

## Source documents

1. Apollo Lessons Learned - Contamination Control for Samples
2. ATLAS Beam Steering Mechanism Lessons Learned
3. Challenger and Columbia Lessons Learned
4. Commissioning MMS - Challenges and Lessons Learned
5. Constellation Program Lessons Learned - Detailed Lessons Learned - Volume 2
6. DSCOVR Contamination Lessons Learned
7. Earthdata Search UX Lessons Learned
8. EMU Lessons Learned Database
9. Human Spaceflight Conjunction Assessment - Lessons Learned
10. JWST Lessons Learned
11. KSC Human Factors Lessons Learned
12. Lessons Learned Briefing PLSS 2.0
13. Lessons Learned from Recent Space Flight Assessments
14. Lessons Learned in Building the Ares Projects
15. Lessons Learned in Engineering
16. Logistics Lessons Learned in NASA Space Flight
17. Lunar Reconnaissance Orbiter - Lessons Learned
18. M-C 1 Engine Valves, Lessons Learned
19. NASA Materials Related Lessons Learned
20. NASA Procurement Lessons Learned
21. Operational Lessons Learned from NASA Analog Missions
22. Orbiter Lessons Learned
23. Personal Air Vehicle Lessons Learned
24. Ramp Traffic Console (RTC) Lessons Learned
25. Solar Dynamics Observatory Lessons Learned
26. SRMS History, Evolution and Lessons Learned
27. Team Collaboration - Lessons Learned Report
28. X-43A Lessons Learned
29. X-57 Maxwell Lessons Learned Report
30. X-Plane Structures Challenges Lessons Learned

All source documents can be found here: <https://ntrs.nasa.gov/>

# Stakeholder Identification & Engagement

## What Works Well

### **Cross-disciplinary and multi-organisation teams**

Several projects brought together stakeholders from NASA centres, contractors, academic institutions, and international partners. For example, the Magnetospheric MultiScale (MMS) commissioning involved NASA Goddard, multiple universities, research institutes, and industry teams. Roles were clearly defined — mission management, payload operations, instrument teams, and integrators — which allowed stakeholders to focus on their expertise while contributing to a coordinated commissioning plan.

### **Defined roles and responsibilities**

The MMS project explicitly documented commissioning roles (e.g., Mission Operations Center, Payload Operations Center, Instrument Team Facilities), showing a conscious effort to map stakeholders to operational responsibilities. This structure appears to have supported coordination despite geographic separation and time zone differences.

### **Independent oversight for safety and quality**

Post-accident responses to Challenger and Columbia included stronger independent program oversight and safety review functions, as well as establishing technical authorities outside of the program's direct chain of command. These mechanisms helped ensure critical safety perspectives were not lost amid schedule or budget pressures.

### **Engagement for requirements verification**

Some technical teams (e.g., ATLAS Beam Steering Mechanism) engaged closely with project science teams early in planning to agree on verification approaches and acceptance criteria. This early sign-off reduced disputes and facilitated smoother testing phases.

---

## What Could Be Done Better

### **Listening to dissenting or cautionary input**

The Challenger and Columbia accident investigations repeatedly found that stakeholder input — especially from engineers and contractors raising concerns — was either not escalated, not acted upon, or diluted in communication up the management chain. Decision-making processes sometimes required engineers to prove a hazard was unsafe, rather than requiring proof it was safe.

### **Maintaining effective communication channels**

Findings pointed to breakdowns in communication between different levels of

management, contractors, and technical teams. In some cases, ambiguous or incomplete recommendations (e.g., Rockwell's input on ice-on-pad concerns before Challenger) led to misinterpretation or inaction.

### **Avoiding over-centralisation under schedule pressure**

Both accident boards noted that schedule pressure could cause decision-making to become insular, with program offices overriding or bypassing independent review. This weakened stakeholder engagement by discouraging full, open discussion of risks.

### **Ensuring safety and quality functions have true independence**

Safety, reliability, and quality assurance offices were sometimes embedded within the same management structures they were meant to oversee, reducing their ability to act as independent stakeholders. This organisational placement limited their effectiveness in challenging unsafe practices.

### **Maintaining stakeholder continuity and capacity**

Budget and workforce reductions, particularly before Columbia, eroded safety and quality teams' ability to engage meaningfully. Reduced staffing meant less capacity for trend analysis, anomaly tracking, and proactive risk engagement — functions critical for informed decision-making.

---

## **Business Case Development**

### **What Works Well**

#### **Clear mission objectives tied to scientific or operational needs**

NASA projects generally had a strong initial case rooted in clear science or mission objectives. For example, the ATLAS Beam Steering Mechanism was justified by its requirement to meet sub-arcsecond pointing accuracy to achieve the ICESat II mission's high-resolution mapping goals. Similarly, the MMS mission configuration flowed directly from the scientific requirement to study magnetic reconnection, dictating a four-spacecraft tetrahedron and specific instrument suites. These technical needs provided a concrete basis for resource allocation and design priorities.

#### **Willingness to reassess sourcing decisions when vendor solutions underperform**

In the ATLAS BSM case, the original plan to use a vendor-supplied mechanism was reconsidered when the vendor's hardware failed environmental testing. Bringing the work in-house eliminated compromises to optical requirements and enabled a better technical outcome, even though it increased internal workload. This shows that NASA can pivot business cases midstream when new information undermines the original value proposition.

## **Integration of safety and risk considerations into project justification**

After Challenger and Columbia, business cases for hardware and procedural changes increasingly incorporated independent safety oversight, updated hazard analyses, and organisational changes as essential project enablers. This indicated a shift toward recognising safety assurance as a critical business driver rather than a separate compliance activity.

---

## **What Could Be Done Better**

### **Full lifecycle cost and schedule realism**

In several cases, the business case underestimated schedule or cost implications of technical requirements. For instance, tight timelines on in-house developments (after vendor failures) created significant schedule pressure. While technical quality was preserved, the cost and time impacts were not fully offset in the original justification. Future cases could better incorporate risk-adjusted cost and schedule forecasts.

### **Maintaining performance requirements under cost pressure**

Before some redesigns (e.g., BSM), compromises to requirements were made to fit vendor capabilities in an attempt to reduce cost. These compromises weakened the original mission case and had to be reversed later, adding rework. This suggests the need to safeguard core performance parameters in the business case and avoid cost savings that degrade mission-critical outcomes.

### **Accounting for organisational and cultural factors in the case**

The Columbia Accident Investigation Board found that budget cuts, workforce reductions, and organisational structures undermined safety oversight and decision-making. These “soft” factors were not always part of the upfront business case, even though they directly affect the ability to deliver the mission safely and on schedule.

### **More robust evaluation of contractor readiness**

Vendor selection sometimes assumed that heritage hardware or prior experience equated to readiness for flight qualification. As seen in the BSM case, design shortcomings only became clear after environmental testing. Business cases could be strengthened by incorporating deeper technical due diligence on supplier capabilities and readiness before committing to procurement.

---

## **Scope Definition & Management**

### **What Works Well**

#### **Requirements driven by clear mission needs**

NASA projects in the binder consistently defined scope starting from concrete mission

objectives. For example, MMS's need to map magnetic reconnection in 3D at multiple temporal scales directly set the scope for a four-spacecraft tetrahedral formation, specific orbital parameters, and a suite of 26 payload components per spacecraft. Similarly, the ATLAS BSM scope was tied to achieving sub-arcsecond pointing accuracy — a requirement dictated by the ICESat II mission's science goals.

### **Early alignment between technical and operational teams**

In several cases, teams engaged early with science leads and technical stakeholders to agree on scope parameters and verification approaches. The BSM project, for instance, had its performance verification plan approved by the science team early, giving clarity on scope boundaries and acceptance criteria.

### **Scope adaptation when technical realities change**

There is evidence that NASA can adjust scope when initial solutions prove inadequate. The decision to bring the BSM in-house after vendor failure expanded scope in terms of in-house workload but restored mission requirements that had been compromised in the original vendor-based scope.

---

## **What Could Be Done Better**

### **Guarding against requirement compromise to fit constraints**

Before the BSM was brought in-house, optical performance requirements were reduced to match vendor capabilities. While this was intended to simplify delivery, it ultimately undermined mission objectives and required scope reversal. This highlights the risk of redefining scope to fit supplier limitations rather than mission needs.

### **Scope realism under schedule and resource pressure**

Projects sometimes faced compressed schedules or staffing reductions that were not fully reconciled with the defined scope. In Columbia-era findings, safety oversight capacity was reduced even though scope still included critical risk management functions. This mismatch created delivery and safety risks.

### **Clearer integration of risk contingencies into scope**

Some programmes underestimated the potential for anomalies, redesigns, or late hardware changes. For example, unforeseen manufacturing tolerance issues with BSM flexures caused significant rework. Scope planning could more explicitly account for contingency activities and buffer time to accommodate such issues without derailing downstream phases.

### **Managing cross-organisational dependencies**

MMS commissioning involved numerous geographically dispersed teams with tightly sequenced activities. While the scope covered all required tasks, last-minute schedule changes due to anomalies or late inputs created strain. This suggests scope definitions

could better model interdependency risks and include built-in flexibility for concurrent operations.

---

## Schedule Development & Control

### What Works Well

#### **Detailed, role-based scheduling in complex missions**

For MMS commissioning, schedule development was highly structured. Long-term (strategic) and day-to-day (tactical) schedules were produced by dedicated planning teams, with clear separation of responsibilities between the Mission Operations Center, Payload Operations Center, Instrument Team Facilities, and mission planners. This structure supported coordination across multiple organisations, facilities, and time zones.

#### **Sequencing to prevent operational conflicts**

The baseline MMS commissioning schedule accounted for strict sequencing of operations to avoid interference between instruments and between spacecraft. This level of forethought helped reduce technical conflicts, especially during deployments and high-voltage activations.

#### **Flexibility to incorporate real-time science opportunities**

While the MMS schedule was tight, teams were able to adjust to include science targets of opportunity and accommodate anomalies. This showed an ability to adapt plans without abandoning core commissioning milestones.

#### **Schedule discipline under safety oversight**

Post-Challenger and Columbia changes emphasised that schedule decisions should not override safety considerations. For example, Shuttle return-to-flight planning included independent technical authority and safety review steps that were built into the timeline, ensuring safety-critical activities were not skipped.

---

### What Could Be Done Better

#### **Avoiding over-compression under pressure**

Several lessons learned caution against cutting corners in the design or testing phases to meet schedule targets. In the BSM case, management pushed for quick hardware delivery, but the engineering team resisted early builds in favour of completing analysis and design properly. NASA's own reflections note that compressing the schedule at the start often leaves no time for rework if problems emerge.

### **Accounting for vendor readiness and rework in schedule**

Vendor failures, such as the BSM supplier's inability to deliver flight-qualified hardware, caused significant rescheduling. Upfront schedule planning could better incorporate supplier readiness checks and contingency time for in-house rework if external hardware does not meet requirements.

### **Maintaining capacity for safety functions under time constraints**

Before Columbia, reduced safety workforce capacity combined with a high flight rate meant anomalies were not always addressed before the next mission. This shows the need to ensure schedule planning includes time and resources for thorough anomaly resolution, even when mission cadence is high.

### **Minimising last-minute concurrency**

During MMS commissioning, delays in instrument inputs and unexpected anomalies increased the number of concurrent operations, putting strain on coordination. Building more buffer into the schedule for critical-path activities could reduce the need for compressed, high-risk overlaps later.

---

## **Cost Estimating & Control**

### **What Works Well**

#### **Strong linkage between technical requirements and resource allocation**

Many projects in the binder, such as ATLAS BSM and MMS, tied cost planning closely to well-defined technical needs. For example, the BSM's sub-arcsecond pointing requirements drove investment in precision manufacturing, multiple verification instruments, and extensive testing — a sign that cost estimates were grounded in mission-critical performance goals.

#### **Willingness to invest in corrective action when risk outweighs savings**

When vendor-supplied hardware failed environmental testing (BSM case), NASA decided to absorb the cost of bringing development in-house rather than risk mission failure. This decision reflects a readiness to reprioritise spending in favour of mission assurance.

#### **Embedding safety and quality into cost justification post-accidents**

Following Challenger and Columbia, cost control strategies increasingly factored in the value of independent oversight, safety panels, and redundancy measures. These investments became recognised as essential cost components rather than “optional” overheads.

---

## What Could Be Done Better

### **Realistic accounting for change-driven cost impacts**

In several cases, the business and cost baselines did not fully anticipate the downstream effects of technical changes. Bringing the BSM in-house restored mission requirements but also increased workload, procurement demands, and schedule length. Future cost estimates could more explicitly model “decision-triggered” rework and transition costs.

### **Avoiding short-term savings that compromise long-term value**

Before the BSM pivot, requirements were relaxed to fit a vendor’s cheaper solution, but this ultimately created redesign and retesting expenses. This suggests cost control should focus on life-cycle value, not just initial outlay.

### **Accounting for safety and oversight resource needs from the start**

The Columbia investigation noted that budget and workforce reductions undermined safety capacity, even though the mission scope still required robust safety engagement. Cost planning should protect these resources as fixed elements rather than adjustable line items.

### **Stronger vendor readiness evaluation to avoid unplanned expenditure**

Vendor shortfalls led to unbudgeted in-house development costs. More rigorous early technical assessments of contractor capability could reduce the likelihood of such cost overruns.

---

## Procurement & Contract Management

### What Works Well

#### **Ability to pivot when supplier performance fails**

The ATLAS Beam Steering Mechanism (BSM) case showed NASA’s capacity to change procurement strategy when vendor-supplied hardware failed environmental testing. By bringing the work in-house, NASA regained control over requirements, eliminated compromises to optical performance, and ensured that mission-critical specifications were met.

#### **Leveraging internal expertise for complex builds**

Once development was brought in-house, NASA was able to apply internal engineering talent to design, analyse, and fabricate the BSM to a higher standard than the vendor-provided alternative. This demonstrated strong capability to reassign work and contracts when it is strategically advantageous.

#### **Clear definition of technical specifications**

Procurement activities typically started from precise technical needs tied to mission



objectives, as seen in MMS instrument acquisition and BSM design requirements. This clarity gave contractors a well-defined target for performance.

### **Post-accident improvements to oversight**

Following Challenger and Columbia, procurement processes incorporated more independent safety and quality review functions, with oversight mechanisms to ensure suppliers met both technical and safety standards.

---

## **What Could Be Done Better**

### **More rigorous vendor capability assessment before award**

The BSM vendor selection assumed that using heritage hardware would reduce cost and risk. However, the hardware had design flaws that prevented flight qualification, requiring a complete pivot. This shows the need for more in-depth technical due diligence before contract award.

### **Avoiding requirement compromise to fit vendor limitations**

In the BSM case, requirements were initially relaxed to fit the vendor's capabilities as a cost-saving measure. This degraded the mission case and led to rework. Procurement processes should protect critical requirements from being traded away too early.

### **Better integration of contingency clauses for non-performance**

When vendors fail to meet specifications, it can cause cost, schedule, and resource ripple effects. Stronger contract terms on penalties, recovery plans, and intellectual property rights could make transitions to in-house development smoother and less costly.

### **Ensuring consistent communication and escalation with suppliers**

Some findings from the Challenger and Columbia investigations show that contractor recommendations (e.g., Rockwell's warnings on ice-on-pad before Challenger) were either misinterpreted or inadequately acted upon. This highlights a need for better structured channels for contractor concerns and risk reporting.

---

## **Risk Identification, Prioritisation & Treatment**

### **What Works Well**

#### **Embedding safety as a core risk domain**

Post-Challenger and Columbia reforms put independent safety and mission assurance functions at the centre of risk identification. These functions were tasked with identifying hazards, tracking anomalies, and maintaining authority over waiver decisions — strengthening the link between safety risk and project governance.

### **Structured technical risk verification**

Projects such as the ATLAS BSM demonstrated rigorous performance verification to manage technical risks. Multiple measurement systems were used to cross-validate results, and the verification approach was agreed with the science team early. This pre-planning reduced the likelihood of undetected performance shortfalls.

### **Ability to adapt controls when risks emerge**

In the BSM case, unexpected structural modes, manufacturing tolerance issues, and actuator design anomalies were identified during testing and addressed through redesigns, part substitutions, or filtering adjustments. This showed a strong capacity to detect and respond to emerging risks during development.

### **Risk-informed decision to change development path**

When vendor hardware failed environmental testing, NASA recognised the risk to mission success and shifted development in-house, accepting schedule and cost impacts to protect core requirements.

---

## **What Could Be Done Better**

### **Ensuring concerns are escalated and acted upon**

Both accident investigations documented cases where technical concerns were raised but not adequately addressed. For example, engineers' warnings about O-ring performance in cold temperatures (Challenger) and foam strike risks (Columbia) were downplayed or reframed as acceptable risk. This points to the need for robust escalation pathways that cannot be bypassed by schedule pressure.

### **Maintaining analytical capability for anomaly trends**

The Columbia findings highlighted that reduced safety workforce capacity led to missed opportunities for trend analysis of recurring issues. Resource planning should ensure enough analytical bandwidth to identify patterns before they become critical.

### **Balancing schedule and budget pressures with risk posture**

In some cases, program management allowed schedule priorities to outweigh proactive risk mitigation, deferring corrective action or relying on partial data. Embedding formal risk tolerance thresholds — with authority outside of the program office — could help counteract this.

### **Protecting against cultural drift toward risk normalisation**

The Columbia investigation showed how repeated acceptance of foam shedding shifted from being a safety concern to “in-family” tolerance, even without a technical solution. Risk registers should include explicit tracking of hazards with no viable mitigation to prevent such normalisation.

---

# Project Team Management & Performance

## What Works Well

### **Highly skilled, cross-disciplinary teams**

Projects like MMS commissioning and ATLAS BSM development relied on multi-organisation, geographically dispersed teams spanning NASA centres, universities, research institutes, and contractors. Each group brought specialised expertise — from instrument design to mission planning — and worked under clearly defined roles.

### **Strong internal capability for complex engineering**

When the BSM was brought in-house, NASA's engineering teams were able to design, analyse, and fabricate hardware to higher standards than the original vendor. This showed that the in-house talent pool could deliver under pressure while maintaining technical excellence.

### **Recognition of individual expertise in critical builds**

In the BSM case, the same highly experienced technician assembled all units, which supported consistency and quality. This kind of continuity in skilled personnel helped maintain performance standards.

### **Post-accident cultural reforms to support openness**

Following Challenger and Columbia, NASA emphasised creating an environment where technical staff could voice dissenting opinions and have them heard by decision-makers, recognising the value of diverse perspectives in complex projects.

---

## What Could Be Done Better

### **Acting on informal observations from skilled staff**

In the BSM case, an experienced technician noticed unusual balance behaviour during assembly but was instructed to ignore it as it wasn't part of the immediate task. This delayed detection of a significant manufacturing tolerance issue. Team management could make better use of informal observations by building them into formal review channels.

### **Avoiding over-reliance on key individuals**

While continuity of skilled staff is valuable, reliance on a single technician or specialist for critical tasks can create bottlenecks or single points of failure. Cross-training could reduce this risk while preserving quality.

### **Maintaining workforce capacity under budget constraints**

The Columbia investigation found that safety, reliability, and quality assurance staffing was cut while operational tempo increased. This reduced the team's ability to engage in risk monitoring and anomaly resolution, undermining performance in critical areas.

### **Encouraging upward flow of technical concerns**

Accident investigations highlighted that team members' technical warnings sometimes failed to reach senior decision-makers or were reframed to fit schedule objectives. Stronger escalation protocols and leadership engagement could ensure that concerns from any level are acted upon.

---

## **Project Governance & Change Control**

### **What Works Well**

#### **Independent oversight structures**

Post-Challenger and Columbia reforms strengthened governance by creating independent technical authority and safety review panels with direct reporting lines outside of program management. This ensured that critical safety and technical decisions were subject to review beyond the immediate project team.

#### **Clear authority and responsibility definitions**

In MMS commissioning and ATLAS BSM development, roles and responsibilities for decision-making were clearly mapped to organisational units — from mission planning to payload operations. This clarity supported structured governance, especially in multi-organisation environments.

#### **Formal requirement verification sign-off**

For the BSM project, the performance verification approach was agreed with the science team early and approved as part of governance. This reduced later disputes over whether requirements had been met.

#### **Willingness to alter governance approach when strategy shifts**

When the BSM development was moved in-house, governance adapted to manage both the increased internal workload and tighter control over technical and quality standards.

---

### **What Could Be Done Better**

#### **Ensuring governance remains strong under schedule pressure**

Accident investigations showed that under high schedule pressure, governance structures could be bypassed or compressed, particularly for safety-related reviews. Maintaining mandatory review gates, regardless of deadlines, would strengthen change control.

#### **Avoiding requirement erosion to fit constraints**

In the BSM case, performance requirements were initially reduced to accommodate a

vendor solution. This change was approved within governance but ultimately proved detrimental, showing the need for governance processes that scrutinise requirement reductions more rigorously.

### **Strengthening change control for cross-organisational dependencies**

MMS commissioning required last-minute changes to sequencing and concurrent operations due to anomalies and late inputs. Change control processes could better anticipate and manage these dependencies to avoid operational strain.

### **Preventing cultural drift toward “acceptable risk” without mitigation**

The Columbia investigation revealed how repeated acceptance of foam shedding shifted it from a hazard to an “in-family” tolerance without a technical solution. Governance should require documented justification and mitigation plans for any change in risk classification.

---

## **Project Delivery & Handover**

### **What Works Well**

#### **Meeting or exceeding technical performance at delivery**

In the ATLAS BSM case, the final flight model met sub-arcsecond pointing requirements with margin, demonstrating that rigorous design, testing, and issue resolution throughout development can lead to high-quality deliverables.

#### **Integration into operational systems**

The BSM flight model was successfully integrated into the ATLAS instrument and located within the spacecraft configuration as planned, showing effective coordination between subsystem teams and the broader mission integration function.

#### **Structured verification prior to handover**

Projects such as BSM followed well-documented verification approaches, combining multiple measurement systems to validate performance across operational ranges and environments before delivery.

#### **Retention of expertise through to integration**

The same technical personnel involved in earlier builds and testing often supported the integration phase, which preserved institutional knowledge and smoothed final delivery.

---

### **What Could Be Done Better**

#### **Capturing lessons learned before final handover**

While individual projects documented technical lessons, the binder shows that some insights (such as informal observations by technicians) were not formally incorporated

until later. A structured “pre-handover” lessons learned process could speed knowledge transfer to other teams.

### **Managing schedule and resource risk close to delivery**

Tight schedules meant that some anomalies and redesigns were resolved late in the development cycle, leaving minimal margin before delivery. Building additional time buffers into late-phase schedules would reduce the risk of rushed fixes.

### **Ensuring operational readiness beyond technical specs**

Accident investigations show that meeting technical requirements at delivery does not guarantee operational safety. Handover processes should explicitly check operational risk status, hazard mitigations, and readiness of safety oversight functions.

### **Improved transition planning for long-term operations**

Some post-delivery operational risks (e.g., normalisation of anomalies such as foam shedding) were not addressed in the handover stage, suggesting the need for clearer ownership of ongoing risk monitoring after project delivery.

---

## **Project Documentation**

### **What Works Well**

#### **Comprehensive technical reporting**

The binder itself demonstrates NASA’s commitment to documenting project history, design decisions, verification processes, and lessons learned in detailed technical papers, conference proceedings, and internal reports. This level of documentation supports future reference and cross-project learning.

#### **Accident investigation records as institutional memory**

The Challenger and Columbia investigation findings and recommendations were thoroughly documented, capturing both technical and organisational failures. These serve as enduring references to inform governance, risk management, and cultural reforms across the agency.

#### **Requirement verification records**

For projects like the ATLAS BSM, the methodology for verifying performance requirements was thoroughly documented, including detailed descriptions of test configurations, measurement instruments, limitations, and data processing methods. This transparency allows replication and audit of results.

#### **Multi-stakeholder input**

In complex projects such as MMS commissioning, documentation incorporated input from NASA centres, contractors, universities, and international partners, ensuring diverse perspectives were recorded alongside technical data.

---

## What Could Be Done Better

### **Closing the loop on informal knowledge**

In some cases (e.g., the BSM technician's early observation of unusual behaviour), useful insights were not documented at the time they occurred, only later after issues emerged. Formal processes for capturing informal observations during development could improve documentation completeness.

### **Ensuring documentation reaches decision-makers**

Both Challenger and Columbia findings showed that critical data and risk analyses existed but were not effectively communicated to higher-level decision-makers. Documentation processes should include escalation protocols for safety-critical information.

### **Maintaining accuracy in manufacturing records**

The BSM case revealed that incorrect inspection reports, caused by improper fixturing and tolerance interpretation, delayed problem detection. Strengthening quality control in documentation of manufacturing and inspection processes would reduce this risk.

### **Sustaining documentation capacity under budget cuts**

The Columbia investigation noted that workforce reductions affected safety and quality functions, which likely impacted their ability to maintain thorough documentation of anomalies and trends. Resource planning should protect documentation functions as a core capability.

---

## Other Lessons Learned

### **What Works Well**

#### **Detailed contamination control procedures**

From Apollo sample handling through later missions, NASA developed and documented highly specific procedures for contamination control — including material restrictions, glovebox environments, organic/inorganic monitoring, and cleanliness standards. These practices preserved scientific integrity for decades and set benchmarks for planetary sample curation.

#### **Technical problem-solving culture**

Across projects, teams demonstrated the ability to methodically diagnose and resolve unexpected technical issues, such as structural mode anomalies, thermal instability, or actuator behaviour. This culture of iterative problem solving allowed projects to meet or exceed demanding specifications.

### **Adaptability in project execution**

Examples like MMS commissioning show that NASA can adjust operations dynamically — incorporating science opportunities, responding to anomalies, and reconfiguring concurrent operations without losing sight of core objectives.

### **Use of multiple verification methods**

In high-precision systems (e.g., ATLAS BSM), combining different measurement instruments (interferometers, autocollimators, theodolites) allowed the team to cross-validate results and overcome limitations in any single method.

---

## **What Could Be Done Better**

### **Avoiding overtesting beyond operational limits**

In the BSM case, a high-potential (hipot) test applied voltages far beyond operational requirements, damaging hardware. This shows the importance of tailoring test parameters to actual use cases rather than blindly following generic specifications.

### **Preventing cultural drift toward risk acceptance**

The Columbia investigation found that repeated exposure to anomalies without incident (e.g., foam shedding) led to their reclassification as “in-family” risks. This normalisation of deviation can erode safety margins over time and should be actively guarded against.

### **Enhancing early detection of manufacturing variances**

The BSM flexure tolerance issue, which caused imbalance, was missed in inspection due to fixturing errors and incorrect interpretation of geometric tolerances. Stronger early-stage inspection and cross-checking could prevent such delays.

### **Protecting safety and quality resources under cost/schedule pressure**

Findings from both Challenger and Columbia noted that safety, reliability, and quality assurance teams lost capacity due to workforce cuts, even as operational demands grew. Sustaining these functions is critical for long-term mission safety and performance.