AIAA SPACE 2010 Conference & Exposition
30 August - 2 September 2010, Anaheim, California

AIAA 2010-8638

Evolving to a Depot-Based Space Transportation Architecture

Frank Zegler and Bernard Kutter
United Launch Alliance (ULA)
9100 E. Mineral Avenue
Centennial, CO 80112

Abstract

Substantial amounts of work have been expended in the past years evaluating the benefits of propellant depots for Lunar exploration. The advantages of these depots are enormous and they provide unmatched performance and reliability benefits once the technology is matured. However the maturation process itself represents a significant challenge. In a cost-constrained environment a step-wise and evolutionary approach to the development of this critical technology is optimal. Such an approach takes on risks in smaller, less costly increments. A foundational functionality is achieved as early as possible- enabling learning by direct experience. Subsequent steps benefit from concrete data and design refinements.

The proposed evolution plan utilizes existing launchers and upper stages coupled with orbital test beds to move from the theoretical position of today to a depot design founded on test data but using existing aerostructures, tanks and equipment. The benefits of this very basic LEO depot as well as subsequent L1 or L2 based depots are shown for missions other than lunar landings. These range from robotic advanced capability probes to virtually anywhere in the solar system to more ambitious manned missions to near earth asteroids and the moon.

With the technique of orbital propellant transfer secured by years of experience a practical architecture for indefinite stay Martian exploration can be readily envisaged as a direct extension of lunar operations. Mars' gravitational environment implies that single-stage-to-orbit vehicles that can reenter without significant aerobraking are practical. Mars orbital propellant depots coupled to these SSTO vehicles enables a strategy of global reconnaissance and exploration.

I. Introduction

The United States is at a point of momentous decisions that will steer the fate of our entire military and civilian space programs for the next 30 years. We are confronted with what appears to be a classical mismatch of needs, desires, and available resources.

Though the policies regarding NASA and exploration dominate the news, there are several problems that must be solved. The military is increasingly dependent on space-based assets that have extremely high value and require the very highest reliability to deliver them to orbit. Simultaneously, there are escalating launch costs due to low launch system utility and increasingly constricted defense budgets. NASA desires to extend crewed operations beyond low Earth orbit (LEO) and perform meaningful exploration activities with relatively flat budgets. NASA must also develop economically viable yet safe transport for crew and cargo to the International Space Station (ISS). At the same time, robotic capabilities are expanding rapidly and there is a strong desire to deploy these increasingly potent technologies throughout the solar system.

Is there a solution that addresses these diverse mission needs while suppressing costs and increasing America's global competitiveness? We believe the answer is yes, but it will require a change of philosophy and policy throughout the government and its contractors. What it does not require is the assumption of excessive risk, drastically elevated budgets, and unrealistic or unsubstantiated assertions as to the mass, safety, reliability or cost of any element of the architecture. We believe that it holds the best hope for a timely, affordable, and flexible space transport system that will be the envy of the world for decades to come.

II. Philosophical Gear Shifting

Before we expand on our solution, let us examine the historical NASA approach. The traditional solutions, effectively summarized in the Constellation program and its variants, envision a transport architecture founded on single-source hardware built using the same engineering and programmatic techniques used for half a century. Implicit in this approach is the thought that the procured item is unique to the government and that no one else would want one. This is the approach taken when buying F-22s or aircraft carriers. Embedded in this is the notion that the U.S. Government (USG) must effectively define the solution, not simply establishish top-level requirements.

Almost from the outset of Constellation was the imposition of minimal innovation. The intent was to effectively leverage Apollo-era design concepts to minimize development risk and bound cost. This approach was touted as the key to success.- “Apollo on steroids” was the motto. However well meaning, this decision had profound impacts.

The desire was to do things bigger than Apollo, including a much larger rocket to achieve success. Indeed, it is seen as a point of national pride that the U.S. should have the biggest lift capability. The purpose or focus of the exploration effort was hardly mentioned. What astronauts and scientists would be actually doing and for how long was almost an afterthought, representing just 5% of Constellation’s budget. The hardware for transportation was at the forefront. Front and center was the insistence that the presence of a heavy lift launcher was mandatory for any useful accomplishment.

As Constellation matured, the heavy lift demand grew steadily until its most valuable feature, being able to use flight proven Shuttle hardware, had entirely vanished. The program demanded new solids, a main propellant tank with a diameter not seen since the 1970s, dozens of new pieces of major primary structure, drastic launch complex changes, and an operational concept requiring unprecedented precision launch timing of two rockets. The number of unique, low-rate components that had to be developed and supported dwarfed the lamentably costly Space Shuttle program leading to program overruns and dramatic schedule delays. One incisive member of the 2009 Augustine Commission summed up this situation by musing that if NASA was to be given the entire Constellation hardware as a gift, it would have to immediately shut it down because operational costs were too high.

Perhaps most concerning was the sheer amount of clean sheet design coupled to low launch rates. Launch vehicle reliability has been attained by recognizing that incremental design changes are the *only* practical way to develop safe and reliable space vehicles and that extensive flight experience is pivotal. There is no debating this or skirting the problem. Nearly every new rocket brought to market in the past years has exhibited either outright failure or very significant anomalous performance during its initial flights. The greater the number of clean drawing sheets at the outset of the design, the greater the probability of malfunction- despite the imposition of costly and time-consuming ground testing. The single-purpose nature of the clean-sheet Constellation vehicles virtually assures that they would see very few flights- a risky proposition.

It is also critical that design and manufacturing be intimately coupled. The Constellation project places the design and manufacturing on separate entities in most cases. This approach is unprecedented in its denial of fundamental systems engineering lessons learned across decades and uncounted programs. It is a recipe for unbounded cost growth, endless delays, poor end performance, and ultimately the elimination or downsizing of the program.

Lost in all this discussion is the fact that not one element of the Constellation development, totaling well over $100B prior to the first lunar mission, could be realistically applied to solve the very serious DoD, NASA science, and commercial launch problems. This fundamental oversight persists in the very latest Constellation incarnations

The often-described conservative approach can thus be seen to be fraught with risk and is postured to generate poor cost, schedule, and technical performance.. The enormous investment fails to address pressing DoD needs, which are of equal or greater criticality. This combination makes the program extraordinarily vulnerable to termination, reconfiguration, or delay, which will inevitably degrade the overall benefit to cost ratio.

III. A New Architecture

In 2009, ULA revealed a transport architecture that embodies an alternative approach that not only matches any traditional heavy-lift based architecture, but also accomplishes tasks that the traditional approach cannot ever perform. This architecture is based on distributed launch using propellant depots and is a direct response to the realization that 70% of the mass we move to LEO orbit is simply propellant. This means that the bulk of the mass for exploration is a commodity substance with little intrinsic value. Its delivery is straightforward and accomplished with the stable of already existing launchers. This means that 70% of the exploration transport task can take place without any significant development or risk. It does not require a preordained solution or substantial government oversight. It is a “go-do” task. In fact, it is a task so simple that it can be subjected to continuous price competition just as everyday earthbound shipping tasks are. The launcher can be engineered to absolute simplicity for minimal cost because the propellant is of low value.

The primary barrier to deployment of this approach is the lack of the orbital depot to receive, store, and forward the propellants. The development of propellant depots has been stymied to date by a combination of grandiose goals and the imposition of unrealistic and unneeded technical demands.

The depots we need must be capable of storing liquid oxygen and hydrogen and the concern has always been that these super-cold propellants would simply evaporate due to the radiant heating from the sun and earth. Extraordinary technology efforts have been underway for years to reduce these “boiloff losses” to zero- to achieve perfect thermal stasis or ZBO. The extensive focus on these ZBO technologies over many years implies to many that a practical orbital depot cannot be emplaced without them. While existing, demonstrated technologies can effectively achieve ZBO for oxygen, it is far more difficult to accomplish this for hydrogen. Hence, there is a perceived technical barrier to orbital depots that constitutes a major risk to committing to depot-based architectures.

Coupled with this artificial technical barrier is the notion that a functional depot must be very large and must be assembled on-orbit from multiple elements. This sizing is associated with historical concepts of mustering an armada of stages and payloads in LEO, fueling them nearly simultaneously including vehicles that will not be used for days or weeks, culminating in a single enormous departure burn to send us on our way. The total propellant masses that would have to be cached in LEO for this concept are indeed large. This implies a costly orbital assembly task likened to the protracted construction of the ISS. With the reluctance to engage in any similar effort anytime soon the depot concept is relegated to the “maybe someday” category. It is perceived as an enormous millstone around the neck of exploration and hence something to be strenuously avoided, leading to the perception that heavy lift is required for exploration.

Both of these perceptions are questionable. To begin with, a depot does not have to be enormous and does not need any orbital assembly or human occupation to be functional. The key to a depot’s effectiveness is to actually use it and that does not mean simply storing propellants indefinitely. There must be a constant flux of mass through the depot to justify its presence, otherwise it is the equivalent of a gas station in the middle of a wilderness. While its capacity might be 120t, the flow through the depot may well be 300t over the course of a year. Importantly even with a depot architecture only 75% of the total propellants lifted to LEO will even be stored in a LEO depot. The balance will be contained in the vehicles receiving propellants from the depot.

From a design standpoint, we want the depot to be as small as possible with the least incident radiation and smallest aerodrag. We want to be able to maneuver the depot so that visiting vehicles require only simple rendezvous and docking systems and so that it can self-deploy to any location we wish without relying on other dedicated vehicles to move it. We want it to be low enough in cost that we do not have to design for a lifetime measured in decades. We want to be able to replace it as technology evolves without an insurmountable cost.

The ZBO barrier disappears when we recognize that a depot is not a device operating in a theoretical, idealized environment. A LEO depot is subject to substantial aerodrag and tidal forces requiring continuous reboost and station keeping propulsion, which requires the expenditure of propellant at the rate of tons per year. Furthermore, the heart of a depot is rendezvous and docking with visiting vehicles. In an economically active depot, these events are numerous and require further propellant. Fortunately, the waste hydrogen that has boiled off happens to be the best-known propellant (as a monopropellant in a basic solar-thermal propulsion system) for this task. A practical depot must evolve hydrogen at a minimum rate that matches the station keeping demands. This places the boiloff rate within the range that existing cryo-storage technology can support. More importantly, the loss rate is not measured against the capacity of the depot:- it is measured against the propellant flux through the depot. Moderate loss rates can be tolerated if you are actually putting the depot to work. Depending on launch vehicle, market pressures and delivery rate the cost of propellant at LEO should fall in the range of 5 to 10 $M/mT. The yearly cost for these losses is small.

So how do we assure that there is propellant demand and that we do not sit with a full LEO depot for years? The key is to look at where exploration efforts want us to go. Every interesting place in the solar system beyond Earth orbit requires us to first push the payload up to escape velocity. This demands a further 3.2 km/sec of energy investment just to get to the starting line. To actually transit to the destination requires further, large investments in energy. This ranges from 1 to 9 km sec depending on destination and trajectory.

We know that we must make these energy inputs. The difficulty is that the timing is extremely erratic due to the occasional nature of interplanetary flight. The advent of an exploration initiative begins to impose more consistent and substantial demands on escape energy but it is still periodic. Somehow we must address the timing issue. With our current non-depot mentality, we are compelled to push all this energy into the payload in a single effort starting from Earth's surface. This is merely an artifact of decades of working with vehicles where we had no other choice. Each individual mission had to supply its own energy. Often this meant invoking the most powerful launchers with additional stages and custom modifications to get to that final peak energy. Spacecraft mass had to be strictly controlled,- driving their designs to exotica or stripping back instruments entirely. The Pluto New Horizon mission with a mass of under ½ mT required the largest Atlas vehicle with five strap-on solid rocket boosters (SRBs) and a Star 48 third stage. . A Mars probe with a transit mass of 4t is considered extremely large. At present, we are effectively saturating the existing launch vehicles. Exploration demands go well beyond these existing capabilities.

What we would like is a way to actually make the investment in earth escape energy gradually and efficiently without actually leaving earth orbit and have it be applicable to any destination. We want to accumulate enough energy to deal with the exploration demands and be useful for smaller missions, a sort of storehouse of "earth escape energy" that we can draw from at need. Can we do this? The answer is yes. We can create an energy savings account by moving propellant to the earth-moon Lagrange points – especially L2. Located 60,000 km beyond the moon, propellant or cargo cached at L2 is very nearly at earth escape energy. It takes only a small nudge to dislodge it from Earth's gravitational grasp. This has been known for decades and L2 is often called a gateway to the solar system.

Just as importantly L2 is in deep space far away from any planetary surface and hence the thermal, micrometeoroid, and atomic oxygen environments are vastly superior to those in LEO. Thermodynamic stasis and extended hardware life are far easier to obtain without these punishing conditions seen in LEO. L2 is not just a great gateway- it is a great place to store propellants.

The position of L2 allows the use of powered gravity assist exit burns; also called Oberth maneuvers that effectively "spend" our accumulated energy savings. A propulsive event performed at high velocity near perigee receives an augmentation of exit velocity. For typical earth departure events, this augmentation is approximately a factor of 2 to 7. A typical Mars departure ΔV of 4.3 km/sec (from LEO) is obtained with a perigee burn of less than 1 km/sec. Because of this we not only cash in our savings but the vehicle doing this last bit of work can be very small. Its dry mass can be much smaller than for a traditional system or conversely our existing vehicles can do proportionately larger tasks. For example, a Centaur class vehicle can move a 7t spacecraft from LEO to L2, refill there, execute an Oberth exit burn, and retain enough propellant to lose over 40 lbs per day during the 8½ month transit and finally execute a 2.7 km/sec insertion burn at Mars. In other words, a depot-augmented, long-duration Centaur can move to Mars orbit what we can barely move to geostationary orbit on the largest rocket in the US inventory using our existing approaches.

Figure 1 illustrates the enormous amplification of vehicle capabilities that comes with depot operations. Existing Atlas/Centaur performance is highly constrained and upgrading to a larger upper stage, while valuable, is only a small incremental benefit. Taking those same stages and departing from a LEO or L2 depot enables the same hardware to deliver far larger payloads to extremely high energies. These are the kind of revolutionary improvements that can change the game for planetary exploration and establish a long-duration human presence on the moon and Mars.

In short, L2 is an ideal location to store propellants and cargos: it is close, high energy, and cold. More importantly, it allows the continuous onward movement of propellants from LEO depots thus suppressing their size and effectively minimizing the near-earth boiloff penalties. Should industrial production of propellants on the moon take place, L2 is optimally placed to put these propellants to work.

Depots are thus seen to provide a continuous demand on launch mass to charge up our "ΔV battery." These propellant delivery launches can be regularly scheduled even though the missions that use the propellant may be very erratic. This is the foundation for an economical space transport architecture. It should be noted that the payloads themselves follow the same mission path: they always go to the same depots. Thus their launch schedule is far more flexible than when we have to execute their departure in a single effort. A Mars-bound spacecraft can be lifted to LEO or L2 at any time and loiter there until its launch window opens. This eliminates typical launch delays such as weather, launch facility, and range issues.

The hardware required to execute the final departure is the simplest possible, and thus least likely to inhibit departure. It has been effectively debugged by prior usage.- We will have seen that stage operate and digested its flight data. Our confidence will be maximized. If the worst happens and we find we have an unhealthy stage, we can replace it. It is, after all, identical to other stages that were expended delivering propellants to the depot and are simply parked there. These are options that simply are not in our playbook at present.

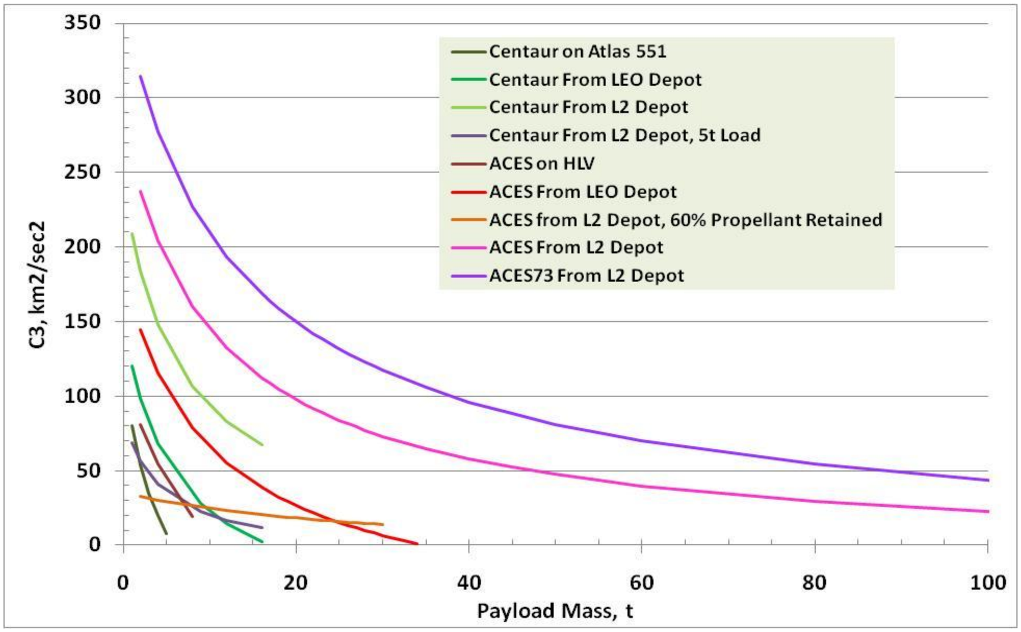

Figure 1. Amplification of vehicle capabilities with depot operations.

IV. The Advanced Common Evolved Stage (ACES)

While basic depot operations occur at smaller scales, at the heart of our architecture is the ACES upper stage. Some background on this hardware concept is required to support depot architecture discussions.

ACES design conceptualization has been underway at ULA for many years. It leverages design features of both the Centaur and Delta Cryogenic Second Stage (DCSS) upper stages and intends to supplement and perhaps replace these stages in the future. The baseline ACES will contain twice the Centaur or 4m DCSS propellant load, providing a significant performance boost compared to our existing upper stages. The baseline 41-mT propellant load is contained in a 5m diameter, common bulkhead stage that is about the same length as ULA's existing upper stages. ACES will become the foundation for a modular system of stages to meet the launch requirements of a wide variety of users. A common variant is a stretched version containing 73t of propellant for use as a tanker and as a component of the depot itself.

ACES is based on a simple modular design (Figure 2). The use of multiple barrel panels, similar to Centaur, provides a straightforward means to building multiple-length (propellant load) stages that are otherwise common. The common equipment shelf accommodates one, two, or four RL10 engines. While ACES can start with existing Centaur and Delta pneumatic, avionics and propulsion systems it is intended to transition to lower-cost and higher capability systems founded on our Integrated Vehicle Fluids (IVF) system concept. IVF eliminates all hydrazine, helium, and nearly all batteries from the vehicle. It consumes waste hydrogen and oxygen to produce power, generate settling and attitude control thrust, and autogenously pressurize the vehicle tanks. IVF is optimal for depot operations since only LH2 and LO2 need be transferred, and it extends mission lifetimes from the present dozen hours to multiple days. Compared to traditional designs it saves more than a metric ton for a vehicle the size of an ACES 41. Importantly it remains a fixed design even as the propellant tank capacities change. Coupled with solar power systems it enables indefinite vehicle operational durations.

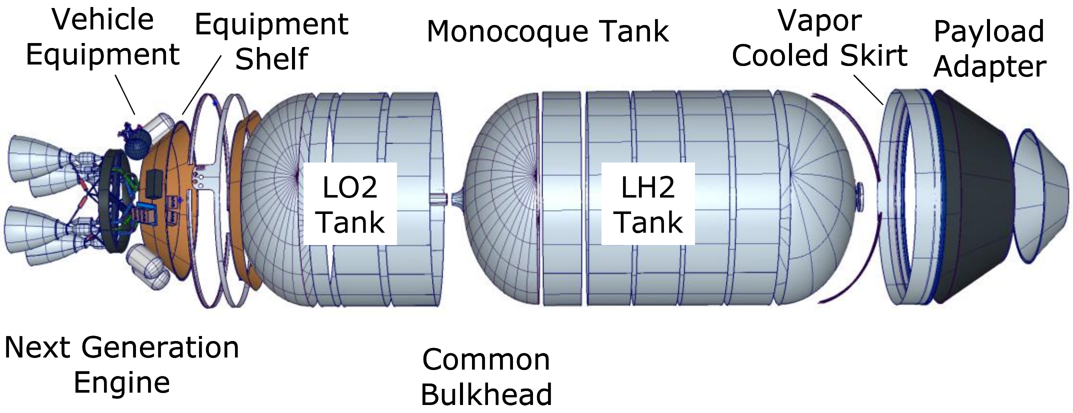

Figure 2. The Advanced Common Evolved Stage

ACES design is optimized with long-duration cryogenic applications in mind. A number of passive-thermal management features are incorporated into the stage at the system level. The tank geometry design minimizes the exposed surface area. Warm equipment is isolated on a separate, thermally controlled shelf, and IVF systems provide a stable thermal environment regardless of vehicle orientation. This is critical for depot applications. Vapor-cooling paths, where vented hydrogen is used to intercept the remaining high-load heat paths, are integrated into the tank structure.

Mission reliability is a key demand and ACES continues to place a top priority on reliability across the entire life cycle, including design, component, and process levels. ACES will retain the present avionics redundancy, and with IVF, it can extend this basic block redundancy to systems that have historically been incapable of it. IVF is modular and it maintains vehicle operations even with a complete module failed. Common cause failures such as leaks, which can disable even traditionally redundant systems, are substantially reduced. So long as propellant remains in the main tanks, the IVF system can maintain systems operation. Significantly, the dual and quad engine ACES variants incorporate engine out. Engine-out capability provides the single largest lever to improving system reliability.

V. How the Architecture is Used

The ACES depot deployment requires four launches: two 554/ACES73 LEO tankers, and the depots themselves which are launched on 511 vehicles. The total time from first launch to initial operational capability is anticipated to take approximately 6 months. The total launcher and depot hardware cost and services are expected be roughly $1B, including orbital checkout.

The depot consists of an ACES 41 main vehicle coupled to an ACES 73 tank (Figure 3). Both the LEO and L2 depots will be subjected to thermodynamic testing to confirm their proper operation prior to final deployment. At the completion of the deployment process, they are not empty but contain approximately 10 tons of propellant at each location. The two ACES 73 tankers are not discarded when empty. They are maintained in readiness to move propellants farther up the gravity well to L2. This permits vehicles delivering propellants to LEO to be very simple. They are only required to fly to LEO and need not be capable of managing the multiday trip to L2.

A model combining NASA exploration missions, interplanetary probes of escalating mass, critical DoD missions and culminating in the first Lunar surface mission is shown in Figure 4. The simulated launch schedule was modulated to create periods of low and high mission activity to illustrate how a depot system would respond. While propellants are anticipated to come from multiple sources for simplicity we assumed that they were all delivered in 26t increments using an ACES 73 tanker.

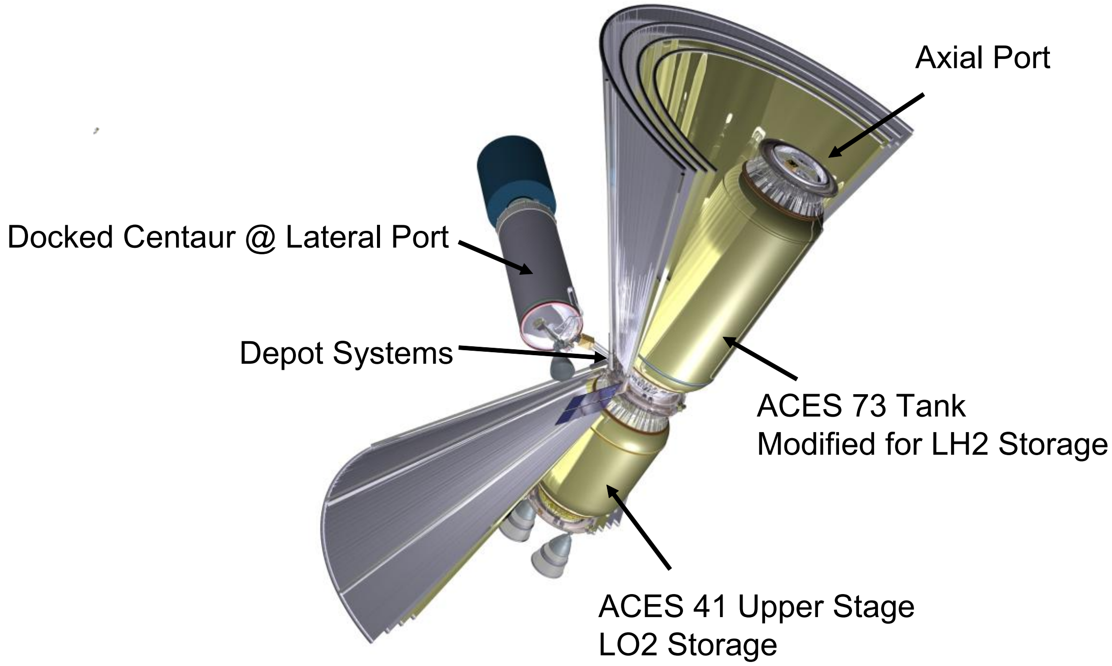

Figure 3. The ACES 121 depot consists of an ACES 41 main vehicle coupled to an ACES 73 tank.

.

Initial exploration activities center on the initial long-duration flights of Orion- likely lunar orbital missions with total ΔV less than 5.4 km/sec. The Orion Command Module and Environmental Control and Life Support System (ECLSS) systems are combined with an ACES 41 stage. These missions are effectively limited by Orion habitation space and food stores. With an L2 depot and routine tanker deliveries, the Orion/ACES could stay in cis-lunar space indefinitely. It is not limited by the present Orion design incorporating low performance hypergolic propellants. The ACES 41 with a 12t Orion/ECLSS module can deliver 5.5 km/sec of ΔV without accounting for boiloff losses, water synthesis, and breathing air demand. This is enough to encompass visitation to any Lagrange point, insertion into low lunar orbit and return, and permits a rapid return to earth with a low entry velocity.

With the Delta IV Heavy Lift Vehicle (HLV) and ACES 41 combination very heavy Orion variants can be lifted to LEO- up to 36t. This would permit larger habitation spaces and additional food consumables. Because of the O2/H2 ACES stage, water and breathing air, the heaviest consumables are produced from the ACES 41. It is anticipated that the Orion extended Command Module with augmented food stores would mass less than 20t. With this mass and using the L2 gateway, we can accommodate a total mission ΔV about 6 km/sec including a 260-day transit time with daily consumables and boiloff of 60 lb/day. This transit time is effectively equivalent to the outbound journey to Mars. This capability can begin to give us access to some near Earth orbits (NEOs), but many remain out of reach due to high energy demands. An Orion/ACES 41 can exchange its ACES stage at L2 for a previously used ACES 73. With this capability, the ΔV is raised to 8.7 km/sec with 340-day transits. This capability envelopes the bulk of NEOs, which are considered for visitation. We typically execute a total ΔV getting to departure energy of less than 1.4 km/sec because we use the L2 gateway and Oberth departure strategy. This preserves the bulk of propellants and directly supports NEO visitation by preserving the maximum propellants for consumables and return delta V.

If a target is especially interesting and we wish to stay there for extended period or if it has extremely high ΔV demands, we can simply apply the same strategy as we would at L2: Send propellant there first and then rendezvous with it. A single ACES 73 tanker can inject in excess of 50t of propellant to a trans-Mars injection energy. The tools we have created should be up to any desired transport task.

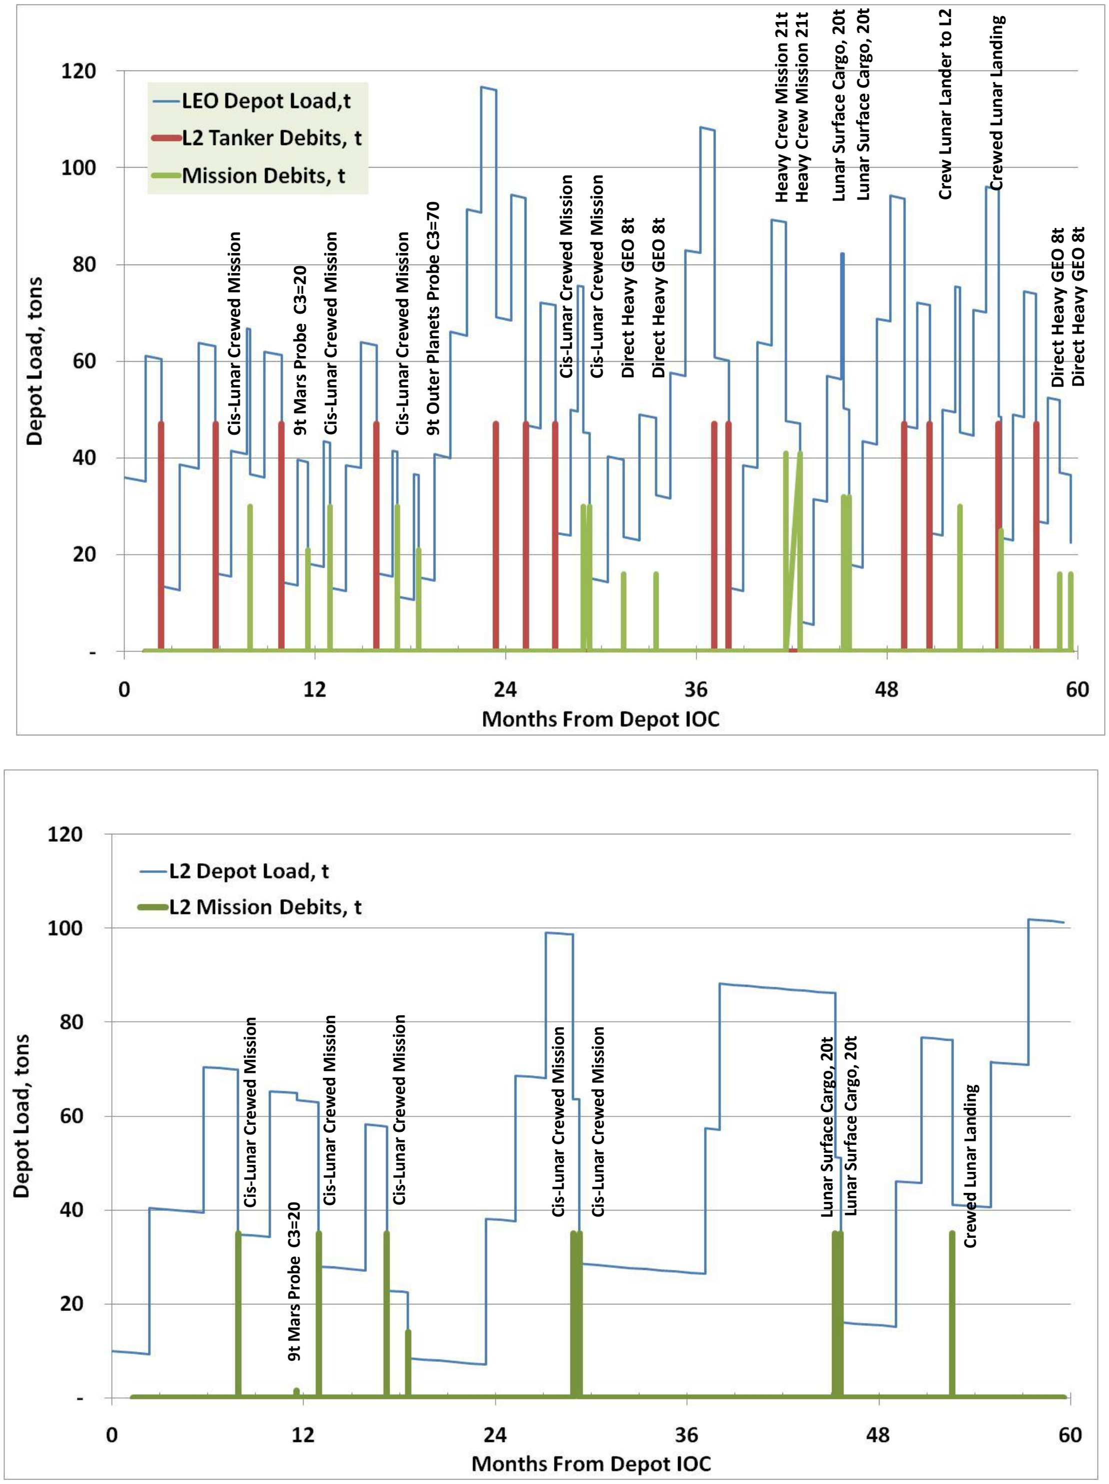

Figure 4. LEO Depot (upper graph) and L2 Depot (lower Graph) work in concert to support varying planetary, crew and Geostationary Orbit missions.

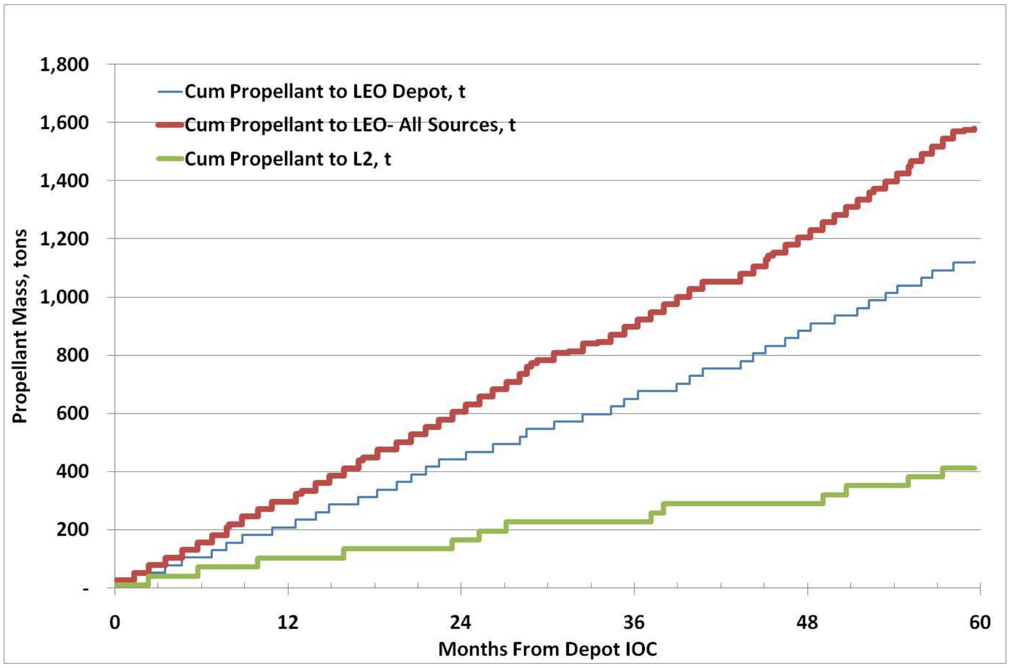

Figure 5. A wide range of missions to are supported by reasonable depot lift capabilities

A basic exploration effort (Figures 3-5) enveloping conceptual crewed and robotic asteroid and cis-lunar missions conducted at a rate of two to four missions per year would require a total lift by propellant suppliers to LEO depot of roughly 200t at an approximate cost of under $2B/year,- the equivalent of nine Atlas 552 vehicles like the one that launched Pluto New Horizons. This type of demand would be an unprecedented spur to industry. Engine production rates would reach levels not seen in half a century. The DoD would be using the identical booster and upper stage hardware and overall production rates would have doubled. The intention of the EELV program, that large-scale commercial space operations would sustain an industry and suppress cost growth for all space users, would be realized. The propellant demand for a twice-yearly lunar landing scenario is approximately 300t/year. Thus, even at the outset of exploration, enormous pressures on launch hardware costs would be brought to bear and overheads would be significantly diluted. Compared to an Ares-style launch approach, the yearly operational savings begin at half a billion dollars and escalate as mission demands increase.

VI. Depot Evolution

We are neophytes when it comes to orbital depots. In keeping with the stark recognition that there is no substitute for flight experience, ULA proposes that a straightforward evolution plan be executed to retire depot risk.

As shown in Figure 6, this evolution plan begins with a testbed, flies this multiple times to mature that hardware, proceeds in increments to a basic depot, and then finally to a depot suited to full-scale lunar and Mars exploration. While this baby-step approach may appear timid, it in fact delivers immediate, valuable, and new capabilities at each step. It allows the maturation of hardware and operational concepts while providing the theoretical underpinnings. It gets our hands dirty early. If our technological noses are to be bloodied we get that experience early on and inexpensively – not at the end of a “big bang” project based on presumption and analysis.

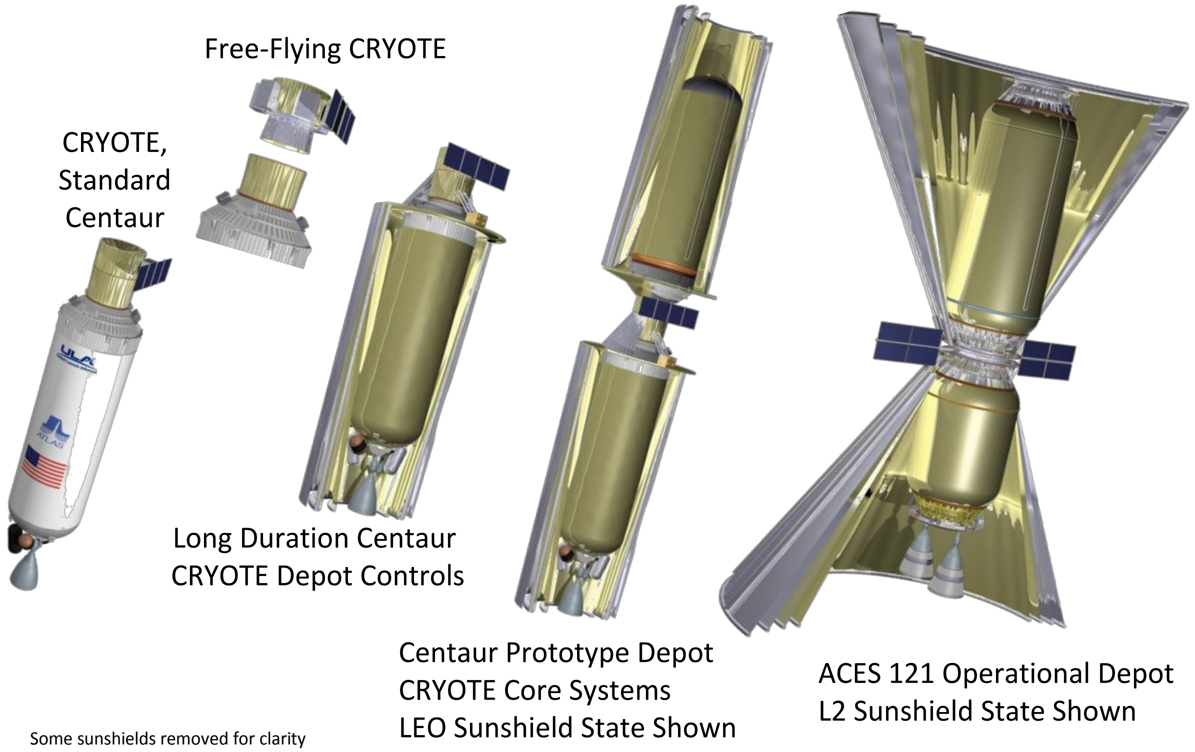

Figure 6. CRYogenic Orbital TEstbed Evolution to Final ACES 121 Depot.

The first step is CRYOTE: the CRYogenic Orbital TEstbed. CRYOTE is a secondary payload that flies on top of any Centaur upper stage. It contains a small, heavily instrumented propellant storage vessel and the complete complement of control valves, disconnects, vents, etc. for conducting an entire range of fill/drain/thermal control/chilldown and vent experiments. Its design is deliberately modular to allow varying suites of instrumentation to be flown. CRYOTE is launched empty so there are no ground cryogenic complexities and it remains inert until the primary payload is deployed. Then CRYOTE, with its own onboard avionics control, fires up and begins experimentation. Cryogenic hydrogen or oxygen are obtained from the copious residuals present in the Centaur main vehicle tanks. These residuals are typically so large that many load/drain or thermal cycles can be performed.

Initial CRYOTE flights use the primary vehicle and then auxiliary batteries, which limits the total on-orbit test time. Follow on flights draw power from a combination of batteries, solar and the IVF system. IVF is a key piece of depot technology that enables generation of peak power for propellant pumping. It also provides attitude control and settling thrust for propellant position control by consuming waste H2 and O2 vapor in a small engine. With these power sources, the CRYOTE test time can be extended to many days and is only limited by the Centaur residuals.

Eventually, CRYOTE becomes a free flying vehicle that is loaded from Centaur, separates from it and can gather data without the influences of Centaur. With its own power and avionics the CRYOTE satellite can operate for months. If desired it can execute rendezvous and docking operations similar to what has been flight proven for hypergolic propellants. With these lessons learned we have the data in hand to assess thermal isolation strategies and propellant management under conditions never before tested. This small satellite will open a large window into this new frontier.

CRYOTE and IVF tests set the stage for more extensive testing involving greater integration with the Centaur upper stage. Early CRYOTE flights of necessity use a standard Centaur that does not have long duration thermal insulation in place. This minimal thermal insulation results in a high propellant loss rate from the Centaur and thus limits orbital test durations. The next step is to add long duration multi-layer insulation (MLI) blankets, improved thermal isolation, and external sunshields. These take advantage of the Atlas 500 design, which encapsulates the entire Centaur and provides a large annular space between the payload fairing and the Centaur for sun shields. The Atlas 500 vehicle also contains a split-disk structure called the Forward Load Reactor Deck that is normally jettisoned right after the PLF. For CRYOTE experiments, we will retain the deck. The FLR deck contains electrical, fluids, and other umbilicals that make it the ideal location for a depot-docking feature. It provides an ideal support for the Centaur sunshields. By leveraging the FLR deck, ULA can obtain an effective test of ultra-high performance

MLI and a full-scale sunshield for a cryogenic tank. With the high residuals provided by adding solids, far lower boiloff rates, and ready sources of power CRYOTE testing can be extended to multiple months in duration.

With the data from the above testing in hand, we can design our first dual-propellant orbital depot. Of course, even simpler, single-propellant depots can be created, but the dual propellant depot offers the greatest forward utility. This depot leverages all the CRYOTE design efforts and simply adds another Centaur-based tank on top of CRYOTE. This tank has its bulkhead shifted to maximize hydrogen storage capacity. Like CRYOTE, the upper tank launches empty so no additional propellant umbilical ground facilities are required. With ULA's proven structural design approach, the payload mass associated with the depot will be less than 2 tons. Launched on an Atlas 551 this means it arrives in LEO with 12 tons of residual propellant. The CRYOTE core continues to provide the propellant handling, power, avionics, and thermal control capabilities. Fluid and structural interfaces are provided at locations on the periphery of the FLR deck. Upon arrival at LEO, the H2 from the Centaur upper stage is transferred to the large forward tank. The Centaur donor LH2 tank is then evacuated. LO2 is then transferred via the CRYOTE pumps to Centaur's larger and more thermally efficient LH2 tank maximizing depot capacity. Nearly 30 tons of propellant can be cached in this prototype depot.

The proto depot is immediately useful for testing the depot concept and for refilling Centaur or Delta upper stages that arrive at LEO. A completely topped off Centaur can inject a 10t satellite directly to geostationary orbit or place the same mass in low lunar orbit, capabilities which do not exist today. This means that DoD payloads that require a dedicated Delta IV HLV can instead be launched on an Atlas 511. Adding in the cost of propellant received at LEO, the total lift cost is tens of millions less than at present. A wide range of missions to Mars, Venus, Mercury, and smaller NEOs can be accomplished with this depot-assisted mission design. Figure 1 shows the range of capabilities that are obtained. If desired the entire depot can self-deploy to L2 to allow data gathering at that location and to provide all the advantages of L2 caching.

The prototype depot sets the stage for the larger capacity depot that is based on larger diameter ACES upper stage. With a capacity in excess of 120t multiple ACES stages can be replenished without a tanker visiting. This permits multiple vehicles to loiter at the depot awaiting departure windows. Like the Centaur based depot the ACES depot is self-deploying to anywhere required and can be maneuvered at need. Any location with sufficient propellant requirements such as LEO, L1, L2, LLO and even inclined orbits can be supported. This depot internalizes the CRYOTE hardware and functionality but augments it by supplying multiple docking locations and fluid transfer ports and is even further optimized for low propellant boiloff. Lessons learned from the Centaur depot with respect to sunshields, MLI placement, vapor cooling, fluids and mechanisms will all be built into this depot. This direct evolutionary extension of prior depot design suppresses costs and risks.

VII. Architecture Extensibility

The deployment and operation of the ACES depots will provide the knowhow and technical confidence to permit serious Mars exploration to occur. The journey to Mars is essentially an exercise in long-duration cryogenic propellant storage. It must be managed on the way to Mars, received and stored from in-situ synthesis, cached in Mars orbit, and finally used for Mars departure back to Earth many years after arrival. With our experience moving hundreds of tons of propellants throughout the Earth-Moon system, we will have proven tools at hand to tackle this critical job.

As shown in Figure 7, the ACES depot can be fitted to act as an in-space stage with a propellant capacity roughly three times that of the basic ACES (ACES 41 contains 41 mT of propellant) and with thermodynamic capabilities that enable flight times measured in months to years. This combination of long duration, high propellant capacity, and low mass enables missions that cannot be otherwise contemplated. Total mission ΔVs in excess of 11 km/sec can be delivered with this vehicle.

To reach our final goal of Mars exploration, this depot/stage can be mated to other identical vehicles- much like what is done today with strap-on solid rockets on boosters. Six vehicles with a combined propellant capacity in excess of 700t enable the direct propulsive insertion of large transit habitats, descent vehicles, and surface installations without staging of tanks or other structure- thus preserving the vehicle configuration for the return to earth. This unprecedented propellant mass combines with an engine complement capable of generating nearly 1.5 million pounds of thrust,- thus enabling reasonable burn durations. Two transfer vehicles can join following TMI to form a system with rotationally induced gravity that further reduces propellant losses and maximizes crew conditioning. Since the departure burn is quite small, the tons of propellants remaining continue to surround the crew's habitat with low molecular weight materials and with relatively large standoff distances to the habitat itself- a reasonable start on radiation shielding.

Fully loaded in Mars orbit, the transit vehicle can depart and retain large propellant reserves for shielding that can then be burned upon reaching earth. By finally staging, the peripheral tanks and habitat sequentially even larger delta V capabilities are obtained. This has the potential for avoiding a direct, high-energy reentry to Earth's atmosphere and instead returning to the L2 depot and a return to Earth via a standard Orion vehicle.

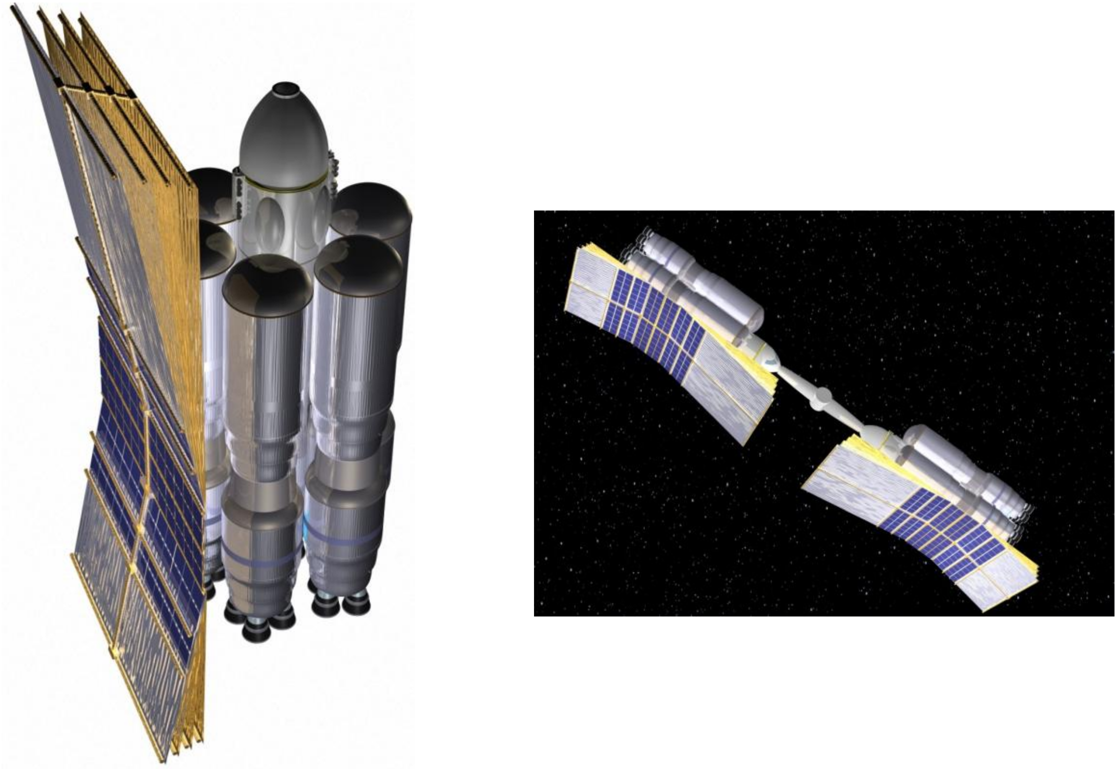

Figure 7. ACES 121 Depots Clustered for Delivery of 16 Crew Transfer Habitat to Mars.

By following this evolution of depots, we mature from the present state up to a basic capability that systematically erases present limitations and ultimately has the interplanetary vehicles needed for large crews to begin to explore a new world. Along the way, we build hardware, mature it, and systematically build upon it. The Mars transit vehicle envelopes the basic ACES stage technology and uses propellant transfer technology pioneered on CRYOTE.

VIII What About the DoD?

Are depots useful for the DoD with its extensive geostationary, polar, and middle-Earth orbit (MEO) assets? It is entirely dependent on usage. Given a low, established propellant delivery cost to LEO, a large GSO satellite can be lifted on a smaller, simpler rocket that can be launched at a higher pace with better payload environments and with a lower total lift cost. In addition, the actual mass of the spacecraft becomes less critical since you would simply buy propellant in LEO to compensate. This may well have positive effects on spacecraft development since the use of high-cost, low-mass exotica could be suppressed.

The reason this is effective is that for direct insertion the upper stage must go all the way to geosynchronous orbit. For most commercial and DoD missions, the upper stage is only placed into a transfer orbit. The depot architecture is not efficient enough to be economically effective for that type of mission unless the spacecraft mass is very high.

However, a new paradigm of spacecraft delivery appears to be settling on us. The days of simply setting an old spacecraft adrift appear to be ending and the permanent disposal of these heavy, delicate objects without creating worse problems due to particulate formation during any recovery process is non-trivial. While there are many proposed approaches for disposing of these obsolete or derelict spacecraft all of them involve the expenditure of substantially more delta V than what has been traditional. It may well be required that old spacecraft be removed at the same time new spacecraft are being emplaced. If this is the case, then depots are highly effective tools. There are various methods proposed for eliminating already derelict spacecraft, but all of them demand that significant amounts of ΔV be available to approach, rendezvous and finally remove the spacecraft. Depots may well be useful tools whenever substantial amounts of ΔV must be applied to heavy objects..

IX Development Costs

So how much does this concept cost and how does it compare to the heavy-lift alternatives? The ESAS or its derivatives have a widely variant price tag of approximately $100B including lunar surface systems. Its recurring launch cost is estimated at $3B/year with an IOC between 2020 and 2025. The cost of the payload elements with their unique low-rate designs is estimated to be approximately $2B/year.

The depot architecture effectively demands only one major development item and that is the ACES upper stage- estimated to be approximately $3.5B including the tanker variant, production rate capability increases and engine upgrades. No completely new upper stage engine is required until Mars operations begin though a line-replacement for the venerable RL10 is advised merely to enable the high production rates that will be demanded. The entire booster system remain as-is for the near-term. The development of a new LO2/kerosene booster engine to supplement the RD-180 is strongly encouraged but is not required to be complete before 2020 at the earliest.

The depot evolution is estimated to cost, including the deployment of operational Centaur and ACES depots at LEO and L2 approximately $3B. LEO propellant donor vehicle development would be borne by the suppliers. The key here is that developmental costs for many elements are diluted across both NASA and DoD budgets.

We continue to believe that a very substantial permanently crewed lunar base is within our grasp and is mandatory before we venture to Mars. With the depot system in place the only remaining transport elements required for lunar surface exploration are the Ascender/habitat and Cargo Delivery modules. With the proposed Dual Thrust Axis Lander concept the lander design uses the ACES tank as the descent system and a substantial portion of lander development costs are avoided- allowing a focus on equipment for long-duration stays on the surface. The development of the lander/ascender elements is estimated to be approximately $5B. The basic surface system development is estimated to be approximately twice that: $10B.

In summary, the entire depot based architecture including all payload elements, payload fairings, habitats, etc. is estimated to cost less than $40B with an IOC for the initial elements in 2016. The recurring lift cost for propellants is approximately $2B/year and this is sufficient to support two major crewed missions or multiple DoD and NASA planetary missions. With the elimination of elements such as the Orion Service module propulsion, lunar descender, and extremely large payload fairings, the cost of the actual payload elements is significantly lower and is principally paced by the cost of the Orion vehicle. Depending on mission mix, the yearly payload costs would range from 1 to 2 \$B/year.

VIII. Summary

Seen from a functional standpoint the exploration of Near Earth Objects, the moon, Mars and high ΔV maneuvering in Earth orbit depends on our mastering the art of moving cryogenic propellants around and keeping them in place. This technology is foundational. No serious Lunar or Mars architecture can avoid depots and indeed the vehicles that will be going to and from Mars are effectively depots in transit.

We believe that the our nation's pride should not be based in having a rocket with the biggest launch mass, but in having an integrated system that can do things no one else can even consider and do it affordably. This is the tradition of our aerospace industry since its inception- to accomplish new things by embracing new approaches irrespective of how strange they may seem relative to the past. By combining proven technology with new thinking, we can accomplish tasks that are presently beyond consideration. If we insist on simplistic, brute-force methods and deliberately avoid innovation we run a very good chance of creating systems which are obsolete on their first day.

The depot architecture is shown to be a flexible and capable system- able to economically address crewed exploration needs including lunar surface habitation, next generation interplanetary probes, direct injection of heavy geostationary satellites and ultimately the crewed exploration of Mars. It anticipates the task of removing derelict spacecraft by providing an infrastructure to permit these high ΔV missions and enables the likely new paradigm of removing a spacecraft for each one deployed.. It is optimal for engaging with new entrants and provides a practical way for all nations to cooperatively support high intensity planetary exploration.

The widely touted heavy lift approach is burdened by an extensive development program whose structure suggests a rocky road ahead. Once completed at a cost over twice that of a depot system, it is relegated to launch rates that are at best 25% of the rate anticipated with depots. All heavy lift missions are unique. These require a substantial overhead for mission design and the costly attentions of thousands of engineers. In a depot system, 80% of launches are identical and contain no unique equipment, lack even a payload, and are bound for the same location. In the end, a heavy lift system can only economically perform the highest-level exploration tasks. It is simply not in the running for addressing any other tasks due to its size, cost, organizational entanglements.

The U.S. industry is increasingly under pressure from international suppliers of launch services. Given that these efforts are hugely subsidized and often a matter of national pride, it is highly unlikely that our industry will be able to compete directly for these entry-level services. By implementing the proposed depot infrastructure, we can actually leverage these new low-cost suppliers as propellant suppliers. Nearly every supplier yearns for business and launch prices are disconnected from actual costs and reasonable business behavior. Delivering propellant creates a market that many entities can supply. It would engage the world community in supporting visionary exploration efforts and the mundane task of cleaning up decades of space junk. The propellant demands are so large that U.S. industry would still be a major player in the LEO delivery task. U.S. launchers with more sophisticated upper stages such as ACES would perform L2 propellant transfers and the high-value satellite emplacement and retirement activities. Effectively, the high-cost U.S. industrial base whould address more sophisticated, non-routine and more mission-critical activities. The depot architecture helps all players become and stay winners.