




Constellation Program Lessons Learned

Volume II
Detailed Lessons
Learned





LESSONS LEARNED

Volume II

Summer 2011

Constellation Program Lessons Learned Volume II

Edited by:

Jennifer L. Rhatigan, PhD, PE
Special Assistant,
Constellation Program

Date

Concurred by:

Deborah J. Neubek
Chief of Staff, Technical
Constellation Program

Date

L. Dale Thomas, PhD, PE
Former Program Manager,
Constellation Program

Date

Approved by:

Charles Stegemoeller
Deputy Program Manger,
Constellation Program

Date

FORWARD

The Constellation Program, as a national-level undertaking, had multiple purposes. These can and will be described from many points of view.

To us Constellation was born of the charred debris, scattered in the Piney Woods of East Texas that were reverently collected as the last remains of that delicate machine Columbia, she that failed to preserve the lives of our friends and colleagues due to our own very human errors. The painful acknowledgement of those errors became foundational principles for what later emerged as the nation's new space exploration program: Constellation.

Those of us who worked in the program were bound by a common belief that America's future lies in the Final Frontier; indeed that America must lead the way there. We joyfully dedicated our working lives and careers to building the ships that would sail those treacherous seas. We had prepared ourselves for this challenge, having learned from the best and honed our skills in the unique arena of human space flight. We were driven by the hope that America would one day step beyond the known, and were blessed when the national will bent toward our dreams for a moment. It was a privilege and an honor to have a hand in this endeavor. Perhaps our greatest shared joy and satisfaction came from the recognition that amid the inherent tumult of a program of this size and scope, we bore the challenges with teammates of superior skill and dedication. We understood that another generation would explore using what we had conceived and built. We rejoiced at our successes, but also knew faults and failures. This, the summation of what we have learned, is our final expression of that cherished hope.

To our families, we beg your forgiveness for the long hours away, and too many distracted conversations. Thank you for indulging our passion.

And, to our generational forebears in human space flight, our hope is that our efforts honor you.

CONTENTS

Forward	iii
Contents	iv
Preface to Volume II	viii
1.0 Introduction	1
2.0 Context — “Where You Stand Depends on Where You Sit”	2
2.1 A Brief Description of the Program	2
2.2 The Program Environment — Key to Understanding the Lessons Learned	4
2.2.1 Program Scope	4
2.2.2 Program Funding	4
2.2.3 The Mosaic of NASA Participation	8
2.2.4 The Generation Gap in Human Space Flight	9
2.2.5 The Size of the Trade Space	10
2.2.6 Phasing of Project/Program Start-up	11
2.2.7 Starting Well vs. Demonstrating Early Success	12
3.0 Key Findings	14
3.1 Robust vs. Optimal Planning — The Only Certainty Is that the Funding Will Not Match the Plan	14
3.2 Schedule Creep and the Fixed Base — The Law of Diminishing Returns	15
3.3 Tailoring and the Design and Construction Standards — Drinking from a Fire Hose	15
3.4 Tailoring Process Simplification — The Law of Unexpected Consequences	17
3.5 Risk-informed Design — Risk as a Commodity	18
3.6 In-house Tasks — Sustaining the NASA Institutional Base vs. Affordably Supporting the Programs — Getting from “or” to “and”	19
3.7 Roles, Responsibility, and Authority — a Non-thermodynamic Application of Entropy	20
3.8 Decision Making — Only as Efficient as Roles, Responsibilities, and Authority Are Clear and Understood	21

3.9	Organization is Organic — You Will Never Get It Right, But You Can Make It Better	23
3.10	Communication Among a Distributed Team — Interpersonal Networks and Information Systems Can Improve Decision Making	24
3.11	Flight Tests – Learning by Doing	26
4.0	Detailed Lessons Learned	29
4.1	Roles, Responsibilities, Authority, and Decision Making	29
4.1.1	Ensure the Decision-making Process Is Aligned with Program, Project, and Center Roles, Responsibilities, and Authority	30
4.1.2	Establish a Cross-level Senior Management Team to Clarify and Improve Roles and Responsibilities	32
4.1.3	Balance Decision Making with the Obligation to Air Dissenting Opinions	34
4.1.4	Implement Communities of Practice to Resolve Integrated Issues	35
4.1.5	Implement a Strong Technical Integration Function	37
4.2	Requirements	38
4.2.1	Reduce Agency Requirements Levied on Contractors	38
4.2.2	Minimize Programmatic Requirements	40
4.2.3	Separate Tailoring of Requirements from the Waiver/Deviation Process	42
4.2.4	Reduce Cost by Maximizing the Use of Industry Standards	43
4.2.5	Use Industry Standards for Data Management	44
4.2.6	Anchor Integration Verification Requirements	45
4.3	Data Management	47
4.3.1	Use a Program-wide Data Systems Architecture	47
4.3.2	Establish a Data-centric (vs. Document-centric) Approach	49
4.3.3	Establish a Model-based Systems Engineering Approach	50
4.4	Schedule and Resource Management	51
4.4.1	Define Objectives and Uses for the Integrated Master Schedule <i>a priori</i>	51
4.4.2	A Bigger Schedule Is Not Necessarily a More Useful Schedule	51

4.4.3	Test Earned Value Management Systems in Phase B for Required Implementation in Phase C	52
4.4.4	Use Standardized Data for Earned Value Management and Cost Accounting	53
4.4.5	Rigorous Baseline Control May Not Be Possible during Early Phases or Rapid Budget Changes	54
4.5	Joint Confidence Level	54
4.5.1	The Current Joint Confidence Level Budget and Schedule Reserve Requirements Are Untenable	55
4.5.2	The Joint Confidence Level Methodology Is a Useful Decision Support Tool	56
4.6	Construction of Facilities	56
4.6.1	Establish Program-level Facilities Expertise during Program Formulation	57
4.6.2	Develop a Facilities Planning Process	58
4.6.3	Accommodate the Unique Risks Inherent in Construction of Research and Development Facilities	59
4.7	Risk Management	60
4.7.1	Ensure Risk Management Is Embedded in the Design Process	61
4.7.2	Good Risk Management Relies on Open Communication	61
4.7.3	Early and Robust Probabilistic Risk Assessment Enables Risk-informed Design	62
4.7.4	Divide and Conquer to Address Complex Risks	63
4.7.5	Link Test Planning to Risk Mitigation	63
4.8	Design and Development	64
4.8.1	Ensure a Common Understanding for Implementation of Human Space Flight System Certification	64
4.8.2	Embed Human-rating Certification in the Core Processes/Tools of the Program	66
4.8.3	Use Risk-informed Design at the Outset of a Project	67
4.8.4	Address Units of Measure Early	67
4.8.5	Design for Operability to Reduce Life-cycle Cost	68
4.8.6	Simulate Processes Early (Run a “Virtual” Mission)	71

4.8.7	Focus the Certificate of Flight Readiness Process	72
4.9	International Partnership Development	73
4.9.1	Engage Potential Partners Early in the Program Life Cycle	74
4.9.2	Develop a Common Understanding of Goals and Objectives to Inform the Architecture	75
4.9.3	Develop an Architectural Framework for Discussion of Partnership Opportunities	77
4.9.4	Build On Current and Past Successes, But Beware of Past Paradigms	79
5.0	Knowledge Capture and Contributors	81
Appendix A. <i>Altair</i> Top Lessons Learned		82
1.	Consider an In-house Skunkworks®-like Team to Initialize New Projects	82
2.	Use Models-based Systems Engineering	82
3.	Start Risk-informed Design during Conceptual Design	83
4.	Involve Industry during Phase A	84
5.	Making Multicenter Teams Work	85
6.	Design during Pre-phase A and Early Phase A	86
Appendix B. Ground Operations Top Lessons Learned		88
1.	Clearly Define Roles and Responsibilities within the Program/Project	88
2.	Streamline Processes to Improve Communication Across Program/Projects	89
3.	Use a Consistent Financial Approach to Ensure Emphasis on Life-cycle Costs	90
4.	Improve Interface Product Integration	92
5.	Address Logistics and Supportability in a Consistent Manner	93
Appendix C. Acronyms		95

PREFACE TO VOLUME II

The prior document (Volume I) provides an executive summary of the lessons learned from the Constellation Program. This companion document, Volume II, expands on this information and provides more detailed analyses for those seeking further background, insight, and information. It contains new subjects that were not covered in Volume I.

In this volume, Section 1.0 introduces our approach in preparing and organizing the content to enable rapid topical location and assimilation of lessons. Section 2.0 describes the contextual framework in which the Constellation Program was formulated and functioned; this is necessary to understand most of the lessons. Section 2.0 supplies an expansion of the briefer contextual description provided in Volume I. As noted in Volume I, context of a former program may seem irrelevant in the heady days of new program formulation. However, readers should take some time to understand the context. Many of the lessons would be different in another context, so readers should reflect on the similarities and differences in their current circumstances. Section 3.0 summarizes key findings, which appear in Volume I, that were developed from significant lessons learned at the program level, and provides a link between the key findings and the topical arrangement of lessons learned in this volume. Readers can use the key findings in Section 3.0 to peruse for particular topics, and will find more supporting detail and analyses in Section 4.0 in a topical format arranged according to traditional NASA program disciplines. The Appendices (A and B) provide insight from the project view of lessons learned.

The reader will no doubt recognize some very similar themes from previous lessons learned, blue-ribbon committee reviews, National Academy reviews, and advisory panel reviews for this and other large-scale human space flight programs; including Apollo, Space Shuttle, Shuttle/*Mir*, and the International Space Station. This could represent an inability to learn lessons from previous generations; however, it is more likely that similar challenges persist in the Agency structure and approach to program formulation, budget advocacy, and management.

Prior lessons learned also drove planning for and formulation of Constellation; the reader may observe attempts, both successful and unsuccessful, to avoid pitfalls encountered by previous programs, and reflect on the outcomes and recommendations.

Perhaps the greatest value of these Constellation lessons learned can be found in viewing them in context with those previous efforts to guide and advise the Agency and its stakeholders.

1.0 INTRODUCTION

“On the actual day of battle, naked truths may be picked up for the asking; by the following morning they have begun to get into their uniforms.”--Gen. Sir Ian Hamilton

These lessons learned are part of a suite of hardware, software, test results, designs, knowledge base, and documentation that comprises the legacy of the Constellation Program.

The context, summary information, and lessons learned are presented in a factual format, as known and described at the time. While our opinions might be discernable in the context, we have avoided all but factually sustainable statements. Statements should not be viewed as being either positive or negative; their value lies in what we did and what we learned that is worthy of passing on.

The lessons include both “dos” and “don’ts.” In many cases, one person’s “do” can be viewed as another person’s “don’t”; therefore, we have attempted to capture both perspectives when applicable and useful.

While Volume I summarizes the views of those who managed the program, *this Volume II* encompasses the views at the working level, describing how the program challenges manifested *in day-to-day activities*. Here we see themes that were perhaps hinted at, but not completely addressed, in Volume I: unintended consequences of policies that worked well at higher levels but lacked proper implementation at the working level; long-term effects of the “generation gap” in human space flight development, the need to demonstrate early successes at the expense of thorough planning, and the consequences of problems and challenges not yet addressed because other problems and challenges were more immediate or manifest.

Not all lessons learned have the benefit of being operationally vetted, since the program was cancelled shortly after Preliminary Design Review. We avoid making statements about operational consequences (with the exception of testing and test flights that did occur), but we do attempt to provide insight into how operational thinking influenced design and testing.

The lessons have been formatted with a description, along with supporting information, a succinct statement of the lesson learned, and recommendations for future programs and projects that may be placed in similar circumstances.

2.0 CONTEXT — “WHERE YOU STAND DEPENDS ON WHERE YOU SIT”

2.1 A Brief Description of the Program

NASA formed the Constellation Program in 2005 to achieve the objectives of maintaining American presence in low-Earth orbit, returning to the moon for purposes of establishing an outpost, and laying the foundation to explore Mars and beyond in the first half of the 21st century. The Constellation Program’s heritage rested on the successes and lessons learned from NASA’s previous human space flight programs: Mercury, Gemini, Apollo, Space Shuttle, and the International Space Station (ISS).

Following the loss of *Columbia*, NASA established the *Columbia* Accident Investigation Board (CAIB) to perform an in-depth review of the Space Shuttle Program. As a result of this review, the CAIB concluded that it was in the best interest of the U.S. to develop a replacement for the Space Shuttle. The CAIB concluded that it should be possible, using past and future investments in technology, to develop the basis for a system “significantly improved over one designed 40 years earlier, for carrying humans to orbit and enabling their work in space.”

In January 2004, The White House issued a new exploration initiative to return humans to the Moon by 2020 in preparation for human exploration of Mars and beyond. As part of this initiative, NASA would continue to use the Space Shuttle to fulfill its obligation to complete assembly of the ISS and then retire the Space Shuttle by 2010. NASA would also build and fly a new Crew Exploration Vehicle (since named *Orion*) by 2014. In 2005, Congress expressly endorsed the President’s exploration initiative and authorized NASA to “...establish a program to develop a sustained human presence on the moon, including a robust precursor program to promote exploration, science, commerce and U.S. preeminence in space, and as a stepping-stone to future exploration of Mars and other destinations.”

In response to the Presidential direction, in 2004 the Agency formed the Exploration Systems Mission Directorate (ESMD) at NASA Headquarters to oversee development of exploration programs. The Constellation Systems Division was initiated to oversee human exploration mission development.

In May 2005, the NASA Administrator commissioned the Exploration Systems Architecture Study (ESAS), a 60-day study to perform the tasks of defining the top-level requirements and configurations for crew and cargo launch systems to support exploration objectives. The study concluded that the launch vehicles should be derived from existing technologies, leveraging the lessons learned from past programs. The ESAS recommended an architecture that formed the basis of the Constellation Program.

Since exploration of the moon and beyond was the overarching goal of the Constellation Program, all elements were designed to perform lunar missions, while also being capable of performing missions to the ISS. The Program was phased as a stepwise capability buildup

largely based on Space Shuttle heritage components. The Initial Capability (IC) comprised elements necessary to service the ISS by 2015 with crew rotations, included the *Orion* Crew Exploration Vehicle, the *Ares I* Crew Launch Vehicle, and the supporting ground and mission infrastructure that would enable these missions. The Constellation Lunar Capability (LC) added the *Ares V* Cargo Launch Vehicle, the *Altair* Lunar Lander, and spacesuits designed for partial-gravity exploration. Lunar outpost elements and capabilities were to follow, including mobility elements such as rovers, permanent or semipermanent habitats, and power and communication elements to support a sustained exploration presence.

The Constellation Program was managed at the Johnson Space Center (JSC) with the following seven “hardware development” project offices:

- Crew Exploration Vehicle Project (the *Orion* spacecraft), at JSC
- Exploration Launch Projects (the *Ares I* and *Ares V* launch vehicles), at the Marshall Space Flight Center (MSFC)
- Ground Operations Project (development of launch and recovery facilities) at the Kennedy Space Center (KSC)
- Mission Operations Project (development of flight and mission control facilities) at JSC
- Extravehicular Activity (EVA) Systems Project (spacesuit and tool development) at JSC
- Lunar Lander Project (the *Altair* lunar descent and ascent module) at JSC
- Lunar Surface Systems Pre-project (advanced planning for lunar habitats and surface support systems) at JSC

Orion, *Ares I*, Ground Operations, Mission Operations, and EVA comprised the IC projects. *Ares V*, the *Altair* and Lunar Surface Systems, and upgrades to the IC to meet lunar interfaces and requirements comprised the LC projects.

By 2010, all elements of the IC had achieved Preliminary Design Review (PDR), and test flights were under way, with the successful launches of *Ares I-X* and the *Orion* Pad Abort-1 tests. The White House announced cancellation of the program on Feb. 1, 2010, on the seventh anniversary of the loss of *Columbia* and her crew.

According to the NASA Authorization Act of 2010 (PL 111-267, Oct. 11, 2010), elements of the Constellation Program will continue in NASA’s current planning. The *Orion* vehicle is redesigned as the multipurpose crewed vehicle (MPCV) and will likely enable crew launch on a variety of launch vehicles. The Space Launch System (SLS), intended as the vehicle to launch crews and cargo beyond low-Earth orbit (LEO), has various configurations under consideration including options based on Space Shuttle and *Ares* heritage designs and hardware. The Ground Operations infrastructure at the Kennedy Space Center, developed for Constellation, will be adapted to the SLS and new commercial launch capabilities.

2.2 The Program Environment — Key to Understanding the Lessons Learned

2.2.1 Program Scope


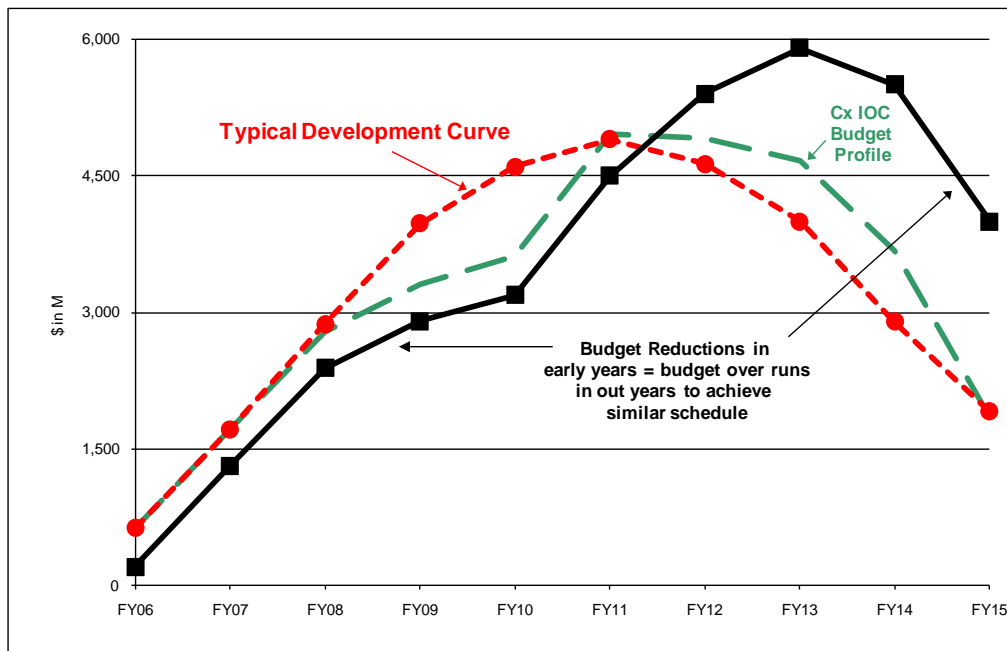
The Constellation Program was conceived as a multi-decadal undertaking, with the primary goal of enabling human exploration beyond LEO. It was to serve as the Space Shuttle successor not only in function, but also in utilization of facilities and workforce. It was to provide the primary mode of delivery of crews to the ISS, enable the construction and operation of an outpost on the moon, and enable human exploration of Mars. As the Agency flagship program, the effort spanned all 10 NASA centers and multiple large-scale acquisitions. It required modernizing an infrastructure designed and sized largely for the Apollo Program in the 1960s. Moreover, it required leading a workforce generationally removed from the previous human spacecraft launch and entry development challenges.

While the near-term focus of the program was providing crew transport to the ISS via the IC, the capability to go beyond LEO drove many technical and programmatic decisions. These decisions were suboptimal when viewed with only near-term objectives in mind, while they were necessary to support the LC phases of the program. They, in some cases, also increased risk to the IC through incorporation of newer technologies with less flight heritage. The need for commonality to decrease life-cycle cost and complexity drove many decisions. For example, the Ares I design used a five-segment, solid-chemical first stage along with the J2-X upper-stage engine. Other simpler, less-costly designs could have been used to optimize for Ares I alone, but in context with the Ares V development, these elements lowered the overall program design complexity, risk, and cost.

Optimization for the long-term program drove not only design decisions, but also infrastructure needs, standards, organizational structure, units of measure (SI [international system of units]), data architecture, etc., that would not have been necessary for the IC alone. The primary objective of stepping beyond LEO with the LC was optimized at the expense of having an IC that was more costly and complex than necessary as a stand-alone development.

2.2.2 Program Funding

Funding for the Constellation Program was inconsistent and unreliable from its initial formulation through its cancellation. Indeed, this inconsistency became the chronic constraint that aggravated all other constraints. Figure 1 illustrates this persistent problem. The program experienced a 10% budget cut prior to PDR. In addition, the funding “profile” (how the money is allocated on a yearly basis) was far less than optimal for a development project. Typical development projects swell in funding as they enter the design development phases. Constellation had to adjust its work to accommodate a funding “notch” in fiscal year (FY)10 as a result of Agency funding limitations due to Space Shuttle retirement. While Constellation was expected to transport crews to and from the ISS soon after Space Shuttle retirement, the funding ramp-up needed for development was not available until *after* the Space Shuttle retired. This, of course, put pressure on program content, schedule, and rationale for the ISS crew rotation mission (Initial Operational Capability [IOC]).



RYS in M	FY06	FY07	FY08	FY09	FY10	FY11	FY12	FY13	FY14	FY15	FY16	FY17	FY18	FY19	FY20	Total
ESAS	2,333	2,976	3,195	4,202	5,034	7,273	7,994	8,324	8,558	8,948	9,547	11,107	11,422	11,744	12,073	114730
PMR09	1707	1779	2514	3085	3454	6085	6346	6991	7145	7856	9145	9983	10294	10582	10854	97820
2010 PBS	1707	1779	2514	3085	3454	5524	5444	5376	5570	6153	6161	6170	6178	6186	6195	71496
PMR09 vs. ESAS Delta	(626)	(1197)	(681)	(1117)	(1580)	(1188)	(1648)	(1333)	(1413)	(1092)	(402)	(1124)	(1128)	(1162)	(1219)	(16910)
Passback vs. ESAS Delta	(626)	(1197)	(681)	(1117)	(1580)	(1749)	(2550)	(2948)	(2988)	(2795)	(3386)	(4937)	(5244)	(5558)	(5878)	(43234)

Fig. 1: Constellation budget profiles before and after budget reductions (ESAS=Exploration System Architecture Study, the initial program budget baseline; PMR=Program Manager’s Recommendation based on “marks” passed down by the Office of Management and Budget [OMB] and NASA Headquarters [HQ]; PBS=President’s Budget Submittal)

Figure 2 illustrates the result of budget pressure on major schedule milestones (PDR, Critical Design Review [CDR], IOC, and human lunar return [HLR]). IOC moved 2 years (2012 to 2014) due to the first budget cut. The program was able to recapture 1 year of that loss by replanning (at greater risk). Subsequent budget cuts pushed IOC further, into 2015. The program was required to hold the IOC schedule commitment to 2015, since this was the “no later than” date that had been committed to the U.S. Congress.




Fig. 2: Consequences of funding cuts on Constellation major milestones

The phasing of the program exacerbated the funding challenges, in a “rob Peter to pay Paul” fashion. The viability of the lunar program phase was under constant threat, as it was the only source to supplement shortfalls in the IC when schedule slips had been exhausted (see effect of the fixed-base costs, below). The lunar phase of the program (e.g., HLR in Fig. 2) bore the brunt of the cuts that could be accommodated in out-years and was pushed even further than the IOC milestones. This is illustrated in Fig. 2 (observe that IOC was “held” in 2015 while HLR slipped almost 2 years).

Constellation was planned with adequate budget reserves, although these reserves were not phased to be most useful to the program (i.e., much of the reserve was phased later in the program schedule rather than earlier). An unfortunately prevalent perspective on program reserves in the federal government is that reserves are seen as “over-funding” that can be depleted to address other issues, or drawn back to the Treasury if unused within an FY. Agency reserves (i.e., allowance for program adjustment [APA] funds) had been eliminated long ago. Constellation reserves were initially planned to fund risk mitigation efforts. Reserves were rapidly depleted due to the combined effects of budget cuts, Agency “taxes,” and efforts to hold schedule without sacrificing content. Subsequently, risk mitigations were reduced and then eliminated, leaving little recourse to resolve the typical technical challenges that crop up in a design of this magnitude and complexity. Risk accrued as a result.

The large acquisitions for design and construction of the major components of the program were planned, sized, and budgeted early in the program, as necessitated to meet schedule objectives.

The prime contractors' costs typically ramped up in accordance with plan. These costs were comprised of both "fixed-base" elements (work force, test facilities, etc.), and variable elements. As the program dealt with subsequent budget cuts, the fixed-base elements were extremely difficult to reduce; cuts tended to be absorbed within variable elements (e.g., deferred hardware buys). This led to a number of problems: schedule risk increased as more activities crept onto the critical path; planned activities were deferred to reduce risk, thus delaying risk mitigation strategies and increasing overall risk; we faced an exercise in diminished returns, in which fewer and fewer items remained in the variable elements of the budget to defer to the next round of reductions.

With the retirement of the Space Shuttle Program looming, there was great pressure to assimilate as much of the shuttle-heritage "fixed-base" (e.g., work force, infrastructure, and industrial base) from that program as possible, per direction NASA received in two Authorization Acts (2005, 2008). While the operational requirements and sustaining budget planned for Constellation were smaller than those of the Space Shuttle Program, nearly every component, building, fleet, process, and cadre of workers had its advocacy group to sustain it. Because of the phasing of the Constellation Program, much of the infrastructure was not needed for the IC but would be needed for the LC. This created an Agency-wide phasing challenge for these assets: Should NASA retire/mothball the asset, or sustain it for many years unused?

The tight coupling of the Space Shuttle retirement to the buildup of the Constellation IC, along with the necessity to plan the IC and the LC with a large degree of commonality, made sense strategically when the plan was developed, but only if all schedule and funding assumptions held. The rationale did not withstand subsequent budget pressure and consequential schedule slips.

The Constellation Program hardware development budget was also, in an absolute sense, large at the Agency level and not yet operational (as Space Shuttle and ISS were), so it was often perceived as the only source for unexpected expenses that cropped up in the Agency. Indeed, expenses were charged to the program that did not benefit the program. This is not atypical for flagship programs in governmental developments.

Cash flow disruptions, resulting from Congressional continuing resolutions (CRs), had been an occasional source of funding problems for the Agency in the past. These became annual events during Constellation's lifetime, and many times lasted well into the new FY. While these are experienced government-wide, they are particularly onerous on developmental programs in early stages when funding levels are increasing toward their peak. A typical program "ramp-up" is nearly impossible during a CR. These funding shortfalls, which sent annual "square-waves" of additional replanning through the program, projects, and contractor communities, were never fully attenuated through the remainder of the FY.

The Constellation Program content, constituted by the elements of the program architecture, along with their associated performance requirements were held as a higher priority than schedule performance by the Agency leadership; therefore, in the budget environment framed above,

schedules slipped, risks accrued, and reserves were depleted before serious consideration was given to descope program content.

2.2.3 The Mosaic of NASA Participation

As Constellation was the flagship program for the Agency, project and program work was distributed across the Agency to support the concept of maintaining “10 Healthy Centers” (Fig. 3). This distribution was planned at the Agency level. The concept was not sufficiently defined to satisfy either centers or the program in light of the Space Shuttle retirement. Details of the implementation were left to the program to sort out in consultation with center management. This resulted in a persistent tension between what was most efficient for the program vs. what was best for a particular NASA center to sustain or grow its current role. The nationwide team was probably larger than necessary, but it also covered “unused capacity” at some centers. The large teams at multiple locations led to instances of overlapping or unclear roles and responsibilities. Decisions taken to perform tasks “in-house” using an on-site work force at a center vs. contracting the task out were often made at the center or project level, and were not optimized for overall program benefit.



Fig. 3: The Constellation distributed team

The involvement of all 10 NASA centers exposed cultural differences that had to be managed by a management team whose members were geographically removed. Each center had different requirements, procedures, and processes that were documented at various levels of

detail and constraint. While the apportionment of responsibility was useful to each center, the Agency had not completed effort to assure that the requirements, procedures, and processes were consistent, non-duplicative, and noncontradictory. This placed a particular burden on the contractor community since it was charged with identification and compliance with duplicative and sometimes contradictory requirements. The mismatch of design and construction standards among centers, along with demands from the “insight” work force at each center, was a worrisome burden for contractors, constantly threatening their ability to meet schedule milestones and cost thresholds. The NASA use of an independent technical authority provided yet another source of direction and, thus confusion, to the contractor community. Although official direction can only flow to contractors via the Contracting Officer (CO), informal discussions and decisions coming from the program and separately from the technical authority sometimes made it difficult for contractors to distinguish a final decision from an interim decision.

The industrial base used by the Space Shuttle Program and, to a lesser extent, the ISS Program was managed by a separate mission directorate at NASA HQ. Overlapping strategic interests among the three programs (e.g., facility usage, floor space, test stands) had to be resolved at the highest levels in the Agency (between mission directorates).

The position of an Agency flagship program carried benefits as well. The program had access to the entire Agency knowledge base, along with skills and infrastructure. The program was viewed as a high priority at all centers, since it represented a path to a sustainable and growing role in most areas of expertise. Engagement of nontraditional human space flight centers allowed the program to tap into unique skills, approaches, and facilities. Contractors benefited from NASA technical depth and expertise, along with access to unique facilities (e.g., engine test stands, large vacuum chambers). The early test flights, *Ares I-X* and *Orion Pad Abort-1*, demonstrated the benefits of multicenter teams. These projects were successful early pathfinders in identifying and resolving the problems that could arise from a conflict of cultures in design, test, and manufacturing.

2.2.4 The Generation Gap in Human Space Flight

The work force that supported Constellation was not only geographically distributed; it was also bifurcated in crucial experience level. Since NASA human space flight developments have occurred on 20- to 30-year cycles, most of the Constellation work force was generationally removed from the previous human spacecraft development, particularly the launch and entry development challenges. The Space Shuttle design was performed in the mid 1970s and tested in the late 1970s and early 1980s, prior to the birth of many of the engineers working on Constellation. ISS design and development (which relied on the Space Shuttle for launch and entry of crew and equipment) took place in the late 1980s and into the mid 1990s. The Constellation Program was replete with ISS experience; however, only senior engineers and leaders in the Constellation Program had been involved in the early development stages and knew the pitfalls and travails in starting up a national-scale program. Most of the work force was

well versed in operating human spacecraft either via Shuttle, the ISS, or both, but only partially aware of how early design compromises affected those very operations.

The Agency took extraordinary steps to deepen the design skills of the work force. The critical thinking necessary to perform systems engineering (SE) design—which was emphasized in formal and informal training, lessons-learned briefings, leadership postings, work force discussions, and work rotations—became a focus at the level of the NASA Administrator. Opportunities for testing smaller development projects early were created to enhance work force experience.

To reach back and capture launch and return vehicle development experience, the Constellation Program created a ready resource known as SAGES [Shuttle and Apollo Generation Expert Services]. This contract provided a simple pathway by which to enlist the aid of retired experts from the NASA past (e.g., George Mueller, Chris Kraft, Glynn Lunney, Dick Kohrs, etc.). Beyond review and advice, SAGES was created to transfer knowledge through mentoring. It was based on relationships between Apollo- and Shuttle-era program managers and discipline experts and the Constellation team. It provided mentors on an as-needed, targeted basis in areas ranging from technical design and analysis disciplines, to ground and flight operations, and to program management.

A perspective drawn from past programs is invaluable in many ways. To individuals, it provides resiliency in the face of challenge, particularly if the challenge has been met before, albeit under different circumstances. To organizations, it lends strength and stability to the organization by having key personnel with this type of resiliency willing to mentor those who do not.

2.2.5 The Size of the Trade Space

An inherent characteristic of the early design phases in a large-scale development is the scope of the “trade space” within which designers and managers can make choices. Fundamental choices at this stage have very long-term consequences. Choices of fuels, engines, outer mold-lines, volumes, materials, etc. have a vast number and sequence of “touch points” in an existing and future infrastructure. In the case of Constellation, these touch points ranged from ground infrastructure already in place from earlier programs, to the lunar and martian surfaces. During maturation of a program, a trade space stabilizes and shrinks as choices are made and designs solidify. As we approach testing and operations, that trade space narrows. While new design work can occur in mature, operational programs, akin to those of ISS or Shuttle, the trade space is typically very limited and choices are few.

Due to the scope of the trade space for Constellation, early choices and decisions were very often brought into question when new knowledge was introduced, as new or overlooked touch points were discovered, or when budget and schedule constraints limited preferred choices. For a work force most adapted to the very limited trade space available in operational programs such as Shuttle and ISS, this cycle of working through options (and changing options as constraints change) can look like churn, indecision, and lack of leadership. And for a leadership team most focussed on socialization of major decisions with Agency management, insufficient attention to this perception leads to frustration at the working level.

This difference in perception is manifest in many of the lessons learned that follow. Here are a few perspectives that illustrate these contrasting points of view. While some may be in opposition, each view is relevant because it affected daily operations and interactions in the work force.

On decision making:

Younger engineer: *“This [changing of decisions] could have been avoided by having a daily tag-up [meeting] as was done....on the ISS Program.”*

Older engineer: *“ISS took much longer to get to PDR....with lots of upheaval getting there. It wasn’t until around CDR that daily tags chaired by a strong leader were added to improve progress and accountability.”*

Younger engineer: *“The program revisits the same decisions over and over. There seems to be no accountability for sticking with decisions that have been made.”*

Older engineer: *“[Space Station] Freedom had three program managers in 4 years. The architecture was constantly changing, from power-tower to a single truss; from solar dynamic power to combined SD [solar dynamics] and PV [photovoltaic] power and back again to SD and back to finally PV. Then segments were cut from the truss, creating the shadowing that we deal with operationally today. And, it took years to get to the 20-kHz power architecture.”*

2.2.6 Phasing of Project/Program Start-up

Two of the primary projects, the *Orion* Crew Exploration Vehicle and the *Ares* Crew Launch Vehicle, were begun well in advance of the Constellation Program.¹ This enabled rapid start-up and early results from each of these projects. They functioned for some time without an integration function between them, which led to an eventual Program Integration (PI) function that was very lean by historical norms.² This rapid start-up enabled rapid contract competition. However, these contracts were not fully scoped for the physical and analytical integration that was necessary. They were focussed on the IC, since trade studies had not been performed that would optimize an integrated architecture. These factors led to costly contract changes when the appropriate level of integration for a complete architecture was better understood.

This approach also provided an additional integration challenge; namely, the projects had formed requirements, budgets, design approaches, acquisition strategies, etc., that needed to be integrated retroactively. The late addition of integrating functions had a long-term effect that was not fully resolved by the end of the program. Strategies that were in a collective self-interest for the program were extraordinarily difficult to implement due to the cost of contract changes.

¹ The most logical approach to starting up a new program is to first understand the “big picture” before decomposing that end state into smaller and smaller elements. This is why it is ideal to first start up a program office with its integration products, then to start up project offices to accomplish the hardware development. This was not the case in the Constellation Program.

² A typical large space flight program historically uses approximately 10% to 15% of its budget for PI. The Constellation Program was relatively lean, spending 5% to 7.5% on these functions.

Examples of these include unit policy (SI vs. heritage English units of design and construction), data architecture, and facilities management. Boards and panels had been formed prior to the program office, which meant these boards and panels had different structures and responsibilities. Integration of these into a program board and panel structure confounded decision making. Moreover, integrated analyses by the program uncovered technical issues that could have been caught earlier (and been resolved in a more efficient manner) if the start-up timing had been different.³

Incentives on contracts were developed primarily to benefit each project rather than the overall program. Missing were incentives for contractors to participate in technical solutions that benefit the integrated program. This drove results that were less than optimum, and sometimes not mutually compatible.

2.2.7 Starting Well vs. Demonstrating Early Success

Many of the following detailed lessons learned describe the repercussions of starting projects before the program is started. The start-up phasing described above was primarily driven by the Agency's need to demonstrate early successes to secure political support.

Experiences from the early stages of ISS may also have contributed to the push for early success. The ISS Program had its roots in the prior Space Station *Freedom* Program, a program that suffered "fits and starts" in the effort to "start well" and proceed to PDR as it underwent several major descopes and redesigns. While it formed in a different political and programmatic environment, it experienced many of the same funding and work force challenges as the Constellation Program did. Consolidation, under the ISS Program in the mid 1990s, alleviated some but not all of the start-up issues.

Much of this was in the minds of the Constellation organizers, who believed that the political will would not be sustained through a *Freedom*-like start-up. Again, we see the perspectives from the work force on this issue:

Program office integration engineer: *"The projects started way ahead of us, so we were always in catch-up mode. They had prime contracts on board, so every integration issue we found was a real headache."*

Project office engineer 1: *"We had to get up and running fast and get our test plans in place. We needed to have some early test successes to keep support for the project. The program office added layers of requirements and reviews. They just didn't get that we were moving and couldn't make changes without big cost and schedule hits."*

³ An example of this was what became known as the "thrust oscillation" issue. Integrated analyses indicated that unacceptably large thrust oscillations from the *Ares I* first-stage burn could be transmitted through the stack to the flight crew in the *Orion* vehicle during launch. Design changes were implemented in the PDR time frame to mechanically dampen the oscillations and sufficiently isolate the crew. We can speculate that earlier identification could have resulted in simpler, lower cost changes.

Project office engineer 2: *“The program was trying to get in our knickers. They wanted to tell us how to design the vehicle. We had a schedule to meet, so this didn’t go down too well.”*

Ground operations engineer: *“We were at the tail end. Projects made changes and didn’t tell us. They didn’t know the change would impact the ground ops hardware. It cost us big time to make late changes.”*

Program control analyst: *“The project procurements were way ahead of the program, and so changes cost big bucks. We didn’t have much flexibility with reserve dollars, but every change we processed to integrate the project hardware and software came down to painful cost and schedule choices.”*

3.0 KEY FINDINGS

Given the challenges described in Section 2.0, the following is a brief listing of the key findings distilled from the Constellation Program. Explanatory information, as well as further detail, can be found in Section 4.0.

The character of these lessons is worth noting. The reader will observe that very few technical issues arose that could not be solved in relatively short order; indeed, the most difficult and most persistent challenges involved cost, schedule, communication, and organization. The reader will no doubt recognize this pattern from previous large-program lessons learned. While the Agency is renowned for technical prowess, senior managers in flagship programs can be faced with a multitude of nontechnical challenges for which they have far less training or preparation. In this respect, lessons learned can be an invaluable source of insight.

3.1 Robust vs. Optimal Planning — The Only Certainty Is that the Funding Will Not Match the Plan

Driving Events

Section 2.1 describes the context and driving events for the funding and planning challenges faced by the Constellation Program. Constellation found itself in a grounds-up re-plan at least annually.

Lessons Learned

The reality is that Agency flagship programs like Constellation must be robustly planned (e.g., “elastic”) vs. optimally planned (“inelastic”). Absent a national imperative akin to the race to the moon, funding will not arrive as planned. Flagship programs are highly visible and stand alone in the national-level budget debates. They are not part of a budget portfolio and, therefore, cannot look elsewhere for funding.

A flagship program will never experience complete alignment of the budget with the plan, either in net or in phasing. Consequently, flagship programs should resist the urge to optimally plan each budget line item.

Recommendations

Plan for Continuing Resolutions CRs for the first quarter of each FY until and unless Congressional processes change substantially.

Decouple programs/projects as much as possible in strategic, programmatic, and technical aspects.

Proactively determine sensitivities and breakpoints in budgets, and communicate these to stakeholders *a priori*. Establish expectations for funding reductions in advance. This is the best hope to influence budget decisions.

Develop scenario(s) to accommodate 5% to 10% funding shortfalls in any given FY. Understand the major drivers to cost and options available (i.e., “knobs to turn”) when unexpected challenges appear.

Supporting Lessons Learned from Section 4.0: 4.2.4, 4.2.5, 4.5.1, 4.9.4

3.2 Schedule Creep and the Fixed Base — The Law of Diminishing Returns

Driving Events

As discussed in Section 2.2, the Constellation Program dealt with funding cuts primarily with schedule slips rather than content reductions (due to NASA Administrator direction). Variable costs are early targets, since they can be deferred (while adding program risk). The fixed-base costs — both within the Agency and for the contractors — become proportionally larger, making schedule slips less and less effective in addressing cost cuts.

Lessons Learned

Each schedule slip resulted in longer periods for which Constellation was expected to maintain the Agency’s Human Space Flight “fixed base” along with the industrial fixed base from its suite of acquisitions. This had the effect of reducing the overall cost savings being sought from the initiating event — the funding cut. A vicious cycle was established that only grew more severe during the life of the program as more and more time supporting the fixed base was added to the program lifetime. Ultimately, there could be no resolution without addressing content reductions.

Recommendations

Understand the inherent limitations of schedule slippage to resolve funding shortfalls, since accrual of the fixed-base portion of the budget erodes the buying power in the out-year funding. Reduction of program content warrants consideration under these circumstances.

Supporting Lessons Learned from Section 4.0: 4.2.4, 4.2.5

3.3 Tailoring and the Design and Construction Standards — Drinking from a Fire Hose

Driving Events

The involvement of 10 NASA centers (described in Section 2.2.3, Mosaic of NASA Participation), along with the overabundance and overly prescriptive nature of design and construction standards led to an avalanche of “shall” requirements on the Constellation contractors. While Constellation was allowed broad latitude to tailor requirements, tailoring required rationale and negotiation with the requirement “holder.” The sheer number of direct and embedded requirements made the task of tailoring intractable.

The following illustrates the extent of the problem: The *Orion* prime contract alone levied design and construction (D&C) specifications from: 34 NASA standards, 10 Johnson (Space Center) Program Requirement (JPR) documents, 11 JSC technical standards, 3 NASA Program Directives (NPDs), 6 NASA Program Requirements (NPRs), and 11 KSC standards. This included the following particularly prescriptive examples:

- NASA STD 8739.3, Soldered Electrical Connections – a 93-page document with 391 “shall” statements inclusive of room temperature/relative humidity and 1077 lumens of light per square meter or 100 foot candles on the work surface.
- NASA STD 8739.4, Crimping/Cables/Harnesses/Wiring – a 66-page document containing 391 “shall” statements requiring (a) a Snell far-vision chart 20/50; (b) a near-vision Jaeger1 at 355.6mm; and (c) reduced Snell 20/20 or equivalent. Requires color vision testing and specifies temperature/humidity and lighting.
- NASA STD 8719.9, Standard for Lifting Devices – a 126-page document containing 996 “shall” statements covering the use of overhead devices as well as powered and motorized devices; approves the use of manufacturer’s procedures and instructs on how to set the parking brakes on mobile and motorized units.

All of the major aerospace contractors, who build systems for multiple government customers (Department of Defense [DoD], National Security Agency [NSA], NASA, etc.), have certified internal processes and suppliers that meet both national and international standards. The set of NASA requirements is often redundant to this.

Lessons Learned

The massive quantity of D&C standards defies tailoring for even a large program such as Constellation. Major aerospace contractors have sufficient processes such that the NASA D&C standards are of questionable value. This is a major driver of program fixed costs.

Recommendations

NASA should assume the burden of “meets or exceeds” evaluations instead of making the contractor prove that it meets NASA requirements. NASA should implement this by establishing an experienced team to audit contractor internal processes at the NASA document level. This process should not burden the contractor with defending the adequacy of each individual requirement. For the most part, the aerospace contractor community is well versed in the appropriate manufacturing and safety measures required to develop space hardware. NASA should build on that expertise.

Currently, this topic (Affordability/Innovation) is the subject of discussion among the Chief Engineer Office, the Chief Health and Medical Office, and Safety and Mission Assurance (S&MA) leadership. A team has been tasked with describing a process and timeline for transforming the Agency approach to Technical Authority requirements and standards. This initiative should establish clear criteria for mandatory (“shall”) requirements. It should also revalidate the current Agency requirements base with the goal of significantly reducing the total

number. Most requirements should become guidelines or best practices, and should take into account the pros and cons of the DoD experience with reduction of requirements.

Supporting Lessons Learned from Section 4.0: 4.2.4, 4.2.5, 4.2.6

3.4 Tailoring Process Simplification — The Law of Unexpected Consequences

Driving Events

One of the many findings of the CAIB was that NASA had institutionalized the waiver and deviation process for requirements such that permanent or semipermanent waivers and deviations to requirements were the accepted norm. The CAIB, as well as other committees and panels studying the NASA human space flight culture, questioned the quality and validity of requirements that need permanent waivers or deviations to remain extant. Consequently, the need to establish a waiver or deviation developed an understandable cultural taint in the work force. Training after the *Columbia* accident reenforced this aversion. Indeed, important NASA stakeholders (e.g., Congress, the Aerospace Safety Advisory Panel, members of the public) were keen to maintain visibility on the NASA use of waivers.

Meanwhile, the Agency did not grapple with the plethora of requirements that drove the practical need for waivers and deviations to requirements. While the Agency approach to requirements was redesigned with the introduction of the independent Technical Authority, the actual number and scope of requirements was not addressed.

This left the Constellation Program in a quandary. It was allowed broad latitude in tailoring requirements to its needs. However, the Agency decided that the program would use deviations and waivers to document tailoring outcomes. Cultural consequences were not evaluated. Moreover, the process for permanent deviations and waivers was formalized by the Agency. Discussions of tailoring were driven down to the level at which they were most productive — to those most familiar with design and implementation; however, waivers and deviations required higher visibility within the Agency to agree on. This cultural inconsistency, coupled with the sheer number of requirements addressed in the previous lesson, made the tailoring process nearly insurmountable.

Lessons Learned

While the use of waivers and deviations was intended to simplify the tailoring process of requirements for the Constellation Program, cultural consequences were not addressed. This had the unintended consequence of confounding the tailoring process rather than assisting it.

Recommendations

Tailoring of requirements is necessary and should not carry a cultural stigma. The process should be revised to remedy the negative implication.

Supporting Lessons Learned from Section 4.0: 4.2.3

3.5 Risk-informed Design — Risk as a Commodity

Driving Events

The broad scope of the Constellation Program permitted some greatly effective experimentation in the design process. The traditional aerospace design process (i.e., “rule-based” approach), wherein the vehicle is designed strictly to requirements (and, by implication, all requirements are equal or nearly so), was used on most elements. Progress in analytical techniques allowed the application of reliability to supersede redundancy in design, strengthening the use of design for minimum risk (DFMR). Two elements used a risk-based approach that resulted in overall better designs.

The *Orion* vehicle was designed initially using the rule-based approach. In meeting all of the requirements, the vehicle was overly constrained and too heavy, and it also had performance challenges. To “close the architecture,” the project implemented a revised design process that prioritized mission-critical, reliability, and safety requirements. A “zero-based” vehicle was designed that met only these requirements. All other design features required justification based on ability to reduce risk, i.e., risk as a commodity, using standard risk assessment tools (e.g., probabilistic risk assessment, hazard analysis, and failure modes and effects analysis) along with the beneficial use of other design commodities (e.g., power, mass, budget). The ability of the *Orion* vehicle to minimize LOC (loss of crew) probability, along with meeting of performance and mass constraints, was improved through application of this process.

The *Altair* vehicle pioneered a similar approach, uniquely applied *a priori*. The Risk Informed Design process used an analysis-based approach to design the absolute minimum vehicle necessary to perform the mission and then added design features by proven value in increasing vehicle reliability and safety. This resulted in more robust design solutions that either eliminated or better controlled hazards.

Lessons Learned

Risk Informed Design (i.e., treating risk as a commodity in design) was a breakthrough for both *Orion* and *Altair*, resulting in designs that “closed” and were safer and more reliable than a rule-based design approach could yield.

Recommendations

A rule-based design approach can remove accountability to understand the implications of design choices. Employ risk-informed design methodology *a priori* to focus design on overall mission success rather than compliance with “design rules.”

Supporting Lessons Learned from Section 4.0: 4.8.1, A.3

3.6 In-house Tasks — Sustaining the NASA Institutional Base vs. Affordably Supporting the Programs — Getting from “or” to “and”

Driving Events

Driving events for this lesson are described in Section 2.2.3 under “Mosaic of NASA Participation.”

While the Constellation Program attempted to distribute work in logical packages, the competing needs of “10 Healthy Centers,” the scope of the program, and phasing considerations (e.g., placement of work in the post-shuttle era) drove instances of incoherent distribution, resulting in muddling of roles, responsibility, and authority.

An example of the *Orion* Launch Abort System (LAS) development further illustrates the wide distribution of roles and responsibilities. This development was a subsystem of *Orion* and is not an extreme example, but it is typical of the rest of the program. Primary responsibilities were divided as follows:

- JSC: *Orion* project management, lead for crew module and vehicle integration, prime contractor oversight and independent analysis
- MSFC: LAS prime contractor oversight and independent analysis
- Langley Research Center: Lead for LAS integration, prime contractor oversight, and independent analysis
- Glenn Research Center: Flight Test Article and pathfinder production for Service Module (SM) and Spacecraft Adapter
- Dryden Flight Research Center: Lead for Abort Flight Test integration and operations

When everybody is responsible for everything, nobody is responsible for anything.

Furthermore, the relatively low cost of most in-house tasks can be deceiving. While sometimes low in cost, the price of confused roles, responsibility, and authority (RR&A) can lead to major headaches, and can flow into the contractor infrastructure. These roles are often coupled with the desire for modernization of associated host-center facilities, which can drive larger costs.

Lessons Learned

Task assignments for in-house work are strategic drivers within the Agency; they are more than a series of make-buy or in-house vs. contractor decisions. Clear RR&As must be a priority. In spite of the relative low cost, the distribution and RR&A for in-house tasks are deserving of management focus to avoid confusion and overlapping roles down the line.

Recommendations

Target in-house tasks for items that will mutually benefit programs (e.g., unique expertise) and centers (e.g., institutional sustainment). Due diligence will be needed to achieve this balance of benefit to the program *and* benefit to the NASA institution. Achieving balance will bring value to

the program as well as provide strategic support for the NASA institution, equipping the Agency to accomplish future programs and missions.

Affordability, institutional benefit, and coherence of RR&A should be three additional strategic considerations in these decisions. Or, simply stated, “Make sure the gain is worth the pain.”

Supporting Lessons Learned from Section 4.0: A.5

3.7 Roles, Responsibility, and Authority — a Non-thermodynamic Application of Entropy

Driving Events

The clarity of RR&A for the Constellation Program was degraded by the combined effects of the wide distribution of program responsibilities via the “10 Healthy Centers” policy, multi-decadal phasing of program development in IC and LC, and assumption of traditionally understood roles from the space shuttle heritage component development as described in Section 2.2.

The Program undertook a number of steps to improve this, including conducting work force surveys and holding regular management retreats. Perhaps the most successful effort was the Program Excellence Team (PET), which was a council of program deputies who frequently met face-to-face and via teleconferencing to work through specific issues related to RR&A.

Lessons Learned

There is no formula or checklist for clear RR&A in an Agency-wide, flagship program. The specific approach on organizational structure taken by Constellation in this regard is addressed in Appendix A, in the paper on formulation of the program. Projects housed at centers where previous or similar work took or takes place will tend to assume these heritage roles, in spite of contrary direction.

Two early test flights provide examples of how strong leadership and meeting imminent schedule milestones can drive improvements in RR&A and decision-making.⁴ The *Ares I-X* test flight and the *Orion* Pad Abort-1 test flight were both pathfinders in how to function and succeed within the broad scope and far-flung Constellation organization. Both were developed as the program relationships were forming, using distributed teams and hardware developments. While they were initially hampered by the same factors described in this section, these teams were able to focus on the most important elements and succeed in meeting all of their test flight objectives.

RR&A can be improved by functional examination, by either combining like tasks or separating functions by need. During development of the Integrated Master Schedule (IMS), implementation was hindered by the conflicting needs of the product/task providers and the schedule assessment

⁴ See also Lesson Learned 4.11 for more benefits of test flights.

team. When these functions were separated, the IMS development quickly gained focus and results improved.

Risk management was challenged because it was separated functionally from cost threat management. When broader aspects of risk (e.g., cost threats, safety, etc.) were included, the results from the system improved.⁵

RR&A for the IC was relatively clear and in daily operation. Roles for the LC, while defined at a top level, were best developed for projects that had “follow-on” content from the IC.

The program was also contemplating the roles the International Partners might take. This formulation is described in Section 4.9.

Recommendations

While it is a good start to define RR&A, periodic examination is required for large-scope, distributed programs. Attention to the effects of heritage RR&A and phasing are particularly important. RR&A should also be planned to adapt to the program phase(s).

Creating a dedicated team of decision makers (e.g., the council of deputies) to critically review and address RR&A can be very effective.

Use of annual employee surveys provides a helpful metric toward gauging the effectiveness of measures taken to address RR&A issues.

Supporting Lessons Learned from Section 4.0: 4.1.1, 4.2.6, 4.4.1, 4.4.2, B.1

3.8 Decision Making — Only as Efficient as Roles, Responsibilities, and Authority Are Clear and Understood

Driving Events

The clarity and effectiveness of the decision-making processes for the Constellation Program were driven by the same events as summarized above in Section 3.7 and as described in more detail in Section 2.2.

The start-up phasing (late start of the program relative to the projects) previously described also confounded decision making. Projects made unilateral changes without considering integrated effects. For example, *Ares* made changes to the “stack height” that impacted the *Orion* and Ground Operations designs; *Orion* deleted a test flight it found unnecessary without understanding the integrated effects. Once the program integration function was operational, decisions were better coordinated but opportunities for early-phase, lower-cost changes were lost.

⁵ NB: Risk management systems quickly degrade to “issues tracking” unless adequate reserves are maintained to address and retire risks at the appropriate time. Lacking this, risks become accepted *de facto*.

The mature decision-making structure affecting Constellation Program decisions is partially illustrated in Fig. 4. Not shown are other organizations having a significant influence on the decisions made: center institutional processes (e.g., Engineering Review Boards, Program Management Councils, facility boards), the Technical Authority (which provides the appeals route for technical decisions), or the internal decision processes for the prime contractors.

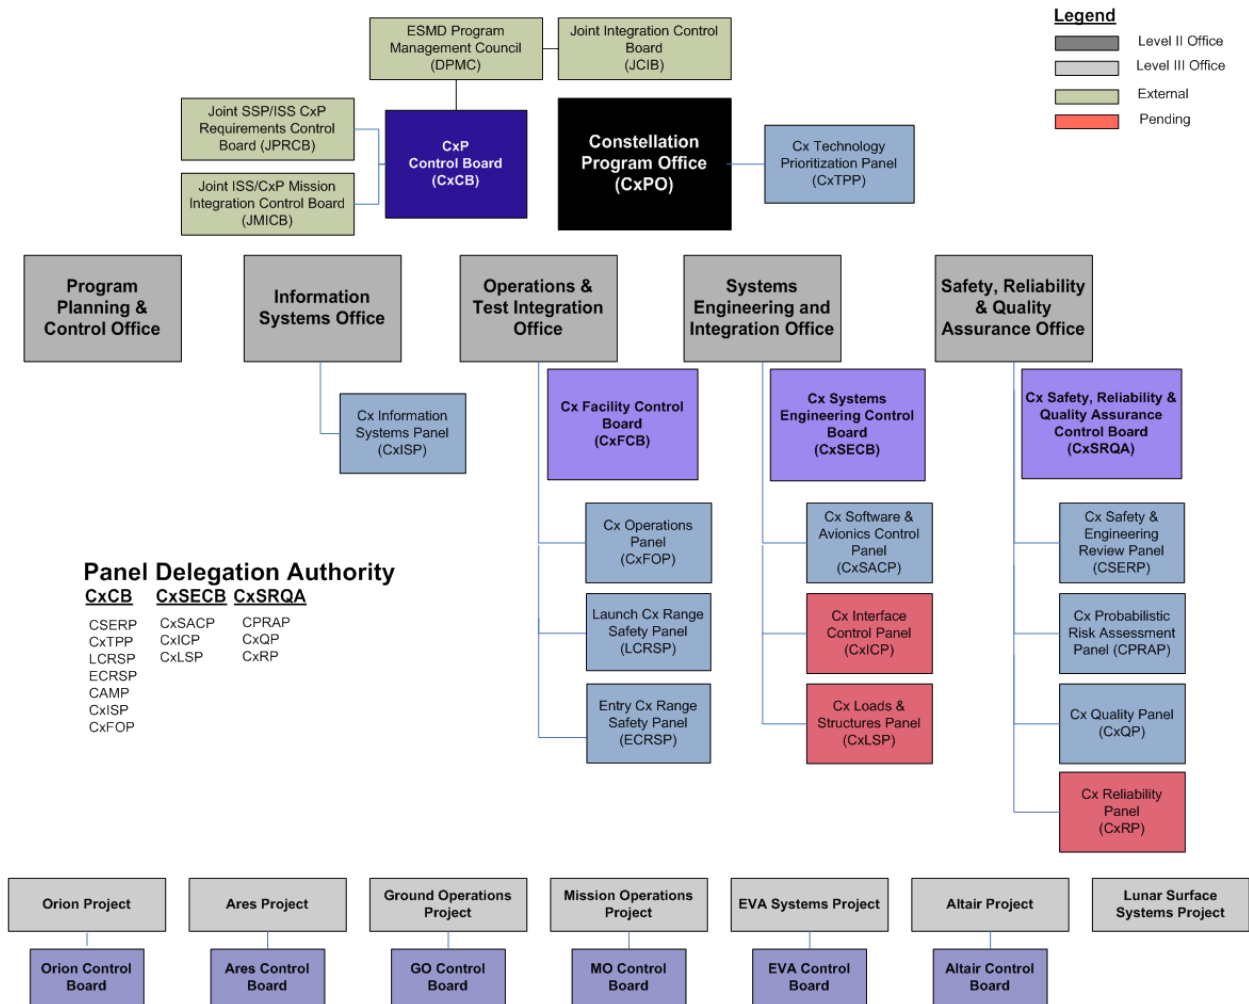


Fig. 4: Illustration of the Constellation Program decision-making board and panel structure. Note what is not shown: center institutional decision-making bodies, Technical Authority (the appeals process), or the prime contractors internal processes. (ESMD=Exploration Systems Assessment Mission Directorate, SSP=Space Shuttle Program, ISS=International Space Station)

Lessons Learned

In spite of constant attention from senior management, the decision-making process remained a persistent issue that only marginally improved over time. In a program of this scope, attempts to balance timely decision making at the appropriate levels, consider strategic and tactical viewpoints, and clearly delineate accountability for execution, while keeping all stakeholders

informed and included, often left someone dissatisfied. The need to maintain the multi-decadal focus was sometimes misunderstood, particularly in projects with immediate, near-term needs.

The need to make decisions regarding NASA center infrastructure that would or would not support Constellation over time was particularly nettlesome, bringing in communities that seldom dealt with program management. These decisions typically necessitated involvement at the Agency level and required extra time and staffing. Major infrastructure changes drew interest and attention from Congress as well, so it was necessary for many entities outside the program to be well informed.

Recommendations

A good start is important to define the decision process; however, constant vigilance is required for large-scope, distributed programs. Invest the time and energy to define a comprehensive decision process that includes all affected parties (Technical Authority, center management, and contractors). Impediments will arise by not anticipating key stakeholder interests.

Attention to the effects of project phasing, particularly for decisions intended to last decades or affect center infrastructure, is very important.

Supporting Lessons Learned from Section 4.0: 4.1.2, 4.1.3, 4.1.4, 4.1.5, 4.4.1, B.1

3.9 Organization is Organic — You Will Never Get It Right, But You Can Make It Better

Driving Events

The organizational structure for the Constellation Program evolved over time. Initially the organization was set up in a model derived from the Apollo Program – an integration function comprised of Systems Engineering and Integration (SE&I); Test and Evaluation (T&E); Operations; Program Planning and Control (PP&C); and Safety, Reliability and Quality Assurance (SR&QA),⁶ along with individual projects including *Ares*, *Orion*, Ground Operations, Mission Operations, and EVA. An Advanced Projects Office was also included to begin formulation of the lunar projects, including the lander and surface systems. Significant changes within the next year included the following:

- Establishment of mission managers for the *Ares* I-X and I-Y flight tests to enable the focus needed for integration.
- Establishment of the Lunar Lander Project (*Altair*) and the Lunar Surface Systems Project, and dissolution of the Advanced Projects Office in response to the completion of early formulation activities for these projects and no immediate need for succeeding early formulation activities beyond the lunar missions.

⁶ J. L. Rhatigan, J. M. Hanley, and M. S. Geyer, *Formulation of NASA's Constellation Program*, NASA-SP-2007-563, October 2007.

In the following year, other organizational refinements were implemented:

- Combination of the T&E Office with SE&I, since the T&E role had become entangled with that of SE&I and the Ares I-X and I-Y mission managers.
- Establishment of the Program Architect to balance the maturation of the Constellation architecture (LC) with near-term (IC) systems development.

In the subsequent year, further organizational refinements were implemented:

- SE&I was reorganized to transform its focus from requirements definition to design integration following completion of the Program Systems Definition Review (SDR).
- The Information Systems Office was formed in response to the focus needed in development and implementation of a program-wide data architecture.

The preceding listing is illustrative, rather than comprehensive, to communicate that the Program adapted the organization over time to both emergent issues and task completion.

Lessons Learned

Organizations should be treated as organic entities that are evolving and adapting to circumstances over time. Needed adaptation can be anticipated according to schedule milestones; the work force and organization should be focused on achieving the next key milestone, not on continuing to operate in the mode that achieved the last key milestone. Key milestones such as design reviews mark natural breakpoints in organizational strategies, and indeed the organizational scheme warrants assessment as part of a key milestone review. Consider that as the Constellation Program moved from requirements definition to development, the needed changes were not uniformly recognized, so the organization was slow to adapt after the SDR to prepare for PDR. Rather than preemptively assessing the organizational model, the program leadership reacted to issues associated with progressing to PDR as these issues emerged. The post-PDR organizational adaptation to CDR objectives was accomplished more quickly since it was anticipated.

Recommendations

Revisit the organization at key milestone reviews to address the changing nature of the work ahead to achieve the next milestone. This should be documented in a revision of the Program Plan and be briefed to senior Agency management at the appropriate milestone briefing. Senior Agency management must be enlisted to mitigate social impacts across NASA centers due to the ebb and flow of assigned work.

Supporting Lessons Learned from Section 4.0: 4.1.4, 4.1.5, 4.4.1, A.5, B.1, B.2

3.10 Communication Among a Distributed Team — Interpersonal Networks and Information Systems Can Improve Decision Making

Driving Events

As discussed in Section 2.2, Constellation's widespread 10-center team created a true communications challenge. While countless assessments and prevailing programmatic wisdom indicate that a small, centrally located team is the most efficient way to build a complex element, Constellation did not have that luxury. This posed many RR&A and decision-making issues, as described above (4.7, 4.8, 4.9), as well as integration challenges. Well-understood RR&A and clear decision-making processes reduce the communication "overhead" for the program team. Beyond that, any and all efforts to improve communications and data management (DM) are beneficial. Constellation used the following to enhance communication and DM over and above holding regular face-to-face meetings:

- A council of project and program deputies, mentioned in 4.7, was formed to resolve issues related to roles, decision making, and communications. The council was formed in early 2007 to tackle issues already apparent in RR&A associated with the formation of the Program Office. This council became an informal communication network; members spent time with each other and gained trust by sharing perspectives and working through problems of mutual interest.
- "Communities of Practice" were formed to foster discussions in particular technical areas and to aid information flow.
- Information technology (IT) tools and applications (teleconferencing, WebEx [Cisco, San Jose, CA], LifeSize® [LifeSize® Communications Inc., Austin, TX], ICE/Windchill [PTC, Needham, Mass.], etc.) were extensively used to enhance the flow of information.
- A data-centric (vs. document-centric) approach to program integration was implemented that enabled linking of configuration-controlled record-level data and automated integration of data direct from authoritative sources, thus eliminating error-prone manual copying of data.

Lessons Learned

This thorny issue was a classic case of "the devil is in the details." Overall RR&As were clear enough, as everyone understood the integrating responsibility of the Program Office with respect to constituent projects; however, translation of this overall understanding into day-to-day practice was needed to address the real issues that cropped up daily.

As such, frequent meetings of the council of deputies were necessary, which gave the informal network an ideal opportunity to form. IT tools, including teleconferencing and WebEx, allowed many virtual meetings to be interspersed with face-to-face meetings, thereby greatly increasing productivity of the group. Over time, as the relationships between group members grew stronger and the team's norms of behavior became more ingrained in team members, more and more of the group's work could be accomplished via these virtual meetings, further increasing productivity.

The council of deputies is only one example; integration groups for each of various engineering disciplines also benefitted from the improved IT collaboration tools over the previous generation, allowing significant improvements in the utility of virtual team meetings.

Due to the late start of the program relative to the start of the projects, incompatible data systems were the norm, so the program had to use a mostly manual approach to integrate data. This often involved use of non-authoritative copies of data, which resulted in decisions being made using out-of-date and inaccurate information (e.g., out-of-date IMS, 40% discrepancy rates in risks). Once a program-wide data systems architecture was established, promising results with core data sets were observed, thus reducing discrepancy rates and communication overhead in support of improved decision making.

Recommendations

Large-scope, geographically dispersed programs can benefit from interpersonal networks established across organizational boundaries.

Affordable, state-of-the-art IT tools can facilitate communications across time and space. They should be available as early and as broadly as feasible.

A large program should be diligent in anticipating the formal and informal networks that will emerge or intentionally be formed, and ensure that these networks have access to the IT tools that will enable these networks to work most productively.

Program-wide data systems architecture for data integration should be used *a priori*; and projects should use software tools that enable data integration and interoperability with published interfaces and, where practical, should also use common software tools.

A data-centric (vs. document-centric) approach for DM should be established at the beginning of the program, with clear and common data requirements between all levels and organizations of the program using formal agreements (e.g., data requirements documents [DRDs]) that include requirements for electronic data.

Supporting Lessons Learned from Section 4.0: 4.1.1, 4.3.1, 4.3.2, 4.3.3, A.5, B.2

3.11 Flight Tests – Learning by Doing

Driving Events

The value of flight tests in program development is undisputed; however, they are inherently resource intensive – which is to say, expensive. As such, they are typically reserved as risk mitigation for significant technical risks and demonstrations of critical capabilities. For the Constellation Program IC, flight tests were primarily focused in two areas: *Ares* launch vehicle ascent performance and *Orion* LAS performance.

Ares I-X was the first test flight in a planned series that would culminate in human rating of the vehicle. It was designed as a crewless development flight test to demonstrate first-stage ascent performance. *Ares* I-X successfully flew in October 2009, achieving all test objectives. One of the primary purposes of *Ares* I-X was to acquire flight data early enough to influence and impact the design and development of *Ares* I. The *Ares* I-X attributes were sufficiently similar to those

of *Ares I* to meet the test flight objectives. The test provided enough data for partial validation of models and processes used in the *Ares I* design, and allowed the test data to improve *Ares I* design and development.

For the *Orion* LAS, a series of flight tests was planned to certify safety and performance in the most challenging launch and flight regimes. The first flight test (Pad Abort-1) demonstrated LAS performance and Crew Module landing systems performance in a pad abort scenario. Launched in May 2010, Pad Abort-1, the second Constellation flight test, achieved all flight test objectives. A primary purpose of the flight, akin to that of *Ares I-X*, was to acquire flight data early enough to influence and impact the design and development of the *Orion* LAS and Crew Module landing and recovery systems. The test provided sufficient data to partially validate models and processes used in the *Orion* design, improving *Orion* design and development.

The *Ares I-X* and Pad Abort-1 efforts were largely independent activities. As such, comparing the two flight test activities reveals fundamental similarities that may characterize successful flight tests and provide insight for future program success.

Lessons Learned

Both flight tests were very successful, achieving all flight test objectives and returning a wealth of technical data early enough in the development cycle to inform design; however, the key lessons learned involved nontechnical returns from the flight tests. In both cases, these complex flight tests stressed the needed organizational constructs and drove clarity into RR&A and the decision-making structures. Had these issues been revealed later and addressed nearer to the first full-system flight (e.g., the planned *Orion-1* mission), significant delays and an increased risk posture would have resulted.

Furthermore, nontraditional approaches can be tried during a flight test, particularly a crewless flight test. For instance, *Ares I-X* employed large structural margins in the mass and dynamics simulators used to emulate the upper-stage and *Orion* hardware; and it used a “ping-test” approach to determine the dynamic characteristics for guidance of the stack rather than a much more expensive and time-consuming structural ground test of the stack. This approach worked very well and, after its successful demonstration on *Ares I-X*, the Program was examining the cost savings vs. risk of adopting this approach for future flights.

Finally it is noteworthy that when cancelled, the Program was in the process of revising the development strategy for the IC to rely more heavily on flight tests, thus saving time and budget by reducing ground tests in response to the increased confidence in the analytical models and simulations that resulted from *Ares I-X* and Pad Abort-1 flight tests.

Recommendations

Plan a series of flight tests to demonstrate key system capabilities and mitigate major technical risks. Ensure that some of the flight tests occur early enough in the development cycle (i.e., PDR time frame) to return data in time to inform the system design. Flight tests have the additional

benefit of providing opportunities to test alternative approaches and discover efficiencies for overall organizational performance.

Supporting Lessons Learned from Section 4.0: 4.2.6, 4.7.5, 4.8.1

4.0 DETAILED LESSONS LEARNED

Section 4.0 comprises the detailed lessons learned. They are arranged in traditional program management and engineering disciplines. The reader should note that many of the lessons cover new topics, i.e., those not addressed in the executive summary lessons learned. Also note that most of these lessons are linked to the Key Findings as supporting information for those seeking further information on a particular topic.

Below is a listing of topical areas contained in Section 4. New topics, not addressed in Section 3.0, are indicated in italics.

- Roles, Responsibilities, Authority, and Decision Making
- Requirements
- *Data Management*
- *Integrated Master Schedule*
- *Earned Value Management (EVM)*
- *Joint Confidence Level*
- *Construction of Facilities*
- Risk Management
- Design and Development
- *International Partnerships*

4.1 Roles, Responsibilities, Authority, and Decision Making

To effectively integrate a multi-element, geographically dispersed human space flight program, roles and responsibilities should be defined early in the start-up process. This should include not only the program and its projects, but also the agency and centers involved. Lack of clarity in this area leads to team tension, inefficient decision making, and wasted efforts. Moreover, as program objectives change with program maturity (Phases A to E), RR&As should be reevaluated and adjusted for relevance to current program milestones and objectives.

4.1.1 Ensure the Decision-making Process Is Aligned with Program, Project, and Center Roles, Responsibilities, and Authority

Supports Key Finding: 3.8

Driving Events

As described in Section 2.0 and Key Finding 3.8, the program decision processes were perceived as duplicative and sometimes inconsistent. Arguably the most important of the factors described in Section 2.2 was the highly constrained budget environment. Budget cuts led to rapidly diminishing reserves. This drove the necessity to retain the remaining reserves at the program level, thus severely restricting decision authority at the project level. Budget cuts drove many revisits to prior decisions in search of remediation. In any decision-making structure, lack of reserves to implement risk reduction, changes, or schedule slips to accommodate reduced funding, etc. compromise the ability of a program to make progress (e.g., reevaluating decisions due to changed circumstances).

Further, the need to hold reserves at a high level directly subverted the best practice, described in 4.1.2, of pushing functions and decisions to the knowledge base (project or lower levels of the program). Lacking resources, boards can solve problems within their scope but often cannot implement the solutions. Rather, the decision must be pushed up through myriad boards to the top, seeking resolution.

The mosaic of NASA participation (Section 2.2.3) and program/project phasing (Section 2.2.6) resulted in a similar mosaic (and phasing) of institutional boards and panels, each with their own style of vetting technical recommendations. These existing institutional processes and procedures were, in some cases, duplicative or contradictory, and were not always clearly linked to program processes that were subsequently established in an effort to implement standardization.

As a result of *CAIB* recommendations,⁷ NASA was in the process of adopting a new governance structure between the institution (i.e., center engineering and program support), Technical Authorities and the Mission Directorates. Constellation was the first program to adopt this new governance structure. While intended to alleviate instances of safety compromises in the face of extreme budget and/or schedule pressures, the new structure added further confusion to decision making at the working-level. As the first major program to operate under these new rules, Constellation had both the benefit (increased safety) and the difficulty (less clarity to decisions) associated with blazing a new trail.

These all contributed to inconsistencies, both apparent and real, in management direction. While projects were under program direction, center boards and the Technical Authority were not under program direction. It became very difficult to converge on a consistent approach once decisions were made by any of these. Further, when an issue arose, there was often confusion

⁷ NASA 2003. National Aeronautics and Space Administration. *Columbia Accident Investigation Board — Report*, Volumes I-VI. August 2003.

about the “entry point” in the decision-making process and who had the ultimate authority. The variety of institutional processes and lack of physical collocation only added to the muddled situation.

While the Deputy-led Program Excellence Team (Section 4.1.2) addressed many of the program-to-project issues, the larger issue of institutional authority vs. program authority was not sufficiently resolved. The frustration at this, which was borne by the working-level engineers seeking resolution, was expressed in the bumpy path to PDR.

For specific task areas, the program and projects conducted multiple Lean Six Sigma⁸ (L6S) events that identified and eliminated thousands of wasted steps. However, much of this “churn” may have been avoided to begin with by establishing a clear program decision-making path with buy-in from all affected institutional organizations.

Lessons Learned

- Boards without resources (budget and reserves) to implement decisions have limited effectiveness.
- Clear lines of authority among the program, projects, and institutions were addressed at the macro level but were not delineated or communicated well to the working level, resulting in tension and inefficiency. The operating model was already in place and was difficult to change.
- The program and institutional/center processes were not well aligned, resulting in inefficient decision making.
- Programmatic decisions on facility construction and modification funding (and the legal necessity not to intermix facility funding) were neither fully understood nor articulated at all levels and institutions (see Section 4.6).
- The need to involve the broadest range of personnel from many centers created a complex communications environment and resulted in frequent revisits of decisions.
- The arrangement of having program managers, deputy program managers, and associate program managers who matrix from different centers enhanced communications, dissemination of information, and team building.
- The concept of the associate program manager (APM) at the centers to facilitate communication between the program and the center was good, but the role of the APM was in some cases unclear and confusing with regard to the projects at that center.
- Communities of Practice (Section 4.1.5) were effective at bridging some of the complexities created by the multiple decision-approval chains.

⁸ L6S is an industry standard approach that combines the statistical analysis of Six Sigma to improve quality with the tools and methods of the Lean methodology, which evaluates processes with a focus on speed and efficiency.

Recommendations

As discussed in Section 2.2, maintain a budget that includes enough reserves for boards to implement necessary decisions.

Create a Control Board flow diagram that documents control board, panel, and working group (WG) relationships and hierarchy. Minimize panels and working groups and clearly define what types of decisions can be made by each. Include relationships that fall outside of the strict program hierarchy (e.g., an institutional Engineering Review Board, Technical Authority, etc.). Publicize these clearly stated authority chains, and hold centers and projects accountable to them.

For widely spread, flagship-scale programs, create a Program Formulation Task Team that includes agency and center representation to integrate a new program structure within existing agency and center structures. Proactively discuss and document what is meant by locating a project at a center initially, then continue the discussion to the task level with the personnel responsible for delivering the product. Clearly defined roles and responsibilities need to be explicitly defined in a Memorandum of Understanding (MOU).

The roles of program-level Technical Authorities and the technical review process must be clearly established and understood by the program, projects, and institutions. As with the technical side of the institution, make the administrative functions (e.g., financial organizations) answerable to the program into which they provide matrix support.

Use methods such as L6S to regularly review processes and organizational structure for duplication and waste.

4.1.2 Establish a Cross-level Senior Management Team to Clarify and Improve Roles and Responsibilities**Supports Key Findings: 3.7, 3.10****Driving Events**

A “best-practice” embraced by the Constellation Program was to push functions and decisions to the lowest possible organizational level, since that is where knowledge resides. This included pushing some roles down to the projects that traditionally had been held at the program level. The desire to implement this best practice led to challenges in RR&A that were exacerbated by the program environment.

As described in Section 2.0, the major projects were established prior to the Program; hence, each entity had its own integration organization, structure, and processes developed independently of the other. The far-flung organization, funding pressures, generation gap in human space flight, and size of the trade-space described in Section 2.0 also contributed to the magnitude of these driving events.

Lessons Learned

In an effort to identify and address areas of tension or inefficiency in RR&A, the Constellation Program chartered the PET to focus on areas for improvement. The PET was a council of deputies, with membership drawn from the program office and project deputies, as well as the Program Chief Engineer (a matrix position from the Technical Authority). This senior leadership team had the latitude to clarify, delete, or modify RR&A and associated processes to achieve better internal operations. Examples of functional areas discussed included Program and project roles and responsibilities for design integration, interface requirements, design and construction standards, schedule management, and vehicle stack integration.

An example was the delegation of vehicle “stack integration” to the Ares project. This is traditionally a program-level role; however, the bulk of the expertise required to perform the analysis resided in the project and delegation eliminated some duplicative efforts. The PET was tasked with clarifying RR&A that had arisen over the performance of the “traditional role” vs. the “new role” in stack integration. After clarifying these roles, the PET oversaw implementation so that the task could proceed more efficiently.

The team met primarily through off-site face-to-face meetings, which proved critical to enabling discussion of difficult and sensitive topics. Thorough and frank discussions resulted in more clarity and improved RR&A, which were documented and disseminated across the program in agreements. Documentation of agreements was instrumental in forming a better integration effort across the various levels of the program.

To quantify the degree of success in the implementation of these agreements, the program used the 4-D survey system[®] (©2010: 4-D Imaging Systems Inc., Oak Ridge, TN) to measure vertical and horizontal team satisfaction in recognized areas for improvement on an annual basis. These progress reports were published program-wide. The surveys indicated that the RR&A areas that had received management attention improved over time. As the program matured, the surveys provided new areas in RR&A that needed attention. While the PET was initially formed as an ad hoc group, it became the resident forum for RR&A improvements.

An added, and extremely important, benefit is that the personal bonds established in this team created a strong sense of trust and teamwork among the members; these relationships later enabled the rapid solution of many problems beyond the PET charter.

Recommendations

A senior-level management team, spanning projects, program, and institutions, is a key element to clarify and improve roles and responsibilities. It should be established from the start of a multi-element human space flight program. This team should meet face-to-face on a regular basis throughout the life cycle of the program and use a quantifiable process to measure success and identify areas for improvement.

4.1.3 Balance Decision Making with the Obligation to Air Dissenting Opinions

Supports Key Finding: 3.8

Driving Events

While decision making was addressed in the prior recommendation, another aspect is worthy of further discussion due to its cultural implications. A finding of the CAIB report (footnote 7) was that NASA culture inhibited frank discussion of dissenting technical opinions. Indeed, this was attributed as one cause of the *Columbia* accident. As a result, the management and work force, particularly at the space flight centers, were robustly trained to allay this cultural inhibition.

Perceptions in the work force indicated that the program was not entirely successful in finding a balance between decision making (as manifest in progress) and the need to find consensus and air dissenting opinions.

Program management was diligent in seeking consensus to assure that all valid stakeholder needs, opinions, and relevant factors were considered in making decisions. The continuously shifting budget situation drove priority changes that, in turn, caused decisions to be reopened. Some working-level personnel perceived this as a management weakness, an inability to “make decisions stick,” and as kowtowing to dissent. What looked like dithering at the working level was, in fact, diligence in identifying cost impacts and seeking cost avoidance for otherwise meritorious or technically necessary changes.

Other factors contributed to this perception, including the program/project phasing (Section 2.2.6), the generation gap in human space flight (Section 2.2.4), and the size of the trade space (Section 2.2.5) in an early-development program. And, for a leadership team most focussed on socialization of major decisions with Agency management, insufficient attention to this perception led to frustration at the working level.

One example was the Integrated Stack Technical Interchange Meeting (TIM), at which the Program sought to synchronize technical baseline requirements for all of the projects involved in the *Ares I* launch vehicle “stack.” The event was intended to include all relevant decision makers in a 2-week, face-to-face period to allow necessary coordination. The outcome was intended as an update to the Constellation Architecture Requirements Document (CARD). While many hard decisions were made and implemented into requirements changes as a result of the TIM, the formal closure of some of these decisions dragged out over a much longer period of time as additional cost or technical issues arose. While this can be attributed to many of the reasons elaborated in Section 4.1.2, the perception of weak leadership, a “back-door” appeals route, and lack of accountability was only reinforced.

On the other hand, Section 3.11 addressed the successful *Ares I-X* test flight. That organization was permitted, indeed supported, to operate with somewhat autonomous authority and was able to ignore many of the drivers described in Section 2.2 due to its ambitious requirements and tight schedule constraints. The flight test leadership stressed the needed organizational constructs

and drove clarity into RR&A and the decision-making structure. Many of these were passed on and implemented into the larger program.

Lessons Learned

Dissenting opinion plays a critical role in any healthy organization. A thorough, thoughtful, and well-understood process must be used for raising, addressing, and resolving dissenting opinions. Otherwise, anyone with a strong opinion can bring progress to a halt.

Program management needs be diligent in socializing decisions not only to Agency management but also to the working level, particularly in early program stages when the trade space is most broad and decisions are often revisited. Accountability, and the perception of such, is critical to morale at the working level.

Recommendations

Establish, communicate, and enforce the open decision-making model and process, including the airing of dissenting opinions. Assure “buy-in” of decisions not only at the Agency level, but also the working level. Do not permit “back-door” appeals as these engender mistrust and frustration.

4.1.4 Implement Communities of Practice to Resolve Integrated Issues

Supports Key Findings: 3.8, 3.9

Driving Events

System Integration Groups (SIGs) were initially established to define and allocate requirements among projects. They were intended to actively include subject matter experts (SMEs) from the projects and contractors, while being led out of the program office. However, due to the pace of the formulation activities, program (Level 2) and project (Level 3) personnel focused on their own level of products, so the vertical integration of the team never completely gelled. In addition, the role of SIGs with respect to program oversight was not clearly defined or communicated.⁹ This resulted in tension between the project-level system leads and the program SIGs over roles and responsibilities, and was a recurring topic at PET meetings.

Conflict around SIG vs. project subsystem lead roles continued until the Program matured, after which significant integrated issues arose that had multiple program, project, and institutional stakeholders. These issues *required* program and project personnel to work better together to find solutions. The Program performed a “reset” by transforming SIGs into communities of practice (CoP), which were defined as groups of technical discipline experts from program,

⁹ Former ISS personnel, who had oversight authority in that program, staffed many of the SIGs. However, Constellation emphasized pushing authority down to the projects (4.1.1). This change in culture was neither immediately recognized by Program management nor made clear in articulation of RR&A between the Program and the projects.

project, and institutional engineering organizations that serve as the authoritative resource for resolving technical integration issues within a specific area of responsibility (e.g., electrical power). These CoPs were directly consulted in technical discussions and decisions made by the new Constellation System Integration Panel (CxSIP) the focus of which was to provide technical integration decisions and recommendations for the Program.

Some examples of integrated issues that were addressed include: launch tower clearance, thrust oscillation, power quality specifications, launch probability, and Human System Integration Requirements interpretations.

Lessons Learned

Communities of Practice proved effective at resolving cross-cutting integrated issues. They helped improve communication within the program/project organization and between centers. The three primary lessons learned in relation to CoPs are identified below.

- They were effective at solving hard problems. Consolidation of program and project technical positions was enabled. Stakeholder buy-in was enhanced since recognized experts in each area were able to achieve integrated solutions.
- They fostered collaboration of the best expertise from across the Agency, industry, and academia. This enabled resolution of complicated technical issues and advanced the volume of knowledge in some disciplines.
- Influence leadership was key to overcoming organizational boundary concerns. *CoPs worked better through leadership facilitation* than direction. This allowed flexibility to resolve issues at the most appropriate point, which could be a program board, a project panel, or an other integration forum.

Recommendations

Communities of Practice should be implemented early with the focus of facilitating integrated issue resolution. Although the pace of formulating a program or a project will encourage organizations to withdraw to focus on their own highest priorities, management and processes should be put in place to incentivize cross-level discussion and decision. Attention should be given to leadership, which is key in determining whether the CoP approach will work within a particular discipline. Specifically, influence leadership is a must, and facilitation rather than direction is needed to allow flexibility. Roles and responsibilities among WGs, panels, and the CoPs should be defined. Management should consistently engage the CoPs directly when solving problems and seek their input when making decisions.

4.1.5 Implement a Strong Technical Integration Function

Supports Key Findings: 3.8, 3.9

Driving Events

Driving events are the same circumstances described in Sections 4.1.1 and 4.1.3; however, aspects of leadership merit a separate discussion. In a program of such breadth, leadership may be strong in some areas and not in others.

Lessons Learned

Pushing traditional program integration roles (Section 4.1.1) to the lowest feasible technical level, coupled with elevating all meager funding reserves to the program level, created significant stressors to the RR&A in resolving technical issues. (Section 4.1.2 describes the forces that drove all reserves to be managed by the program manager rather than at lower levels.) Communities of Practice (see Section 4.1.4) helped, but could not resolve, the fundamental lack of integrated technical leadership.

The Standing Review Board (SRB), which was charged with oversight of the program, found that the program was deficient in technical integration and recommended that a Technical Deputy be appointed to alleviate this issue. Program management was evaluating the implementation of this recommendation at the time of Program cancellation.

The cultural implications of this are partially addressed in Section 2.2.5. A strong and prevalent belief resided at the working level that program leadership was consumed with political, budget, and schedule issues, and that insufficient attention was applied to the many technical integration issues. Moreover, no one was delegated the authority to clear the backlog of technical decisions that lingered. A yearning for leadership, particularly in the program (Level 2) integration function, was a prevalent theme around the water coolers.

The Program took steps in strengthening technical integration by refocusing the SE&I office toward design integration, and creating a CxSIP that served as a clearinghouse for key integration issues.

As noted in Key Finding 3.9, it is critical for an organization to recognize its weaknesses and adapt as its functions change over the life cycle of the program. Once these changes occur, it is just as critical to fully communicate to the working level why these changes have been made.

Recommendations

In a program of the size and scope of Constellation, authority, responsibility, and accountability should be delegated to a clearly defined technical integration function to assist the program manager in efficient resolution of technical issues. It is best to consolidate this responsibility earlier, rather than later, to avoid morale issues related to inattentive leadership.

4.2 Requirements

A long-term trend in the aerospace industry has been the proliferation of requirements and specifications. The forcing function for this has been attributed to the well-intentioned desire to “regulate” good behavior and practices. Failures involving loss of life and loss of mission (LOM) understandably drive care and detail into specifications. NASA has been no exception to this trend.

As noted in Section 2.2.3, the different cultures that have developed at each NASA center have driven 10 different sets and approaches to D&C standards and specifications. Although the Agency has been systematically trying to establish Agency-wide D&C standards, there are still a multitude of center-specific standards in practice. The collective effect on the contract environment (and cost) is illustrated in the lessons in this section.

Notably, this effect drives the size and scope of the contractor work force. Each requirement in each standard and each specification must be verified.¹⁰ This task comprises a large part of the fixed-base costs, addressed in Key Finding 3.2 that can greatly hinder the flexibility of a program to respond to changing circumstances.

Recommendations in this section call for a balanced, measured approach in application of “shall” statements to contracts.

4.2.1 Reduce Agency Requirements Levied on Contractors

Supports Key Findings: 3.2, 3.3, 3.4

Driving Events

The procedural and policy directives (NPRs and NPDs) for the Agency contain tens of thousands of programmatic requirements and standards. While many Agency requirements properly address the outcomes sought (“what to do”), the vast majority attempt to incorporate and standardize best practices (“how to do”) to maximize the likelihood of achieving those outcomes. Each NASA center layers its own set of requirements and standards on the set that exists for the Agency, as described in Section 2.2.3.

These requirements drive cost in large ways. Contractors must verify every requirement, and the program must ensure that every requirement is auditable. NASA represents a small portion of the overall contracted effort for most aerospace contractors. DoD, which typically represents the largest source of work, maintains its own set of standards and specifications, most of which are duplicative of NASA standards and specifications. There are costs associated with retraining portions of a contractor work force to use a different process for the relatively small portion of NASA work. Risks of errors are introduced because of switching processes for similar tasks

¹⁰ A requirement is an individual “shall” statement within a document that may contain hundreds of individual “shalls.”

between contracts. This drives costs up and lowers safety. Costs are associated with documenting a requirement and assessing the impact of a potential requirement. Finally, potentially hundreds of labor-hours of debate are spent just to tailor a single requirement.¹¹ Confusion ensues when the contractors are expected to sort through and understand the differing NASA requirements among NASA centers.

Design and construction standards were the greatest challenge due to their sheer number and scope at the designer level. Over prescription is very costly and time consuming because it severely limits the ability of the designer to balance risks, and often results in suboptimal outcomes for the intended system.

One example would be standards for lifting hardware. Each NASA center has a standard for this, and the various standards are not identical. This necessitates a contractor dealing with multiple centers to, depending on the locale, know, understand, train, and staff to implement multiple ways of lifting hardware. This approach increases both risk and cost.

Since D&C standards are usually written by the functional discipline experts, it is only natural for them to set standards that minimize risks in their particular area. But, sometimes adherence to these drive unwise system design decisions that actually *increase* risk overall and lead to a generally poorer design. It can also drain resources from the project to do the “perfect” design rather than the “good-enough” design. Many of the design and construction standards are also written to cover a wide array of products and do not necessarily fit any one product well.

It is important to acknowledge that experienced contractors are also motivated to apply best practices and document them as corporate processes, if they do not already use Voluntary Consensus Standards.

The NASA Administrator stated that for Constellation, “all requirements need to fight their way on to the table.” Agency senior leadership embraced this approach (the Office of Chief Engineer [OCE] and Office of Safety and Mission Assurance [OSMA] in particular). However, reality proved to be that Constellation had to fight the working-level “requirement holders” over each requirement, and the burden was on the Program to document agreements regarding applicability and tailoring. On program cancellation, the number and scope of Agency NPRs and NPDs had not decreased.¹² Constellation did negotiate dramatic reductions in the number of applicable requirements and tailored the remaining documents into significantly fewer requirements.

¹¹ While the projects were allowed to do substantial project-specific tailoring, the process was time consuming and costly. It was not uncommon to have running dialogues of weeks to months for the applicability or meets/exceeds evaluation of a single requirement within a document containing hundreds of requirements. Often this was outside of the purview of the project manager, since the requirement “owners” or institutional representatives would defend the need for the requirement. In some cases, resolution had to take place at senior management levels.

¹² As of this writing, focus has been renewed at the Agency level to dramatically decrease the number of requirements within NPR 7120.5—a major parent document for programs and projects. The authors hope this renewed focus is managed to a successful conclusion and is promulgated to other Agency documents.

While this was a great help, the numbers were still too high. The biggest challenge was the sheer volume of “shalls.”

Lessons Learned

The “blanket approach” to requirements used by the Agency is untenable. While minimizing requirements on contractors was a priority for the Constellation Program, the resulting requirement base was, and still is, excessive. This drove a large portion of the contractor fixed-base costs.

This is a problem not imposed from outside the Agency, but is purely within the scope of the Agency to remedy.

Recommendations

Proactively address the issue of requirement over prescription in traditional acquisitions that are driven by Agency- (and multiple center-) level requirements and standards. Establish clear criteria for mandatory requirements at the “shall” level and revalidate the current Agency requirement base. The Agency, led by OCE and OSMA, must establish a process that will:

- Identify the criteria that must be met before “shalls” can be used in any Agency-level document.
- Capture best practices as “shoulds” in these documents to provide guidance for providers who do not have extensive space flight experience or documented standards.
- Reduce verification documentation and closure overhead by identifying criteria that can allow the “signature warranty” provider and process sampling by the customer to suffice for verification of targeted “shalls.”
- Review all Agency-level requirements documents with the intent of reducing the number of “shalls” by 75% to 90%.
- Reintroduce the use of NPGs to provide flexibility in managing NASA contractors to reduce the cost of leveraging valuable experience.
- In parallel, address the proliferation of center-level, non-unique documents with the same criteria.

4.2.2 Minimize Programmatic Requirements

Supports Key Findings: 3.2, 3.3, 3.4

Driving Events

The Constellation Program also fell victim to the Agency trend described above (Section 4.2.1), by generating requirements documents that were not wholly driven by Agency needs. Contributing factors include:

- A “system of systems” was being built that would span decades to develop and was intended to be used for generations. The program needed to ensure that the systems would integrate

well downstream, without needing to have the trailing systems fix the problems left by the early developments. However, after the scope of the program was reduced, the Program was slow to reassess and reduce the scope of the requirements.

- The program team bore the scars of the *Columbia* accident and lengthy ISS development and integration. Team members needed to avoid the mistakes that had cost so much.
- The Program was driven to reduce life-cycle cost. From a life-cycle cost perspective, having standard processes and requirements span across multiple systems seemed to be the most effective approach.
- There were perceived deficiencies in existing standards such as MIL-STD-1540 when applied to human-rated elements intended to operate reliably for decades with minimal logistical support capacity.

The Program was also guilty of writing requirements where guidelines or other less-prescriptive communication would have provided more flexibility. Although it was recognized that “one size does not fit all” when it comes to requirements, it was believed that through simple tailoring and allocating processes, the burden of requirements would be minimal. This proved to be incorrect.

An additional contributing factor was the use of requirements to define “design solutions” well above the designer level. Unfortunately, this practice began with ESAS and was evident all the way up to Level 0 (at NASA HQ).

The program, in particular SE&I and T&E, were in catch-up mode (see Sections 2.2.6 and 2.2.7), having started late with respect to the projects and their associated contracts. SE&I staffed up very quickly and began pumping out requirements and plans at rapid pace. Although program-level boards attempted to keep the number of requirements to a minimum, individuals and institutions applied continual pressure to add more requirements to address perceived deficiencies or noncompliances.

There were, however, attempts to reduce the number of requirements and their cost impacts. One example was the successful Safety, Reliability, & Quality Assurance Requirements Technical Forum. All projects were requested to work with their contractors to identify candidate requirements having significant cost or that were not enhancing mission assurance. Because the program had seen that many of the cost impacts were driven by misinterpretation of requirements, the projects/ primes also provided their understanding of the intent of a costly requirement. Over 3 days, with the participation of 90 empowered project and contractor staff, the team deleted 43 requirements, reduced the scope of 52 requirements, and modified 56 requirements to provide more clarity or flexibility in implementation. Program representatives also emphasized a willingness to engage in further tailoring and project representatives left with a better understanding of previously misinterpreted requirements.

Lessons Learned

Even programs committed to ensuring the implementation of a small, cost-effective requirement base can end up with much more than expected.

A repeated theme in debates on requirements revolved around the competing needs of minimizing life-cycle cost (particularly the need for operability) vs. near-term cost impacts in an ever-dwindling budget situation. Programs should carefully evaluate whether “consistency across the program” is necessary, and target for such consistency where it brings true value.

Recommendations

Challenge all self-generated program-level requirements to provide lower-level implementers as much flexibility as possible. Do this early and more than once, particularly in de-scope situations.

4.2.3 Separate Tailoring of Requirements from the Waiver/Deviation Process

Supports Key Finding: 3.4

Driving Events

With the advent of Constellation, the Agency OCE revised the Agency tailoring process. This activity occurred in parallel with Constellation requirements development, creating confusion within the work force. Notably, this even included definition of terms.

As explained in Key Finding 3.4, unintended consequences arose when Constellation applied the Agency requirements tailoring process.

Within the DoD (the U.S. Air Force in particular), the acquisition community has long had proven configuration management (CM) processes for the application of deviations and waivers that could be traced to MIL-STD-480, MIL-STD-973, and MIL-HDBK-61. *These are used primarily for noncompliances during production and delivery.* Likewise, tailoring was a separate CM process that was selectively used within the acquisition community, mainly in the early stages of program start-up, i.e., Phases A and B. Any higher-level document (i.e., MIL-STD, MIL-SPEC, etc.) that was authorized for tailoring could be tailored by a program in part or entirety. The document owner is always consulted if the document is tailored. If the entire document is tailored, the tailored version is “owned” by the organization and typically given a unique identifier that associates the document to an organization or a program. Deviations and waivers are given the proper attention and scrutiny with the approval level limited to the program director in his/her role as chairperson of the configuration control board. Approval of deviations/waivers is not delegated to lower-level boards/panels. Deviations/waivers are tracked and maintained until they are closed. No “permanent” deviations/waivers are allowed. Furthermore, the traditional definitions of deviations and waivers and tailoring are used as follows:

Deviation – A specific written authorization, granted prior to the manufacture of an item to depart from a particular requirement(s) of the current approved configuration documentation of an item for a specific number of units or a specified period of time.

Waiver – A written authorization to accept an item that, during manufacture or after having been submitted for government inspection or acceptance, is found to depart from specified

requirements but nevertheless is considered suitable for use “as is” or after repair by an approved method.

Tailoring – The process used to adjust or seek relief from a proscribed requirement to accommodate the needs of a specific task or activity (i.e., program or project specific).

NASA, in an attempt to standardize the tailoring concept across all centers, decided to use deviation/waiver definitions, processes, and documentation to tailor outcomes. This brought several unintended consequences that made tailoring problematic. These included: formalizing the concept of permanent deviations or waivers, the use of waiver documents for tailoring, and dealing with the cultural aversion to waivers explained in Section 3.4. In addition, the program desire to delegate tailoring to the lowest level (where the knowledge resides) conflicted with the higher-level visibility required for a waiver.

Lessons Learned

Some may argue that NASA can make the work force more comfortable with the fact that waivers are only form of paperwork to ensure proper documentation. However, the cultural taint is not only a NASA-internal problem. Outside stakeholders (Congress, the Aerospace Advisory Panel [ASAP], National Academies, etc.) have preconceptions about the waiver process. Repeated reviews with ASAP can attest to that fact. Processing waiver paperwork unnecessarily complicates tailoring. And, as seen in the two prior lessons, the amount of tailoring required is massive for a program akin to Constellation. This is also a large driver of fixed costs in the contractor base.

Recommendations

Reduce the cost of the requirements tailoring process by returning the definitions of deviations and waivers to their traditional role (akin to the DoD process) of temporarily dealing with nonconformances during production.

Tailoring should not carry a cultural stigma. Allow tailoring to be documented by the simplest method possible.

4.2.4 Reduce Cost by Maximizing the Use of Industry Standards

Supports Key Finding: 3.1, 3.2, 3.3

Driving Events

The Constellation Program, supported by the Technical Authorities, went to great effort to evaluate contractor (industry) standards. It is important to recognize that experienced contractors are also interested in applying best practices. They document them as corporate

processes, if they do not already use Voluntary Consensus Standards (VCS). While NASA has embraced the use of VCS,¹³ the Agency has been slow in its implementation of these standards.

As noted earlier, major contracts were competed prior to program office formation, so the opportunity to perform program-wide analysis of appropriate standards was missed.

Lessons Learned

Contractor-based standards that are compliant with VCS were found to be mostly consistent with Agency standards. The exceptions that were not already fully compliant could be easily assessed at lower cost to programs and NASA.

Recommendations

NASA should maximize the use of contractor standards deemed “good enough” by Agency technical experts and only levy specific requirements to fill significant gaps. This would be best done on an Agency-wide basis to minimize impact on new programs. However, until this is done, programs and projects should consider performing this assessment before contracts are let.

4.2.5 Use Industry Standards for Data Management

Supports Key Findings: 3.1, 3.2, 3.3

Driving Events

Akin to the prior recommendation (4.2.4), this recommendation specifically addresses DM standards. Many of the data systems at the projects and within program integration offices were chosen prior to Constellation start and before the Program had identified requirements for data interoperability. A data system always requires some cost for interoperability, and the variables for that cost depend heavily on standardization. Most data systems used by the projects and program offices did not support or adhere to industry standards for data interoperability (e.g., Web services); and since many of them were closed-source systems, it was both difficult and expensive for the Program to achieve its DM goals. In addition, the Program spent significant effort developing various new standards rather than using existing industry standards or working with appropriate governing bodies to evolve those standards (if needed).

¹³ The OMB Circular requires that agencies report their use of standards on either a “categorical” or a “transactional” basis. Those agencies that report on a categorical basis are not required to list each instance in which a government-unique standard is used in lieu of a private-sector standard in procurement actions. However, such agencies are required to have a system in place to ensure that government-unique standards are developed only when suitable private-sector standards are unavailable for use. At present, only the DoD and NASA consistently report on a categorical basis. For all agencies, in those cases in which government-unique standards are required because private-sector standards do not exist, use of the government-unique standard is not subject to reporting. From *The Thirteenth Annual Report on Federal Agency Use of Voluntary Consensus Standards and Conformity Assessment*, NISTIR-7718.

Lessons Learned

Developing independent, program-specific standards has many challenges that can be mitigated by using existing standards and working with appropriate governing organizations to evolve those standards to address specific needs (NASA Standards, Federal Standards, or International Standards).

Recommendations

Industry standards for DM should be adopted and promoted in compliance with the National Technology Transfer and Advancement Act. These standards should be made explicit, and named as dependencies with the purchase or development of systems and services supporting agency data. Furthermore, contracts and MOUs should be used to enable the adoption of specific criteria.

Program standards (including data interoperability) should be based on industry standards (e.g., American National Standards Institute, International Organization for Standardization [ISO], World Wide Web Consortium, and Organization for the Advancement of Structured Information). These rules for DM operations should be mapped to project or center standards, and levied on NASA projects and contractors as early as possible in the life cycle of a program.

Procurement of data systems should take into account the ability of those systems to support the standards or to be modified by NASA to become compliant. Deviation from program standards should be addressed via a named governing board in charge of waivers, enforcement, and data architecture. It is also essential that data systems reflect an architectural framework rather than defining one. An enterprise framework for data or data systems is necessary for global adoption, governance, and standardization across the agency to reduce costs, improve integration efficiencies, and provide effective traceability.

4.2.6 Anchor Integration Verification Requirements**Supports Key Findings: 3.3, 3.7, 3.11****Driving Events**

As described in Section 2.2.6, the phasing of Program start behind the project starts led to many out-of-phase functions. One out-of-phase function was the development of system integration verification requirements prior to the establishment of certification criteria. This resulted in verification requirements that were not properly anchored to the performance and mission certifications. Another result was the integrated test and evaluation needs for certification, above and beyond requirements compliance verification, were not well understood across the stakeholder community.

Verification is the work done to create sufficient evidence to determine that, on delivery, the system will meet specified performance requirements and include required capabilities that will accomplish the design reference mission (DRM) in all applicable natural and induced

environments. Verification should, but frequently does not, include the work necessary to validate that component procurements, acceptance testing, manufacturing process controls, and assembly checkout will consistently deliver flaw-free products.

Nevertheless, verification work alone does not produce sufficient evidence to certify the system because it does not include all of the work necessary to validate that the behavior of the system is sufficiently well described and understood, that well-understood operating rules and procedures can accomplish operational missions within design constraints, and that the system is suitable for its intended purpose. Verification is a one-time-only “snapshot” of requirements compliance, while certification is the living, sustained, “as-certified baseline” definition of what the system is and what its capabilities and constraints are. The organization responsible for “sustainment” sustains the system certification, of which verification data make up only a historical portion.

Lessons Learned

The lack of a strong relationship between the verification requirements and the certification criteria did not allow for effective assessment of impacts resulting from reductions of verification events. Verification events, such as flight and ground test, were reduced without consistently assessing the impacts against baseline certification criteria, which resulted in increased risk postures. These reductions in test events were typically driven by cost and schedule constraints.

RR&A were not well defined. As a result, the T&E integration effort was less than effective. The appropriate, integration-level verification products were being developed, but products needed for interface and integration T&E were out of sync with the various project production schedules (likely resulting in cost and schedule impacts).

Integration-level verification needs were not identified and communicated in time to negotiate and acquire the deliverables necessary to support integration T&E events. The program and projects were not in sync with regard to integration software, analyses, and simulator delivery dates; and the integration effort required to resolve known disconnects across the independent planning for each of the stakeholders was not completed prior to the PDR.

Recommendations

Define the human-rating integration performance and mission certifications early in the program to influence development of integrated system verification requirements. Integration certification and verification deliverables should be defined no later than the PDR to effectively negotiate with multiple stakeholders. Example integration deliverables include: models, analyses, software loads, simulators, facility assets to support integrated testing, and stakeholder verification test results required to update integrated models and analyses.

Decision makers must take corrective action when the satisfaction of verification requirements is the sole or overriding consideration when building and implementing the T&E campaign. This frequently happens for two reasons. First, the understanding that verification is only a subset of a necessary and sufficient T&E campaign is not widely held; in other words, too few people fully understand and appreciate the difference between verification and certification. Second,

verification is a contract artifact, and “closure” of specifications “checks a box” for contract Statement of Work (SOW) closure, therefore management frequently emphasizes that T&E work must be linked to a driving requirement – sometimes even making policy or building processes that allow only work that can be linked to specification closure, considering any other T&E work to be “out of scope.” Similarly, tracking only closure of specified requirement “shalls”, as opposed to tracking the completion of all the T&E campaign work requirements can result in unbalanced emphasis on verification over other certification work.

4.3 Data Management

This section addresses DM and data systems for managing, defining, building, and operating large space systems. Proactive and systematic DM is key to enabling review and management of the overall integrated vehicle and program (e.g., integrate authoritative requirements with computer-aided design [CAD] models with Parts List with Problem Reports), mitigating risks and reducing costs associated with DM.

Lack of data integration leads to cost and risk increases due to use of non-authoritative/out-of-date data, miscommunication of data, and missing data. Key to DM is the use of a data-centric approach, as distinct from a document-centric approach, that involves managing electronic data effectively for data consistency and integration with: 1) structured format (e.g., XML [extensible markup language] or database) accessible at the field level (e.g., specific shall statements, not the whole requirements document); 2) searchable form (e.g., full text and field/value-based queries); and 3) linkable format that requires a unique identifier and machine/programmatic accessibility from external systems.

As discussed in Section 4.2.5, many data systems were chosen by the projects and program offices prior to Constellation start and/or before the Program had identified requirements for data interoperability. In many cases, these data systems did not support or adhere to industry standards for data interoperability.

In addition to the recommendations in this section, the reader will note that lack of DM drove problems in other systems, such as the IMS (Section 4.4), EVM (Section 4.5), and risk management (RM) (Section 4.8).

4.3.1 Use a Program-wide Data Systems Architecture

Supports Key Finding: 3.10

Driving Events

The Program did not have a data system architecture deployed across all of the projects to support integrated DM at its inception. For example, no support was given to link authoritative data sets such as: Requirements (component must sustain X impact), CAD models (component model), Parts List (components and subcomponents), and Problems (specific problem with component).

Projects and centers used a diverse mix of existing and new data systems with diversity of data formats (as described in Section 4.2.5), and limited or no support for interoperability. As a result, the Program used a manual approach for data integration that resulted in decisions being made using out-of-date and sometimes inaccurate information (e.g. out-of-date IMS, 40% discrepancy rates in risks). In addition, the lack of common software tools led to

- Additional cost and time expended to build/maintain multiple systems for similar data (e.g., shadow document and product structures).
- Increased cost and time to find and integrate data.
- Limited capability for efficient trending/data mining.

Program-wide data systems architecture was developed, deployed, and used 4½ years into the Program, with promising results for some core data sets.

Lessons Learned

Program-wide data systems architecture is needed to mitigate DM risks. The architecture should show clear traceability to the program DM requirements, and should specify an approach for automated data integration (aka “mash-ups”), which is a common practice in industry and a best practice in government.

The approach should

- Identify program data owners and their data responsibilities for key data sets (e.g., integration office manages requirements, safety office manages problem reports).
- Manage a definitive list of authoritative systems (tools) for authoritative data sets (e.g., requirements, risks, hazards).
- Ensure that new tools comply with program integration requirements for data integration using Web services (with fields and formats defined in DRDs using a data-centric approach).
- Implement updates or workarounds (e.g., a cache for a system that lacks integration capabilities) for legacy tools that do not meet requirements.
- Manage data interoperability agreements between system owners.
- Adequately address information security constraints from inception.
- Consider selecting or migrating to common open-source or source-available software tools that can be readily modified to align with the architecture.

Recommendations

A program-wide data systems architecture for data integration should be used with data owner identification and responsibilities, a managed list of authoritative data sources (tools), managed requirements for data integration using Web services, and managed data interoperability agreements between system owners. Projects should use software tools that enable data integration and interoperability with published interfaces and, where practical, common software tools.

4.3.2 Establish a Data-centric (vs. Document-centric) Approach

Supports Key Finding: 3.10

Driving Events

Constellation projects were established before the Program (as described in Section 2.2.6), and used numerous legacy processes and data systems that did not support data integration. Once the Program was established, data integration became critical for overall SE, Program management, and sustaining support, yet no clear and common data requirements existed to support data integration between levels and organizations (e.g., program-to-project, project-to-project). This resulted in misunderstandings (since groups had to translate between data produced in different formats), significant manual rework of data products that took time and introduced errors, and use of out-of-date and incomplete data for decision making (that resulted in increased cost). For example, different groups came to meetings with different conclusions based on different data, and argued which data were authoritative to attempt to reach agreement.

Lessons Learned

A data-centric, model-based approach to DM needs to be established at program inception, with all levels of management involved in implementation of the approach. Without this, the program incurs increased integration costs, additional risk from underspecified data requirements, and manual DM using documents.

To implement this approach, clear and common DRDs for specific data that are focused on sustaining engineering efforts should be developed for internal NASA and NASA-to-contractor data exchange, including the following:

- Requirements for defined attributes (fields, values, metadata, etc.)
- Electronic data formats for computerized data integration (e.g., XML)
- The NASA data systems where the data are to be delivered
- The delivery mechanism (Website upload vs. system-to-system integration push after every change)

A data-centric approach such as this enables linking of configuration-controlled, record-level data and automated integration of data direct from authoritative sources, eliminating error-prone manual copying of data. With this approach, CM can be done at the level of individual fields and data sets rather than of whole documents.

Recommendations

A data-centric (vs. document-centric) approach for DM should be established at the beginning of the program, with clear and common data requirements between all levels and organizations of the program using formal agreements (e.g., DRDs) that include requirements for electronic data. Key to this is to have a management structure for the approach that interfaces with key stakeholders (projects, program offices, and prime contractors).

4.3.3 Establish a Model-based Systems Engineering Approach

Supports Key Finding: 3.10

Driving Events

A comprehensive and clearly articulated Model-based Systems Engineering (MBSE) process was not established at the beginning of the Program and was not effectively communicated to the project offices until the *Altair* lunar lander office was initiated. As a result, consistent models were not established throughout the Program; and, when improvements were offered later in the program life cycle, there was little desire or funding provided to enhance the SE process. Resulting issues include: Program and project offices had operational concepts with substantial synchronization issues; SE practices were sidestepped to accommodate schedule pressure and avoid perceived costs (e.g., system functional analyses were not consistently performed or prematurely terminated); and the architecture was not ultimately traced to DRM objectives.

Lessons Learned

The benefits of using MBSE on the *Altair* design are articulated in Appendix A.2. Its use resulted in streamlined engineering requirements development and validation by developing integrated products among the Requirements, Design, and Operations communities.

Establishing requirement traceability to program figures-of-merit and mission and operational capabilities with proper functional analysis was an important step in ensuring requirement quality and completeness in the *Altair* implementation of MBSE. The definition of MBSE requirements and process standards as well as up-front education and broader work force experience greatly assisted the *Altair* project in the paradigm shift associated with MBSE. *Altair* successfully established traceability of the MBSE process to decision gates and design review criteria, and identified organizational responsibilities to implement the MBSE process. Configuration DM processes to efficiently and affordably manage the MBSE models and the corresponding data, in addition to managing documents generated from the models, were also key factors in the success of *Altair* and also provided a common focal point for SMEs.

Recommendations

Model-based Systems Engineering processes, with the ability to govern integration activities, should be established during program formulation to create an integrated set of models that all projects will use. Adoption of MBSE industry-wide standards integrated with NASA standards and policies is essential to MBSE success and should be strongly supported by program leadership throughout all organizations. MBSE training should be provided to model developers and SMEs to ensure consistent model development and traceability up to program decision criteria.

4.4 Schedule and Resource Management

This section outlines the primary lessons learned for the program IMS and EVM. This section should be read in context with the DM limitations described in Section 4.3. Lack of a consistent DM structure greatly hindered the development of an IMS and an earned value system.

4.4.1 Define Objectives and Uses for the Integrated Master Schedule *a priori*

Supports Key Findings: 3.7, 3.8, 3.9

Driving Events

Program development of an IMS was hindered by a number of factors. Primary among these was a lack of a common DM structure, described in Section 4.3. The late start of the Program with respect to the projects (Section 2.2.6) was also a contributing factor. Finally, the lack of clarity as to the purpose and use of the IMS delayed its implementation and use further.

Lessons Learned

Projects and program viewed the purpose of the IMS in fundamentally different ways, for high-level reporting vs. supporting detailed analysis, respectively. The Program did not explicitly define, or internally agree on, the purpose of the IMS. Lack of a single, clearly communicated vision from the Program resulted in the Program calling for multiple sets of data from the projects in various formats. Some projects objected to the increased effort being levied on them when the value of the end product was unclear and, in some cases, counterproductive. This lack of mutual understanding and intent, exacerbated by poor communication, led to development delays. The IMS architecture drove cases in which the Joint Confidence Level (JCL) analysis efforts (Section 4.5), detailed tracking of SE flows, and calculating an analytical critical path were mutually exclusive. When the Program set and communicated the clear goal of achieving a cross-program, analytical, critical path in the IMS, cooperation improved and implementation became easier.

Recommendations

The objectives and uses of the program IMS should be defined and clearly communicated prior to scheduling development activities, as these will drive the DM structure and cost to support the IMS.

4.4.2 A Bigger Schedule Is Not Necessarily a More Useful Schedule

Supports Key Findings: 3.7, 3.9

Driving Events

Due to the lack of clarity of purpose and conflicting opinions on purpose, the IMS development team, working with project counterparts, attempted to develop an IMS that included many activities at Levels 4 and 5 (e.g., contractor and subcontractor activities). As a result, the IMS

contained a huge data set that was difficult to maintain and use effectively. The value of the resulting schedule was questionable.

The Program also developed a high-level Program summary schedule that incorporated all key Program events on about 3 pages. This high-level schedule included an assessed critical path, which was developed by considering project-calculated critical paths, risk areas, and key technical and programmatic integration points. This summary schedule reflected the Program content in a more manageable package, and proved to be a much better management tool.

Lessons Learned

For large, multi-project NASA programs, a summary schedule approach with assessed critical path and key integration points can provide a more useful management tool with lower development and maintenance costs in comparison to a highly data-intensive schedule database.

Recommendations

Assess the depth and amount of information necessary to effectively manage the program. Consider that a summary schedule integrated from project detailed schedules may comprise the fidelity useful for program management at a lower overall cost.

4.4.3 Test Earned Value Management Systems in Phase B for Required Implementation in Phase C

Supports Key Findings: 3.1, 3.2

Driving Events

NPR 7120.5D requires the use of EVM in Phase C of developmental programs such as Constellation. While EVM is taught in NASA program/project management training, the NASA work force does not have extensive experience on its purpose and best use through successful application. Moreover, with the exception of the Jet Propulsion Laboratory, NASA centers do not have validated EVM systems in place.

In early 2009, the Program began to implement earned value reporting at Monthly Program Reviews (MPRs) in preparation for Phase C readiness (post PDR). Problems in understanding, readiness, data discrepancies, and perceptions of the value added from EVM were encountered.

While on-line and on-site training was available, this proved inadequate. One-on-one training was implemented to assure that all managers responsible for EVM reporting were able to manage the systems. This type of training encourages questions and creates a relaxed, open atmosphere for learning and will prevent planning errors that corrupt EVM data and damage EVM system credibility.

Lessons Learned

Early implementation of EVM in Phase B was difficult because of unrealistic expectations and insufficient training among the Constellation Program and project management team. One-on-one training was the most effective technique in the EVM training effort.

Recommendations

Programs should plan to provide EVM status at the MPR in Phase B and indicate this via the program plan. The EVM team should schedule “one-on-one” EVM training with managers very early in the program. This is necessary to have all managers engaged in dialogue prompted by EVM usage for cost control, schedule management, and corrective actions planning.

Note the following cautions in application of EVM described in Section 4.5.3.

4.4.4 Use Standardized Data for Earned Value Management and Cost Accounting**Supports Key Finding: 3.10****Driving Events**

Due to the start-up phasing described in Section 2.2.6, the involvement of every NASA center (and associated contractors) described in Section 2.2.3, and the lack of an adequate data system architecture (Section 4.3), the earned value system was very difficult to implement. Not all processes were documented among the Program, projects, and centers, and no standardized reporting formats or templates were used NASA wide. Additionally, inconsistent rates and full-time equivalent conversion factors were applied by each center and project as well as the Program to calculate budgets and performance. These inconsistencies made reconciliation of products a Herculean task.

Further, cost accrual personnel were using a completely different financial system than that used by the earned-value community. This was primarily driven by the need to interface with heritage systems at NASA centers (i.e., NASA cost accounting is center oriented). This difference consumed much time and effort in reconciliation, and was particularly onerous because of the constant need to re-plan as budgets were cut. There was also no valid system in place to manage or track procurement costs (a major portion of the Program budget).

Lessons Learned

When the program and its projects do not have standardized data, integration and data quality suffer. Budget estimates and performance calculations are never truly consistent due to use of inconsistent data.

Recommendations

Data gathering and reconciliation needs to be accurate and consistent in concert with Recommendation 4.3.1.

4.4.5 Rigorous Baseline Control May Not Be Possible during Early Phases or Rapid Budget Changes

Supports Key Findings: 3.1, 3.2

Driving Events

As described in Section 2.2.2, the Program encountered inconsistent and unreliable funding through its lifetime. Budget changes directly affected the major schedule milestones and the ability of the Program to maintain schedule integrity during rapid budget changes. The effort to establish a working baseline in Phase B to execute early use of EVM was extremely difficult in this environment. In addition, the processes used for baseline CM (e.g., baseline change requests and a baseline change log) were still in development and not ready to track the numerous changes.

Moreover, the broad trade-space described in Section 2.2.5 is a natural driver of uncertainties in an early-phase program.

Lessons Learned

Phase B EVM requires a rigorous CM system in place to allow for the program formulation turbulence inherent in this phase. Without a well-documented and current baseline, quality EVM data generation is problematic.

Recommendations

Although the baseline is still being established during the formulation phase, an initial baseline change control log and baseline change documentation should be developed to trace budget and funding changes and communicate those changes regularly. Logs should be developed to enable the staff to have an historical trace.

4.5 Joint Confidence Level

This section outlines the primary lessons learned for the JCL, which is a methodology pioneered by Constellation at the behest of NASA HQ that can be used by a program to provide confidence levels for the achievement of program objectives based on budget and schedule information. It attempts to predict the confidence in success given a budget phased against schedule milestones. NASA HQ set a 65% confidence level for Constellation.

This topic is technical in nature, and the interested reader is encouraged to review the complete report on the Constellation development and usage of this technique.¹⁴

¹⁴ Kuo, F. and Wilson, S. "Joint Confidence Level Requirement: Policy and Issues." NASA TM-2011-216154.

4.5.1 The Current Joint Confidence Level Budget and Schedule Reserve Requirements Are Untenable

Supports Key Findings: 3.1, 3.2

Driving Events

Requirements have been documented in NPD 1000.5 and NPR 7120.5 directing programs to budget to a 70% joint cost and schedule confidence level. Although the Constellation Program was permitted latitude to plan to a 65% JCL, analyses showed that this requirement, when coupled with the inherent uncertainties in large-scale, tightly coupled programs associated with human space flight, drives the need for high budget and schedule reserves. This is driven from the fundamental nature of portfolio management. There is a beneficial effect to cost confidence provided by portfolio diversification because cost and budgets are fungible; program managers can transfer money from a project in surplus to another in shortage. However, extending the requirement to include schedule, as in JCL, reverses the benefits of the portfolio effect. This is because, unlike money, one cannot transfer time from a project that has been completed early to one that is late.

The budget status of the program is described in Section 2.2.1 as is the first key finding, 3.1. The current 65% JCL demands budget and schedule reserves that are untenable for most projects and programs. The Agency and government budgeting process (along with many historical examples) precludes the types of reserves necessary to maintain this level of confidence.

Lessons Learned

To meet the 65% JCL requirement, budget and schedule reserves of 30% to 50% are needed for a typical, tightly coupled human space flight program. Using the minimum 30% reserve requirement and the approximate \$24B budgeted for IOC, this translates to a reserve of **\$7.2B for IOC alone**.

For the entire Constellation Program through FY20, the budget was approximately \$100B (including \$12B in reserves). Using the minimum 30% on program content (\$88B) renders **\$26.4B** that would have had to be set aside for reserves for the full program. This is more than double the actual reserves held, and does not include the required schedule reserve of 30%. Clearly, this amount of funding is excessively large with respect to the overall NASA funding and would not be allowed to sit idle in the Constellation budget in light of other Agency needs.

Recommendations

Implement an approach to stipulate requirements for cost confidence alone, but still enforce an emphasis on schedule uncertainty as an inherent element within cost estimates. This approach is described in more detail in the citation above.

4.5.2 The Joint Confidence Level Methodology Is a Useful Decision Support Tool

Supports Key Findings: 3.1, 3.2

Driving Events

The implications of a JCL target requirement notwithstanding, the analysis provided invaluable insight into the health of programs and projects and served as a vital portal into schedule, cost, and technical issues. The most significant contribution of JCL is its ability to provide information to program managers about the *comprehensive* impact of risks to schedule and cost, focusing on key issues at the macro level. It has proven to be a valuable management tool that merges the risk, cost, and schedule stance and also captures the dynamics of the interrelationships.

For the Constellation Program, the JCL process generated new strategic dialogue and understanding among program and project stakeholders — and even within the project control team. Analytical cooperation, as well as traceability, understandable methodology, identifiable input sources, and analytically rigorous results, are all hallmarks of successful JCL in action.

Lessons Learned

JCL analysis is highly useful in providing information to program managers about the comprehensive impact of risks to schedule and cost, thus enabling focus on key issues that most greatly effect the plan.

Recommendations

Continue to encourage (or require) programs to conduct JCL analysis to drive out the most relevant issues and focus on those that have the greatest impact.

4.6 Construction of Facilities

Management of facilities construction and modification has inherent complexities that were not fully appreciated at the outset of the Program. Since Constellation was the first major NASA development requiring new and modified facilities at multiple locations in many years, oversight skills were initially lacking. As described in Section 2.2, the Program was inheriting much of the vast space shuttle infrastructure, and was viewed as a source to solve pent-up facility maintenance demands. Moreover, management of facility construction is not a subject typically covered in program management training, so much of the management team was unfamiliar with some of the unique challenges. These include the following:

- *Facility use.* Some specialized facilities relied on the space shuttle (and, hence, Constellation) as the single, or nearly single, user. Centers desired to retain capabilities, but had difficulty either finding other users or justifying the costs to a single program (or both). Some facilities were nearly duplicated at centers, so securing a program commitment for usage was competitive and highly sought. Once the commitment was secured by a particular venue, cost estimates tended to increase.

- *Construction of facilities (CoF) funding.* Congress requires higher visibility for funding for new facilities, and funds must be managed separately from other program funds, reducing flexibility. Conversion of CoF funds to research and development (R&D) funds (and vice versa) is extremely difficult. Increased Congressional visibility increases visibility at NASA HQ and resident centers, necessitating additional reviews.
- *Center infrastructure.* Since most facilities reside at NASA centers, senior center management expected to be a part of the decision chain, regardless of funding source.
- *Cost control.* While not flight hardware, CoF is fraught with historical cost overruns and overestimations. This creates additional management difficulties in a budget that is highly visible and inflexible.
- *Facility development in parallel with hardware development.* New and modified facilities for use by Constellation were under design in parallel with hardware design in many cases. In some instances, hardware design changes drove changes in facility designs. The inflexibility of the CoF process made this a particular challenge.

To avoid these issues in future programs, flagship programs should establish an experienced facility team, conduct independent cost estimation on construction projects, and develop a facilities construction planning process that will guide selection and development of facilities and ensure due diligence is performed.

4.6.1 Establish Program-level Facilities Expertise during Program Formulation

Supports Key Finding: 3.1

Driving Events

The Constellation Program experienced all of the issues related to construction of facilities listed above. Moreover, there were multiple CoF projects under way at several NASA centers. Most of the issues resulted from one or more of the following:

- Inadequate front-end planning (FEP) and due diligence
- Poorly defined, immature, and changing facility requirements
- Vehicle requirements changes that affected ongoing CoF (since the vehicle designs were immature at the time that CoF had to begin)
- Limited funding and no contingency
- Lack of CoF experience in programs and projects
- Complexities and confusion in determining which parts of the project required CoF vs. R&D funding
- Late engagement of center facility expertise

Lessons Learned

The program and projects must have CoF expertise in their construction of technical facilities (manufacturing, test and operations facilities).

Recommendations

Establish a small, experienced facilities team early at the formulation of the program. This team should include facility cost estimators and facility development experts, as described below. This team oversees facility selection and development as well as aids in determining and tracking program facility use according to the facility planning process described below.

4.6.2 Develop a Facilities Planning Process**Supports Key Findings: 3.1, 3.7****Driving Events**

Most facilities issues experienced by the Program could be linked to inadequate FEP, in part driven by the late Program start relative to projects (Section 2.2 6).

Moreover, the starting state of the facility was not necessarily understood for the facilities that were being modified. That is, the facility condition report was not detailed enough to accurately identify all problems, and estimates were based on the building only and not on unique vehicle requirements. Initial proposed budgets for some facilities did not account for the immaturity of the facility requirements. Facility requirements were often immature at the time at which the facility project needed to begin. Since CoF must start years before vehicle design requirements are solidified, vehicle design requirements were changing that would impact a facility that was already under construction. Some projects did not have project management plans. Trade studies and cost estimates often did not consider full life-cycle costs.

Lessons Learned

Appropriate due diligence should be followed in selecting and developing facilities needed to support the program. The program documented the CoF processes once these unique aspects were understood.

Recommendations

During program formulation, develop a facility selection and development process that is consistent with NASA processes, laws, and regulations, and that reflects the organization of the program. This FEP process must have built-in flexibility so it can be tailored for the unique facility complexities. The process should ensure appropriate due diligence is done throughout the development cycle: from initial trade studies through construction and commissioning. Project milestones should be linked to the existing budget processes to provide the most current and accurate facility cost estimate available. The process should ensure that environmental impacts, energy efficiency, and historic preservation are considered in development. The process should lay out roles and responsibilities regarding who is developing the various due diligence products. It should clearly define the role of the centers, the projects, the program, and other stakeholders that should be involved. Milestone approval requirements should be clearly defined (Reference:

CxP 70700 Constellation Program, Constellation Management Plan, Annex 10 Facility Planning Process.)

4.6.3 Accommodate the Unique Risks Inherent in Construction of Research and Development Facilities

Supports Key Findings: 3.1, 3.7

Driving Events

Many of the facilities that are needed to support a new program are technically complex (unlike an office building); therefore, these projects have significant technical risks. In Constellation, the Reverberant Acoustic Test Facility at the Plum Brook Station Space Environment Test facility and the Stennis Space Center A-3 rocket test stand had requirements that were beyond any facility in existence. These one-of-a-kind facilities led to technical risks that required technology development and testing to mitigate them. Technical risks can lead to additional cost or schedule risks; these technical risks should thus be accounted for in CoF budgets and schedules. Accounting for these risks also requires flexibility in budgeting and schedule.

A related factor is the immaturity in the requirements for a facility in a development program such as Constellation. Facilities, such as those to support testing, manufacturing, or operations, are required early in the program life cycle. Often the manufacturing, testing, and operations requirements are being developed at the time the facility construction must start, which leads to unavoidable changes to the construction (driving costs up).

In other cases, significant new requirements are levied on a facility. For example, after construction start, the Program required the Mechanical Vibration Facility to accommodate not only the IC hardware, but also LC hardware – namely *Altair* vibration testing.

The existing agency CoF process supports building administrative-type facilities that have low technical risks. To accommodate the challenges of building an R&D facility, a program must have large reserves set aside for the CoF projects and the flexibility to change funding levels without requiring an Operating Plan change, separate management, or Congressional approval. With the current CoF system, once a program sets aside reserves in the CoF account, it cannot reclaim that funding for R&D accounts should the facility reserve not be required. Changing Congressional operating plans just to manage risk in a facility development project is not good fiscal practice.

Lessons Learned

Construction of facilities processes should be modified to better accommodate R&D facility development and provide funding flexibility.

Recommendations

Modify existing CoF process to address the risk associated with the building of R&D facilities. Flexibility could be provided that allows the recapture of a CoF budget that has been used for

construction or to increase funding due to construction changes. A second tool for gaining limited flexibility is for the program to group construction projects at a geographic location into a larger project. Grouping would allow the program the flexibility to move funds from one subproject to another at the same location without requiring an operating plan change. For example, the numerous construction projects required at Michoud Assembly Facility (MAF) for the upper stage could be strategically grouped so that money could be transferred from one CoF project at MAF to another.

4.7 Risk Management

The Constellation Program implemented an integrated RM process compliant with NPR 8000.4, Continuous Risk Management, and consistent with experience on the ISS Program and SSP. Successful RM faced significant challenges, as described in Section 2.2. Chief among these was the dwindling budget reserve that in large measure consumed program flexibility to address risks. Risks were diligently tracked and reported, but many lingered and, indeed, accrued due to budget limitations.

As with Section 4.4, this section should be read in context with the DM limitations described in Section 4.3. Implementation of RM was partially hindered early in the Program by the lack of a common, requirements-driven data system. An integrated RM process for systems that must interact physically or operationally is essential to discovering, characterizing, and reducing integrated risks. To support these activities, it is important to maximize similarity in common standards for defining, comparing, and tracking risks. Using a common risk database can substantially improve the effectiveness of this activity by providing common reports, fields, definitions, evaluation criteria, and electronic sharing capability.

4.7.1 Ensure Risk Management Is Embedded in the Design Process

Supports Key Finding: 3.5

Driving Events

The complexity, size, and limited budget described in Section 2.2 were circumstances that engendered a difficult RM environment. While a rigorous RM system was implemented, using the best assets and knowledge of the Agency, management processes can develop a life of their own, yielding a myopic view and confusion about the value of the process vs. the product. The Program was most successful with RM when risks were identified, analyzed, and brought to resolution as part of the design process rather than independent of the design process.

Lessons Learned

This team-centered risk identification process worked especially well as part of the Integrated Design Analysis Cycle (IDAC) process. IDAC provided a venue for iteratively improving analysis tools, integrating analysis disciplines, and investigating design decisions and operational assumptions. It brought stakeholders from the directorate, Program, projects, and elements together on a regular basis to define assumptions, perform integrative analysis, interpret results, and draw conclusions. IDAC organized the design process and enabled systematic risk identification. As a result, many key technical issues were identified during early design cycles rather than later. Examples include: achieving acceptable probabilities for LOC/LOM, structural loads on *Orion* and *Altair* during trans-lunar injection, first-stage thrust oscillations, and mass growth during design maturation. This supported early action to reduce or eliminate risks, and enhanced integrated understanding of design and performance.

The Program experience illustrated that many risks require multiple design cycles to be fully addressed via research, analysis, option development, and, ultimately, testing.

Recommendations

Implement common risk standards, risk system, risk training, and joint risk review process on programs and projects that have significant interdependencies. Put this system in place early in development to ensure that architecture and mission-driving risks are understood early in the design process. Embed the RM process with the IDAC process to ensure maximum impact on the design.

4.7.2 Good Risk Management Relies on Open Communication

Supports Key Findings: 3.7, 3.8, 3.10

Driving Events

Several communication issues hindered implementation of an RM and tracking process. First, the purpose and use of the RM system was not well understood throughout the Program and projects. Some viewed the database of risks as too large and cumbersome; others viewed it as

a centralized tool to capture all risks, with simple processes to limit what was seen based on user need. Second, projects were reluctant to provide detailed risk insight because they viewed such insight as micromanagement. The Program needed broad risk insight due to the small and dwindling reserve level; projects viewed the level of insight needed as a lack of trust. Third, cultural issues were at play. While some centers saw a high number of potential risks as a good thing (folks are paying attention), others saw it as a bad thing (there is something terribly wrong if all these risks exist). These thorny issues hindered discussion and resolution of Program risks.

Lessons Learned

As with the IMS (Section 4.4.1), the purpose and use of the RM system was neither fully communicated nor embraced by all key stakeholders. This led to distrust and a feeling by some parties that the risk system was overly burdensome and of limited value.

Recommendations

In establishing RR&A for programs and projects, the amount of insight necessary at each level for proper RM should be established and communicated *a priori*. Expectations of any risk system should be established early; do not assume they are understood because the system was used on a prior program.

4.7.3 Early and Robust Probabilistic Risk Assessment Enables Risk-informed Design

Supports Key Finding: 3.5

Driving Events

The Program built early probabilistic risk assessments (PRAs) to inform early design studies. As described in Key Finding 3.5, two projects used risk-informed design to great benefit. This was enabled by having a PRA available so that trade studies could proceed that weighed options for risk reduction.

Lessons Learned

The two projects (*Orion* and *Altair*) were able to use a risk-informed design process to improve mission success probabilities and reduce mass. These analyses relied on having early PRA available, and could not have proceeded successfully without this analysis capability in place.

Recommendations

Risk-informed design (treating risk as a commodity) can only be performed if good PRA data are available. Assure that an early and robust PRA is established to support risk-informed design.

4.7.4 Divide and Conquer to Address Complex Risks

Supports Key Finding: None

Driving Events

Complexities and interdependencies in design and, thus, risks drove complexity in solving those risks. Appropriately delegating aspects of risk mitigation to address complex risks arose as a management challenge during early design.

Lessons Learned

Division of risk mitigation was most effective if portions were defined within the scope of, and delegated to, individual teams. For instance, design risk for the liquid oxygen/hydrogen common bulkhead was broken down into many child risks that could be managed by individual teams. Each child risk contained a portion of the detailed mitigation plan developed to address every aspect of the overarching issue related to manufacture and production of the common bulkhead stage. This risk illustrated that it can be helpful to establish parent and child risks for complex or integrated issues. This approach makes work more manageable while retaining an integrated understanding of the risk and strong risk communication between stakeholders. The parent risk established a context, while the child risks focused on mitigation of specific risk aspects.

While the establishment of parent and child risks was found to be very helpful in managing complex integrated risks, it was also found that care should be taken to limit the establishment of “parent risks” to cases in which there is a unique higher-level mitigation activity. Simply establishing parents for all risks that are common across systems and elements inserts unneeded complexity into the risk process and clutters up the risk system. In the case of mass growth, for example, the Program did not track an integrative parent risk because it was felt that mitigations for potential mass growth would simply reflect standard SE practice. Risks were established on those systems that had mass issues, such as *Orion*. These teams conducted additional analysis and design activities to achieve their mass targets.

Recommendations

Use parent-child risk relationships to document where multiple stakeholders are performing significant mitigation responsibilities for the same integrated risk. However, limit the parent-child structures to risks with unique and significant mitigation contributions to avoid overpopulation of the risks system with less important risks.

4.7.5 Link Test Planning to Risk Mitigation

Supports Key Finding: 3.11

Driving Events

Choosing how best to test, and how many tests are needed, prior to human space flight is a debate explored more fully in Section 4.9. Testing is expensive in hardware development, and is

most often first on the chopping block when budgets are squeezed. It is therefore necessary to develop a solid and compelling rationale for testing, or elements of testing, to form cogent program and project plans.

Lessons Learned

To better focus limited test resources, the Program developed a Risk Informed Test (RIT) approach that was aimed at identifying the most significant operational and developmental risk drivers and ensuring that tests appropriately explored these through analysis, ground test, and flight test. RIT relied on evaluating both qualitative and quantitative systems safety assessment, historical failure history, historical test and verification experience, and expert elicitation to help identify test priorities, test fidelity, and frequency.

A qualitative approach such as RIT, which collects what is known about failure modes and the effectiveness of analysis and test to mitigate and understand these risks, can be very helpful when informing these decisions.

Recommendations

While more testing provides greater assurance of safety and maturity growth, there are clearly cost and schedule limits on how much testing can be done. An approach such as RIT can be helpful in discerning the best test approaches.

4.8 Design and Development

4.8.1 Ensure a Common Understanding for Implementation of Human Space Flight System Certification

Supports Key Finding: 3.11

Driving Event

As the Program matured, the Agency was unable to clarify requirements to govern the first flight of humans on a new spacecraft.¹⁵ This was, in part, due to the generation gap in human space flight discussed in Section 2.2.4. The space shuttle was designed in the mid 1970s, and tested in the late 1970s and early 1980s. Few, if any, senior leaders in the Agency had dealt with the issues surrounding vehicle certification for human space flight, particularly for vehicle launch and reentry.

The NASA human-rating criterion¹⁶ is stated as: “Given that the human [space flight] system is designed for human [space flight], it is accepted that the objective is to fly humans when risks to

¹⁵ Such requirements exist for payloads and reflect varying levels of test of new commercial flight vehicles based on the level of government insight and importance of the payload.

¹⁶ NESC-RP-10-00619. “Readiness for First Crewed Flight.” April 12, 2011.

the crew safety have been mitigated to the point where the need or benefit is worth the residual risks.”

This criterion, while sensible and easy to understand in principle, is extremely difficult to implement as certification guidance. It is instead left to program management to interpret into a design, development, analysis, and test plan that leads to certification for human space flight. And, while program staffs are skilled in this, the breadth of interpretation leaves a program extremely vulnerable in the face of budget challenges once that plan is formulated. Moreover, as some have observed,¹⁷ the current strategy does not adequately address “unknown-unknowns,” or the unexpected things the hardware can “tell” the designers during test.

Section 2.2.1 described the budget challenges faced by the Program during its tenure. A primary target during multiple rounds of budget cuts was the test program leading to human certification. This test program became difficult to defend in the face of myriad interpretations of what it should take to certify for human space flight. The definition of “residual risk” was widely debated and never clearly understood, such that any critic could challenge the test plan. The interested reader is referred to a summary accounting of changes to the certification test plan during the life of the Program, cited in footnote 17, below.

The Constellation Program was advancing with necessary tasks that would define the certification requirements for a human space flight system. The integrated certification plan was not complete as required for the PDR in February 2010. This delay was due to the interpretation challenges noted above, along with the mismatch between the project procurement contract requirements vs. the integrated requirements produced by the Program SE&I organization (see Section 2.2.6). After PDR, work on the plan was halted as more cost/schedule-friendly variant strategies were studied. While an operational certification plan was not completed, a number of lessons learned were identified that addressed several inefficiencies that originated from these driving events.

Lessons Learned

A cogent, integrated test plan cannot withstand budget challenges if the underlying requirements and rationale (certification criterion, residual risk) are left open to broad interpretation.

The definition of “certification” (in terms of content and in relation to content definitions of terms such as “verification” and “validation”) at the element and integrated hardware/system level were unclear (see Section 4.2.6). The Master Integration and Verification Plan (MIVP) was released at the Systems Requirements Review (SRR) but did not adequately define the certification process or the plan. It provided a high-level definition to Program elements but not to the level of detail necessary for certification planning (e.g., defining the role of the Systems Acceptance Review, Design Certification Review, and Certification of Flight Readiness in the overall certification plan). MIVP Revision A was released for the PDR and included an updated certification plan.

¹⁷ Noriega, C. et al. “First Crewed Flight: Rationale, Considerations and Challenges from the Constellation Experience.” In preparation as a NASA TM.

Work should be performed early in the life cycle of a program to develop an integration and test strategy (including flight tests) as a companion to design definitions, starting with system concept development and maturing with the design through certification completion. The early availability of this strategy would enable timely contribution of key driving test and evaluation figures of merit to support making better-informed decisions, such as the “what and when” of procurement content. Based on this information, “right to left” integration can occur to identify key integration points driving (hardware, software, data objects, etc.) product deliveries and critical paths through the certification campaign. The integration and test strategy should accompany design descriptions and be controlled in the same manner. The strategy should be used as a planning benchmark to provide a clearer and considerably more valid basis of content, cost, and schedule estimation.

Recommendations

The Agency should interpret its current certification criterion to a level such that a program plan for design, analysis, and test can be clearly seen to either meet or not meet that criterion, particularly in the face of budget pressures.

Programs should

- Clearly define what “certification” means prior to awarding contracts.
- Clearly define the test and evaluation roles and responsibilities prior to awarding contracts.
- Develop an integration and test strategy for accomplishing certification concurrently (and to a comparable extent and fidelity) with initial system concept development.

4.8.2 Embed Human-rating Certification in the Core Processes/Tools of the Program

Supports Key Finding: None

Driving Events

The Constellation Program was the first program to implement NPR 8705.2B, which was approved coincident with the Program SDR. A draft human-rating certification package (HRCP) was developed for the SDR. The NPR 8705.2A and B versions both require implementation of a safety process; one key difference between them is that Version B places more emphasis on performing the analysis to drive the design vs. using a rule-based approach to drive the failure tolerance for catastrophic events, which is the emphasis in Version A. Since the NPR was approved after contracts were put in place, needed NPR data had to be accommodated with the currently planned contract deliverables. Some potential gaps existed for the new NPR-required specific design assessment data.

The Constellation human-rating approach was to embed human rating into Program standard procedures. This placed responsibility for compliance with the designers rather than establishing a separate process. It also leveraged the NASA human space flight heritage and culture for implementation of these procedures.

A team, which was staffed by technical experts responsible for ensuring the satisfaction of the NPR requirements, project representatives responsible for ensuring project compliance, technical authorities, and astronauts, was responsible for ensuring that the evidence provided met the intent of the NPR requirements.

The team demonstrated to human-ratings stakeholders that the required content was covered (and augmented, if necessary) within the Program products. The HRCP provided pointers (via hyperlinks) to the evidence contained within the planned Program documentation. The Program used the requirements software tool to demonstrate compliance with NPR technical requirements. Fault tolerance and the use of the hazards database is a key area in which this type of negotiation with stakeholders is necessary.

Lessons Learned

Build human-rating precepts into the program teams along with standard processes and tools. Do not establish human-rating activities as a separate process unto itself.

Recommendations

For future human-rated procurements, include the process-related requirements from the NPR, including access to all human-rating data in the contract SOW.

Establish a team with the implementers and technical experts to develop the HRCP. Include the technical authorities and astronaut crew on that team to assure the certification package content satisfies the NPR.

4.8.3 Use Risk-informed Design at the Outset of a Project

This lesson learned is addressed in whole in Key Finding 3.5.

4.8.4 Address Units of Measure Early

Supports Key Findings: 3.1, 3.2

Driving Events

NASA has attempted to move the Agency and its supporting contractor infrastructure from use of English units of measure to the more widely used and accepted SI units of measure. Most notably, the ISS Program attempted this in the late 1980s and early 1990s, particularly since the International Partners use SI units. The ISS Program abandoned the effort when costs of implementation became prohibitive. Human space flight programs have lagged in this effort, primarily due to the reliance on heritage infrastructure that was built in English units (e.g., launch pads, test stands, transport ships and crawlers, railcars, etc.). NASA science missions have enjoyed the best success at this conversion, since they are less reliant on heritage infrastructure.

The Agency attempted to implement SI units once again, using the Constellation Program as leverage to incentivize the change. The CARD established SI as “primary units of measure,”

stating: “All Level 1, 2, and 3 integrated analyses, performance and verification systems will be completed in SI units.” However, the CARD allowed “English units” for Level 4/5 (contractor and subcontractor) if justified. The CARD was ambiguous on “operational data.” Yet as of 2007, the Program was not compliant – there was a units “Tower of Babel” in presentations and analysis. This was primarily due to a lingering heritage infrastructure issue that drove a similarly unacceptable cost overhead for implementation. The risk of a units-produced error was clear and present.

A directive was approved in late 2007 that reinforced the use of SI units and provided details on where and how to use a “hybrid” units system based on proven practice in robotic missions. It also addressed the primary constraint that “flying” in SI units requires testing in SI units (i.e. “fly like you test”).

Lessons Learned

While the Program and projects were willing to implement this new directive, cost estimates provided by the major contractors (who were well into design due to early contract starts) were very high, approaching \$370M.¹⁸ Funding was unavailable, even at a low level, to start implementation and negotiate the contractor costs. Ultimately, due to these funding limitations, the Program had no choice but to issue a new directive in June 2009 that specified U.S. customary (English) units across the Program.

Recommendations

Establish adequately specific units requirements *a priori*, including a well-thought-out and executable implementation approach that addresses heritage infrastructure and contractor conversion costs. Unless SI units are specified prior to acquisition, the cost to change will be prohibitive.

Assign a champion to enforce implementation at all levels until use of SI units is standardized in the human space flight community.

4.8.5 Design for Operability to Reduce Life-cycle Cost

Supports Key Finding: 3.1

Driving Events

One of the many observations of the CAIB was that the space shuttle was expensive to operate. The need to use Apollo-era infrastructure, along with early design compromises that saved then-year dollars as opposed to out-year dollars, led to an expensive operational program. During formulation, Constellation was challenged to design in lower operational costs while using and upgrading the shuttle-era (and, in some instances, Apollo-era) infrastructure and “fixed-base” (see Section 2.2). All involved recognized that program sustainability depended on lower life cycle and,

¹⁸ Notably, these costs were primarily due to changes in documentation, rather than hardware.

thus, operational costs. Moreover, lower spending on IC operational costs would allow the development of the LC to proceed.

An example concept was the “clean pad” approach, which recognized that every attachment, umbilical, structure, etc. on the launch pad that touched the vehicle drove an interface that needed a design, an ICD, operational requirements, procedures and launch constraints, and a team of operators. This, in turn, drove operational costs and complexity. Designers were challenged to drive interfaces to an absolute minimum to achieve the “cleanest” pad possible.

Another launch-related concept for the LC was to shorten the time between launches of the crew vehicle (*Ares I*) and the cargo vehicle (*Ares V*) to drive minimal launch preparations into designs. A target of 1.5 hours between launches challenged designers to reduce both launch preparation and processing times.

An activity was initiated, during the SRR in late 2006, to formalize these “stretch requirements” for inclusion into the CARD. Teams from the operations centers (JSC, KSC) brainstormed an initial set of candidates that would help the Ground Operations and Mission Operations projects to reduce their costs to operate the integrated program systems. The teams developed a set of proposed requirements that were assessed by the projects and elements for cost and schedule impact. Roughly half of the proposed requirements were implemented into Revision A of the CARD. Those that were not implemented as requirements became a set of Constellation Program “design guidelines” that were scored for their implementation progress at subsequent Program milestone reviews.

Many of the design guideline items found their way into the architecture where there was performance or schedule margin to accommodate them. Some examples of these design changes for operability include the use of a single common fuel for all *Ares* reaction control systems (hydrazine), non-pyrotechnic separation systems for all flight-to-ground umbilicals, and locating all vehicle-to-ground pad interfaces within the tower-facing arc of the flight system.

Several initial stretch requirements were deleted from Program specifications during the season of SDRs. The most significant of these deletions was the requirement for meeting targeted operation cost caps, by system, at full-operational capability. Others were tailored or ameliorated based on budgetary and performance limits of the architecture.

The most successful of the stretch requirements dealt with minimizing the critical path of integrated launch operations processing. An integrated WG was established to estimate and track the impact of vehicle and infrastructure design changes to the targeted 45-day vehicle processing flow. By the time of Program PDR, the architecture was showing the capability to meet this aggressive timeline requirement. While working this ongoing assessment, strong working relationships between the Ground Operations Project team and the vehicle design elements were established in which the common goal was developed and shared.

The Program recognized that the stretch requirements were only able to address the operations phases of the life cycle (post-handover from the contractor). To attack life-cycle cost challenges

inherent in production and sustaining engineering, the Program established an operability “czar” to focus on the fixed and recurring costs of the flight elements.

Lessons Learned

Design for operability must be a driving metric in the initial architectural trades, particularly for programs planned with significant operational lifetimes. Operability requirements imposed after the major architectural decisions are made must compete for budget and performance margins that, in the case of Constellation, were extremely limited.

Operability must be defined uniformly across all program and project elements. Ground Operations saw operability as a measure of cost/complexity to perform launch system integration and checkout; *Ares* saw it through the lens of the ease of maintenance of its designs; Mission Operations aligned it as being associated with operational robustness and flexibility of in-space elements; *Orion* assessed it as the reliability and robustness of its subsystems. While these definitions had some overlap, they often led to competing interpretations of the intent.

Focus on design for launch site operability improvement provides some cost-reduction benefit, but is only a small contributor to the total life-cycle cost of complex space systems. Operability improvement is further constrained by incorporation of legacy systems and infrastructure.

Many stretch requirements were defined with both “threshold” and “objective” values. Threshold requirements were minimum targets to be met with objective requirements as goals to continue to work toward as systems matured. In practice, while projects did do assessments to reach the objectives, many objectives were not implemented because the needed contract changes were cost prohibitive.

Recommendations

Programs should establish an initial cost model, which includes operability, and maintain that model as systems evolve from SRR through PDR to CDR. Consider using the best practices of “design-to-price” and “cost as an independent variable” as implemented in DoD acquisitions.

Use of “stretch requirements” that incorporate both “threshold” and “objective” specifications can be maximally effective if they are included up front in acquisitions or procurements through performance incentives for contractors. Programs should consider these measures to assist contractors in focusing on design innovations that provide operability improvements of significant life-cycle cost drivers.

4.8.6 Simulate Processes Early (Run a “Virtual” Mission)

Supports Key Finding: None

Driving Event

The Program freshly developed mission integration and certificate of flight readiness (CoFR) processes needed to be assessed for efficiency and productivity. Typically, these processes are “shaken out” during early missions or mission preplanning. Since the technology was available, the Program undertook an early mission simulation, prior to PDR, that included prelaunch, in-flight, and post-flight planning and operations.

An ISS crew rotation mission was chosen. The first simulation was termed Virtual Mission-1 (VM1), and a second simulation, VM2, was planned prior to Program cancellation.

Lessons Learned

Gaps and areas of improvement were discovered in the mission processes and, in one case, an *Orion* vehicle design deficiency was discovered.

Process improvements: Several critical mission milestones that had schedule or assignment disconnects were identified. An example was the Mate Review: There was some difference of opinion among the *Ares* project, Ground Operations Project, and ATK [Alliant Techsystems, Inc., Minneapolis, Minn.] (the *Ares* first-stage contractor) as to who was responsible for the stacked first stage. While *Ares* assumed responsibility for VM1, the issue remained open and the plan was to simulate this activity during VM2 to resolve it. The VM1 team also identified schedule disconnects with the Mission Control Center and Constellation Training Facility reconfiguration. Other mission milestones needed processing improvements. For example, the CoFR was a manual process that was very labor intensive. VM1 helped confirm that mission stakeholders needed a continuous status capability that could be provided by a software tool (e.g., the CoFR dashboard).

Design improvements: The management team challenged the virtual team during VM1 with an ISS control moment gyroscope (CMG) failure such that a new CMG would need to be processed quickly for launch in the *Orion* SM. The VM1 team identified the need to access the SM unpressurized cargo later than planned (L-270). The command module (CM)/SM stack had to be accessed before leaving the KSC operations and checkout (O&C) building because no ground support equipment existed to access the SM in the KSC Multipurpose Processing Facility. For VM1, the team was able to load the CMG later in the flow because the CM/SM stack was still in the O&C and the Ground Operations team worked around the clock. Another simulated emergency, which coincided with an SM redesign study to accommodate unpressurized cargo, resulted in several new *Orion* recommendations for the SM design (e.g., overhead mounting using a lunar propellant tank mounting point, radiator removal). This underscored the concept that design decisions made early in the program can drive operational capabilities later.

Recommendations

The virtual mission is an excellent vehicle with which to exercise planned mission management processes. In the case of VM1, the team refined many of the processes and began developing needed software tools.

Additional virtual missions are recommended with lessons learned from VM1. A virtual mission simulation to exercise an initial (non-operational) ISS rendezvous flight would be very helpful in identifying and resolving any process problems that might be present in the longer mission management template planned for these initial flights.

4.8.7 Focus the Certificate of Flight Readiness Process**Supports Key Finding: None****Driving Event**

Perceptions that the current human space flight CoFR process was inefficient and needlessly burdensome led the Program manager to charter a CoFR L6S composed of representatives from all Program organizations, as well as the ISS and Space Shuttle Programs, and an external senior advisor.¹⁹ The L6S team charter was to develop a CoFR process that was “integrated, efficient, and streamlined.”

Lessons Learned

The existing Program CoFR (flight preparation processes) was people, meeting, and paper intensive. The CoFR lean event led to a philosophy of a continuous (anytime) status capability that was to be facilitated by a CoFR tool (the CoFR “dashboard”). In addition, the process was redesigned to be product focused vs. organizationally focused. That is, for a group to be integral to the CoFR process, it must provide a concrete product such as an assessment or a report that contributes to certification. This eliminated several “checkers checking the checkers” that have been added to the CoFR process over the years. A similar activity regarding WGs revealed inefficiencies with unchartered WGs using precious resources.

Recommendations

Continue to use L6S or similar methodologies to assess process efficiency and effectiveness. Continue to use applications/dashboards, “roll-ups,” etc. to get a continuous status of not only CoFR but of any work process that spans a considerable time duration. This should be done in lieu of infrequent and large status events, and to avoid the generation and use of paper products for such events. This is not only inefficient and time consuming, but leads to “stale” information.

¹⁹ The ISS and Space Shuttle Programs were included because the Constellation CoFR process had been largely based on prior program experiences.

The full set of recommendations was implemented in CxP 70175, the Constellation Program Certification of Flight Readiness Process Plan (CPP).²⁰

4.9 International Partnership Development

These lessons learned related to international partnership development are focused specifically on topics relevant to NASA and Constellation internal program planning and development. Two recent and very relevant documents provide broad consensus views of international partnership lessons learned. The first of these, “International Space Station Lessons Learned as Applied to Exploration,”²¹ was developed by the ISS partnership to provide timely guidance in planning the Constellation Program. These lessons learned were widely shared and frequently discussed in the activities described below. The second, “The ISECG [International Space Exploration Coordination Group] Reference Architecture for Human Lunar Exploration—Summary Report”²² contains consensus lessons learned from the perspective of all international agencies involved in the development of the ISS. Both of these documents are valuable tools for planning; however, since they are consensus-based documents, they do not provide a view internal to the Agency.

As Constellation was forming, the ISS International Partners expressed an interest in continuing to pursue human space flight challenges to expand frontiers together with NASA. As Constellation represented the next step in human space exploration, consulting International Partners could have been an important step in strengthening our existing partnerships. It also brought us the opportunity for expanding international partnerships to additional space agencies beyond the ISS partnership. It is widely recognized that international agreements enhanced the sustainability of the ISS and, thus, could do so for future human space flight efforts.

However, NASA policy and federal regulations restricted the scope of potential international partnership opportunities available for the Constellation Program to projects and hardware not designated as critical path items (see Lessons Learned below). As a practical matter, export control regulations exclude international partnerships on launch vehicle development and any shared technologies that are defense related. Thus, the focal point for potential partnership opportunities occurred during the later phases of the Program—namely, the lunar exploration (LC) phase.

²⁰ Although this document was not baselined prior to Program cancellation, a full draft was produced in support of the Program PDR and may be found as an official record within that review.

²¹ ISS Multilateral Coordination Board. *International Space Station Lessons Learned as Applied to Exploration*. 2009.

²² International Space Exploration Coordination Group. *The ISECG Reference Architecture for Human Lunar Exploration—Summary Report*. July 2010.

The primary mechanism for discussion of partnership opportunities was the ISECG, which was comprised of 14 space agencies from around the world.²³ ISECG discussions were led by the NASA HQ Exploration Systems Mission Directorate and supported by personnel in the Program.

4.9.1 Engage Potential Partners Early in the Program Life Cycle

Supports Key Finding: None

Driving Events

The Program prepared its first *integrated* lunar-phase budget estimate in 2008. It was at this time that initial concept designs and potential architectures were mature enough to frame an integrated discussion. This integrated framework provided enough knowledge to initiate understanding and planning of the scope and roles that potential partners could take.

A small team of Constellation personnel began activities within the ISECG to explore potential avenues that could lead to partnerships and build relationships that would be key to this understanding. This team had the following objectives:

- To understand national priorities. An understanding of the national priorities of each agency, with respect to human exploration, was critical to gauge whether, when, and to what extent a partnership opportunity might evolve.
- To develop a realistic planning framework. A realistic understanding of the extent and nature of potential or contemplated contributions from each agency was necessary for the Program to plan a viable architectural framework.
- To inform the NASA budget process with realistic assumptions. This understanding allowed the development of a set of budget assumptions based on current and projected international states of affair based on national priorities.
- To build a sustainable framework for the lunar phase of the Program. As the sustainability of the ISS was achieved via international partnerships, the team envisioned a lunar partnership that would anchor a long-term program within participating agencies and governments.

Understandably, each of the agencies with ISS partnership experience viewed future opportunities in the context of ISS participation. The ISS was nearing the “assembly complete” milestone, and agencies were readying for full utilization of the ISS research capabilities. Moreover, much of the Partner hardware on ISS was deployed in later program phases; so large budget expenditures

²³ In alphabetical order: ASI [Agenzia Spaziale Italiana] (Italy); BNSC [British National Space Centre], now UKSA [United Kingdom Space Agency] (United Kingdom); CNES [Centre National d'Études Spatiales] (France); CNSA [China National Space Administration] (China); CSA [Canadian Space Agency] (Canada); CSIRO [Commonwealth Scientific and Industrial Research Organization] (Australia); DLR [Deutsches Zentrum für Luft- und Raumfahrt] (Germany); ESA [European Space Agency]; ISRO [Indian Space Research Organisation] (India); JAXA [Japanese Aerospace Exploration Agency] (Japan); KARI [Korea Aerospace Research Institute] (Republic of Korea); NASA (U.S.); NSAU [National Space Agency of Ukraine] (Ukraine); and Roscosmos (Russia). “Space agencies” refers to government organizations that are responsible for space activities.

had been recently spent and/or committed. Full utilization of agency investments in ISS was a high priority and was just being realized. Non-ISS partners involved in ISECG were either seeking ISS participation (e.g., South Korea) or contemplating space stations of their own (e.g., China). Recent global economic challenges had left most agencies challenged to meet ISS commitments and left little to no room in governmental budget planning to accommodate lunar exploration on an equivalent scale.

Lessons Learned

NASA led two major WGs—an objective WG and an architecture WG—that sought to develop realistic plans in light of the national constraints described above. Results of these efforts were quite striking as they greatly influenced NASA internal and external planning for lunar exploration. Prior to these interactions, NASA had envisioned a lunar outpost as a base from which to explore the moon; something akin to a “space station on the moon.” In working with potential partners, the international group conceived a new architectural plan that was much more flexible and responsive to national priorities and budget realities. This is described in more detail below (Section 4.9.2).

Recommendations

Include potential International Partners in architectural discussions as early as possible to achieve realistic planning.

4.9.2 Develop a Common Understanding of Goals and Objectives to Inform the Architecture

Supports Key Finding: None

Driving Events

After conducting initial formative architectural discussions via ISECG, it became apparent that an understanding of national goals and objectives was needed to drive formulation of architectural scenarios. NASA proposed a WG within the ISECG to develop this understanding. This group, which was formed in early 2009, was chartered to pursue a common understanding of agency objectives for human lunar exploration.

NASA had spent the previous 2 years working with a series of external and internal groups to develop a set of objectives for lunar exploration. These objectives spanned scientific, technological, economic, cooperative, and educational objectives that could be made possible, and provided benefit via human lunar exploration. While not yet prioritized, these objectives had been used to drive the initial NASA architectural concept of an “outpost,” described above. An outpost from which to explore was believed to satisfy the largest number of NASA objectives. Therefore, NASA was able to articulate goals and objectives for human lunar exploration towards a common international set.

Lessons Learned

Timing was important here. In 2006,²⁴ NASA led an initial effort to develop exploration themes and objectives in collaboration with the international community as well as other stakeholders,²⁵ collecting nearly 200 potential objectives. The apparent similar motivation and rationale for exploration by worldwide space agencies led to the development of the Global Exploration Strategy (GES) framework document and the creation of ISECG. This document articulates five exploration themes shared by participating space agencies. However, meaningful consensus beyond these five themes was not attainable because no other agency at that time had done enough preparation with national stakeholders to engage in discussion to sufficient depth.

By 2009, ESA had conducted a series of workshops with external stakeholders to understand lunar exploration objectives and was also in a position to discuss commonality among interested parties. ESA member agencies (ASI, CNES, UKSA, DLR) had participated in the stakeholder reviews and, thus, were able to provide input to discussions based on the particular interests and capabilities of their agency. CSA was conducting a similar review of exploration objectives. JAXA was in the midst of a national review of lunar objectives and was able to provide insight from a series of external stakeholder discussions during 2009. Other agencies, while not having conducted formal national reviews by 2009, wished to participate in the formulation of a common set of objectives for human lunar exploration. The time was ripe for a group of nine agencies to embark on discussions.

The WG collected more than 600 national objectives ranging from broad, large-scope statements to detailed, investigation-level statements. It became clear that a distillation of these was necessary to achieve an understanding of where commonality among agencies would lie. Something between the five GES themes and these 600 national objectives, traceable to each, was necessary to define the level of commonality achievable at this early, pre-program stage. The WG agreed these statements should be called goals and also agreed on a definition. NASA led the development by providing a set of 10 summary-level statements of goals. These goals were traceable to the full set of NASA stakeholder-articulated objectives as well as the five GES themes articulated by ISECG.²⁶ This set of goals provided a framework within which to discuss commonality. Other agencies followed the NASA lead by providing summary statements of their own. It was from these statements that a common set of 15 goals for human lunar exploration was developed. ISECG endorsed these in late 2009, and they became the drivers for the architectural work to follow.

From the NASA perspective, these goals achieved the international commonality necessary to frame architectural concepts and became an important step forward. Moreover, these common goals did not provide sufficient support to the concept of a permanent outpost on the moon that

²⁴ NASA. [Online] 2006. http://www.nasa.gov/exploration/news/GES_FAQ.html.

²⁵ NASA. [Online] 2006. http://www.nasa.gov/exploration/home/why_moon_objectives.html.

²⁶ Rhatigan, J., Matsumoto, K., Carey, W., Hipkin, V., Kim, H-D., *Developing a Common Set of Human Lunar Exploration Goals and Objectives*, IAC-10.B3.1.11, Sept. 2010.

NASA had espoused. This drove NASA to engage the ISECG agencies in the architectural discussions described below.

Meaningful architectural discussion could be held with a common understanding of goals and objectives in hand. NASA learned that development of common goals and objectives depends on sufficient stakeholder engagement by participating agencies. Most importantly, the common goals allowed NASA to perceive that an “outpost” architecture was not necessary to achieve most agencies goals (see next lesson learned).

Recommendations

Before engaging in architecture discussions, develop a common understanding of goals and objectives among interested agencies. This common understanding will drive the architecture.

4.9.3 Develop an Architectural Framework for Discussion of Partnership Opportunities

Supports Key Finding: None

Driving Events

An architectural WG was chartered by ISECG, at the behest of NASA, shortly after ISECG was chartered. In our view, this WG struggled in its early efforts because there was no common understanding of, or framework for, its discussion. Early efforts focused on developing common international standards and interfaces; but without an architectural framework, these discussions faltered.²⁷

In 2008, NASA announced its intent to build an outpost for human exploration of the moon. During the 2009 ISECG discussions and, in particular, the formulation of the common set of human lunar exploration goals (Section 4.9.2), NASA learned that most agencies would not be able to endorse this type of architecture in the near term, and possibly never.

There were several compelling reasons for this. First, as discussed in Section 4.9.1, agency budgets were highly constrained due to both ISS commitments and the stalled global economy. Commitments in the foreseeable future would be limited by budget realities. Second, most of the ISS partner agencies envisioned expanded roles for themselves in the next major international undertaking.²⁸ Third, the ISS experience had led to an understandable aversion among partner

²⁷ NASA may have been inhibited by its previous experience with ISS, in which NASA determined an architecture and invited other agencies to contribute elements. NASA may have tacitly assumed that the next program would proceed in the same way; i.e., along the “outpost” path. This is explored in greater depth in the next lesson learned.

²⁸ While partner agencies had taken on dependent roles in the ISS Program, that experience had developed extensive new capabilities within these agencies, and they no longer saw themselves as dependent on NASA to provide the same level of leadership and integration expertise.

agencies to a high degree of dependency on NASA.²⁹ This high degree of dependency on NASA, exacerbated by the delay imposed by the loss of *Columbia*, drove a desire to seek a participatory role unlike the “ISS on the moon” concept proposed in the NASA outpost vision. Fourth, and perhaps most importantly, because of the development of the common goals described above, an outpost was not necessary to achieve common goals. Indeed, these goals called for a much different architectural approach.

Via the ISECG, NASA sought to lead the development of an architectural framework, later called the Reference Architecture. Through a series of workshops, 10 agencies participated in its development. The architectural framework was endorsed by ISECG in June 2010.

This architecture is described in detail elsewhere.³⁰ Key elements of the Reference Architecture that were important to NASA include the following:

- It represented a flexible, phased approach for lunar exploration that demonstrated the importance of agencies working together early in program formulation. Neither a lunar outpost nor a series of Apollo-style missions, the Reference Architecture reflected anticipated global realities and challenges while enabling significant scientific and exploration accomplishments.
- It would accommodate innovations in technologies, international priorities, and programmatic constraints over time.
- It enabled opportunities for multiple partnerships coupled with a phased deployment for partners at varying levels of maturity. Thus, it was much more attractive to ISECG agencies.
- It could use launch vehicle assets from many agencies, not just those able to launch humans.
- It was composed of largely independent phases, employing an international range of human-rated and robotic assets over time to explore the moon. These could be started, paused, and/or restarted over time, depending on budget resources.

The Reference Architecture strategy provided continuous human and/or robotic science and exploration activity, spanning multiple locations on the moon, starting from at least 1 year before the first flight crew arrives.

Lessons Learned

The development of an architecture framework, in collaboration with potential International Partners, resulted in a highly viable architecture for NASA. It was much more responsive to both NASA and international community goals and objectives than the initial outpost concept. It was flexible and could accommodate programmatic and budget changes over time. Early engagement of the international community resulted in a much more robust concept for collaborative human

²⁹ International modules were deployed late in the assembly sequence of ISS and could only be launched by the space shuttle. Partner astronauts were transported to and from the ISS by either the NASA space shuttle or the Roscosmos Soyuz.

³⁰ International Space Exploration Coordination Group. *The ISECG Reference Architecture for Human Lunar Exploration—Summary Report*. July 2010.

exploration of the moon. Indeed, this method provides a cooperative framework for any human exploration endeavor in the solar system.

Recommendations

Partnership discussion should be founded on a common architectural understanding, informed by national priorities, budget realities, and common goals and objectives.

4.9.4 Build On Current and Past Successes, But Beware of Past Paradigms

Supports Key Finding: None

Driving Events

Much of the NASA senior leadership engaged in the ISECG activities had spent large parts of their careers on the ISS Program, and were experienced in dealing with the very successful ISS partnership. This provided an invaluable experience base from which to build future partnerships. It also led to some initial limitations in what NASA was able to envision as a future with international partners. The following discussion, while not exhaustive, illustrates some of these hazards.

Lessons Learned

Stability: Perhaps the greatest benefit of the ISS international partnership is the stability it has provided over more than 20 years for the governments and agencies involved in it. Having interdependency among nations obligated long-term funding support among partners and led to a stable program that was sustained through unprecedented political change. Of course, this benefit was not lost on those planning the future program. Stability was a primary reason, along with the benefit of sharing resources, that the Constellation Program sought partnerships.

Maturity: NASA had engaged the ISS partnership at a time, in the mid 1980s, when most other space agencies were in a developmental mode of operation. One attractive reason to join NASA in the ISS partnership was to use the partnership as a mechanism to develop a skill base and learn from NASA. After Roscosmos joined the ISS partnership in the mid 1990s, the benefit was doubled for other agencies—to learn from the two organizations capable of launching humans into space. CSA, ESA (along with its member-state agencies), and JAXA all developed substantial space-faring capabilities over the length of the ISS partnership.

By 2009, and with a fully active ISECG, it was clear those agencies participating in ISS viewed themselves as mature space-faring agencies, no longer in need of the strong leadership role NASA provides on ISS. Indeed, many of these agencies were contemplating the costs and benefits of continued reliance on NASA and Roscosmos vs. developing their own human launch capabilities. And yet for some time, many within NASA assumed the Agency would take the same role it does currently on ISS. The ISECG members expressed the desire for NASA to continue to lead, but in a different way.

This new style of leadership required an internal challenge to NASA: to define a leadership role that recognized current global capabilities and strengths rather than leaning on past paradigms. Being open to this new way of leading resulted in collaborative definitions of goals, objectives, and a robust and flexible architecture described in the previous lessons learned (4.9.3). While the small community of NASA employees engaged with ISECG quickly adapted this style, it represented a cultural change within the Agency. It was not fully articulated and adapted within the broader community in the short time before the Constellation Program was cancelled.

Limiting Opportunities: The programmatic decision, described in the opening to Section 4.9, limited partnership discussions to noncritical path items. This policy resulted in partnership discussions on only the later lunar exploration phase of the Program. It also created offense with long-term and loyal International Partners, that NASA would not discuss future partnerships in the initial planning of its new exploration program.

In retrospect, we can pose a number of questions that challenge this policy. We do not know the answers, and perhaps never will, but the questions concerning this policy are relevant to planning any subsequent human space flight program.

- Did it come about under the ISS paradigm? Did it too strongly reflect past knowledge of global space-faring capability rather than current capabilities?
- Did it weigh too strongly the problems that might arise from “critical path” partnerships at the expense of benefits? Could more have been done to avoid offending long-term international partners?
- Could stability have been added to the Program by allowing “critical path” partnerships on near-term development elements?

Granted, these are retrospective questions and could be viewed as unfairly challenging what was widely believed to be a wise policy at the time. We offer these as food for thought to planners of the next program.

Nevertheless, the Program did not achieve the benefit of stability that an international partnership can provide. NASA may have been hindered by both the late recognition of the contemporary views other space agencies held of themselves and the policy limiting partnership opportunities to noncritical-path items.

Recommendations

Do not “manage to the last program.” Welcome participation by mature international partners in early program formulation and seek roles for less mature partners that are consistent with their capabilities and long-term interests.

5.0 KNOWLEDGE CAPTURE AND CONTRIBUTORS

The entire Constellation Program team contributed to these lessons learned, in both the content of Volume I Executive Summary and that of Volume II.

An executive review committee assisted in synthesizing the lessons learned. This committee was comprised of: Bob Ess (*Ares I-X Mission Manager*), Jeff Hanley (Program Manager 2005-2010), Deborah Neubek (Chief of Staff—Technical), Charlie Stegemoeller (Deputy Program Manager and former Program, Planning and Control Director), Dale Thomas (Program Manager and former Deputy Program Manager), and Teresa Vanhooser (*Ares Project Manager*).

Dale Thomas, Deb Neubek, and Jennifer Rhatigan synthesized and structured these contributions into the Lessons Learned documents (Vol. I and Vol. II).

The Program and its projects conducted knowledge-capture activities throughout its duration. These included many L6S events, a multitude of technical reports, conference papers, and final lessons learned activities. All formal program documents, reports, review materials, lessons learned reports, etc. will be archived with NARA and made available in the ESMD ICE system.³¹ Those lessons learned that are deemed key risks are being fully documented via text and video interviews as knowledge-based risks (KBRs). In addition, hundreds of lessons learned are being entered into the NASA Engineering Network's Lessons Learned Database (NEN LLDB). To make finding all of these materials easier, the ICE system will include a Constellation Program Knowledge Management Webpage that will link to the various topics.

Current URLs for these resources are below:

Exploration Systems Risk and Knowledge Management Portal:

<https://ice.exploration.nasa.gov/ice/site/km>

NASA Engineering Network (Click on the Lessons Learned Tab): <https://nen.nasa.gov/web/nen>

Public NASA Engineering Network Lessons Learned: <http://llis.nasa.gov/llis/search/home.jsp>

In addition to the hundreds of Constellation team members who contributed to the content of this report, the following played a significant role in synthesizing that data in the development of these lessons learned:

Jose Amador, Kathy Andrews, Bill Arceneaux, David Bell, Joe Caram, Sean Carter, R. H. Coates, John Connolly, Alan Dickey, Dave Hall, Linda Ham, Fred Kuo, Kathy Laurini, Burt Laws, Jon Lenius, Charlie Mallini, Deb Neubek, Carlos Noriega, Ned Penley, Jennifer Rhatigan, Doug Sander, Beto Sanchez, Debbie Stapleton, John Turner, and Keith Williams.

³¹ In researching lessons learned from previous programs, primarily the Apollo and Space Shuttle Programs, a number of key documents were scanned into electronic form for wide distribution within the Program and archived in the Constellation database. These documents are available to interested readers as well.

APPENDIX A. *ALTAIR* TOP LESSONS LEARNED

1. Consider an In-house Skunkworks®-like Team to Initialize New Projects

Supports Key Finding: 3.1

Driving Events

A study was performed in late 2006 to identify options for initiating the development of a lunar lander to meet the Human Lunar Return requirements in the National Space Exploration Policy. That study determined insufficient budget was available to execute a lunar lander development project using a standard NPR 7120.5 process with contractors performing pre-phase A and Phase A/B studies. The Lunar Lander Project Office (LLPO) was formed to implement a more streamlined and efficient path through the project formulation phase. Key tenets of the approach included the following:

- NASA Administrator buy-in to the approach and broad latitude to deviate from NASA policies, as needed
- A small, handpicked team made up of spacecraft and subsystem design experts from across the Agency
- A set of guiding principles that included a small set of hard requirements and a clear goal of developing a spacecraft specification based on an in-depth understanding of every additional requirement levied
- Project start-up by collocating the multicenter team in a single facility for the first 2 months of the project, which maintained that cohesiveness with regular, short-term colocated meetings

Lessons Learned

A small, focused multidisciplinary in-house NASA team is a very effective way to initiate new projects. The small team was able to provide *Altair* project insight to impact integrated lunar-phase Constellation activities at very low relative cost.

Recommendations

New projects should begin with a small, focused team that uses the LLPO/*Altair* experience as a basis for initial project formulation.

2. Use Models-based Systems Engineering

Supports Key Finding: None; supports recommendation 4.3.3

Driving Events

The Constellation Program initiated the requirements development process by determining the capabilities needed from each of the systems in the architecture to accomplish the mission, and levied requirements through various documents, including the CARD. As a project within a large

program, it is the responsibility of the project to implement a more detailed assessment of the mission to validate requirements levied on it from the program and determine whether these requirements are necessary and achievable, as well as to scrutinize the mission for latent requirements. The *Altair* project used MBSE as the approach to requirements development. MBSE focuses on building data models that clearly depict the operational flow of the mission. These hierarchical models manage the complexity of the mission and functional allocations. They make excellent integration products for a group review between different organizations for consistency in assumptions. The operational models can then be analyzed to determine the functionality the vehicle must possess to enable these operations. These models are kept in the requirements database for the program, and cross-references between the vehicle functions and mission phases in which the functions are used serve as validation of the function. Further, linkages between functions and requirements that enforce them on the design provide complete traceability from requirements to the concept of operations. Finally, by assigning durations to the activities within the operational models, activities can be simulated to provide a complete timeline of the mission from the same models that are being used to establish vehicle functionality. By developing these common products, a single model set becomes the authoritative source for the requirements, functionality, and operations, thereby improving the quality of the requirements, integrating the various organizations within the projects, and minimizing disconnects.

Lessons Learned

Model-based Systems Engineering streamlines requirements development and validation by cultivating integrated products among the Requirements, Design, and Operations communities.

Recommendations

Model-based Systems Engineering should be used to provide an integrated set of products on which architecture, element, system, and subsystem requirements can be sequentially decomposed.

3. Start Risk-informed Design during Conceptual Design

Supports Key Finding: 3.5

Driving Events

The initial lunar design analysis cycle (LDAC-1) for the *Altair* lunar lander provided a spacecraft with “minimum functionality.” Minimum function was defined to mean that the lander was designed to meet the primary key driving requirements, not the complete set of Level II Lander requirements. Primarily, it meant that the vehicle and its subsystems were designed with zero fault-tolerance and with no consideration for contingency operations. It was understood that this design did not represent a flyable vehicle; rather, this provided a theoretical vehicle that could perform a lander mission if everything worked with 100% reliability, obviously an unrealistic design point. However, during LDAC-2, the risks that could result in LOC failures were identified and prioritized, and specific vehicle, subsystem, and component trade studies were performed to

identify the amount of reliability improvement needed as a function of mass of the alternatives. The resulting vehicle concept was referred to as the Increased Safety vehicle. This process was repeated during LDAC-3 for the LOM risks to provide the Enhanced Mission Success vehicle. A risk assessment tool was developed using component failure data to quantify risk associated with subsystem and component alternatives. This included reliability data from the space shuttle and ISS databases, the lander operational timeline to define mission phases and component operating times, and a vehicle preliminary hazard assessment to determine which components could lead to a LOC. The term “risk-informed design” is used because the design decisions were literally “informed” by the quantified risk information instead of using a rule-based decision process. This gave the project team much greater insight into risk drivers at an early stage of project development.

Lessons Learned

A risk-informed design process provides the project with early, in-depth insight into risk drivers so that safety, reliability, and capability functions can be systematically added to a minimum function design baseline. This allows a more complete understanding of integrated effects on other primary constraints, such as spacecraft mass.

Recommendations

A risk-informed design process should be used early in the project formulation phase to achieve efficient design and understand major design drivers.

4. Involve Industry during Phase A

Supports Key Finding: None

Driving Events

One of the early tenets of the LLPO was to solicit input from industry throughout the project formulation phase. In June 2008, a little over 1 year after the LLPO was created, the office awarded broad area announcement (BAA) contracts to five companies: three major traditional aerospace companies and two small, independent companies.

The BAA had two primary objectives. The first asked the companies to review the government LDAC-1 minimum-function lunar lander conceptual design and plans for implementing risk-informed design, and to provide suggested alternative design concepts. The second asked companies to provide recommendations for the roles and responsibilities of government and industry for lunar lander development. A process was then set up whereby the companies that participated in the BAA and other entities meeting International Traffic in Arms [ITAR] export control regulations could continually obtain updates on the government design.

This two-way interaction was extremely valuable to both government and industry. It provided the government some important alternative perspectives and industry with information that helped focus internal investment funding. The input from industry also helped shape the roles of

the government and industry that were incorporated into the *Altair* acquisition strategy. As the *Altair* project proceeded into Phase A concept exploration and refinement, and while preparing for the SRR, this two-way interaction developed a process whereby multiple contractors would be integrated into the *Altair* team to support the work.

This process was the basis of a BAA contract that was in preparation for release at the time Constellation was cancelled. This process would have maintained government leadership of the design until a prime contractor was selected for flight vehicle development, enabling the government team to develop a better Request for Proposals (RFP) solicitation and allow all contractors greater insight into the information that the government was using to shape the requirements.³²

Lessons Learned

Industry should be brought into the formulation process of a project early to ensure the NASA project team is aware of alternative ideas and perspectives, and help the contractor community better understand the basis of the requirements.

Recommendations

While the acquisition model that *Altair* was developing may not be appropriate for widespread application, the primary recommendation is that the project team should explore unique and creative ways for incorporating the industry community into the project as early as possible.

5. Making Multicenter Teams Work

Supports Key Findings: 3.6, 3.9, 3.10

Driving Events

When *Altair* began as the LLPO, it collocated approximately 35 people in a single, small building at JSC for a period of 8 weeks. This period allowed the team to get through most of the “forming, storming, norming,” and begin performing as a high-performance team before returning to home centers and organizations. The project maintained team cohesiveness by planning a travel budget that allowed a significant part of the team to get back together for a minimum of 3 to 4 days once a quarter to work together. These periodic collocations were rotated through the centers; and, to the extent possible, the team stayed in the same hotels, dined together, and spent time socializing. This approach ensured that the personal relationships that were initially formed were maintained, and that as new members joined the team, they were more quickly assimilated. The interpersonal bonds were tested as inter-center institutional competition for roles and responsibilities emerged; however, the strength of the team prevented impacts to progress.

³² The *Altair* approach to NASA and industry collaboration was specifically cited as a good model by two contractors responding to the Heavy Lift and Propulsion Technology Request for Information (HLPT RFI).

Lessons Learned

Diversity provided by a multicenter team is worth extra management effort. While teleconferences and Internet-assisted virtual meetings are required, the key to operating with a multicenter team is to periodically bring the team together to work in a single location. Regular collocations/face-to-face meetings with a significant portion of the team, including all key system, subsystem, and discipline leads, are essential to maintain team cohesiveness when using multicenter teams.

Recommendations

Specific considerations that enabled the LLPO/*Altair* multicenter team to be highly effective included: 1) leadership of key portions of the spacecraft design and analysis was assigned to different centers; 2) initial and subsequent collocations all included social events designed to allow people to get to know each other; 3) collocations were scheduled, to the maximum extent possible, to avoid conflicts with family vacations—i.e., consideration was given to when K-12 school breaks; and 4) collocation meetings were well planned to make maximum use of the time together.

6. Design during Pre-phase A and Early Phase A

Supports Key Finding: None

Driving Events

One of the key initial tasks for the LLPO was to develop a preliminary in-house design within 6 to 9 months, with emphasis on performing significant design; i.e., focus on the “D” in DAC [design analysis cycle]. The design work allowed the team to validate and identify weaknesses in the parametric modeling. A few notable examples, as follows can illustrate this point:

- The mass of the landing legs could not be easily scaled from the Apollo lunar module due to the dramatic size difference between it and the hydrogen-fueled *Altair*. By performing detailed analysis to determine the required size for stability at the bounding landing site terrains and then developing and analyzing a design, we were able to improve mass confidence.
- As the team used the design and began identifying assembly, integration, and test during both development and operations, it identified the Constellation Program decision, which was made early in LDAC-2 to increase the descent module diameter to take advantage of the 10-m *Ares V* shroud, had significant implications for transportation and thermal-vacuum propulsion system testing. That assessment found the *Altair* descent module could not be transported in the NASA Guppy aircraft (the largest available) but would require barge transportation. The barge would require an icebreaker to clear a path through Lake Erie during winter to perform an acceptance test of each descent module propulsion system in the Plumbrook Station B2 test facility.

These examples represent some of the more dramatic items that were revealed by allowing the NASA in-house team to perform significant design work during the project-formulation phase.

Lessons Learned

Design work during pre-phase A and early Phase A can identify important issues that cannot be determined until a design concept is of sufficient maturity to evaluate a more complete set of development, test, and evaluation considerations.

Recommendations

A design concept should be developed, as early as possible in the project formulation phase, to allow a more complete understand of programmatic implications. The concept needs to be more than an artist rendering or a parametric design. It needs to have engineering design analysis substantiating it so that a design team is made of SMEs with the experience to foresee potential issues during later project phases.

APPENDIX B. GROUND OPERATIONS TOP LESSONS LEARNED

This Appendix contains the top lessons learned from the Constellation Ground Operations Project. These are linked to main volume key findings as applicable.

1. Clearly Define Roles and Responsibilities within the Program/Project

Supports Key Findings: 3.7, 3.8, 3.9

Driving Events

As discussed in the main volume, the late start of the Program relative to that of the projects led to unclear RR&A. Many WGs were solving the same issues in parallel with other WGs across the Program (e.g., logistics). Furthermore, requirements and expectations were unclear due to lack of formal documentation. Some examples are provided for illustration:

- Both the Ares project and Ground Operations Project believed they were responsible for the processing timeline for operations at KSC because both had CARD requirements that were to be verified via the KSC processing timeline. Therefore, both organizations developed and maintained what each considered the “Constellation Program Ground Operations Timeline.” These timelines were not consistent; each required a significant effort to maintain. Resolution of discrepancies therefore required lengthy discussion.
- Management of integrated and stand-alone flight hardware testing at KSC (which involved/impacted multiple projects) was not actively controlled at the Program. There was no consistent definition of which flight hardware testing was going to be performed at KSC.
- While there was sufficient participation by the projects to define what testing was required within the launch vehicle and the spacecraft during development of integrated test objectives by the Flight Element Integrated Test WG, there was very little (if any) analysis of the integrated system to determine testing between the elements and the effects of testing on one element with the others (e.g., integrated vehicle test, end-to-end testing).
- Maturing requirements caused significant rework and delays. For example, vehicle changes drove costly changes to already designed and fabricated ground systems. Different organizations were working to different vehicle configurations and technical baselines at the same time.

Lessons Learned

Unclear accountability led to duplication of effort among projects. Integration of multiple projects must remain the responsibility of the overarching program and not be delegated to one (or more) supporting projects. The specific responsibilities of each of the projects must be clearly defined and socialized among and between all stakeholders to minimize confusion and inefficiency across the program.

Engineering analysis of the integrated launch vehicle and spacecraft must be performed and used to define integrated test objective details for testing at KSC. Program baseline synchronization activities should be done periodically and often to assure that all are using the same baseline.

Recommendations

The program should clearly define roles and responsibilities of the individual projects as soon as possible and reinforce those roles regularly, revisiting as necessary. WGs involving multiple projects should be chartered at the program level with clearly defined roles and responsibilities to avoid duplication of effort and maintain clear accountability.

Maintain at the program level complete responsibility for performing analytical integration for the entire launch vehicle and spacecraft, and determining test objectives/requirements for the overall integrated vehicle.

Synchronize vehicle and ground system design to the same baseline. Program synchronization with all systems (mission systems, EVA, etc.) should be performed quarterly, at a minimum, to ensure goals and objectives are met. No metal should be cut until simulation models and schedules are coordinated.

2. Streamline Processes to Improve Communication Across Program/Projects

Supports Key Findings: 3.9, 3.10

Driving Events

Two different processes were levied for records management for the Ground Operations Project: the KSC Business Records Template, governed by NPD 1441.1D; and the program “catalog” (i.e., access database) with many additional attributes for program records.³³ This resulted in an unstructured change management process that drove costs and schedules. Because of change management process immaturity, unfunded requirements were levied on the Ground Operations Project and costs were not adequately quantified. The Program created change processes (Change Requests, Document Change Notices [DCNs]) to expedite and document changes. The DCN, for instance, was intended to push through Program-level changes quickly. Instead, the change processes merged into one cumbersome process that required weeks or months before changes could be instituted.

During change evaluations, multiple document revisions appeared during the evaluation period. This caused confusion and wasted technical resources. While the Program had an approved formal Technical Coordination Meeting process, its application varied between Change Package Managers (e.g., virtual, email only, telephone only). As it was difficult to give direction to Change Package Engineers, complaints were often elicited from them.

³³ Similar situations existed at each NASA center supporting the Constellation Program.

An organization-wide common process and tool for agreements regarding exchange of hardware, software, data, etc. was not implemented until late in the Program lifetime. The Constellation Program bilateral exchange agreement (BEA) tool was useful and user friendly with very responsive developers; however, there was much controversy over the use of the term “bilateral exchange agreement” (originated for agreements between countries). Further issues involved the process itself, where the confusion was over whether a BEA needs to be approved prior to inclusion in the schedule and enforcement of the agreement (e.g., at what point and by whose authority is a providing organization directed that due dates cannot change).

Lessons Learned

The use of new and/or different processes can add churn and reduce effective communications. It is best to adopt known and proven processes for change management and standardization of agreements within a program. Better processes and communication will eliminate confusion and frustration on agreements.

Recommendations

Develop and document a robust change management process from inception that precludes levy of unfunded requirements.

The Agency should develop NASA-wide processes for records management. Implement this at the start of any new program, with continuous monitoring for any new additional requirements or improvements.

Create one change control process that can be used to make changes to project/program-level documents (minimize the number of change control processes). This will eliminate unnecessary overlap and confusion in the change control process. Both vehicle and ground should proceed from an agreed baseline such that all stakeholders are expeditiously notified when changes occur.

Establish clear and concise processes up front pertaining to agreements.

3. Use a Consistent Financial Approach to Ensure Emphasis on Life-cycle Costs

Supports Key Finding: None

Driving Events

The acquisition strategy did not address program life-cycle supportability due to the late start of the Program relative to that of the projects (Section 2.2.6). Each prime contractor defined life-cycle supportability strategies that were inconsistent across projects. This resulted in different expectations for maintenance concepts, and support infrastructure required at the launch site.

Further, the involvement of all NASA centers (Section 2.2.3) and multiple contractors introduced confusion, since each center/contractor used somewhat different definitions for general accounting terms (e.g., WYE [work year equivalent]) to characterize work content.

Program requirements (formats, actions) tended to focus on cost data, while the Ground Operations Project focused more on obligations (which are more representative of work/content in process). This difference in focus impacted the timely receipt of funding. Program shifts/adjustments mandated on project cost could not be reconciled/traced to NOA [new obligation authority]/obligations. Projects performed a thorough bottom-up estimate, significantly focusing on obligations; but the same emphasis on cost is not undertaken, which leads to suboptimal communication with stakeholders.

Liens were captured as cost threats in the same manner as potential risk impacts.

Lessons Learned

Supportability strategies and concepts for flight and ground systems must be defined early in the program, as part of the acquisition strategy, and incorporated into the RFPs to influence designs and have a positive effect on program life-cycle costs.

The Agency, program, and projects depend on the information provided to make prudent decisions. Consistent and well-defined terminology between all levels and projects, as well as an understanding of these definitions, would increase accuracy and efficiency in providing the data and enhance analysis. Clear definitions and consistency should be used in reporting WYE numbers. A properly developed and used work breakdown structure (WBS) will enhance efficiency, minimize errors, and lead to better analysis of work content as related to budget and schedule.

An increased emphasis on NOA/obligation requirements by the program would lead to a better understanding of the work/content being undertaken in relation to monetary requirements and schedule. An increased emphasis on cost requirements by the project would lead to a better ability to communicate with external stakeholders by providing a more thorough and complete understanding of technical work and milestones accomplished in relation to budget and schedule.

The Program mixed potential risk impacts with liens/shortfalls. This turned the risk system into a budget tool, which did not allow analysis of potential cost threats vs. unfunded mandates.

Recommendations

Develop an acquisition strategy early in the program that includes life-cycle supportability strategies and concepts prior to release of RFPs. Include strategies and concepts in the RFPs and ensure that if contract negotiations make any changes to one contract they are changed the same way in all the contracts, as applicable. Ensure that guidance conferences emphasize the need to keep supportability, logistics, and maintenance concepts the same.

Clearly define all important financial terms across the Agency and assure consistent use. Develop and use WBSs at a level at which financial actuals are received, characterize the work content (regardless of center), and reflect the planning and execution structures (e.g., by phase, etc.).

Program PPBE and phasing plan requirements should include NOA/obligation requests, as appropriate. Project stakeholders supporting PPBE and phasing plans should forecast cost requirements more thoroughly.

Capture true risk-related cost threats as such, separating liens (i.e., change request impacts) or remove liens from the risk system. Financial impacts, as part of a change process, should be addressed on approval rather than driven into the risk system.

4. Improve Interface Product Integration

Supports Key Finding: Supports the data management recommendations, Section 4.3

Driving Events

The following are examples of difficulties encountered in negotiating Interface Control Documents/Interface Requirements Documents (ICDs/IRDs):

- Flexibility was not permitted in product templates for inclusion of data as tables or figures in lieu of verbiage, thus introducing confusion.
- The same requirements were allocated to multiple ground and/or vehicle elements, creating confusion.
- Required drawing exchanges were not defined early in development work for the interfaces.
- It was difficult to navigate and search interface data in the chosen document format.
- It was difficult to obtain quick updates of approved interface data or to easily communicate interface data to designers or new employees.
- Interface change impacts were not well balanced between projects. The *Ares* project was often able to implement requirements changes without showing detailed rationale illustrating that several concepts had been explored.

Lessons Learned

Flexibility in product templates (ICD/IRD), tailored for each interface, can improve the outcome of the interface.

Real dimensional measurements are needed in the ICDs to control the design on both sides of the interface properly.

A single neutral party should manage the project-to-project interfaces. Each project should be responsible for bringing forward full and consistent change packages to the Interface Manager for discussion prior to implementation.

Recommendations

Allow the interface team to define the most effective template (this may be unique for each interface). Organize interfaces by project, and allow projects to sub-allocate requirements.

A Website site navigable by a graphical representation of the vehicle and ground should be used to specify all data for each interface, with requirements generated from a requirements database.

Define standard drawing/model exchanges early on in the development work for the interface. Use a controlled-access Website to specify all data for each interface (similar to system used by shuttle). Provide easy navigation to interface data: requirements, drawings, etc.

Provide real dimensional measurements in the ICDs.

Each project should have an office dedicated to interface management. The project-to-project interface requirements baseline should be controlled by a neutral board/unbiased interface management group, and not at the project-to-project level.

5. Address Logistics and Supportability in a Consistent Manner

Supports Key Finding: None

Driving Events

As discussed in the main volume, projects initiated contracts prior to formation of the Program. These organizations were not mature enough to adequately specify logistics and supportability requirements. Logistics tools for the *Orion* and *Ares* projects were developed independently from the Ground Operations Project with the end result being different final products. The Ground Operations Project participated in the WGs for both *Orion* and *Ares*. It became evident, in working with the contractors, that the supportability requirements and direction levied on the contractors was not the same. This caused both groups to rework some products to allow them to interface with each other.

A comparison was done between the Program supportability plan and the project supportability plan to map requirements between documents. Gaps and disconnects were uncovered between the plans.

Lessons Learned

Critical logistics data items such as source maintainability and recoverability code charts need to be common across projects to facilitate parts transfer.

Ensuring continuity between parent-child documents is important to prevent confusion about requirements.

Mismatched logistics requirements included in projects RFPs caused rework later while trying to align acquisition/performance guidance between projects.

Additional supportability requirements that affect life-cycle costs need to be defined in program-level requirements documentation. Life-cycle costs are affected by supportability factors such as

commonality and interchangeability. If these factors are not stated as requirements, the design organization is not compelled to incorporate revisions or suggestions (guidance) into the design.

Recommendations

Identify essential logistics infrastructure data items that need to be common across projects, and specify data items in a program requirements document.

Ensure that the program-level supportability plan is used as a parent document when creating project-level plans, and routinely perform a gap analysis between the parent-child documents.

Develop a standard set of supportability requirements to be applied to all projects prior to contract initiation.

Performance requirements can be traded against other logistics and supportability requirements for the system so the program will have a more accurate assessment of life-cycle costs vs. system performance.

APPENDIX C. ACRONYMS

APA	allowance for program adjustment
APM	associate program manager
ASAP	Aerospace Advisory Panel
ASI	Agenzia Spaziale Italiana
BAA	broad area announcement
BEA	bilateral exchange agreement
CAD	computer-aided design
CAIB	<i>Columbia</i> Accident Investigation Board
CARD	Constellation Architecture Requirements Document
CDR	Critical Design Review
CM	configuration management
CMG	control moment gyroscope
CNES	Centre National d'Études Spatiales
CO	Contracting Officer
CoF	construction of facilities
CoFR	certificate of flight readiness
CoP	community of practice
CR	continuing resolution
CSA	Canadian Space Agency
CxSIP	Constellation System Integration Panel
D&C	design and construction
DAC	design analysis cycle
DCN	Document Change Notice
DES	discrete event simulation
DFMR	design for minimum risk
DLR	Deutsches Zentrum für Luft- und Raumfahrt

DM	data management
DoD	Department of Defense
DRM	design reference mission
DRD	data requirements documents
ESA	European Space Agency
ESAS	Exploration Systems Architecture Study
ESMD	Exploration Systems Mission Directorate
EVA	extravehicular activity
EVM	earned value management
FEP	front-end planning
FY	fiscal year
GES	Global Exploration Strategy
HLPT RFI	Heavy Lift and Propulsion Technology Request for Information
HLR	human lunar return
HQ	Headquarters [ie, NASA Headquarters]
HRCP	human-rating certification packet
IC	Initial Capability
ICD	Interface Control Document
IDAC	Integrated Design Analysis Cycle
IMS	Integrated Master Schedule
IOC	Initial Operational Capability
IRD	Interface Requirements Document
ISECG	International Space Exploration Coordination Group
ISO	International Organization for Standardization
ISS	International Space Station
IT	information technology
ITAR	International Traffic in Arms
JAXA	Japanese Aerospace Exploration Agency

JCL	Joint Confidence Level
JPR	Johnson (Space Center) Program Requirement
JSC	Johnson Space Center
KBR	knowledge-based risk
KSC	Kennedy Space Center
L6S	Lean Six Sigma
LAS	Launch Abort System
LC	Lunar Capability
LDAC	lunar design analysis cycle
LEO	low-Earth orbit
LLPO	Lunar Lander Project Office
LMI	Logistics Management Information
LOC	loss of crew
LOM	loss of mission
MAF	Michoud Assembly Facility
MBSE	Model-based Systems Engineering
MIVP	Master Integration and Verification Plan
MOU	Memorandum of Understanding
MPCV	multipurpose crewed vehicle
MPR	Monthly Program Review
MSFC	Marshall Space Flight Center
NOA	new obligation authority
NPD	NASA Program Directive
NPR	NASA Program Requirement
NSA	National Security Agency
O&C	operations and checkout
OCE	Office of Chief Engineer
OMB	Office of Management and Budget

OSMA	Office of Safety and Mission Assurance
PBS	President's Budget Submittal
PDR	Preliminary Design Review
PET	Program Excellence Team
PI	Program Integration
PMR	Program Manager's Recommendation
PP&C	Program Planning and Control
PPBE	planning, programming, budgeting, and execution
PRA	probabilistic risk assessment
PV	photovoltaic
R&D	research and development
RFP	Request for Proposals
RIT	Risk Informed Test
RM	risk management
RR&A	roles, responsibility, and authority
S&MA	Safety and Mission Assurance
SAGES	Shuttle and Apollo Generation Expert Services
SD	storage device
SDR	Systems Definition Review
SE	systems engineering
SE&I	Systems Engineering and Integration
SI	international system of units
SIG	System Integration Group
SLS	Space Launch System
SM	service module
SME	subject matter expert
SOW	Statement of Work
SR&QA	Safety, Reliability, and Quality Assurance

SRB	Standing Review Board
SRR	Systems Requirements Review
SSP	Space Shuttle Program
T&E	Test and Evaluation
TIM	Technical Interchange Meeting
UKSA	United Kingdom Space Agency
VCS	Voluntary Consensus Standards
VM1	Virtual Mission-1
WBS	work breakdown structure
WG	working group
WYE	work year equivalent
XML	extensible markup language