-time analytics

Data Lakehouse

• **Data Catalog (Nessie)**	Git-like versioning for data
• **Distributed Query (Trino)**	Federated SQL across source
• **Tabular Data (Iceberg)**	Time travel queries, schema evolution, it's a format
• **Object Data (S3 compatible)**	Cost-creative storage

Stream Processing

• **Persistent Streams (Kafka)**	Event backbone with ordering guarantees
• **Stream Processing (Flink)**	Stateful real-time analytics

Data mesh provides Lakehouse analytics + Hybrid SOA & Event Driven architecture

Platform: Infrastructure Automation

RA231

Capabilities: Self-healing infrastructure with Kubernetes & GitOps

GitOps & Continuous Delivery

• **FluxCD**	Declarative GitOps for Kubernetes
• **Crossplane**	IaC across clouds
• **Harbor**	Helm charts & container images
• **Bootstrap IaC**	Initial cluster provisioning

Core Services

• **Service Mesh**	Istio with Envoy proxy
• **Auto Scaling**	Karpenter for node provisioning
• **Cluster DNS**	CoreDNS/external DNS (Route53)
• **Observability**	Grafana, Mimir, Loki, Tempo

Policy & Security

• **Policy as Code (Kyverno)**	Declarative validation and mutation
• **Software Bill of Material**	Automated TLS certificate management
• **Scanning**	Static (Trivy) & Dynamic analysis (Falco) of workload behavior
• **IAM (Keycloak)**	Cloud provider integration

Control plane provides zero-touch operations through GitOps reconciliation and policy enforcement.

What is PlanDev?

A multi-mission **activity planning, resource simulation, and command sequencing software framework** developed by the Advanced Multi-Mission Operations System (AMMOS) team at JPL.

https://nasa-ammos.github.io/plandev-docs/

How It All Flows Together

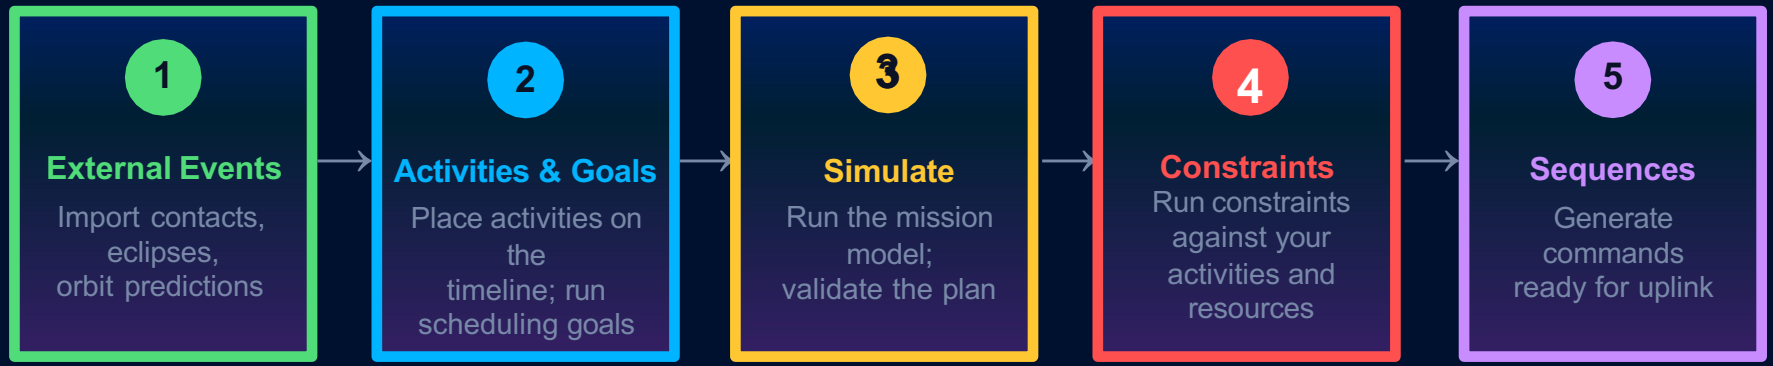

Solution

RA231

RA231 SEXTANT

An AI intelligence layer that continuously closes the loop between simulation and reality.

Intelligent Model Authoring

Natural language → validated PlanDev models.
Weeks → hours.

Automated Plan Evaluation

Pre-uplink constraint checking via real PlanDev simulation.
Instant feedback.

Continuous Model Calibration

Telemetry vs. predictions, continuously.
Quarterly → daily.

This is not an autonomous planner. Sextant doesn't generate or uplink plans on its own. It's an intelligent assistant — it drafts, it flags, it recommends. The operator decides at every step.

ARCHITECTURE

How It Works: LLM-Modulo + Human-in-the-Loop

Tools, Not Magic

Sextant calls real tools — PlanDev for simulation, YAMCS for telemetry, Trino for lakehouse queries. When Sextant says "battery margin drops to 38%," it ran a real PlanDev simulation to get that number.

Human-in-the-Loop, Always

The AI suggests. The operator decides. Every model, plan change, and parameter update requires review and approval. All actions are logged and auditable under the operator's identity.

Gets Smarter Over Time

Operator interactions become training data. Models learn your mission's patterns. Knowledge transfers across missions — Mission B starts from where Mission A left off.