Organization of Findings and Recommendations

Findings and Recommendations

- F1 Collecting the Right Samples
- F2 Communicating the Importance of MSR
- F3 Overall Organizational Structure
- F4 Agency-level Leadership and Engagement
- F5 ERO/CCRS and the NASA/ESA Partnership
- F6 OS Impact Across MSR Elements
- F7 UV Decontamination of Possible Biohazards on the OS Exterior
- F8 NASA Coordination with US Regulatory Agencies on Backward Planetary Protection
- F9 Architectural Robustness and Resiliency
- F10 Programmatic Assessment
- F11 Independent Review Structure
- F12 Culture and Communication

Additional Important Findings and Recommendations

- F13 Verification and Validation
- F14 Cross-Organization Engineering Management
- F15 Telecommunications Infrastructure
- F16 Helicopter Accommodation Risk Balance
- F17 Technical Baseline Management and Change Control
- F18 Launch Vehicles
- F19 Workforce Capacity and Expectations Post COVID-19
- F20 Supply Chain

The slight distinction between “Findings and Recommendations” versus “Additional Important Findings and Recommendations” is as follows:

- Findings and Recommendations are those entries with broad, all-encompassing implications to the MSR Campaign and inter-Center/Inter Agency agreements.
- Additional Important Findings and Recommendations are those entries with system-level design or execution implications. These entries includes findings such as Verification and Validation, Launch Vehicles, Workforce Capacity and Expectations post COVID, and Supply Chain.

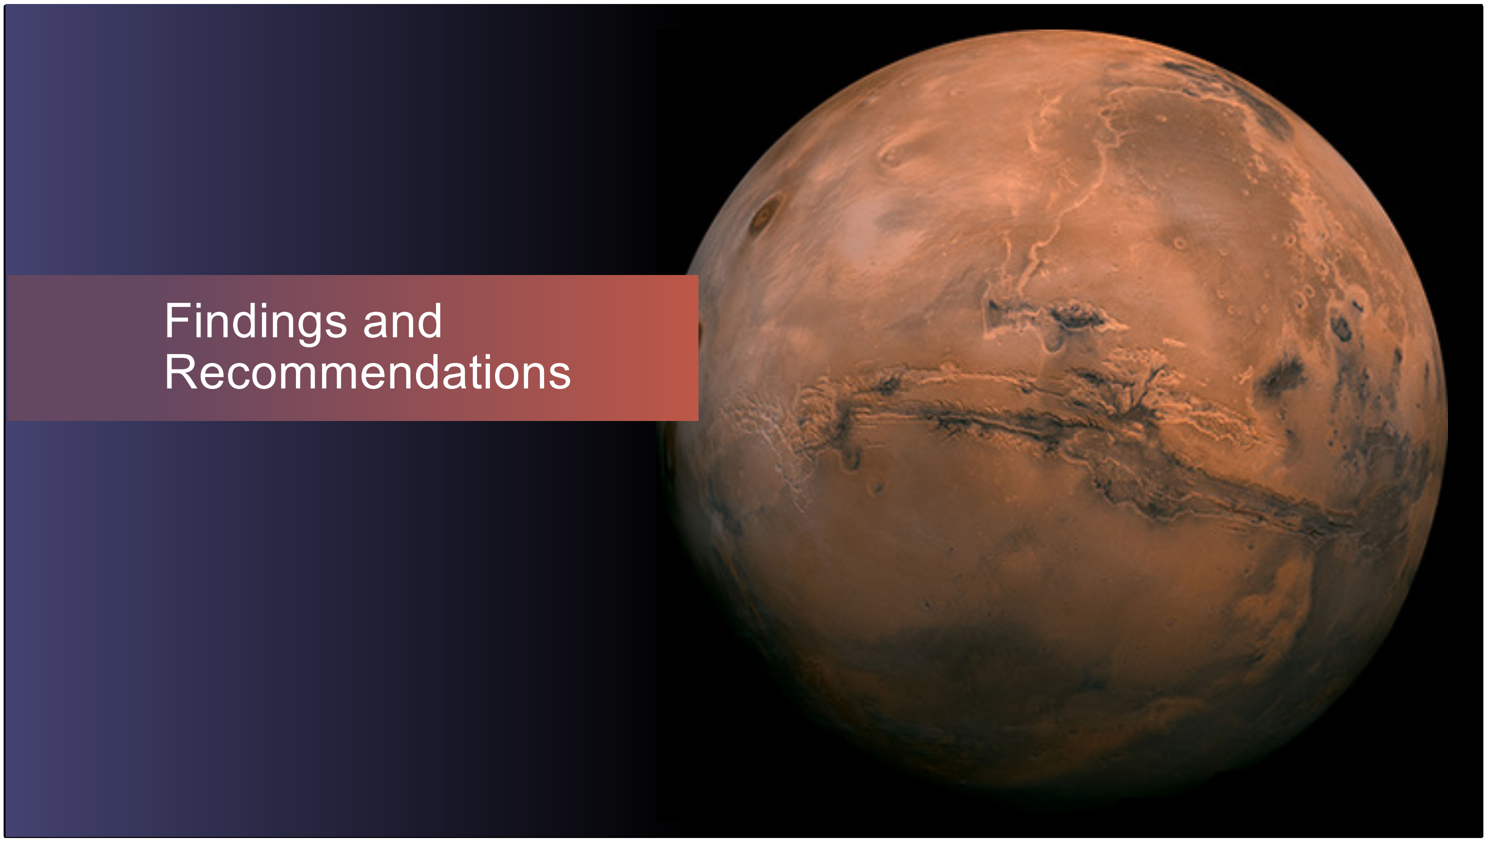
Findings and
Recommendations

F1: Collecting the Right Samples (1 of 2)

The cache of samples deposited at Three Forks is return worthy but is not an optimal sample set because it does not represent the full diversity of geologic environments along the rover's traverse that could preserve signs of life. The 20 samples now collected by and carried on Perseverance are of very high scientific value – higher than the cache at Three Forks. The science value is expected to increase even more as Perseverance collects 18 additional samples from the crater rim and beyond.

- A backup set of 10 return-worthy samples has been dropped at the Three Forks depot. However, we do not know in advance how evidence for Martian life could be preserved in the geologic record. The diversity, number, and value of additional samples increases substantially as Perseverance continues its traverse.
- Driving up and out of Jezero Crater will allow sampling of highly diverse geologic units on the margin, rim, and possibly beyond, enabling the return of an optimized sample suite.
- Such a traverse requires that there be a certifiable site for landing SRL relatively nearby the samples to be returned. Also required is a clear plan for M2020 to operate with a focus on direct transfer to SRL or placing one or more additional depots.
- Drivers on the OS design include the number of samples to be returned and the surface timeline for retrieval of samples from Perseverance or a depot.
  - A case can be made that the specific number of samples returned is less important than the diversity and scientific richness of those samples, if the returned samples expand the MSR scientific objectives that can be realized.
  - Any reduction in returned sample mass will limit the number and, possibly, types of analyses that can be performed on Earth using state-of-the-art laboratory instrumentation, potentially deferring the realization of the full scientific return of MSR.
  - It is not evident that deliberations around a reduction in the number of tubes carried by the OS have accounted for the above considerations or have taken place via an independent assessment of the risk to the mission versus the impact to MSR science.

A carefully considered, science-based rationale for the return of 500 g (~1.1 pound) of Martian sample material has been documented in the peer-reviewed scientific literature several times over more than a decade. The scientific justification for returning 500 g of Martian sample material is based in part on in the types of expected analyses and the quantities of sample material needed to conduct these analyses. The sample tubes are estimated to hold at the most 15-16 g each. This design results in the planned capacity of the OS to return ~32 tubes. Not all of the tubes will be filled to capacity, and only a portion of the returned material will be made available to researchers in the years immediately following the tubes' arrival on Earth. A yet-to-be-determined quantity of the returned samples will be placed in long-term storage for future generations, as was done for Apollo lunar samples and will be done for the asteroid sample planned for return by OSIRIS-REx.

Perseverance's cache is the primary cache for return to Earth, representing the broadest diversity of samples and geological diversity of Mars. It is therefore important to strike the right balance between the largest number and the diversity of samples as the OS design is settled.

A reduction in the capacity of the OS to only 18 tubes could yield approximately 288 g (~0.64 pounds) of the desired 500 g. Given that some of this mass will go into long-term storage for future generations, even less than ~288 g would be available initially for analysis at US institutions and across the international scientific community.

R1: Collecting the Right Samples (2 of 2)

Recommendations:

- (R1.1) The relative science merit of the 10 samples on the ground at Three Forks should be independently evaluated against additional samples currently on Perseverance and samples that may be collected during continued science operation of the Perseverance rover. Science merit should be a key input to landing site selection for MSR and operational strategies for the Perseverance rover, Sample Retriever Lander (SRL), and Mars Ascent Vehicle (MAV).
- (R1.2) NASA should develop and implement a formal, metrics-based and timeline-based plan for operating Mars 2020 with a focus on direct transfer to the SRL or placing one or more additional depots.
  - Whether Perseverance 1) remains at the second certified landing site, 2) conducts sorties from that site, or 3) drops a second cache and is permitted to continue along its traverse towards a third certifiable landing site should depend on factors including the science value of the samples collected to date, the health of the rover, the mission architecture, and arrival dates of future flight elements.
- (R1.3) NASA should direct Mars 2020 to certify a landing site outside Jezero Crater as soon as feasible.
- (R1.4) Mars 2020 and MSR should develop a joint plan that balances the risk of degradation or failure of Perseverance against obtaining the highest value samples near the crater rim and beyond. The plan should be appended to the existing memorandum of agreement between MEP and MSR.
- (R1.5) NASA should revise the MSR Program requirement (6.1.1) related to the capability to return 30 sample tube assemblies and implement a requirement that is based on science value, the MSR timetable, and other factors relevant to successfully achieving the optimum science from MSR, including total sample mass.

F2 and R2: Communicating the Importance of MSR

The societal, technological, and scientific significance of MSR makes it a mission of the highest importance in NASA's long-term exploration strategy, and this is not being communicated consistently and clearly with the public and stakeholders .

- A successful MSR campaign will revolutionize our understanding of the history of Mars, the Solar System, and the potential for life beyond Earth.
- The exploration of Mars and the return of samples are integral parts of NASA's Moon-to-Mars initiative and our commitment to US space leadership in partnership with ESA.
- The technical challenge and audacity of MSR will inspire the next generation.
- NASA is not sending a consistent and unified message to Congress, the scientific community, and to the public regarding MSR's scientific and strategic importance to the nation.

Recommendation:

- (R2.1) NASA should develop a strategy and implement a compelling communication plan that reflects MSR as an Agency priority and as a priority for the nation that historically has dared to do the seeming impossible in space. NASA's senior leadership should vigorously participate in the execution of that plan.

The planetary science community is substantially larger and more diverse now than when MSR was first conceived decades ago. We have made many new discoveries throughout the Solar System in those intervening years. The number of high-priority scientific targets of exploration has justifiably expanded, particularly since the inception of the Decadal Surveys.

However, the growth and ambition of planetary science does not change the fact that MSR remains a longstanding, two-time Decadal priority for legitimate scientific, technological, and geopolitical reasons. Now that this mission has begun with the collection of scientifically selected samples, and NASA/ESA are on the verge of being able to return those samples, **it is important to establish deeper, meaningful communication about the scientific value of this longstanding goal as well as the ways in which it will benefit NASA beyond planetary science and as part of NASA's Moon-to-Mars strategy.** A particular and understandable concern from the community is the impact that MSR may have on the rest of the Planetary Science Division and SMD flight mission portfolio. NASA must address this concern while clarifying and dispelling the notion that cancelling MSR necessarily means greater budgets for everybody else in the Planetary Science Division or even the Mars Exploration Program. Cancellation may also call into question the feasibility of other ambitious sample return efforts envisioned by the planetary science community.

It is worth adding that leadership in space exploration is a hallmark of USA's soft power in the world. Peaceful exploration of space serves to demonstrate United States' technological expertise to deliver, no matter how difficult. These efforts inspire the nation and the world community in the journey, opening the door to significant international collaborations in space. Communicating the importance of MSR as a priority for US scientific and technological reasons is important, particularly in light of China's plans for preeminence in the world stage.

F3: Overall Organizational Structure (1 of 2)

The success of Mars exploration is critically dependent upon the technical and programmatic success of the MSR campaign. MEP and MSR are currently operated and managed separately, even though MSR depends on MEP assets. This arrangement is creating unnecessary silos, and potentially pits MEP against MSR in terms of attention given by senior HQ management. Furthermore, the current MSR program structure and management results in an unnecessarily complex organization. This structure results in multiple communication pathways and unclear lines of accountability and authority between HASA HQ, its field centers and labs, and ESA.

- Mars science and management advocacy at HQ is split between the Planetary Science Division/MEP and MSR programs. This arrangement lacks the single champion that is required at this critical phase of the program.
- NASA has not established an appropriate programmatic reporting structure as required for a program of this significance and complexity.
- MSR has not demonstrated an adequate level of maturity in its execution of programmatic management at JPL and the Marshall Space Flight Center (MSFC) due to inadequate staffing, processes, and tools that are specific to programmatic analysis.
- JPL's role of complete programmatic management over multiple and extremely complex activities and NASA field centers is an enormous task that distracts from developing SRL and managing critical MSR campaign (Level-1.5) interfaces.
- The current MSR management structure and approach does not enable the very different engineering and management cultures of the participating organizations to collaborate for efficient resolution of issues.

The current programmatic workforce at NASA HQ consists of a senior civil servant Deputy Program Director and one support contractor.

The MSR HQ office has provided guidance to the JPL MSR Program office regarding programmatic analysis and reporting and is being reflected in including the draft program and business plans. However, much work remains to implement these plans consistently with standards that are commensurate with a flagship mission.

The need for experienced programmatic staff increases and acquisition of assessment tools for MSR has been acknowledged by the respective institutions that are contributing to MSR.

The cultures at JPL, MSFC, GSFC, and ESA differ in multiple ways, particularly in the degree of structure in their decision, documentation, and Project control processes. A key part of NASA HQ's Program Office role is to be able to normalize the inputs and metrics against plans.

R3: Overall Organizational Structure (2 of 2)

Recommendations:

- (R3.1) Combine the current MEP and MSR programs and establish an integrated Mars Exploration Program office at NASA HQ, reporting directly to the SMD AA and encompassing all the activities within the Mars enterprise.
- (R3.2) Unify and strengthen the scientific and programmatic leadership and advocacy within this new, integrated Mars Exploration Program in order to enhance communication with senior NASA leadership, the planetary science community, and other stakeholders.
- (R3.3) Establish full programmatic responsibility and authority within the Mars Exploration Program office at HQ, separating the budget from the Planetary Science Division and creating a properly-staffed Program Planning and Control (PP&C) group within the new HQ program office.
  - All Projects within the integrated MEP should report programmatically to the HQ MEP PP&C group.
  - The HQ MEP PP&C should set project analysis and reporting standards and should conduct independent monthly programmatic assessments based on Project-provided data.
- (R3.4) Have JPL retain its MSR campaign-level technical role of leading cross-functional technical teams that provide integrative support to the MSR campaign.
- (R3.5) Give full technical responsibility and authority for the individual MSR projects to their respective labs/field centers with the directors held fully accountable for their organization’s contribution.
- (R3.6) Enhance project business office PP&C staffing and capabilities (e.g., tools, processes) in order to ensure the implementation of the depth and breadth of programmatic management necessary for a large complex mission (e.g., planning, maintenance, tracking, assessment, analysis/forecasting, schedule cost/risk analysis, control).

NASA would need to augment the integrated MEP programmatic staff above the combined levels in the current MSR and MEP in order to ensure that there is sufficient program planning and control (e.g., cost and schedule planning, maintenance, tracking, assessment, analysis, forecasting and risk analysis).

To enable SRL and MAV to analyze and report cost and schedule status effectively, both of their programmatic teams will need to be augmented with 3-4 highly qualified FTE/WYEs.

The JPL campaign-level technical role includes providing key technical support to the MEP office at HQ.

Definition of full responsibility for the individual MSR Projects entails the following:

- JPL retains full responsibility and authority for SRL.
- Goddard Space Flight Center assumes full responsibility and authority for the CCRS payload contribution to the ESA-provided ERO mission, along with authority for JPL-provided hardware.
- Marshall Space Flight Center assumes full responsibility and authority for the MAV payload contribution to the JPL-provided SRL.
- Push responsibility for programmatic controls to the particular MSR project.

F4 and R4: Agency-level Leadership and Engagement

MSR follows many decades of strategic investment. The mission has been identified as the top priority of the Decadal Survey. However, leadership at the SMD AA level and above has not been adequately engaged at a level that is commensurate with a mission of such importance to NASA, the nation, and NASA's international partners.

- MSR has not been treated as a mission of such technical boldness and international visibility that requires a highest-priority mindset within the top leadership at NASA in conjunction with its partner ESA.
- MSR Program leadership lacks sufficient access to NASA executive leadership.
- NASA has not established an appropriate organizational and reporting structure at the highest levels as required for a program of this significance and complexity.

Recommendations:

- (R4.1) NASA senior leadership should manage and advocate for MSR as a priority for the nation's space exploration goals.
- (R4.2) The director of the integrated MEP must be a world-class Science Program Directorship and Community leader who is responsible for prioritizing sample return from Mars utilizing current assets, coordinating with ESA, and meeting overall science goals that are consistent with the recommendations of the Planetary Science and Astrobiology Decadal Survey.
- (R4.3) The director of the integrated MEP should report on the status of MSR to the NASA Associate Administrator at no less than a bi-weekly cadence.

MSR represents a greater than ~$20 billion investment to date and a very challenging mission in the $8-11B class. All the components must work with narrow margins for error. The end-to-end effort requires a heightened level of attention, risk management, and reassurances.

Programs of this magnitude and complexity, particularly those with multiple centers and international partnerships (e.g., ISS, JWST, Shuttle, Hubble, Artemis) have historically had direct and regular access, interaction, and accountability with the HQ A-suite. Each of these programs have had very challenging phases where top-level attention and actions were key to their success.

This finding also relates to agency-level proactive advocacy in terms of public interactions (speeches, social media, outreach events, conferences) and congressional interactions (testimony, Hill visits, talking points, industry alignment). These same programs have also faced difficult gauntlets in terms of public and administration/congressional support. During these challenging times, advocacy made the difference between success and failure.

MEP deserves world-class leadership with proven records of success. Strong leadership is required to:

- Create a team culture that is aligned around the mission objectives and mission success
- Communicating strategy and vision, keeping stakeholders onboard
- Keep all of the partners and projects aligned
- Making and communicating well-considered, integrated, and sometimes difficult decisions

F5: ERO/CCRS and the NASA/ESA Partnership (1 of 2)

The tight technical and programmatic coupling of ERO to MSR and the early schedule phasing for ERO has created unique and untenable organizational, schedule, and technical issues for both the ESA/Airbus ERO team and the NASA Goddard Space Flight Center (GSFC) CCRS team.

- Both SRL and ERO are ambitious and technically challenging standalone Projects with clean interface boundaries, yet they are being technically managed as coupled elements of a single mission.
- The organizational construct between JPL, ESA/ERO, and GSFC/CCRS leads to confusion regarding responsibility and accountability for overall management of the ERO mission.
- JPL personnel supporting CCRS and providing CCRS hardware report to the MSR program office, not to the CCRS Project Office. This arrangement contributes to the confused lines of authority and communications between the GSFC and JPL CCRS elements, as well as between GSFC management of CCRS and the Program.
- The ESA contract for the spacecraft adds a program-level interface that has complicated ERO/CCRS/MSR Program Office interactions and has created tensions between the teams.

The MSR campaign architecture requires tight operational, scheduling, and design coordination between JPL, Goddard, ESA, and ESA's spacecraft vendor Airbus. This arrangement introduces notable communication and coordination challenges, especially when addressing the intricate engineering issues tied to the ERO mission. ERO's design is highly impacted by the MSR campaign design, which subsequently affects the CCRS design. CCRS has significant componentry provided by JPL, which is also in charge of the overall campaign design. These dynamics create circular interconnected dependencies on top of difficult engineering challenges.

The timing of ERO's Mars orbit arrival to support SRL's EDL and surface operations results in an intricately linked development and launch schedule dependency between ERO and SRL. This dependency influences costs and introduces complexities into engineering design decisions. When evaluating the primary objectives of each mission, there are opportunities to streamline these interactions at the highest level, particularly in relation to the recovery of the OS in Mars orbit. This recovery could occur at an unspecified time, obviating the necessity for a specific order and timing of SRL and ERO launches.

The fixed-price and contract management approach with Airbus requires ESA to be understandably guarded regarding interactions between CCRS and Airbus engineering. With the spacecraft bus design well ahead of CCRS, the GSFC team has had to design CCRS to accommodate the spacecraft bus and a highly prescribed set of interfaces. The complexity is further compounded by uncertainty regarding the choice of the launch vehicle for ERO. Interactions and technical interchange between ESA, CCRS, and Airbus have not been adequate given the complex nature of the ERO mission.

R5: ERO/CCRS and the NASA/ESA Partnership (2 of 2)

Recommendation:

- (R5.1) Decouple the ERO and SRL development schedules by transferring full responsibility for the integrated ERO mission to ESA. This arrangement will allow GSFC to work with ESA and Airbus without going through JPL.
  - Focus the NASA/JPL efforts on the elements and infrastructure that are required to get the OS into a defined and stable Mars orbit through all available assets (in collaboration with ESA for Trace Gas Orbiter (TGO) relay communications), while ESA focuses on OS retrieval and OS return to Earth.
  - JPL should continue to lead overall campaign technical integration and management of Level 1.5 requirements including backward planetary protection, in proper coordination with ESA.
  - Accept ESA's ability to successfully accomplish ERO, recognizing ERO's importance to ESA and its goals in Mars exploration.
  - Give GSFC full responsibility and authority for leading the CCRS development (including the JPL components) as a deliverable payload to ESA.
    - ESA should provide adequate visibility into driving technical parameters and margins and should help facilitate closure of the CCRS design.
    - NASA and ESA should collaborate more closely in order to better integrate the ERO spacecraft and CCRS teams into a one-team approach that incorporates contractual considerations into a modified interagency agreement.

The MSR campaign should be approached as two distinct missions. Both missions rank among the most challenging endeavors ever undertaken by NASA and its partner ESA. SRL's mission involves recovering samples from Perseverance and placing them into a well-defined and stable Mars orbit. ERO retrieves these samples from Mars orbit and returns them to Earth. Consequently, the primary interplay between the two missions revolves around the OS design, encompassing considerations such as mass, volume, shape, Mars orbit, and reverse planetary protection for all phases of the MSR campaign. By centering the interaction between the two missions on OS recovery and planetary protection, each organization can devote its resources entirely to its MSR campaign contribution and can assume complete accountability for its success.

Given the intricacies of both the ERO spacecraft and CCRS, coupled with the separation of the ERO and CCRS missions, it becomes imperative for NASA and ESA to closely align their Airbus and CCRS teams. This alignment is crucial to provide each team with essential insights into the other's designs, allow collaborative solutions to each other's issues, and to establish a unified approach to agreed-upon verification methods.

F6 and R6: OS Impact Across MSR Elements

The OS is the critical element that is central to the entire MSR mission. Failure to converge upon an acceptable and stable OS design will continue to produce significant design uncertainty throughout the elements of any MSR Program architecture.

- The OS problem is uniquely complex because the OS is not only highly constrained due to mission requirements, but the OS also touches multiple systems in every MSR project element.
- Finalizing the OS design is critical to establishing a technical and programmatic baseline, as well as for helping settle backward planetary protection issues.
- Key partners at NASA Centers and ESA report cost impacts, schedule delays, and the stress of personnel that are caused by repeated redesign of OS accommodations and interfaces.

Recommendation:

- (R6.1) Close the OS design and all associated interfaces before the Program's preliminary design review season. Even if there are changes in the overall MSR Program plans, and these changes cause changes in the OS, the necessity of closing the OS design as soon as possible (and before other elements' designs are finalized and reviewed) remains undiminished.

The OS is a critical element of any Mars Sample Return architecture, serving as the primary containment vessel for the samples and their launch into Mars orbit. The OS must be discoverable in Mars orbit for an extended period, captured, contained, decontaminated, and secured into the Earth return vehicle for landing after the return journey to Earth.

Various systems and organizations interface through the OS. OS design interfaces with, and thus is constrained by and impacts, all parts of the MSR system.

- The OS is implemented by the Jet Propulsion Laboratory.
- The mass and volume of the OS is driven by the number of sample tubes it needs to contain, sample integrity and planetary protection requirements, as well as robotic manipulation related interfaces with the Sample Transfer Arm (STA, supplied by the European Space Agency)
- In turn, the mass and volume of the OS drives the size of the Mars Ascent Vehicle (MAV; managed by Marshall Spaceflight Center). MAV in turns drives the size and mass of the Sample Return Lander (SRL; managed by the Jet Propulsion Laboratory) and its launch vehicle.
- The exterior optical properties of the OS are critical for detectability by the Earth Return Orbiter (ERO; supplied by the European Space Agency).
- The OS also drives the Capture Containment and Return System in several ways (CCRS; managed by Goddard Space Flight Center with some elements contributed by the Jet Propulsion Laboratory). The CCRS is accommodated on the ERO.

F7 and R7: UV Decontamination of Possible Biohazards on the OS Exterior

Over the past year, efforts to simplify the CCRS design and to save mass resulted in changes to the plan for backward planetary protection. Thermal sterilization or fitting of the seal on the container enclosing the OS was replaced by use of UltraViolet (UV) illumination to decontaminate possible Martian biohazards on the OS exterior.

- In December 2022 at the CCRS PDR, a baseline design with heat sterilization of the weld for the Primary Containment Vessel was presented as compliant to requirements and feasible to implement.
- In February 2023, ESA and NASA OSMA expressed concerns about the risk of going forward with an emerging proposal for UV as an alternative to thermal methods.
- In June 2023, a restructuring agreement was signed by ESA for Level-1.5 requirements. This agreement made clear mention that it is not an endorsement of UV treatment but allows NASA to proceed at risk.
- JPL's proposed UV treatment dose and wavelength for decontamination of the OS exterior are currently being quantified, but much work remains. Independent validation of the UV design is planned during the next year.

Recommendation:

- (R7.1) Results of JPL's ongoing microbiological testing of the proposed UV treatment should be reviewed by NASA's Office of Safety and Mission Assurance as well as relevant US regulatory agencies. If the results are accepted as preliminary validation of the UV design, then the reviewing organizations should select an independent laboratory for additional decontamination testing and final validation of compliance with backward planetary protection requirements. NASA should work closely with ESA on review and validation of the MSR design for UV decontamination.

At the CCRS PDR in December 2022, thermal welding to close and sterilize the seal on a container enclosing the OS was presented as meeting success criteria. The IRB did not hear a compelling explanation for why subsequent to PDR, the thermal plan was deemed too complex and was replaced with a plan for UV treatment. Being new, the UV treatment plan requires extensive testing of effective biological kill and the flight readiness of the UV source.

F8: NASA Coordination with US Regulatory Agencies on Backward Planetary Protection (1 of 2)

The Office of the Chief Scientist (OCS) coordinated an Independent Scientific Review of the potential use of an active UltraViolet (UV) treatment to minimize or eliminate possible Martian biohazard contamination on the exterior of the OS. This study was an important step toward interagency coordination on MSR. The review recommendations include substantial additional work to confirm dust-load models, to evaluate the effectiveness of UV for microbiological and macromolecular decontamination, and to develop a mechanism for verification that the UV illumination system will operate as planned during capture and containment of the OS in orbit at Mars.

- There is a risk that without an interagency agreement on the subjects of effective biohazard decontamination and intact containment, there will be a failure to approve transfer of the Martian samples to NASA.
- Planetary Protection coordination for Apollo was accomplished by establishing the Interagency Committee on Back Contamination (ICBC).
- MSR currently has no similar interagency coordinating organization or process. The IRB is not aware of any formal contact between NASA and the White House Office of Science and Technology Policy as needed to begin work on the protocols for decision-making and the interagency chain-of-command on safety decisions during MSR landing at the Utah Test and Training Range (UTTR).
- OCS has proposed lunch-and-learn activities with key interagency representatives for the coming year.

R8: NASA Coordination with US Regulatory Agencies on Backward Planetary Protection (2 of 2)

Recommendations:

- (R8.1) All aspects of the MSR design for compliance with backward planetary protection requirements should be reviewed and accepted by NASA's Office of Safety and Mission Assurance as well as the relevant US regulatory agencies.
    - In addition to UV decontamination of possible Martian biohazards, safety reviews should cover models and testing for breaking the chain of contact and robust containment of non-sterilized material from Mars.
- (R8.2) NASA should initiate engagement with the White House Office of Science and Technology Policy and relevant regulatory agencies in order to establish protocols for decision-making and chain-of-command structure during safety inspection of the Earth Entry System at the UTTR landing site.

When the samples from Mars' Jezero Crater land on Earth, multiple regulatory agencies will participate in deciding if all unsterilized Martian material is contained and safe for transfer to a receiving facility. Although Jezero Crater is not considered to be currently habitable, there are other locations on Mars where environmental conditions could possibly enable carbon-based life to survive and replicate. Consequently, any samples returned to Earth from Mars must adhere to backward planetary protection requirements that are focused on breaking the chain of contact and utilizing robust containment of all unsterilized material. NASA's technical standards are rigorous and compliant with international standards for planetary protection as established by the Committee on Space Research.

F9: Architectural Robustness and Resiliency (1 of 2)

The current MSR architecture is highly constrained and is not sufficiently robust or resilient to delays in the launch periods of the main architecture elements.

- SRL and CCRS have inadequate technical margins and have limited opportunities to find additional technical resources.
- A delay of SRL past the 2030 launch period greatly reduces the likelihood of mission success using the present architecture because of the aging Perseverance rover and the surface mission timeline challenges.

High-level analysis suggests that other architectures may be more robust and more resilient to schedule risk. These architectures may fit within yearly funding constraints, although at higher lifecycle costs due to the element schedules being extended with less-than-optimal funding profiles.

- Such architectures might include a separately-landed fetch rover/MAV lander, nuclear power, and additional telecommunications assets.

The current outsourcing strategy is consistent with the proven level of contractor expertise.

- Rearrangement of responsibilities among these proven NASA, NASA partner ESA, and industry players might be warranted if the architecture is changed.

The current architecture is intelligently constructed but is also highly constrained. The current architecture is based on an assumed budget and schedule plus other parameters including technical margins that do not seem realistic at this point in the lifecycle. The architecture should be reexamined with respect to all parameters and constraints in order to assure that the mission is both robust and resilient to schedule and other mission constraints.

Resilience, in particular schedule resilience, must be a fundamental part of the architecture assessment. Resilience also comes through agility within each element and between elements. This agility requires enough decoupling such that the individual projects are functionally coupled at the mission level but are otherwise fundamentally independent from a developmental perspective.

Robustness should follow from the architecture. Architectural robustness comes through making each system element as individually robust as practical within the constraints imposed by the system-level requirements. Robustness does not necessarily require a blind application of a Class A approach or redundancy mandates. The goal should be a design that is well matched to the systems' requirements and is focused on resilience that is specific to the mission goals and mission success without added features or complexity.

R9: Architectural Robustness and Resiliency (2 of 2)

Recommendations:

- (R9.1) NASA should examine other potential architectures or variants of the current architecture in order to determine whether there are options that offer greater technical robustness and schedule resilience to launch phasing or launch delays, while fitting within annual budget constraints and lowering programmatic risk.
  - Actively seek and analyze alternatives that might provide larger technical margins than the current architecture while also being technically simpler at the individual project, system, and subsystem levels.
  - The mindset should be simplicity in design, taking advantage of heritage design and approaches, examining mission classification specific to each project element, and execution of a robust cross-cutting V&V program.
  - Enlist independent programmatic analysts to assess each technically viable alternative in order to determine whether the alternative can be implemented with high confidence and lower cost and schedule risk within the annual budget constraints imposed by NASA.
  - Examine which path forward best balances technical and programmatic constraints, given the current status of the Program and the recommendations made in this report.
  - Alternate architectures should be carefully compared to variants of the existing architecture in terms of technical, schedule, and cost risk.
- (R9.2) If a significantly different alternate architecture is determined to be the best way forward for the Program, the MSR Program should consider reallocating responsibilities among participants or adding other experienced and qualified sources when developing an acquisition strategy for the alternatives.

F10: Programmatic Assessment (1 of 3)

The management of MSR as a hybrid Single-Program/Tightly-Coupled Program impedes effective management, limits insight, and introduces inefficiencies and inaccuracies.

- The management structure requires excessive effort to establish and maintain programmatic insight into the individual hardware elements, where most of the budget resides.
- Element managers have limited control over budgets and Unallocated Future Expenses (UFE).

The program is not ready to be baselined, technically or programmatically.

- From the start, the program plan has been over-constrained in terms of budget, schedule, and workforce.
- Program performance to date has been inadequate and has been hampered by unclear responsibility, accountability, authority, and organizational complexity.

The IRB assessment shows that both lifecycle cost and annual funding requirements exceed current estimates in the 2024 President's Budget Request.

- The Program's anticipated annual funding requirements are significantly higher than the PBR (FY24-FY28).
- Analysis shows that the earliest feasible launch opportunities are in 2030, but that assumes an annual allocation of more than $1B during FY25-FY28.
- The current reference architecture is not viable within the likely available funding profile.

Analysis shows that an MSR Program for the present or alternate architectures within the launch dates under consideration, and with the current division of work between NASA and ESA, will:

- Cost NASA at least $8B and as much as $11B.
- Require on the order of $850M-$1B per year during the development period.

The MSR Program is currently managing to one Management Agreement (MA) and one Agency Baseline Commitment (ABC) as part of the hybrid Single-Program/Tightly-Coupled Program structure. This structure complicates the management of programmatics by limiting responsibility, accountability, and authority for budget and schedule performance to only the program level (not the project level), thus limiting visibility into cost and schedule performance of the projects against budget and schedule commitments recognized by SMD and the Agency. Given the size, complexity, and cost of each project, SMD and the Agency need separate budget and schedule commitments at the project level.

The program is not ready to be baselined due to continued technical issues that have yet to be resolved and anticipated fiscal year funding streams that are out of line with Program-estimated funding requirements. The earliest feasible launch opportunities will occur in 2030. However, utilizing these opportunities is only possible with adequate fiscal year funding and a total lifecycle cost on the order of $8-11B.

Programmatic Assessment (2 of 3)

- The IRB's independent programmatic assessment shows that **$8-11B** is the probable range (50%-80% confidence level) for the total MSR Program Life Cycle Cost (LCC) - *Range for IRB ICE ERO (2030)/SRL (2030) and Various Alternate Architectures*
  - The IRB's independent programmatic assessment shows that **$8.0-9.6B** is the probable range (50%-80% confidence level) for the total MSR Program LCC with the earliest probable Launch Readiness Dates (LRDs) for both ERO and SRL in 2030.
  - There exist a variety of potential alternate architectures that the program may choose to consider in order to add robustness and resiliency to the Program and/or to operate within the constraints of a fiscal year budget cap. IRB analysis suggests that the alternate architecture solution space has an LCC range of **$8.4-10.9B** with various LRD combinations in the 2030-2035 timeframe.
- The probable cost range demonstrates the uncertainty based on the various scenarios estimated by the IRB.

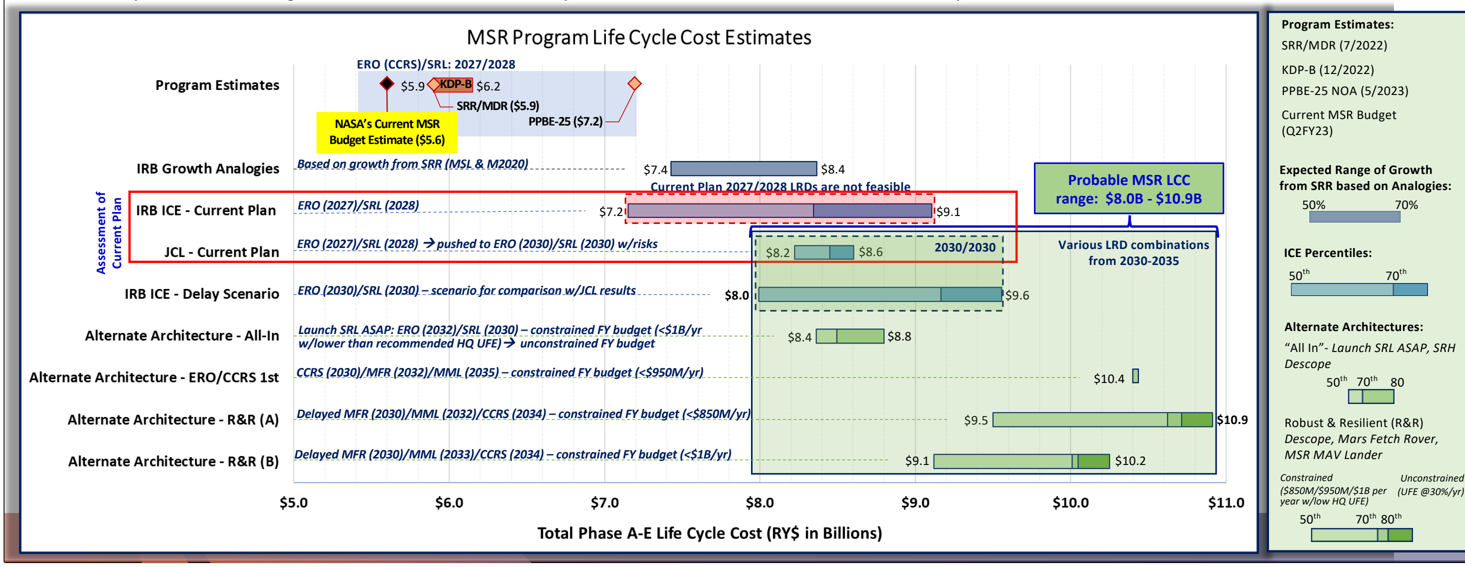

The IRB developed a programmatic assessment of potential MSR Program cost and schedule scenarios, including the program's current plan. While the ICE for the current plan shows a range of costs for the ERO/2027 and SRL/2028 LRDs, the JCL indicates a very low probability of achieving the current LRDs. The earliest likely LRDs for both ERO and SRL are in 2030. The subsequent IRB ICE for the delayed (2030/2030) scenario illustrates a rephased cost plan for comparison. This scenario, combined with potential alternate architectures, provides a probable lifecycle cost range of approximately $8-11B.

The alternate architectures considered were not detailed concepts and were intended only to represent a potential solution space for programmatic analysis. The alternate architectures included decoupled projects (where launches are not dependent on one another) with fiscal year budget scenarios that are either constrained (up to $850M, $950M, or $1B/fiscal year) or unconstrained. Scenarios included:

- The "All In" approach, which prioritized SRL to launch as soon as possible with helicopters removed but retaining the arm.
- The "Robust and Resilient (R&R)" approaches, which disaggregated the landed systems into two missions: 1) MSR Fetch Rover (MFR), which includes a 2xMER rover carrying the OS, scaled-down Skycrane EDL, and STHS, and 2) MSR MAV Lander (MML), which includes the MAV and a MAV landing/launch system using a scaled-up Skycrane EDL. The helicopters and the arm were removed from the MAV lander due to the addition of the MFR. One option launched ERO/CCRS first, with a constrained funding level of $950M/year. Another set of options launched the landed systems first and examined the impact of two constrained funding levels – $850M/year and $1B/year.

Overall, the IRB analysis suggests the alternate architecture solution space has an LCC range of $8.4-10.9B with various LRD combinations in the 2030-2035 timeframe.