Abstract

The growing number of spacecraft operating in Earth orbit is creating an increasing need for capabilities that extend beyond launch and initial deployment. Satellites may require inspection, relocation, anomaly assessment, hosted payload support, servicing, or other forms of physical interaction after reaching orbit. These operations are typically addressed through mission-specific spacecraft, dedicated transfer vehicles, or other specialized systems, which can limit flexibility and require substantial resources for individual missions.

This whitepaper proposes the concept of the **Reusable Orbital Drone (ROD)** as a class of uncrewed, autonomous or highly automated vehicles designed to repeatedly operate in orbit and perform multiple types of post-launch missions throughout their operational lifetime. Rather than designing a separate spacecraft for each mission, an orbital drone is intended to combine reusable vehicle infrastructure with modular mission capabilities, allowing a common platform to support different operational requirements over multiple flights.

The paper develops a systems-level framework for reusable orbital drones, including their defining characteristics, vehicle architecture, mission lifecycle, autonomy requirements, reusability considerations, mission interfaces, and potential operational applications. It examines the technical functions required for orbital transit, rendezvous, proximity operations, mission execution, departure, recovery, refurbishment, and subsequent redeployment. The paper also proposes a preliminary classification of orbital-drone missions spanning orbital transfer, inspection, hosted payload operations, emergency response, servicing, and future orbital sustainability applications.

The underlying approach is to treat orbital vehicles not solely as mission-specific spacecraft, but as **reusable operational infrastructure** capable of providing services across multiple missions and customers. This shift introduces new considerations in vehicle utilization, turnaround, modularity, autonomy, and fleet operations, while also creating technical challenges in propulsion, guidance and navigation, proximity operations, fault management, safety, and long-duration reuse.

Orbit Keeper, a platform under development by Graviron Aerospace, is presented as an implementation of this broader architecture. The objective of this whitepaper is not to define a single vehicle configuration, but to establish a systems framework for understanding and developing reusable orbital drones as an emerging class of post-launch space infrastructure.

1. Introduction

Space operations have historically been organized around individual missions and the spacecraft designed to perform them. A satellite is launched to a particular orbit, configured for a defined mission, and expected to operate with limited interaction with other spacecraft throughout its lifetime. When a mission requires an object to be inspected, relocated, serviced, refueled, upgraded, or removed, a new vehicle or a dedicated servicing mission is often required.

This approach has been effective for many traditional space missions, but the operating environment is changing. The number of spacecraft in orbit is increasing, commercial constellations are becoming more complex, and the value of maintaining and interacting with space assets after launch is becoming increasingly important. At the same time, capabilities such as autonomous rendezvous, proximity operations, docking, robotic manipulation, and in-orbit transportation are moving from experimental technologies toward practical operational capabilities. NASA and ESA are both investing in these capabilities as part of the broader development of in-space servicing, assembly, manufacturing, and logistics.

This creates an opportunity to reconsider how orbital vehicles are designed. Instead of developing a separate spacecraft around every individual mission, a reusable vehicle could be designed from the beginning to perform multiple missions over its operational lifetime. Such a vehicle would not be defined primarily by a single payload or destination, but by its ability to repeatedly enter the orbital environment, maneuver between locations, interact with other space assets, perform useful work, and support subsequent missions.

This whitepaper explores that concept through the proposed architecture of the **Reusable Orbital Drone (ROD)**. The concept treats the orbital vehicle as reusable operational infrastructure rather than as a single-use or mission-specific spacecraft. The objective is to examine the technical and operational principles required to make such a vehicle practical and to establish a framework for understanding its potential role in post-launch space operations.

1.1 The Post-Launch Problem

Reaching orbit is only the beginning of a spacecraft's operational life. Once a vehicle has been deployed, its ability to change its orbital state, interact with other spacecraft, or receive physical assistance is generally constrained by the capabilities that were designed into it before launch. A spacecraft that was not designed for servicing, docking, or relocation may have few practical options when its mission requirements change or when an unexpected problem occurs.

This creates a gap between **access to orbit** and **operations in orbit**.

The launch industry has made substantial progress in reducing the cost and increasing the frequency of access to space. However, putting an object into orbit does not automatically provide a way to move, maintain, inspect, upgrade, recover, or otherwise interact with that object after deployment. Those activities require additional vehicles, propulsion, navigation, sensing, communications, robotic or mechanical interfaces, and increasingly autonomous operations.

The need for these capabilities is already recognized across the space sector. NASA's work in on-orbit servicing has demonstrated the technical requirements associated with autonomous rendezvous, relative navigation, robotic operations, and interaction with client spacecraft. ESA is similarly developing an in-space transportation and logistics ecosystem that includes rendezvous and docking, orbital transportation, refueling, servicing, and standardized interfaces.

The challenge is therefore not simply to develop a vehicle capable of reaching another spacecraft. The larger challenge is to make **post-launch operations repeatable, flexible, and economically useful**.

A conventional approach often treats these activities as individual missions. An inspection mission may require one spacecraft. An orbital relocation mission may require another. A servicing mission may require a vehicle specifically designed around the target spacecraft and its interface. A debris-removal mission may require yet another architecture. Although these vehicles can share technologies, their hardware is frequently optimized around a particular mission profile.

This creates several limitations.

First, the cost of developing and launching a dedicated vehicle can be difficult to justify for relatively small or infrequent orbital tasks. Second, mission-specific designs can limit flexibility when the operational requirement changes. Third, a vehicle that performs a single mission and is subsequently retired does not fully take advantage of the hardware, propulsion system, avionics, and other infrastructure that were required to place it in orbit.

The problem becomes more significant as the number and diversity of orbital assets increase. Future space operations are likely to involve not only the deployment of new spacecraft, but also the movement, maintenance, inspection, assembly, servicing, and eventual removal of existing assets. ESA has explicitly identified in-orbit transportation and logistics as an emerging ecosystem, including services such as payload transfer, life extension, maintenance and repair, end-of-life management, and in-space infrastructure.

This suggests a shift in how post-launch operations can be approached.

Rather than designing the vehicle around the question:

“What spacecraft is this mission going to launch?”

the design process can instead begin with:

“What work needs to be performed in orbit, and what reusable vehicle can perform that work repeatedly?”

That distinction is central to the concept explored in this whitepaper.

A reusable orbital vehicle designed around repeated operations could provide a persistent layer between launch systems and the spacecraft that operate in orbit. It could be deployed when required, travel to a target or destination orbit, perform its assigned task, and then continue to another mission or return for refurbishment and redeployment. The same underlying vehicle could therefore support different customers and mission types over its operational lifetime.

Such an architecture introduces its own engineering challenges. Repeated orbital operations require reliable propulsion, guidance and navigation, autonomous rendezvous and proximity operations, robust communications, fault management, thermal and structural design for reuse, and interfaces that allow different mission payloads or customers to be accommodated. NASA identifies rendezvous, proximity operations, docking, relative navigation, and safe spacecraft maneuvering as key technical capabilities for future in-space operations.

The **post-launch problem**, therefore, is not simply a lack of individual servicing technologies. Many of the required technologies already exist in various stages of development. The broader challenge is how these technologies can be integrated into a reusable, autonomous, multi-mission vehicle architecture that can provide useful services repeatedly.

The following sections examine this transition from mission-specific spacecraft toward reusable orbital infrastructure and introduce the Reusable Orbital Drone as one possible architecture for addressing it.

1.2 From Mission-Specific Spacecraft to Reusable Orbital Infrastructure

Traditional spacecraft are generally designed around a specific mission, payload, orbit, and operational lifetime. This works well when requirements are known in advance, but makes many post-launch activities dependent on dedicated vehicles or mission-specific designs.

An alternative is to separate the **vehicle from the mission**. A reusable platform can retain a common core while its payloads, tools, and operational objectives change between missions. This approach is already reflected in the broader development of in-space transportation, where ESA is pursuing reusable vehicles, rendezvous and docking, standardized interfaces, refueling, and orbital logistics.

This represents a shift from **mission-specific spacecraft toward reusable orbital infrastructure**: instead of developing a new vehicle for every orbital task, a reusable platform can repeatedly provide different services to assets already in orbit. The Reusable Orbital Drone concept explored in this paper is one proposed architecture for this model.

1.3 Objectives

The primary objective of this whitepaper is to establish a systems-level framework for **Reusable Orbital Drones (RODs)** as a potential class of multi-mission vehicles for post-launch space operations.

The paper aims to achieve the following objectives:

1. **Define the concept of a Reusable Orbital Drone.**
   Establish a clear definition and identify the characteristics that distinguish an orbital drone from conventional satellites, spacecraft, orbital transfer vehicles, servicing vehicles, and reusable spaceplanes.
2. **Establish a common systems architecture.**
   Identify the major subsystems and functional elements required for a reusable orbital drone, including propulsion, guidance, navigation and control, avionics, communications, power, thermal management, autonomy, and mission interfaces.
3. **Describe the operational lifecycle of an orbital drone.**
   Develop a representative mission flow covering deployment, orbital transit, rendezvous, proximity operations, mission execution, departure, recovery, refurbishment, and subsequent redeployment.
4. **Identify potential mission classes.**
   Examine how a common reusable vehicle could support different post-launch activities, including orbital transfer, inspection, hosted payload operations, emergency response, servicing, and future orbital sustainability applications.
5. **Examine the role of reusability and multi-mission operation.**
   Consider how designing a vehicle for repeated missions changes requirements for vehicle durability, modularity, maintenance, turnaround, interfaces, and operational planning.
6. **Identify key technical challenges.**
   Highlight the engineering challenges that must be addressed for reusable orbital drones to become practical, particularly in autonomous navigation, rendezvous and proximity operations, propulsion, fault management, safety, and repeated orbital operations.
7. **Establish a basis for further development and research.**
   Provide a systems-level foundation that can be used to guide future engineering studies,

technology development, simulations, demonstrations, and refinement of specific orbital-drone architectures.

The objective of this work is therefore not to present a single finalized vehicle design. Rather, it is to establish a structured way of thinking about reusable orbital drones and the technical requirements associated with making them useful, repeatable, and adaptable across multiple missions.

1.4 Scope

This whitepaper focuses on the **vehicle and mission architecture of reusable, uncrewed orbital drones intended to perform useful operations after reaching orbit**. The discussion is primarily concerned with Earth-orbiting applications and with the technologies and operational concepts required to support repeated missions.

The scope includes the following areas:

- Definition and classification of reusable orbital drones;
- System-level vehicle architecture;
- Propulsion and orbital maneuvering;
- Guidance, navigation and control;
- Autonomous rendezvous and proximity operations;
- Avionics, communications, power, and thermal systems;
- Mission payloads and standardized interfaces;
- Vehicle reusability, refurbishment, and turnaround;
- Multi-mission and modular vehicle operations;
- Orbital transfer, inspection, hosted payload, emergency response, and servicing missions;
- Operational safety and fault management;
- Fleet utilization and the transition toward reusable orbital infrastructure; and
- A preliminary development framework for demonstrating and maturing the required capabilities.

The whitepaper does not attempt to provide a detailed design for a specific flight vehicle, nor does it establish final subsystem specifications, detailed component selections, or flight-qualified requirements. Those decisions depend on vehicle configuration, mission profile, target orbit, payload, launch architecture, regulatory requirements, and subsequent engineering analysis.

The scope also deliberately distinguishes reusable orbital drones from adjacent vehicle classes. Reusable launch vehicles, conventional satellites, crewed spacecraft, dedicated orbital transfer vehicles, space tugs, and reentry vehicles are considered where they provide relevant technical or architectural context, but they are not themselves the primary subject of this work.

Similarly, this paper does not attempt to provide a complete treatment of orbital servicing, active debris removal, space manufacturing, or broader space-logistics markets. These applications are discussed primarily to demonstrate the types of missions that a reusable orbital-drone architecture could support.

The framework presented here is intended to remain **vehicle-agnostic at the architectural level**. Individual implementations may differ significantly in size, propulsion, payload capacity, recovery method, autonomy, orbital regime, and mission duration while still conforming to the broader principles described in this paper.

Orbit Keeper, developed by Graviron Aerospace, is referenced where useful as an example of a vehicle being developed within this architectural framework. The paper, however, is intended to describe the

broader concept of reusable orbital drones rather than serve as a detailed technical specification or product document for Orbit Keeper.

2. Defining the Reusable Orbital Drone

The Reusable Orbital Drone (ROD) is proposed in this whitepaper as a class of uncrewed orbital vehicles designed to perform repeated post-launch operations across multiple missions.

The concept combines capabilities that already exist independently across several spacecraft classes, orbital maneuvering, autonomous rendezvous, servicing, modular payloads, and reusability, around a common vehicle architecture. The key distinction is therefore not a single new subsystem, but the intended utilization model: one reusable vehicle supporting different orbital missions over its operational lifetime.

An ROD is designed around a reusable core containing systems such as structure, propulsion, power, avionics, communications, and guidance, navigation and control. Mission-specific capabilities are accommodated through configurable payloads, tools, and capture or docking interfaces.

The architecture is intentionally vehicle-agnostic. Individual implementations may vary in size, propulsion, payload capacity, orbital regime, mission duration, and recovery method while retaining the same underlying principles.

2.1 Definition

For the purposes of this whitepaper, a Reusable Orbital Drone (ROD) is defined as:

> *An uncrewed orbital vehicle designed to perform repeated post-launch missions using a common reusable vehicle platform, with modular mission capabilities and sufficient onboard autonomy or automation to safely execute its assigned operations in orbit.*

Five elements are central to this definition:

- Uncrewed: No human crew is required onboard; ground operators may supervise or command the vehicle.
- Orbital: The vehicle is designed to maneuver and operate within an orbital environment.
- Reusable: The same vehicle is intended to support multiple mission cycles rather than being retired after one mission.
- Multi-mission: The vehicle can support different operational tasks through configurable mission capabilities.
- Autonomous or automated: The vehicle can perform defined functions onboard, with the level of autonomy varying by mission phase.

Reusability in this context is more than surviving a previous flight. The vehicle must be designed around cumulative mission cycles, including subsystem life, propellant requirements, mechanism wear, health monitoring, and refurbishment.

The ROD is therefore best understood as an architectural class within the broader spacecraft category, rather than as a replacement for the term spacecraft.

2.2 Core Characteristics

The ROD architecture is characterized by the following properties:

1. Reusability
The vehicle is designed for repeated mission cycles, with its operational lifetime treated as a system-level design parameter.

2. Multi-mission capability
The same vehicle can support different mission classes rather than being optimized around a single orbital task.

3. Modularity
Mission-specific payloads, tools, and interfaces can be changed while retaining the common vehicle core.

4. Orbital mobility
The vehicle can perform orbital transfers, phasing, rendezvous, station-keeping, and proximity operations as required by its mission.

5. Autonomy and automation
Time-critical functions can be performed onboard, while mission planning and higher-level decisions may remain ground-supervised.

6. Interoperable interfaces
Defined mechanical, electrical, data, payload, and docking interfaces allow the vehicle to accommodate changing mission requirements.

7. Operational persistence
The vehicle is treated as an orbital asset that can be assigned to multiple missions throughout its useful life.

8. Safe spacecraft interaction
Where required, the vehicle must safely approach, inspect, capture, dock with, or otherwise interact with other orbital objects.

These characteristics are intended to work together. Reusability without multi-mission capability produces a reusable vehicle; mobility without servicing or mission flexibility produces a transfer vehicle. The ROD architecture combines these capabilities around repeated orbital utilization.

2.3 Orbital Drones vs. Existing Vehicle Classes

The ROD overlaps with several established vehicle classes. The distinction is primarily architectural rather than technological.

Vehicle class	Primary role	ROD distinction
Conventional satellite	Persistent operation of its own payload/mission	ROD is designed as a mobile service asset

Space tug / OTV	Move spacecraft or payloads between orbital locations	Transportation is one ROD mission class, not its sole purpose
Servicing spacecraft	Inspect, maintain, refuel, or manipulate another spacecraft	Servicing is one ROD mission class within a broader platform
Reusable spaceplane	Provide transportation to/from orbit and return	ROD emphasizes work performed within orbit
Reusable Orbital Drone	Repeated, multi-mission orbital operations	Common reusable platform + configurable mission capabilities

These categories are not mutually exclusive. A single vehicle can combine transportation, payload hosting, rendezvous, servicing, and reentry capabilities. ESA's Space Rider, for example, is explicitly described as a reusable uncrewed transportation system whose concept can also extend toward in-orbit services and close-proximity operations.

The purpose of the ROD classification is therefore not to create rigid boundaries between spacecraft. It identifies a particular design philosophy: a reusable orbital vehicle whose core remains in service while its missions and operational capabilities can change over time.

3. State of the Art

The technologies underlying reusable orbital drones are already being developed across several areas of the space industry. Autonomous rendezvous, orbital servicing, orbital transportation, reusable vehicles, and modular payload systems each provide pieces of the broader architecture.

The current state of the art can therefore be understood through four adjacent vehicle classes: reusable orbital vehicles, servicing vehicles, space tugs, and spaceplanes.

The purpose of this section is not to provide an exhaustive industry survey. It is to establish the relevant technical precedents and identify where the ROD architecture differs.

3.1 Existing Reusable Orbital Vehicles

The Boeing X-37B is the strongest operational example of a reusable uncrewed orbital vehicle. It has demonstrated repeated orbital missions, long-duration operation, autonomous reentry, and recovery for subsequent flights. Its primary role, however, is orbital experimentation and government missions rather than commercial multi-customer orbital services.

ESA's Space Rider represents another approach. It is being developed as an uncrewed reusable orbital transportation and laboratory platform capable of carrying different payloads, operating in LEO, and

returning payloads to Earth. ESA also identifies potential future applications involving close-proximity operations and in-orbit services.

Together, these systems demonstrate that reusable uncrewed orbital operation is technically achievable. The ROD concept differs primarily in making repeated orbital work and multi-mission utilization the central purpose of the vehicle.

3.2 Orbital Servicing Vehicles

Servicing vehicles provide much of the technical heritage required for ROD operations. Their capabilities include relative navigation, rendezvous, proximity operations, inspection, capture, docking, robotic manipulation, and spacecraft relocation.

NASA's OSAM-1 architecture, for example, incorporated autonomous real-time relative navigation, servicing avionics, robotic arms, and systems for capturing and refueling a client spacecraft. Although NASA ultimately cancelled OSAM-1 following technical, cost, and schedule challenges, the program illustrates both the capabilities and integration challenges involved in robotic servicing.

Commercial demonstrations such as Astroscale's ELSA-d have also advanced rendezvous and capture technologies. NASA continues to develop autonomous relative-navigation systems because final rendezvous and capture can require onboard decision-making that cannot be performed reliably through continuous ground control.

For the ROD architecture, servicing is therefore best viewed as one major mission class built on an established technical foundation, rather than the entire definition of the vehicle.

3.3 Space Tugs

Space tugs provide another important precedent: orbital mobility as a service.

Their primary role is to move payloads or spacecraft from one orbital location to another. ESA's current in-space transportation program explicitly identifies reusable space tugs, rendezvous and docking, refueling, onboard intelligence, and standardized interfaces as elements of a future orbital logistics ecosystem.

The ROD architecture extends beyond transportation. An ROD can use the mobility of a space tug while also supporting inspection, servicing, hosted payloads, emergency response, and other orbital operations.

The distinction can therefore be summarized as:

Space tug: mobility platform.
Reusable Orbital Drone: multi-mission orbital work platform.

The two architectures can coexist within the same future orbital logistics ecosystem.

3.4 Spaceplanes

Reusable spaceplanes demonstrate another important element of the ROD concept: repeated access to and return from orbit.

The X-37B provides extensive flight heritage in reusable autonomous orbital operations, while Space Rider is being developed as a reusable uncrewed transportation system for LEO.

The principal difference is the center of the mission. A spaceplane generally treats transportation to and from orbit as a central function. An ROD treats operations within orbit as the primary mission environment.

Spaceplanes therefore provide important technological heritage for RODs, particularly in reusability, autonomy, and recovery, but do not by themselves define the multi-mission orbital-service model.

3.5 The Unfilled Architectural Space

The current landscape demonstrates most of the individual capabilities required by an ROD:

- Reusable orbital vehicles → repeated flight and long-duration orbital operation;
- Servicing vehicles → rendezvous, capture, inspection, and physical interaction;
- Space tugs → orbital mobility;
- Spaceplanes → reusable access to and return from orbit.

What is less clearly established is the combination of these capabilities around a common reusable vehicle whose primary purpose is repeated, multi-mission work within the orbital environment.

The proposed architectural intersection can therefore be expressed as:

Reusability + Orbital Mobility + Proximity Operations + Modular Mission Capability + Repeated Utilization

This is the architectural space addressed by the Reusable Orbital Drone concept.

The distinction is not that existing vehicles cannot perform some of these functions. They increasingly can. ESA's current programs, for example, are explicitly moving toward reusable vehicles, standardized interfaces, refueling, onboard intelligence, and multi-mission orbital logistics.

The ROD concept instead organizes these capabilities around a different utilization model: Rather than building a spacecraft for every orbital task, operate a reusable vehicle capable of performing many of them.

4. Reusable Orbital Drone Architecture

A reusable orbital drone cannot be designed the way a conventional, mission-specific spacecraft is designed. A conventional spacecraft is optimized around a single payload, a single orbit, and a single operational lifetime. A reusable orbital drone must instead be optimized around a common vehicle infrastructure capable of supporting many different missions, payloads, and customers across repeated flights. This distinction shapes every subsystem decision described in this section.

The architecture proposed here is organized around a **shared core vehicle bus** , the structural, power, propulsion, avionics, and communications infrastructure that remains largely constant across missions , combined with **modular mission interfaces**, primarily the capture/docking interface and payload bay, which vary depending on the mission being performed. This separation is what allows a single reusable vehicle family to serve orbital transfer, inspection, hosted payload, and servicing missions without requiring a new spacecraft design for each mission type.

The subsections that follow describe each major subsystem group individually. It is important to note, however, that no subsystem in this architecture can be designed in isolation. Guidance and navigation outputs become propulsion commands. Power budgets constrain how often avionics and navigation sensors can operate. Thermal design is shaped as much by orbital regime and mission duration as by

onboard electronics. Communication availability determines how much autonomy a given mission phase actually requires. The interfaces between subsystems are, in many respects, as architecturally significant as the subsystems themselves.

4.1 System-Level Architecture

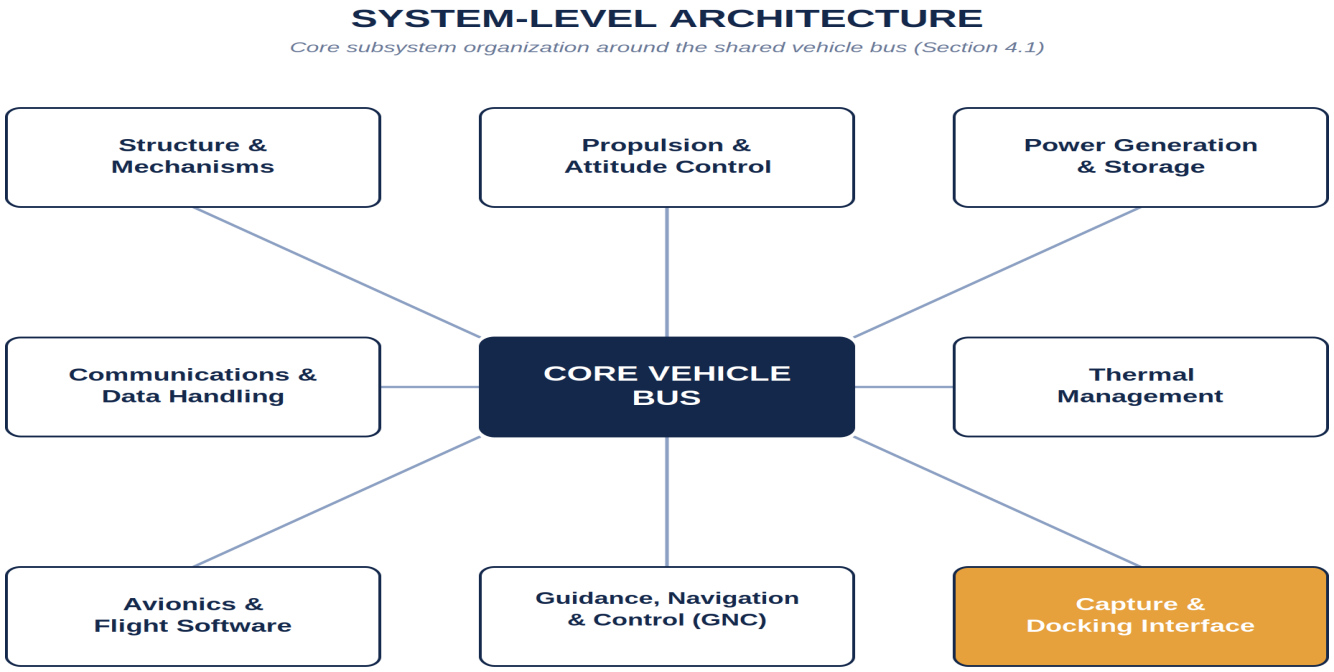

Figure 1 - Orbit Keeper conceptual system architecture. Subsystem groupings are illustrative of design scope, not final subsystem specifications.

The vehicle is organized into eight functional subsystem groups around a common structural bus: structure and mechanisms; propulsion and attitude control; power generation and storage; avionics and flight software; guidance, navigation and control; the capture and docking interface; communications and data handling; and thermal management.

Two architectural principles distinguish this arrangement from a conventional, single-use spacecraft bus.

First, subsystems that determine mission life , propellant capacity, mechanism actuation cycle life, battery cycle life , must be sized for **multiple missions**, not a single mission duration. A conventional spacecraft bus is typically sized to survive one mission plus margin. A reusable orbital drone bus must be sized to survive a defined number of complete mission cycles, each of which independently stresses propulsion, mechanisms, and power systems.

Second, the architecture deliberately concentrates variability in a small number of interfaces rather than distributing it across the whole vehicle. The core bus , structure, propulsion, power, avionics, and communications , is intended to remain functionally constant across mission types. The capture/docking interface and the mission payload bay are the primary points where a given mission's specific requirements are accommodated. This is not simply a cost-saving measure; it is what makes the reusability claim technically coherent. A vehicle that requires substantial bus-level redesign for

each new mission type is not meaningfully more reusable than a fleet of purpose-built spacecraft , it has simply moved the non-recurring engineering cost from "per vehicle" to "per mission." Subsystem sizing and budgeting throughout this section follows standard systems-engineering practice for spacecraft design [1].

4.2 Propulsion

The propulsion subsystem serves two distinct functions: primary orbital maneuvering (orbit raising, phasing, plane changes, and transfer between mission targets) and fine attitude control during proximity operations, typically provided by a separate reaction control system.

Propellant budgeting is the single most consequential constraint on the reusability claims made elsewhere in this paper. A propulsion system sized only for one mission's delta-v requirement does not produce a reusable vehicle; it produces a single-use vehicle that happens to be structurally capable of a second flight. The propulsion architecture must instead be sized against a cumulative delta-v budget spanning the target number of missions per vehicle, with appropriate margin for each mission's phasing, rendezvous, and disposal or repositioning requirements.

The selection between chemical, electric, and cold-gas or hybrid propulsion concepts remains an active design trade rather than a settled decision, and is driven primarily by the delta-v and response-time requirements of near-term LEO logistics and inspection missions. Each option carries different implications for reusability specifically:

- **Chemical propulsion** offers higher thrust and faster response, which is advantageous for time-critical maneuvers and proximity operations, but is comparatively propellant-inefficient , meaning multi-mission delta-v budgets require proportionally more propellant mass carried at launch.
- **Electric propulsion** offers substantially higher propellant efficiency (specific impulse), allowing more cumulative delta-v to be extracted from a given propellant mass , directly supporting a higher number of missions per vehicle , at the cost of low thrust and correspondingly longer transfer times, which affects mission cadence and time-to-rendezvous.
- **Hybrid or cold-gas approaches** may offer a middle path for specific mission phases, such as fine proximity-operations control, where precision matters more than raw efficiency.

This trade is not purely technical. It is directly coupled to the mission classes described in Section 6 and the reusability targets described in Section 7: a vehicle intended primarily for slower-cadence, high-delta-v missions favors electric propulsion, while a vehicle intended for rapid-response or time-critical servicing favors chemical propulsion, with corresponding effects on how many missions a single vehicle can realistically perform before propellant exhaustion becomes the limiting factor.

4.3 Guidance, Navigation & Control

The guidance, navigation, and control (GNC) subsystem provides relative navigation and maneuvering authority throughout the mission, but its requirements change substantially across mission phases, and the architecture treats it accordingly.

During long-range transit and initial rendezvous, relative navigation can rely on coarser techniques , GNSS-relative positioning, ground-based tracking and monitoring - with sufficient time margin for ground-supervised decision-making. As range to the target closes, navigation requirements shift toward higher-precision, higher-update-rate sensing: RF ranging at medium range, and optical or LIDAR-based relative pose estimation at close range, typically combined through sensor fusion to maintain a reliable navigation solution as lighting conditions, target geometry, and relative dynamics change.

This progression is matched by a corresponding shift in autonomy, discussed further in Section 4.8: ground-supervised operation is appropriate where communication latency does not meaningfully constrain safety, while proximity operations , where communication round-trip time may exceed the time available to respond to an anomaly , require the vehicle to carry its own fault detection, collision-avoidance, and abort logic onboard.

Collision-avoidance logic in particular must be pre-planned rather than computed live under fault conditions. Rendezvous and proximity-operations trajectories should be designed from the outset with defined keep-out zones and pre-committed retreat trajectories, so that a sensor fault, missed capture attempt, or loss of navigation confidence results in a known, safe vehicle behavior rather than a real-time decision made under degraded information.

GNC output is the primary driver of propulsion commands during proximity operations and is correspondingly one of the tightest-coupled interfaces in the overall architecture, requiring close coordination with both the propulsion subsystem (Section 4.2) and the avionics subsystem responsible for executing GNC computations in real time (Section 4.4).

4.4 Avionics & Flight Computing

The avionics subsystem provides onboard computing, fault detection and management, and telemetry handling, and functions as the vehicle's primary point of subsystem integration.

Computational demand on the avionics subsystem is not constant across the mission. During orbital transit, processing needs are largely dominated by housekeeping, health monitoring, and telemetry , relatively low-rate, low-urgency tasks. During proximity operations, the same avionics architecture must support real-time GNC computation, sensor fusion, and fault response at update rates and latencies that leave little margin for error. The avionics architecture must therefore be designed for its peak demand phase, not its average demand, even though that peak may represent a small fraction of total mission duration.

Fault detection and management is treated as an onboard-resident capability rather than a ground-supervised one for any mission phase where communication latency could compromise safety. This includes defined "safe mode" behaviors , degraded but stable vehicle states that can be entered automatically in response to sensor faults, navigation uncertainty, or loss of ground contact , allowing the vehicle to preserve itself and await further instruction rather than requiring immediate, correct autonomous decision-making under every possible fault condition.

Because the avionics subsystem interfaces with essentially every other subsystem , commanding propulsion, reading GNC sensor inputs, managing power allocation across subsystems, and scheduling communications windows , its architecture has an outsized influence on overall vehicle reliability and is treated as a primary integration point during system-level testing (see Section 9).

4.5 Power & Thermal

Power. The power subsystem consists of solar array generation combined with battery storage, and , consistent with the reusability principle established in Section 4.1 , is sized to support the vehicle's full multi-mission operational life rather than a single mission's duration. This is a meaningful departure from conventional single-use spacecraft power system design, where battery sizing need only account for one mission's charge/discharge history.

Battery degradation across repeated deep charge/discharge cycling and thermal cycling is a genuine limiting factor on vehicle lifetime, independent of and in addition to propellant availability (Section 4.2) and mechanism wear (Section 7). A full treatment of vehicle reusability must therefore account for

power-subsystem degradation alongside propulsion and mechanical wear, rather than treating propellant margin as the sole reusability constraint.

Thermal. Thermal management combines passive and active control, sized for the thermal cycling characteristic of the vehicle's operating orbit across an extended, multi-mission operational life. Thermal design in this architecture is driven as much by orbital regime and cumulative mission duration as by onboard electronics heat load , a vehicle intended for multiple missions across its lifetime experiences substantially more thermal cycling than a single-mission spacecraft of comparable electronics load, and the thermal subsystem must be designed against that cumulative exposure.

The power and thermal subsystems are tightly coupled: power budget directly constrains how much duty cycle is available to active thermal control, avionics, and GNC sensors, particularly during eclipse periods or extended proximity-operations phases where power generation may be limited by vehicle orientation.

4.6 Communications

The communications subsystem provides the ground link and manages onboard data handling across all mission phases, but its role changes significantly as autonomy increases during proximity operations (Section 4.8).

During long-range transit and ground-supervised rendezvous, reliable, low-latency communications availability is a functional requirement , ground-based decision-making depends on it. During proximity operations, however, the architecture is designed under the explicit assumption that communications may be **intermittent or unavailable** rather than continuously reliable, whether due to link geometry, target-vehicle interference, antenna pointing constraints during maneuvering, or simple range and timing limitations. This assumption is deliberate: designing proximity-operations safety logic (Section 4.3, Section 8) around the availability of a continuous ground link would create a single point of failure in exactly the mission phase where failure is least recoverable.

Onboard data management therefore includes not only telemetry formatting and downlink scheduling, but also local storage and prioritization of health and mission data during periods when ground contact is unavailable or degraded, ensuring that a complete mission record is recoverable even if real-time downlink is interrupted during the highest-risk mission phase.

4.7 Mission Payload & Interfaces

The mission payload and interface subsystem is the primary point of architectural modularity in the vehicle and is what allows a single reusable core bus to support the range of mission classes described in Section 6.

Two interface types are relevant. The **capture and docking interface** provides the mechanical means of rendezvous, berthing, and physical interaction with a target spacecraft, and is designed first for compatibility with cooperative targets , spacecraft built with a compatible docking interface , with capability against less-prepared or non-cooperative targets identified as a longer-dated, higher-technical-risk capability rather than a near-term claim. The **payload bay** provides a standardized mechanical, electrical, and data interface allowing different mission-specific payloads , inspection sensors, hosted-payload hardware, servicing tooling , to be integrated onto the same core vehicle without bus-level redesign.

The capture and docking interface represents the architecture's highest-risk and lowest-technology-readiness subsystem, and correspondingly the area of greatest engineering

investment and architectural differentiation. Specific mechanism design, actuation approach, and redundancy strategy are outside the scope of this paper. What is architecturally significant here is the design principle rather than the mechanism itself: by isolating mission-specific physical interaction into a single, well-defined interface, the remainder of the vehicle , propulsion, power, avionics, GNC , can remain common across mission types, which is the structural basis for the platform economics argument developed elsewhere in this paper.

Precedent for cooperative-target capture and docking already exists in the industry. Northrop Grumman's Mission Extension Vehicle (MEV) program successfully docked with a live Intelsat satellite in geosynchronous orbit in February 2020, and again with a second Intelsat satellite in 2021, demonstrating that remote, uncrewed docking with a prepared client spacecraft is an operationally proven capability rather than a purely theoretical one [2]. Astroscale's ELSA-d mission similarly demonstrated magnetic-capture rendezvous and docking with a cooperative client in low Earth orbit [3]. Most notably, DARPA's Robotic Servicing of Geosynchronous Satellites (RSGS) program , built on the same flight-proven MEV bus , launched in July 2026 and is expected to begin on-orbit servicing demonstrations in 2027 following transit to geosynchronous orbit [4], representing the most current operational validation of the capture-and-servicing architecture described in this section at the time of writing.These precedents also show that no single capture approach has become an industry standard: mechanical docking probes, pre-installed magnetic plates, and robotic-arm or electrostatic capture are all being pursued in parallel, each with different assumptions about target cooperation and pre-installed client hardware. Orbit Keeper's capture interface remains an open architectural trade for the same reason. The field itself has not yet converged on one method.

4.8 Autonomy

Autonomy in this architecture is not a single, binary capability but a staged progression that varies by mission phase, and it functions as the integrating theme across the guidance and navigation (4.3), avionics (4.4), and communications (4.6) subsystems described above.

The architecture targets **ground-supervised operation for long-range rendezvous**, where communication latency does not meaningfully constrain the time available for a human-in-the-loop decision, transitioning to **increasingly autonomous operation during proximity operations and capture**, where communication delay and the pace of relative-motion dynamics make ground-based intervention impractical within the time available to respond safely.

This staged approach reflects a deliberate risk-management choice rather than a technical limitation to be minimized as quickly as possible. Claiming full end-to-end autonomous operation without ground supervision , across all mission phases, on a vehicle without flight heritage , would overstate a capability that has not yet been demonstrated. Full autonomous operation is treated in this architecture as a longer-term capability goal, developed incrementally as flight experience accumulates, rather than a near-term operational claim.

Autonomy architecture at this stage is oriented primarily around reliable, well-bounded behaviors: defined fault responses, pre-planned abort and retreat trajectories (Section 4.3), and safe-mode fallback states (Section 4.4), rather than more general autonomous planning or decision-making capability. This is a conservative starting point by design , the priority in early operational missions is predictable, verifiable vehicle behavior under fault conditions, with more sophisticated autonomous capability introduced as a later-stage development goal once foundational behaviors have been validated in flight.

5. Mission Architecture

The architecture described in Section 4 exists to perform a repeatable sequence of mission phases, each of which places different demands on the subsystems described there. This section describes that sequence , from launch through mission execution to subsequent redeployment , and identifies where each phase's requirements connect back to the vehicle architecture.

It is worth stating explicitly at the outset that this mission flow is not evenly weighted. Some phases , orbital transit, for example , are comparatively well-understood, low-risk, and draw heavily on existing flight heritage across the industry. Other phases, particularly rendezvous, proximity operations, and capture, are where the majority of this architecture's technical risk and differentiation are concentrated. The subsections below are organized sequentially, but the level of architectural attention given to each is intentionally uneven, and Section 5.4 in particular reflects that emphasis.

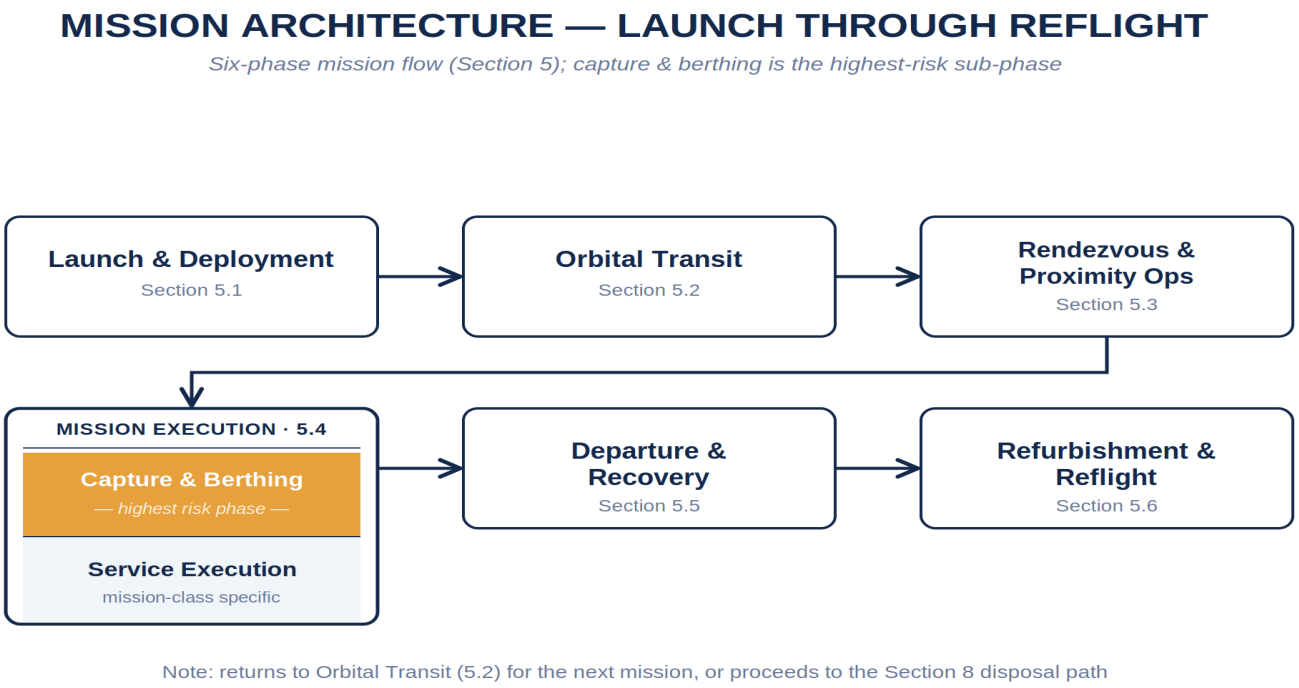

Figure 2 - Illustrative mission flow, launch through reflight. Exact phase durations and delta-v budgets are mission- and orbit-dependent.

5.1 Launch & Deployment

The mission begins with launch, either as a rideshare payload or via dedicated launch, to an initial parking orbit. Deployment from the launch vehicle marks the transition from a passive launch payload to an operationally active spacecraft: initial telemetry establishment, subsystem checkout, and verification that the vehicle survived launch environments , vibration, acoustic, and thermal loading , within expected parameters.

This phase is procedurally similar to deployment for any conventional spacecraft. Its significance for a reusable orbital drone is primarily that it represents the first of what will be repeated deployment-equivalent events across the vehicle's operational life , each subsequent mission begins from an analogous state, whether that state is reached via fresh launch or via repositioning from a prior mission (Section 5.6).

5.2 Orbital Transit

Following deployment or the conclusion of a prior mission, the vehicle executes the orbital maneuvers required to reach its next target: orbit raising or lowering, plane changes, and phasing maneuvers to align with the target's orbital position and timing.

This phase consumes the largest single share of the mission's propellant budget (Section 4.2) under most mission profiles and, depending on the propulsion concept selected, may extend from hours to days or weeks , chemical propulsion enabling faster transit at higher propellant cost, electric propulsion trading transit time for propellant efficiency. Because the time margin during transit is comparatively large relative to proximity operations, this phase operates under ground-supervised guidance and navigation (Section 4.3, Section 4.8): communication latency does not meaningfully constrain decision-making at this range, and ground-based mission planning retains primary authority over trajectory decisions.

5.3 Rendezvous & Proximity Operations

As the vehicle approaches its target, mission operations transition from orbital transit to rendezvous. Long-range rendezvous continues under ground-supervised navigation, typically relying on GNSS-relative positioning and ground-based tracking to close the remaining distance to the target's vicinity.

As range continues to close, the mission enters proximity operations, and both navigation and autonomy requirements shift accordingly. Navigation transitions toward higher-precision, higher-update-rate sensing , RF ranging at medium range, optical or LIDAR-based relative pose estimation at close range , consistent with the sensor progression described in Section 4.3. Autonomy shifts correspondingly toward onboard-resident decision-making, since communication round-trip time at this stage may exceed the time available to safely respond to an anomaly using ground intervention alone.

Proximity operations is also the phase in which the architecture's safety design assumptions (Section 4.6, Section 8) become operationally active: communications availability is treated as potentially intermittent, and the vehicle is expected to carry pre-planned, onboard-resident abort and retreat logic rather than depending on a continuous ground link to remain safe.

5.4 Mission Execution

Mission execution comprises two distinct sub-phases , capture/berthing and the mission-specific task itself , which are treated separately here because they carry substantially different risk profiles and architectural implications.

Capture and berthing. This is the phase in which the vehicle physically engages the target through the capture and docking interface described in Section 4.7. It is, without qualification, the highest-risk and highest-differentiation phase of the entire mission architecture. Unlike the preceding phases, which draw on rendezvous and proximity-operations techniques with meaningful precedent elsewhere in the industry , precedent that now includes DARPA's RSGS Mission Robotic Vehicle, launched July 2026 on the same flight-proven bus used by Northrop Grumman's MEV-1 and MEV-2 docking missions [2][4] , capture and berthing depends on the vehicle's own capture mechanism , the architecture's lowest-technology-readiness subsystem and its primary point of engineering differentiation.

For this reason, capture and berthing should not be treated as one phase among several in an evenly-weighted mission sequence. It warrants its own dedicated failure-mode analysis, its own abort

criteria distinct from the general proximity-operations abort logic described in Section 5.3, and its own redundancy considerations specific to the capture mechanism and its actuation sequence. A bounded number of capture attempts, clearly defined criteria for aborting to a safe standoff distance, and pre-committed retreat trajectories specific to whatever relative geometry exists at the moment of an aborted attempt are treated as baseline requirements for this phase rather than optional refinements.

Mission-specific execution. Once berthed or otherwise engaged with the target, the vehicle performs the task specific to the mission class being flown , inspection, payload transfer, servicing, refueling, or another operation as described in Section 6. Unlike capture and berthing, which is common architecture across all mission classes, this sub-phase is where mission-class-specific requirements are actually exercised, and its duration and complexity vary substantially depending on mission type: an inspection pass may require only minutes of sensor data collection, while a servicing operation may require an extended, multi-step sequence of mechanical or fluid-transfer operations.

5.5 Departure & Recovery

Following mission execution, the vehicle disengages from the target and repositions , either toward its next assigned mission or to a safe standoff orbit pending further tasking. The term "recovery" here refers to the vehicle's return to a stable, healthy operational state following mission execution, not to physical return to Earth; the vehicle remains on-orbit throughout its operational life, and departure is better understood as the transition back to the orbital-transit phase (Section 5.2) for a subsequent mission than as a terminal event.

Departure includes a post-execution health assessment: verification that the capture interface, propulsion system, and other subsystems stressed during mission execution remain within operational parameters before the vehicle is committed to its next assignment. This assessment feeds directly into the turnaround and reflight considerations described in Section 5.6 and Section 7.

5.6 Refurbishment & Reflight

Refurbishment, for a reusable orbital drone, does not resemble refurbishment of a reusable launch vehicle. The vehicle does not return to Earth between missions, and ground-based inspection, part replacement, or physical rework are not available options. Refurbishment in this architecture instead refers primarily to two activities: replenishment of consumables , principally propellant, potentially via a refueling rendezvous with a depot or tanker vehicle rather than return to a ground facility , and assessment of subsystem health via onboard telemetry and diagnostic data accumulated across the preceding mission.

It is worth noting explicitly that on-orbit refueling remains a technically difficult, unproven capability at scale. NASA's own attempt to demonstrate autonomous robotic refueling of an unprepared client satellite, the OSAM-1 mission (formerly Restore-L), was cancelled in 2024 after nearly a decade of development, citing continued technical, cost, and schedule challenges [5]. This is not cited as a precedent to build on, but as a cautionary data point: refueling architecture for this vehicle should be treated as a longer-term capability to be validated incrementally, consistent with the conservative reusability framing developed in Section 7, rather than assumed as a near-term operational baseline.

This distinction matters architecturally. Because the vehicle cannot be physically inspected or repaired between flights, confidence in its readiness for a subsequent mission depends entirely on the fidelity of its onboard health-monitoring capability (Section 4.4) and on subsystems being designed, from the outset, for a known and bounded number of reuse cycles rather than for reuse in general. The number of missions a given vehicle can support is therefore not a fixed architectural constant but

a function of accumulated wear against the specific limiting subsystems discussed in Section 7 , principally propellant margin and capture-mechanism actuation cycle life.

Where a vehicle's health assessment indicates it remains within operational parameters, it proceeds to a subsequent mission via orbital transit (Section 5.2), closing the mission loop described in this section. Where a vehicle approaches or exceeds its design reuse limits, it instead proceeds to the end-of-life disposal path described in Section 8 , passivation and controlled deorbit or repositioning to a disposal orbit , rather than continued reflight.

6. Mission Classes

The vehicle architecture described in Section 4 and the mission flow described in Section 5 are, by design, common across mission types. What actually varies between one mission and the next is concentrated in two places: the mission-specific execution sub-phase (Section 5.4) and the payload or interface configuration required to support it (Section 4.7). This section describes six representative mission classes that a reusable orbital drone architecture is intended to support, illustrating how that shared architecture is applied differently depending on the task.

These six classes are not exhaustive, and a given vehicle implementation may support only a subset of them, particularly early in its operational life. They are presented here to demonstrate the range of post-launch operations a single common platform can plausibly address without bus-level redesign, and to make explicit which subsystems each mission class stresses most heavily.

Mission	Description
6.1 Orbital Transfer	Repositioning satellites between orbits - orbit raising/lowering,plane changes, and relocation of client spacecraft.
6.2 Inspection	Close range visual and sensor-based assessment of a target's structural,thermal and operational health without physical contact
6.3 Hosted Payloads	Carrying third-party sensors or experiments to orbit to validate hardware ahead of a dedicated mission.
6.4 Emergency Response	Contingency operations for disabled or anomalous spacecraft including rapid-response rendezvous support.
6.5 Servicing	Physical interaction with target vehicle refueling,component replacement and life extension tasks requiring capture/berthing.
6.6 Recovery & Sustainability	End of life operation including deorbiting, re-orbiting to disposal orbits and active debris removal.

Table 1 - Mission classes overview.Description are illustrative; a given implementation may support a subset of these classes.

6.1 Orbital Transfer

Orbital transfer involves rendezvous and capture of a client spacecraft (Section 5.3, Section 5.4), followed by use of the vehicle's own propulsion (Section 4.2) to reposition the combined stack , orbit raising, lowering, plane change, or relocation to a new operational orbit , before release.

This mission class is architecturally distinct from inspection or hosted-payload missions in that it requires sustained mechanical engagement with the target throughout the maneuvering sequence, not merely brief contact. The propulsion subsystem must be sized to account for the added mass of the client spacecraft during the transfer maneuver, and the capture interface must maintain a stable mechanical connection under the loads imposed by active maneuvering rather than a static, unpowered dock.

6.2 Inspection

Inspection missions involve close-range visual and sensor-based assessment of a target's structural, thermal, or operational condition without physical contact. The vehicle performs proximity operations (Section 5.3) , potentially including a controlled flyby or circumnavigation of the target to obtain sensor coverage from multiple aspects , and mission execution (Section 5.4) consists of sensor data collection rather than any mechanical interaction.

Because no capture or berthing occurs, inspection missions draw primarily on the guidance, navigation and control subsystem (Section 4.3) and avoid engaging the vehicle's highest-risk subsystem entirely. This makes inspection architecturally the lowest-risk mission class among those described here, and a natural candidate for early operational missions ahead of capture-dependent mission types.

6.3 Hosted Payloads

Hosted-payload missions differ structurally from the mission classes above in that they do not necessarily involve rendezvous with another spacecraft at all. The vehicle instead carries a third-party payload , a sensor, experiment, or technology demonstration , in its payload bay (Section 4.7) and operates it on behalf of a customer, drawing on the power, thermal, and communications subsystems (Section 4.5, Section 4.6) to support the payload rather than on capture or close-proximity navigation capability.

This is architecturally the simplest mission class described in this section, since it may require neither proximity operations nor the capture and docking interface. It is correspondingly well suited to early operational missions, allowing flight experience and reliability data to accumulate on the vehicle's core bus before capture-dependent mission classes are attempted.

6.4 Emergency Response

Emergency response covers contingency operations in support of a disabled, anomalous, or otherwise degraded spacecraft , for example, a target that has lost attitude control or is tumbling. This mission class draws on the same rendezvous, proximity-operations, and capture architecture as orbital transfer and servicing missions (Section 5.3, Section 5.4), but under materially harder conditions: the target's relative motion may be unpredictable, and the timeline available to plan and execute a safe approach may be compressed relative to a routine, cooperative servicing mission.

Emergency response places correspondingly higher demands on the fault management and collision-avoidance logic described in Section 4.3, and on the vehicle's onboard autonomy (Section 4.8), since the assumptions that hold for a stable, cooperative target , a fixed docking interface, predictable relative motion , may not hold here. This mission class also has implications for fleet operations more broadly: providing a credible emergency-response capability requires a vehicle to be

available for rapid tasking rather than fully committed to a scheduled mission queue, a consideration addressed further in Section 7.4.

6.5 Servicing

Servicing is the most architecturally demanding mission class described in this section. It requires full capture and berthing (Section 5.4) followed by an extended period of physical interaction with the target , refueling, component replacement, or other life-extension tasks , while the two vehicles remain mechanically joined.

Where inspection missions avoid the capture interface entirely and orbital transfer missions engage it only for the duration of a maneuver, servicing missions exercise the capture and docking interface for the longest duration and under the widest range of operational conditions, including whatever mechanical or fluid-transfer operations the specific servicing task requires. Servicing missions are correspondingly the mission class most directly limited by the capture-mechanism reusability considerations discussed in Section 7 , each servicing mission represents a disproportionate share of the interface's cumulative actuation-cycle wear relative to other mission classes.

6.6 Recovery & Sustainability

This mission class covers end-of-life operations: deorbiting, re-orbiting a defunct object to a disposal orbit, and active debris removal. It is grouped separately from servicing because its targets are frequently uncooperative , a defunct satellite or debris object was not designed with a compatible docking interface and cannot assist in its own capture.

As noted in Section 4.7, capture of non-cooperative targets is a longer-dated, higher-technical-risk capability relative to capture of cooperative, docking-interface-equipped targets, and recovery and sustainability missions are the primary mission class where that capability gap is operationally relevant. This mission class also connects directly to the end-of-life and passivation considerations described in Section 8: a vehicle performing debris removal or deorbit operations must itself eventually be disposed of safely, and the disposal path for the target and the vehicle's own eventual disposal draw on the same underlying safety principles.

7. Reusability & Orbital Utilization

Section 4.1 established that a reusable orbital drone must be architected around multiple mission cycles rather than a single mission duration. This section makes that principle concrete: what actually determines how many missions a given vehicle can perform, what "turnaround" means for a vehicle that never returns to Earth, and how reusability is honestly represented before it has been flight-proven.

The framing throughout this section is deliberately conservative. Reusability is a genuine, meaningful technical and economic advantage over expendable, single-use architectures , but it is a bounded advantage, limited by specific, identifiable subsystems, not an open-ended capability. This section names those limits explicitly rather than treating reusability as a general property of the vehicle.

7.1 Single-Mission vs. Multi-Mission Architecture

A single-mission spacecraft is sized to survive one mission plus margin; every subsystem's life requirement is defined once and does not need to be revisited. A multi-mission architecture changes this calculus for a specific subset of subsystems , principally propulsion, the capture and docking interface, and power storage (Section 4.2, Section 4.5, Section 4.7) , whose life-limiting parameters

(propellant mass, mechanism actuation cycles, battery charge/discharge cycles) must instead be sized against a target number of complete mission cycles.

This is the technical basis for the platform-economics argument introduced in Section 4.1: reusability only produces a genuine cost advantage if the non-recurring engineering and manufacturing cost of the vehicle is amortized across more than one mission. A vehicle that is structurally capable of a second flight but not actually architected , propellant, mechanisms, batteries , to reliably complete it does not realize this advantage; it merely carries the mass and complexity of reusability without the economic benefit.

7.2 Turnaround

Turnaround describes the period between the conclusion of one mission (Section 5.5, Departure & Recovery) and the start of the next (Section 5.2, Orbital Transit). For a reusable orbital drone, turnaround does not involve any of the activities typically associated with reusable vehicle refurbishment on Earth , physical inspection, part replacement, or ground-based rework are not available, since the vehicle remains on-orbit throughout its operational life.

Turnaround instead consists of two activities: replenishment of consumables, principally propellant, and a diagnostic health assessment based on onboard telemetry accumulated during the preceding mission (Section 4.4). The fidelity of this assessment is the primary factor determining how quickly and how confidently a vehicle can be committed to its next mission; a vehicle with limited onboard diagnostic capability must either accept greater uncertainty about its own condition or apply more conservative operating margins to compensate.

It is worth noting, by way of an imperfect but useful analogy, that even reusable vehicles with full ground access are not maintenance-free between flights. SpaceX's Falcon 9 booster, despite returning to Earth and undergoing physical inspection between launches, still requires genuine refurbishment work rather than simply refueling and reflying. A reusable orbital drone, which has no ground access at all, should be expected to require at least as much operational conservatism between missions; turnaround is not a zero-maintenance event simply because no physical servicing is possible.

7.3 Vehicle Lifetime

Three subsystems independently limit how many missions a given vehicle can perform, and the vehicle's actual operational lifetime is set by whichever of the three is exhausted first:

- **Propellant margin** (Section 4.2) , the cumulative delta-v budget consumed across repeated orbital transit, rendezvous, and maneuvering phases.
- **Capture-mechanism actuation cycle life** (Section 4.7) , the docking and capture interface undergoes the most repeated, high-stress mechanical cycling of any subsystem per mission, and is likely the tightest constraint on vehicle lifetime among the three.
- **Battery cycle life** (Section 4.5) , degradation across repeated deep charge/discharge and thermal cycling, which limits vehicle life independent of propellant or mechanism wear.

A reusability target of 2–3 missions per vehicle, referenced elsewhere in this paper, should be understood explicitly as a design target derived from these three constraints, not as a number that has been demonstrated in flight. Validating it requires two distinct steps: ground-based life-cycle testing of the capture mechanism and other actuated components , following established mechanism verification practice such as ECSS-E-ST-33-01, the space engineering standard governing mechanism design, life testing, and verification [6] , ahead of flight, and subsequent confirmation and refinement against real degradation data once the vehicle has actually flown multiple missions.

Until that flight data exists, this paper treats the 2–3 mission figure as a bounded, propellant- and mechanism-limited estimate rather than a proven capability, consistent with the conservative posture adopted throughout this paper regarding claims that have not yet been flight-validated.

7.4 Fleet Utilization

At the level of a fleet rather than a single vehicle, reusability introduces a scheduling and retirement question that does not exist for expendable spacecraft: whether individual vehicles are retired on a fixed schedule, or on a condition basis determined by measured wear against the limiting subsystems described in Section 7.3.

A fixed-schedule retirement policy is operationally simpler but economically conservative , it necessarily assumes worst-case wear accumulation for every vehicle in the fleet, regardless of the specific missions any individual vehicle actually performed. A condition-based approach, informed by the onboard health-monitoring data discussed in Section 7.2, allows each vehicle's remaining operational life to be assessed individually, extracting more total service life from the fleet at the cost of requiring higher confidence in onboard diagnostic data.

Fleet utilization also has a direct connection to mission-class demand. As discussed in Section 6.4, a credible emergency-response capability requires that at least some fleet capacity be held available for rapid tasking rather than fully committed to a scheduled mission queue. This is fundamentally a fleet-level reusability question: it requires knowing, with reasonable confidence, which vehicles in the fleet have sufficient remaining propellant and mechanism life in reserve to be tasked on short notice, which is only possible under a condition-based view of fleet-wide vehicle lifetime.

8. Autonomy & Safety

Orbit Keeper is designed around **bounded autonomy**: safety-critical behaviors such as fault detection, collision avoidance, aborts, and safe-mode transitions are performed onboard, particularly during proximity operations where continuous ground intervention may not be practical. Full end-to-end autonomy remains a longer-term development objective.

During rendezvous and capture, **keep-out zones, approach corridors, and retreat trajectories are defined before proximity operations begin**. Communication loss is treated as a possible operating condition, with predefined safe responses. Capture attempts are bounded, with mandatory aborts after predefined failure conditions rather than unlimited retries.

Safety extends through end-of-life. Vehicles will be **passivated and disposed of through controlled reentry or an appropriate disposal orbit**, consistent with applicable orbital-debris mitigation requirements. NASA-STD-8719.14 provides technical requirements for limiting orbital debris, while ISO 24113 establishes international debris-mitigation requirements for unmanned space systems.

The vehicle designed to support orbital sustainability must itself remain sustainable throughout its operational life.

9. Development Framework

Development of a Reusable Orbital Drone requires progressive validation of the vehicle, its subsystems, and the operational processes required for repeated missions. The objective is not to demonstrate the complete mission architecture in a single step, but to progressively reduce technical risk as increasingly representative environments and mission scenarios are introduced.

A representative development framework is:

Phase 1 - Ground Validation
Validate structural systems, propulsion, avionics, power, communications, software, interfaces, and mission operations through laboratory testing and integrated ground tests.

Phase 2 - Atmospheric and Suborbital Testing
Where required, validate relevant guidance, navigation, recovery, communications, and environmental-response capabilities in increasingly representative flight conditions.

Phase 3 - Orbital Demonstration
Demonstrate the core vehicle architecture in orbit, including orbital maneuvering, communications, navigation, and basic mission operations.

Phase 4 - Proximity Operations Demonstration
Progressively introduce rendezvous, relative navigation, proximity operations, inspection, and eventually capture or docking. These capabilities should be demonstrated under controlled conditions before being applied to more complex targets.

Phase 5 - Multi-Mission Operations
Demonstrate the ability of the same vehicle architecture to support different mission configurations and repeated mission cycles, including refurbishment and redeployment.

This progressive approach allows individual technologies and interfaces to mature before they become dependencies for higher-risk missions. ESA's current InSPoC programme follows a similar philosophy, developing and testing rendezvous, docking, refilling, onboard intelligence, and logistics capabilities through incremental stages leading to in-orbit demonstrations.

The final objective is not simply a successful first flight, but a vehicle architecture capable of transitioning from **technology demonstration to repeatable operational use**.

10. Technical Challenges

The Reusable Orbital Drone architecture combines several technically demanding capabilities that are individually complex and become more challenging when integrated into a reusable, multi-mission system.

The principal challenges include:

Rendezvous and Proximity Operations.
Safe navigation around another spacecraft requires accurate relative-state estimation, controlled approach trajectories, collision avoidance, and reliable abort behavior.

Capture and Docking.
Physical interaction introduces uncertainties in relative motion, contact dynamics, target condition, and interface loads. These challenges become particularly significant when interacting with uncooperative or non-functional spacecraft. ESA's current servicing work identifies capture strategy, approach corridors, safety and abort logic, and contact-dynamics modelling as key areas requiring further maturation.

Autonomy and Fault Management.
Increasing onboard autonomy requires reliable sensing, decision-making, fault detection, isolation, and recovery. The system must also remain capable of transitioning to safe states when mission conditions fall outside expected parameters.

Vehicle Reusability.
Repeated missions introduce cumulative wear and degradation across propulsion systems, mechanisms, batteries, thermal protection, structures, and other components. Demonstrating useful vehicle life requires ground life-cycle testing followed by flight experience.

Interfaces and Modularity.
A multi-mission vehicle depends on interfaces that allow different payloads, tools, and client spacecraft to interact with the vehicle safely and predictably. Standardized interfaces are increasingly recognized as an important enabler for interoperable in-space transportation and servicing systems.

Orbital Safety and Operations.
As orbital activity increases, reusable vehicles must operate within increasingly complex traffic environments while managing conjunctions, debris, communication constraints, and end-of-life requirements.

These challenges do not represent independent engineering problems. They are closely coupled: for example, greater mission flexibility increases autonomy and interface requirements, while repeated operations increase the importance of reliability and health monitoring. The development process must therefore evaluate the vehicle as an integrated system rather than optimizing individual subsystems in isolation. This systems-level, iterative approach is consistent with NASA's systems-engineering framework for developing and operating complex systems across their full lifecycle.

11. Future Applications

The Reusable Orbital Drone architecture is not limited to the initial mission classes described in this paper. As vehicle capability, autonomy, interfaces, and orbital infrastructure mature, the same underlying architecture could support a broader range of activities.

Potential future applications include:

- **Orbital logistics:** transportation and repositioning of spacecraft, payloads, and components between orbital locations.
- **Satellite life extension:** inspection, maintenance, relocation, and other operations intended to extend the useful life of orbital assets.
- **Orbital infrastructure:** deployment, maintenance, and inspection of larger infrastructure that cannot be economically treated as a single self-contained spacecraft.
- **In-space assembly and manufacturing:** transporting tools, components, and payloads to support construction and manufacturing activities in orbit.
- **Orbital sustainability:** end-of-life management, debris recovery, and other activities aimed at reducing the long-term accumulation of objects in orbit.
- **Responsive space operations:** rapidly deploying a vehicle to inspect, assess, or assist an asset when an unexpected event occurs.

These applications are consistent with the broader direction of the space industry. NASA identifies servicing, assembly, manufacturing, maintenance, and infrastructure development as capabilities that can increase the performance and lifetime of space systems, while ESA is developing technologies for orbital transportation, refilling, servicing, logistics, and fleet operations.

Over time, individual vehicles could also become part of larger **orbital service networks**, where multiple reusable vehicles are scheduled according to mission requirements rather than being assigned permanently to a single customer or spacecraft. Such a model would shift orbital operations further toward infrastructure and fleet utilization.

The practical scope of these applications will ultimately depend on vehicle capability, mission economics, regulatory frameworks, standardized interfaces, and the development of supporting orbital infrastructure.

12. Conclusion

Space operations are moving beyond a model in which spacecraft are launched, deployed, and expected to operate largely on their own. Growing orbital activity and the increasing value of existing space assets create a need for capabilities that can move, inspect, maintain, service, and otherwise interact with those assets after launch. NASA and ESA are both developing technologies and programs aimed at making such in-space operations increasingly routine.

This whitepaper has proposed the **Reusable Orbital Drone** as an architectural approach to that problem: an uncrewed orbital vehicle built around a reusable core, configurable mission capabilities, orbital mobility, and sufficient autonomy to perform repeated operations across multiple missions.

The concept does not depend on a single new technology. Its significance comes from integrating several capabilities, reusability, maneuverability, rendezvous, proximity operations, modularity, autonomy, and mission flexibility around a common vehicle intended for repeated use.

The central proposition is therefore simple: **Orbital vehicles can be designed not only to operate in space, but to repeatedly perform work in space.**

Turning that proposition into an operational capability will require further development in propulsion, autonomous navigation, capture and docking, vehicle reusability, interfaces, safety, and fleet operations. The architecture presented in this paper provides a framework for pursuing that development.

Orbit Keeper, currently under development by Graviron Aerospace, represents one implementation of this broader concept. Its development and future demonstrations can provide an opportunity to validate, refine, and expand the principles described in this paper.

The long-term objective is not simply to build a reusable spacecraft. It is to establish a **reusable layer of orbital infrastructure** through which post-launch space operations can become more flexible, repeatable, and accessible.

References

[1] Wertz, J. R., Larson, W. J., and D'Souza, B. (eds.). *Space Mission Analysis and Design*, 3rd ed. Microcosm Press, El Segundo, CA, 1999.

[2] Northrop Grumman SpaceLogistics. ***Mission Extension Vehicle (MEV) Program.*** Northrop Grumman.
MEV-1 successfully docked with Intelsat IS-901 in February 2020, followed by MEV-2 docking with Intelsat IS-1002 in April 2021.

[3] Astroscale. ***ELSA-d Mission: End-of-Life Services by Astroscale*** *– Demonstration*. Astroscale, 2021–2024.
ELSA-d demonstrated rendezvous, relative navigation, and repeated magnetic capture of a cooperative client spacecraft.

[4] Defense Advanced Research Projects Agency (DARPA). *Robotic Servicing of Geosynchronous Satellites (RSGS)*. DARPA, 2026.

RSGS integrates DARPA's robotic servicing payload with SpaceLogistics' Mission Robotic Vehicle for GEO servicing operations.

[5] National Aeronautics and Space Administration (NASA). *On-orbit Servicing, Assembly, and Manufacturing 1 (OSAM-1)*. NASA. NASA's OSAM-1 program developed technologies for autonomous rendezvous, robotic servicing, refueling, and in-space assembly before being discontinued following technical, cost, and schedule challenges.

[6] European Cooperation for Space Standardization (ECSS). *ECSS-E-ST-33-01C Rev. 2: Space Engineering, Mechanisms*. ECSS, 2019. Provides requirements and design guidance for spacecraft mechanisms, including design, sizing, interfaces, and verification.

[7] Boeing. *X-37B*. Boeing.
Technical and operational information on the X-37B uncrewed autonomous reusable spaceplane, including its repeated orbital missions, autonomous reentry, and recovery.

[8] European Space Agency (ESA). *Space Rider Programme*. ESA Space Transportation.
Technical and operational information on Space Rider as an uncrewed reusable orbital transportation and laboratory platform, including its potential extension toward in-orbit services and close-proximity operations.

[9] European Space Agency (ESA). *Building the European Space Logistics Ecosystem for In-Space Transportation*. ESA Commercialisation Gateway. Describes ESA's development of reusable space tugs, rendezvous and docking, refueling, standardized interfaces, and broader in-space logistics infrastructure.

[10] National Aeronautics and Space Administration (NASA). *NASA Systems Engineering Handbook*, NASA/SP-2016-6105 Rev. 2. NASA, 2017. Systems-engineering framework referenced for lifecycle development, system design, integration, verification, and technical management.

[11] National Aeronautics and Space Administration (NASA). *NASA-STD-8719.14C: Process for Limiting Orbital Debris*. NASA, 2019. Technical requirements and assessment methods for limiting orbital-debris generation, collision risk, and end-of-mission risks.

[12] International Organization for Standardization (ISO). *ISO 24113:2023, Space Systems: Space Debris Mitigation Requirements*. ISO, 2023. International requirements addressing debris mitigation for unmanned systems operating in or passing through near-Earth space.

Published by

The Team,
Graviron Aerospace, Inc.
Beyond Launch