

Proposed \$1B Stratospheric Aerosol Injection Experimental Research Fund

By ChatGPT 5.6

Anyone can copy, modify and rename any or all material
at no charge via the [CC BY 4.0](#) license. For original files, click [here](#).

August 10, 2026 | v1r | Copyright © 2026 by Glenn Weinreb

Table of Contents

Table of Contents.....	2
What is Stratospheric Aerosol Injection (SAI)?.....	3
SAI Field Experiments - \$1B Fund	4
C1, C2 and C3 Experimental Pathways	14
SO ₂ versus Direct H ₂ SO ₄ Aerosol.....	20
Boeing 777 Sulfur Feedstock and Aerosol-Generation Energy	22
Reflectivity per Dollar and per Kilogram of Sulfur	24
Project Phases (i.e. what happens first)	27
Purchase Orders	32
Recipients of Purchase Orders.....	38
Project Management	43
Funding Source Requirements	51
Top SAI Papers	59
Top Researchers	66
Scientific Advisory Panel.....	71
Management Advisory Panel.....	75
Agreements.....	81
Termination Shock.....	90
SAI Aircraft, Instrumentation and Plume Tracking.....	94
Further Reading	99

What is Stratospheric Aerosol Injection (SAI)?

Prompt to AI: Summarize Stratospheric Aerosol Injection (SAI) in roughly 300 words.

ChatGPT Conversation

<https://chatgpt.com/share/6a6781b7-3710-83ea-a9f5-7316c51a65f4>

Response From AI:

Stratospheric Aerosol Injection (SAI) is a proposed form of **solar radiation modification** intended to cool Earth by increasing the amount of sunlight reflected back into space. The leading concept would use specialized high-altitude aircraft—potentially supplemented by balloons or other platforms—to release sulfur dioxide or aerosol particles into the lower stratosphere. Sulfur dioxide would react to form tiny sulfate droplets, producing a cooling effect broadly analogous to that observed after major volcanic eruptions.

Climate models indicate that sustained, carefully distributed injections could lower global average temperature relatively quickly and potentially reduce some heat-related climate damages. The cooling could be adjusted by changing the quantity, latitude, altitude and timing of injections. However, SAI would **not remove carbon dioxide**, stop ocean acidification or eliminate the need to reduce greenhouse-gas emissions.

The central concern is that SAI would create a different climate rather than precisely restoring the previous one. Possible effects include changes in regional precipitation and monsoons, heating of the stratosphere, ozone depletion, altered atmospheric circulation, whiter skies, and changes in the balance between direct and diffuse sunlight. Benefits and harms could also be distributed unevenly among countries and regions.

SAI would require continuous operation because aerosols eventually fall out of the stratosphere. If a large cooling program were suddenly stopped while greenhouse-gas concentrations remained high, temperatures could rise very rapidly—a danger known as **termination shock**.

SAI has not been demonstrated as a global climate intervention and remains scientifically, politically, ethically and legally controversial. Unresolved questions include environmental monitoring, international decision-making, compensation for damage, control over deployment, geopolitical misuse and whether reliance on SAI might weaken emissions-reduction efforts. Major scientific assessments therefore describe SAI as a possible risk-reduction tool worth researching—not a substitute for emissions reductions or a technology ready for deployment.

SAI Field Experiments - \$1B Fund

Prompt to AI: If a funding source wants to spend \$1B on Stratospheric Aerosol Injection (SAI) experimental research over 5 years, what experiments should it fund? Assume the reader has not read any other atmospheric-experiments document. Organize the program around C1 (SO₂ gas), C2 (direct H₂SO₄/H₂O aerosol), and C3 (alternative candidate particles). Use release size based on observability and scientific value rather than an arbitrary low mass ceiling.

Response From AI:

A \$1 billion, five-year SAI experimental research program should be designed to answer a practical question: can a specified injection pathway be predicted, measured, controlled, and evaluated well enough to inform a later decision about whether SAI research should continue, change direction, or stop? The program should not attempt to cool a region or the planet. Its purpose is to produce decision-grade knowledge about the physical chain from release to radiative effect and side effects.

The program should be built around three core experimental pathways: C1 for SO₂ gas; C2 for direct sulfuric-acid pathways, including mechanically generated H₂SO₄/H₂O aerosol and H₂SO₄ vapor generated from liquid acid using aircraft exhaust or a dedicated vaporizer; and C3 for alternative candidate particles. Laboratory work, near-field plume facilities, natural analogues, monitoring infrastructure, engineering-economic modeling, and global modeling should support these pathways rather than operate as separate research silos.

Design principle

To minimize harm, we intend to:

Use the minimum atmospheric perturbation that is large enough to produce the measurement needed for a specific decision.

A release that is so small that the plume becomes unobservable before the important chemistry, microphysics, or radiative effect can be measured can waste aircraft time and money. Conversely, increasing release mass solely to create a larger signal can also distort plume microphysics. Release quantity therefore has to be chosen together with release geometry, duration, concentration, altitude, and the measurement objective.

Before each controlled release, the program should perform an observability and value-of-information analysis. That analysis should determine the release mass required to obtain a statistically useful signal and should show why the proposed experiment is preferable to existing observations, laboratory studies, modeling, or a natural volcanic analogue.

The three core SAI experiment families

C1 — Inject SO₂ gas and follow the complete sulfur pathway

C1 tests the simplest sulfur-based SAI pathway: release sulfur dioxide gas into the lower stratosphere and measure what the atmosphere actually produces. The physical chain is:

SO₂ release -> dilution and transport -> oxidation -> H₂SO₄ vapor -> nucleation and/or condensation -> sulfate particle growth -> coagulation -> optical properties -> radiative effect -> heating and ozone-relevant chemistry

C1 is valuable because SO₂ is comparatively straightforward to store and release, and it is the sulfur form used in many climate-model studies. Its weakness as an experiment is time. The sulfate aerosol of greatest interest does not appear instantly. The plume stretches, folds, and dilutes while the chemistry proceeds, so C1 requires repeated plume interception over days and weeks.

- Primary questions: How rapidly does SO₂ convert? Does sulfur form new particles or condense onto existing particles? What particle-size distribution results? How does that distribution evolve? How much scattering is obtained per kilogram of sulfur?
- Measurement window: minutes and hours for near-field dilution; days to weeks for oxidation and particle growth; longer where detectability permits.
- Core instruments: SO₂ and H₂SO₄ gas measurements, nanoparticle and accumulation-mode sizing, aerosol composition, ozone and reactive chemistry, aerosol optics, lidar, radiometry, winds, water vapor, and trajectory data assimilation.
- Success criterion: the program can predict the evolving sulfur mass, particle distribution, optical properties, and chemistry within decision-relevant uncertainty bounds.

C2 — Engineer H₂SO₄ aerosol by competing direct-delivery architectures

C2 bypasses the slow SO₂ oxidation step, but it should no longer be treated as a single injection method. The program should compare at least three direct-acid architectures:

- C2A** Mechanically generated H₂SO₄/H₂O droplets.
- C2B** Concentrated liquid H₂SO₄ pumped into a hot aircraft-exhaust/vaporization region so that H₂SO₄ vapor nucleates and coagulates in the wake; and
- C2C** H₂SO₄ vapor produced by a dedicated heater or heat exchanger independent of the engine exhaust.

The objective is to learn which architecture produces the most useful, reproducible sulfate particle population after the first seconds and minutes of plume processing.

- Primary questions: Which generation mechanism produces the desired particle-size distribution after 1 second, 10 seconds, 1 minute, 10 minutes, 1 hour, and one day? How do mass flow, acid concentration, exhaust temperature, plume dilution, nozzle geometry, compressor or pump power, and ambient water vapor affect nucleation and coagulation? How much radiative forcing is obtained per kilogram of sulfur, per aircraft-hour, and per dollar of total system cost?
- Measurement window: milliseconds and seconds in ground/plume facilities; seconds to hours in atmospheric releases for near-field microphysics and optical closure; optional revisits over one to several days for aging, transport, and persistence.
- Core instruments: particle size and number from nucleation through accumulation mode, composition and acidity, H_2SO_4 vapor, exhaust temperature and dilution, mass-flow and energy measurements, optical scattering/absorption/extinction, lidar plume geometry, above/below radiometry and polarimetry, meteorology, and ozone-relevant chemistry.
- Success criterion: the program can predict delivered particle properties and radiative forcing from the generation mechanism and release conditions, and can rank C2A/C2B/C2C against C1 using standardized radiative-forcing-per-sulfur and radiative-forcing-per-dollar metrics.

C2 should begin earlier than the most expensive C1 campaigns because direct-acid signals are available immediately. Year 1 should emphasize ground and near-field plume experiments that compare C2A/C2B/C2C before large outdoor releases. This allows the program to validate instruments, plume models, radiative closure, and the measurement architecture while eliminating weak aerosol-generation approaches early.

C3 — Inject alternative candidate particles only after laboratory prequalification

C3 evaluates whether non-sulfate particles could provide useful sunlight reflection with lower ozone, heating, persistence, or other side effects. Candidate materials sometimes discussed in the literature include calcite, alumina, silica, and titanium-containing particles. The program should not assume that a material that looks attractive in a radiative model will behave safely or predictably in the real stratosphere.

- Stage 1: laboratory optical measurements, surface chemistry, aggregation, aging, toxicology, and ozone-relevant reaction studies.
- Stage 2: chamber and ground-dispersion tests to determine whether particles can be generated without uncontrolled agglomeration or contamination.

- Stage 3: kilogram-scale atmospheric tests only for candidates that pass predefined laboratory criteria.
- Stage 4: larger closure tests only if the material remains scientifically competitive and the observability analysis shows a clear need.

C3 should receive less early funding than C1 and C2 because sulfate has the strongest natural analogue in volcanic eruptions and the largest existing scientific base. C3 is an option-preserving program, not a reason to delay sulfate work.

Release-size framework

Release size should be selected from the measurement requirement, not from a universal ceiling. The following bands are useful for planning; they are not automatic approvals.

Experimental scale	Illustrative release mass	Primary purpose	Typical observation window
Process-scale	~1–100 kg	Near-field wake physics, nucleation, delivered particle distribution, instrument validation	Seconds to hours
Closure-scale	~100–10,000 kg (0.1–10 tonnes)	Measure plume optics, radiative flux, chemical evolution, and model closure with adequate signal	Hours to days
Pathway/transport campaign	Multi-ton campaign total when justified	Follow deployment-relevant plume evolution and reacquire material over days to weeks	Days to weeks

The upper end of a mass range should be used only when modeling shows that smaller releases will not produce the required signal. A multi-ton campaign does not imply a single concentrated dump. For many questions it is better to distribute material along a long flight path or among repeated sorties so that local plume concentration remains representative of a plausible SAI release geometry.

Illustrative \$1 billion portfolio

These are planning estimates in 2026 dollars, not vendor quotations.

Program	Five-year budget	Main purpose
C1 — SO ₂ pathway experiments	\$270M	Multi-day and multiweek experiments measuring conversion, particle growth, transport, radiative efficiency, heating, and ozone-relevant chemistry

Program	Five-year budget	Main purpose
C2 — H ₂ SO ₄ delivery-architecture and aerosol experiments	\$220M	Compete mechanical aerosol, exhaust-vaporized H ₂ SO ₄ , and dedicated-vaporizer approaches; measure near-field particle formation, optics, radiative forcing per kg-S, and radiative forcing per dollar
C3 — Alternative particle experiments	\$60M	Laboratory-first screening followed by tightly staged field tests for promising non-sulfate candidates
Laboratory/chamber chemistry, microphysics and toxicology	\$70M	Ozone kinetics, aerosol aging, surface chemistry, particle generation, and candidate-material screening
Volcanic/wildfire natural-analogue rapid response	\$60M	Observe large natural stratospheric perturbations that deliberate experiments should not reproduce
Shared aircraft, instruments, lidar, balloons and monitoring infrastructure	\$75M	Reusable monitoring capability for C1/C2/C3 and natural events
Modeling, observability, plume prediction and data assimilation	\$75M	Predict detectability, guide flight targeting, reconstruct plumes, and convert observations into model improvements
Targeted baseline and SAI detection system	\$25M	Measure only background quantities that materially limit interpretation; test detection of meaningful SAI-scale perturbations
Downstream deposition, ecology and cloud-interaction studies	\$25M	Track descent/deposition and study relevant ecological and cloud pathways without pursuing climate-scale cooling
Program management, governance, international participation, open data and independent review	\$70M	Fiduciary control, stage gates, environmental review, participation, data publication, audits, and independent evaluation
Contingency and rapid-response reserve	\$50M	Weather losses, instrument failures, unexpected natural events, and high-value follow-up work

About \$550 million is explicitly centered on C1, C2, and C3. Most of the remaining budget exists to make those experiments measurable, interpretable, and independently credible. Thus, roughly three-quarters of the program is directly connected to the question: what happens when a candidate SAI material is actually released into the stratosphere, and can that outcome be predicted?

Phase 0 — Observability and value-of-information gate

No major outdoor release should be approved until an independent design review receives a written report answering the following questions:

1. What quantity or process is uncertain today?
2. What is the present best estimate and uncertainty range?
3. What later engineering, environmental, or policy decision depends on reducing that uncertainty?
4. Could existing data, a laboratory experiment, a chamber experiment, a volcanic plume, or additional modeling answer the question more cheaply?
5. What release mass and geometry are expected to produce a measurable signal?
6. What concentration and particle signal are expected after 1 minute, 1 hour, 1 day, 1 week, and other relevant times?
7. Which instruments can detect the signal and with what uncertainty?
8. What is the probability of finding the plume again at each planned revisit?
9. Would 100 kg work? Would 1 tonne work? Would 10 tonnes materially improve the answer?
10. At what point would additional material stop adding useful information or begin to distort the process being studied?
11. What environmental exposure and deposition are expected?
12. What result would cause the program to stop or abandon the pathway?
13. What result would justify replication or a larger follow-on experiment?
14. What independently testable predictions will be made before the data are collected?
15. What engineering decision will the result resolve—for example, whether to retain C2A, C2B, or C2C?
16. How will the result change the estimated radiative forcing per kilogram of sulfur, per aircraft-hour, and per dollar?
17. Could a ground engine/exhaust-plume experiment answer the question before any outdoor stratospheric release?

This gate directly addresses the risk of spending tens of millions of dollars on an exquisitely instrumented experiment whose plume is too weak to answer the stated scientific question.

Four-aircraft experimental and radiative-closure architecture

For the most important C1 and C2 closure experiments, separate four functional roles: a release aircraft that creates and precisely logs the perturbation, an above-plume aircraft, a through-plume aircraft, and a below-plume aircraft. The three monitoring aircraft provide complementary remote-sensing, in-situ, and radiative measurements; the release aircraft should normally remain separate so its hazardous-material systems and possible contamination do not compromise clean measurements.

Release/injector	Released mass and mass flow, feedstock state, particle-generation settings, exhaust/vaporizer conditions, aircraft position/attitude and ambient meteorology	Defines the controlled input and provides the source record needed for mass balance, plume modeling and reproducibility
Above plume	Up/down shortwave and longwave radiometry, spectral radiometry, polarimetry, lidar and multispectral imaging, with precision GPS/INS	Measures reflected/absorbed radiation, plume horizontal structure and vertical geometry from above
Through plume	SO ₂ /H ₂ SO ₄ , aerosol size and number, composition/acidity, scattering/absorption, ozone-relevant chemistry, water vapor, winds and turbulence	Provides direct physical and chemical measurements of the material being modeled and anchors remote-sensing retrievals
Below plume	Up/down radiometry, spectral attenuation, UV-visible spectroscopy/DOAS, upward-looking lidar, column optical depth and background context	Provides the second side of the radiative-flux difference and measures what radiation passes through the plume

Ground or mobile lidar, balloon profiles, targeted satellite observations and a real-time plume-location/navigation system should supplement this stack. Not every release requires all four roles simultaneously, but a campaign claiming radiative closure should explain any missing layer of observation. For multi-day C1 work, plume reacquisition should be treated as a core experimental capability rather than an informal flight-planning task.

Illustrative flagship C1 campaign

One useful C1 design is a high-latitude lower-stratospheric SO₂ campaign in a dynamically coherent late-winter air mass. An illustrative campaign might release about 10 tonnes of SO₂ total, distributed over multiple flight legs and sorties rather than concentrated at one point. For example, approximately 20 release sorties could each place about 500 kg along long flight tracks over roughly ten days. These numbers are illustrative; the Phase 0 calculation should determine the actual mass and geometry.

The scientific objective is not to create a surface-temperature signal. It is to generate a sulfur perturbation that can be reacquired often enough to measure the conversion from SO₂ to sulfate, particle-size evolution, optical efficiency, chemical surface area, and transport. Sampling might be concentrated on release day and then repeated around days 1, 2, 4, 7, 10, and 14, with later flights only when the plume remains measurable.

The principal products should be measured curves such as:

- SO₂ mass remaining versus time.
- Sulfate mass formed versus time.
- Particle number and size distribution versus time.
- Aerosol surface area versus time.
- Optical scattering and absorption per kilogram of sulfur versus time.
- Observed plume location and geometry versus pre-registered model forecasts.

If a 10-tonne campaign cannot be reacquired reliably, that is itself a major result: the program would know that multiweek C1 field measurements are less useful than expected and should redirect resources toward different designs, larger natural analogues, or different measurement methods.

Illustrative flagship C2 campaign

C2 can use a faster experimental rhythm, but the first major investment should be a competitive engineering down-selection rather than an immediate large atmospheric release. The program should first test C2A mechanical atomization, C2B exhaust-assisted vaporization, and C2C a dedicated vaporizer in laboratory and full-scale or near-full-scale exhaust/plume facilities. Atmospheric closure events should then use only the architectures that demonstrate reproducible particle generation and acceptable safety.

The strongest science occurs during the first milliseconds through hours. Ground and engine/plume tests should resolve vaporization, nucleation, initial coagulation, and energy demand. Airborne near-field passes should measure the particle distribution after seconds to minutes; mid-field passes should measure continued coagulation, dilution, and water equilibration; and long straight legs above and below the plume should measure radiative flux. A subset of cases should be revisited the following day to measure persistence and continued evolution.

Important Consideration

C2 should produce a map of generation mechanism + feedstock + energy source + mass flow + release geometry -> particle properties -> optical properties -> radiative effect -> radiative forcing per kg-S and per dollar. That map can then be compared directly with C1 (SO₂) to determine whether the theoretical controllability advantage of direct acid survives real aircraft-plume physics and total-system economics.

C3 candidate-material screening

C3 should proceed more conservatively because the atmosphere has far less experience with deliberate stratospheric releases of many proposed solids than it has with sulfate. Each candidate should pass laboratory gates for optical performance, heterogeneous ozone chemistry, aggregation, persistence, toxicology, and degradation products before an outdoor test is considered.

Initial C3 releases should normally remain kilogram-scale and focus on whether the material can be dispersed, identified, and measured without uncontrolled agglomeration. A larger experiment should require a positive scientific case that the candidate could materially outperform sulfate and that the larger release is necessary for optical or chemical closure.

Five-year sequence

Year	Priority	Illustrative activities
Year 1	Build and down-select generation systems	Phase 0 modeling; laboratory/chamber work; build near-field plume and exhaust test capability; compare C2A mechanical aerosol, C2B exhaust vaporization, and C2C dedicated vaporization; build the standardized sulfur/cost-efficiency model; instrument integration and natural-analogue readiness
Year 2	Small C2 closure first; first serious C1 test	Conduct small atmospheric C2 tests of the surviving generation systems with intensive seconds-to-hours measurements; validate radiative-closure geometry; conduct the first serious C1 release only after observability modeling and plume-reacquisition rehearsals
Year 3	Compare sulfur pathways	Major C1 campaign; matched C1/C2A/C2B or C2C experiments under similar conditions; equal-sulfur, equal-aircraft-time, and equal-cost comparisons; selected C3 laboratory-to-field transitions
Year 4	Broaden conditions and scale	Test different H ₂ SO ₄ flow rates, dilution geometries, altitude/latitude/season cases, and background aerosol conditions; second major C1 campaign; 777-class reference-scale mass/energy/cost validation; tightly screened C3 field work

Year	Priority	Illustrative activities
Year 5	Replicate, challenge, and down-select	Independent model challenges and replication; integrated ranking of SO ₂ , mechanical H ₂ SO ₄ aerosol, exhaust-vaporized H ₂ SO ₄ , dedicated-vaporizer H ₂ SO ₄ , and surviving C3 candidates; publish requirements for –1 W/m ² scenarios without authorizing deployment

Program boundaries

The program should impose boundaries based on climate consequence, environmental review, and scientific necessity rather than a blanket prohibition on multi-ton experiments.

- No experiment should be designed to produce detectable regional or global climatic cooling.
- No deployment-scale fleet should be built under the experimental research fund.
- No C3 material should be released outdoors before laboratory toxicology, ozone-chemistry, and persistence review.
- No major release should proceed without the Phase 0 observability/value-of-information gate, environmental review, scientific approval, and independent governance approval.
- Release mass, geometry, and timing must be documented publicly in advance to the extent compatible with safety and regulatory requirements.
- All atmospheric data, model code, calibration records, and negative results should be released on a defined open-science schedule.
- Completion of the five-year program should create no presumption that deployment or larger experiments are authorized.

Primary program deliverable

The most valuable product of the \$1 billion should be an open, validated predictive and technoeconomic system. Given a sulfur precursor, aerosol-generation architecture, release rate, particle specification, aircraft wake, altitude, latitude, season, and meteorological state, the system should predict plume dispersion, chemical conversion, particle evolution, radiative forcing, stratospheric heating, ozone-relevant effects, feedstock and energy requirements, aircraft utilization, and total cost—with quantified uncertainty. The field experiments are valuable to the extent that they make that system demonstrably better.

C1, C2 and C3 Experimental Pathways

Prompt to AI: How should the \$1B program divide effort among C1, C2, and C3? Explain the experiments from first principles, including why C2 should generally be tested before the most expensive C1 campaigns and how paired C1/C2 experiments should be used.

Response From AI:

C1, C2, and C3 are complementary pathways, but C2 should now be treated as an engineering competition rather than one fixed aerosol-release method. C1 asks what the atmosphere makes from SO₂. C2 asks whether engineers can make a more useful sulfate aerosol directly—and which of several physical methods for doing so is best. C3 asks whether a non-sulfate material can outperform sulfur pathways without introducing worse problems.

The core-pathway budget is: \$270 million for C1, \$220 million for C2, and \$60 million for C3 (\$550M total). Shared aircraft, instruments, laboratory work, plume modeling, and governance add substantial support beyond those line items.

Pathway	Five-year core budget	Field emphasis	Why this share
C1 — SO ₂	\$270M	Fewer, longer and more expensive multi-day/multiweek campaigns	Most deployment-relevant simple sulfur pathway; expensive because the plume must be found repeatedly while chemistry evolves
C2 — direct H ₂ SO ₄ pathways	\$220M	Ground-heavy architecture competition followed by repeatable same-day/short-duration closure campaigns and selected aging revisits	Fast signal; tests mechanical aerosol, exhaust vaporization, and dedicated vaporization; directly measures the cost-versus-particle-size tradeoff
C3 — Alternative particles	\$60M	Laboratory-heavy screening; limited staged outdoor tests	Preserves options but should not delay sulfate research; chemistry and toxicology are less established

C1 and C2 test different parts of the same sulfur problem

Feature	C1 — SO ₂ gas	C2 — engineered H ₂ SO ₄ aerosol
Starting material	SO ₂ gas	Concentrated liquid H ₂ SO ₄ and/or H ₂ SO ₄ /H ₂ O solution; final aerosol may be mechanically atomized or formed after vaporization
What must happen before strong sulfate optics appear	Oxidation to H ₂ SO ₄ , then nucleation/condensation and growth	Depends on architecture: preformed droplets equilibrate and mix; vaporized H ₂ SO ₄ nucleates and coagulates in the wake before aging and transport
Strongest experimental window	Minutes to weeks	Milliseconds to hours; optional days
Main uncertainty	Conversion, nucleation, condensation, growth and long-range plume evolution	Generation mechanism, vaporization/atomization, wake dilution, nucleation, coagulation, energy demand, persistence, and cost-normalized radiative efficiency
Operational difficulty	High: repeated plume search and reacquisition	Moderate-to-high: corrosive feedstock plus specialized aerosol-generation hardware; same-day closure is possible
Potential deployment advantage	Simple storage and release hardware	Potentially higher radiative efficiency and particle-size control, if useful size distributions can be produced at practical mass flow and cost

Why begin with C2 engineering and ground tests

C2 gives the program immediate feedback and can answer many of its most important questions on the ground. Mechanical atomization, exhaust-assisted vaporization, dedicated vaporization, corrosion, pumping, heat transfer, and the first milliseconds-to-seconds of particle formation should be challenged in instrumented plume facilities before expensive field campaigns. The surviving systems can then be tested in the atmosphere, where the aerosol can be measured almost at once.

C2 also isolates a central uncertainty in the global-model literature. Models that show an advantage for accumulation-mode H₂SO₄ generally prescribe a post-plume particle distribution; they do not resolve whether a real aircraft can create that distribution at economically useful mass flow. The program should therefore treat generation method and mass flow as primary experimental variables, not merely assume that an ideal particle population leaves the aircraft.

Accordingly, Years 1–2 should emphasize C2 process and closure experiments. C1 modeling, engineering, and rehearsals proceed in parallel, but the first expensive multiweek C1 campaign should use lessons learned from C2.

Illustrative campaign counts

Campaign counts are less important than information value. The following is a reasonable planning architecture if major monitoring hardware is reusable:

Pathway	Illustrative five-year campaign architecture	Illustrative release scale
C2 (H ₂ SO ₄ architectures)	Ground/plume screening of C2A/C2B/C2C; then 6–10 major atmospheric campaigns using only surviving approaches, with paired and replication cases	Ground tests from gram/s to multi-kg/s equivalent flow; 1–100 kg atmospheric process tests; ~0.1–10 t closure events only when needed
C1 (SO ₂)	4–6 major campaigns, including at least two multiweek pathway campaigns and several smaller process/paired experiments	1–100 kg process tests; ~0.1–10 t closure events; multi-ton campaign totals when needed for repeated tracking
C3	Laboratory screening of many candidates; perhaps 4–8 outdoor experiments across only the best candidates	Usually kg-scale at first; larger only after explicit material-specific review

Matched C1/C2A/C2B/C2C experiments

The highest-value sulfur experiments should compare C1 with the best surviving C2 architectures under closely matched atmospheric conditions. A campaign can release comparable sulfur mass on nearby flight tracks as SO₂, mechanically generated H₂SO₄ aerosol, and/or H₂SO₄ vapor formed with exhaust or a dedicated vaporizer, using the same monitoring system to compare them.

A paired campaign should compare at least:

- Delivered and newly formed particle-number distributions during the first hour.
- Particle size and sulfate mass after one day.
- Optical scattering and absorption per kilogram of sulfur.
- Aerosol surface area relevant to heterogeneous ozone chemistry.
- Stratospheric heating per unit radiative forcing.

- Fraction of sulfur remaining in the target altitude range.
- Predictability of dispersion and reacquisition.
- Aircraft payload and operational complexity per unit of useful radiative effect.

The comparison should not ask simply which method leaves more aerosol aloft. It should report three standardized comparisons: equal sulfur, equal aircraft time, and equal total cost. The final metric is cooling-relevant radiative forcing per kilogram of sulfur and per dollar, adjusted for aerosol lifetime, stratospheric heating, ozone-relevant surface area, sulfur deposition, safety, and controllability.

C1 (SO₂) release design: increase total signal without creating an unrealistic local plume

If plume dilution is the problem, increasing total released mass is a rational response. But total mass and local plume concentration should be treated as separate control variables. A campaign can increase total detectable sulfur by lengthening the flight track, extending release duration, or using repeated sorties instead of injecting the entire mass into a short, highly concentrated segment.

This distinction matters because high local concentrations can increase nucleation and coagulation and thereby change the particle distribution the experiment is trying to study. A good C1 design therefore asks two questions at once: How much total sulfur is needed for reacquisition, and what release rate per unit length best represents the pathway being tested?

C2 aerosol-generation experimental matrix

C2 can exploit controlled variation much more aggressively than C1. A campaign can divide a release into distinct segments and vary:

- Aerosol-generation mechanism: mechanical atomization, exhaust-assisted vaporization, or dedicated vaporization.
- Energy source and measured parasitic energy: compressor, pump, electrical/hydraulic power, engine-exhaust heat, or dedicated heater.
- H₂SO₄ vaporization fraction and unevaporated droplet/residue fraction.
- Near-field dilution rate and plume residence time before mixing with background air.
- Initial particle-size distribution.
- Acid/water concentration at release.

- Mass flow per second.
- Release duration and geometry.
- Altitude.
- Ambient temperature and water vapor.
- Background aerosol loading.
- Continuous versus segmented release.

Each segment should be linked to a pre-registered prediction for particle size, number, sulfur mass, radiative properties, energy use, and cost. That makes C2 a calibration laboratory in the actual stratosphere: the program can determine which combinations of generation architecture, engineering settings, and atmospheric conditions produce stable and efficient particle populations, and whether the optimum in forcing per kg-S is also competitive in forcing per dollar.

C3 down-selection

C3 should use a tournament structure. Many candidates can be studied in the laboratory, fewer should reach chamber testing, and only the best should reach the atmosphere. A candidate should be dropped when it fails a major criterion rather than kept alive simply because money was budgeted for it.

Gate	Minimum evidence before advancing
Laboratory optics	Measured spectral scattering/absorption and realistic particle-size assumptions
Chemistry	Acceptable heterogeneous reaction behavior under stratospheric temperature, pressure, humidity and trace-gas conditions
Toxicology/environment	Characterized toxicity and plausible degradation/deposition products
Engineering	Demonstrated generation and de-agglomeration with a known size distribution
Small field test	Measured plume dispersion, particle integrity, composition, and detectability consistent with predictions
Larger closure test	A credible case that the candidate materially improves a decision-relevant metric relative to sulfate

End-of-program decision

At the end of five years, the program should be able to state, with quantified uncertainty, whether SO₂, mechanically generated H₂SO₄ aerosol, exhaust-vaporized H₂SO₄, a dedicated H₂SO₄ vaporizer, or any C3 candidate remains scientifically and economically credible. It is completely acceptable for one or more pathways to be abandoned. The program succeeds by reducing uncertainty and eliminating bad options, not by proving that SAI must advance.

SO₂ versus Direct H₂SO₄ Aerosol

Prompt to AI: One SAI option is to inject SO₂ and then let it oxidize into H₂SO₄ aerosol over several weeks. Another option is for an airplane to emit an H₂SO₄ aerosol directly. If one creates an H₂SO₄ aerosol with an airplane, is that the same as what you get from SO₂ that has naturally converted to an H₂SO₄ aerosol over many weeks? If not, how do these differ? What are the advantages and disadvantages of each option?

Response From AI:

No. The mature aerosol can have similar basic chemical composition—sulfuric acid and water—but the aerosol population is not determined by chemistry alone. Particle number, size distribution, spatial concentration, mixing with background aerosol, and the history of nucleation, condensation, and coagulation determine radiative efficiency and chemical surface area. Two sulfate populations with the same total sulfur mass can therefore behave quite differently.

For C1, SO₂ is released as a gas and is dispersed while oxidation proceeds. The pathway is SO₂ -> oxidation -> H₂SO₄ vapor -> nucleation and/or condensation -> particle growth -> coagulation -> mature sulfate aerosol. Because oxidation and transport occur over days to weeks, much of the sulfate can condense on existing particles or form new particles after the original aircraft wake has diluted.

For direct H₂SO₄, there are at least two physically distinct possibilities. One is to create droplets mechanically and inject an H₂SO₄/H₂O aerosol. Another is to carry liquid H₂SO₄, vaporize it near the aircraft, and let the H₂SO₄ vapor nucleate and coagulate in the rapidly expanding wake. Pierce et al. (2010) proposed the vapor route; Weisenstein et al. (2022) then tested the global consequences of injecting a prescribed accumulation-mode H₂SO₄ distribution, while explicitly noting that the first plume-scale seconds are unresolved and uncertain.

Feature	SO ₂ -> atmospheric sulfate	Direct engineered H ₂ SO ₄
Initial material	SO ₂ gas after release	Preformed droplets or H ₂ SO ₄ vapor generated from liquid acid
Time to sulfate aerosol	Days to weeks for oxidation and growth	Essentially immediate for preformed aerosol; seconds for vapor nucleation and early coagulation
Particle-size control	Indirect; strongly controlled by atmospheric chemistry and background aerosol	Potentially higher, but only if wake microphysics preserves the intended size distribution
Spatial control	SO ₂ can travel and dilute before conversion	Aerosol formation begins close to the release path

Feature	SO ₂ -> atmospheric sulfate	Direct engineered H ₂ SO ₄
Near-field coagulation risk	Lower before oxidation because the precursor is gaseous	Potentially high if direct particles or H ₂ SO ₄ vapor are too concentrated in the wake
Natural analogue	Strong volcanic analogue	Much weaker natural analogue; requires engineering validation
Operational complexity	Relatively simple release hardware, but chilled/liquefied handling may be needed	Corrosive liquid handling plus atomizer, vaporizer, exhaust interface, or other generation hardware
Potential radiative efficacy per kg-S	Lower in current multi-model comparisons at the same sulfur injection	Higher in current models if accumulation-mode particles can actually be produced

At 5 Tg-S/year, the Weisenstein et al. (2022) three-model intercomparison found net top-of-atmosphere forcing of about -0.9 to -2.5 W/m² for SO₂ and -1.6 to -3.8 W/m² for prescribed accumulation-mode H₂SO₄. The corresponding radiative-efficacy ranges are approximately 0.18–0.50 and 0.32–0.76 W/m² per Tg-S/year. The models also produced smaller effective particle radii for accumulation-mode H₂SO₄ than for SO₂ at the same sulfur injection.

The important qualification is that these global models prescribe the H₂SO₄ particle population after plume processing. They do not demonstrate that a real aircraft can create it at the required mass flow. This is why C2 should focus heavily on the first milliseconds, seconds, and minutes after release.

If two pathways somehow produced exactly the same particle number, size distribution, composition, altitude, spatial distribution, and mixing state, the atmosphere would have little reason to care how the H₂SO₄ molecules were produced. The practical difference is that SO₂ and direct-acid architectures usually do not arrive at that same aerosol state.

Key references: Pierce et al. (2010), *Geophysical Research Letters*, DOI 10.1029/2010GL043975; Weisenstein et al. (2022), *Atmospheric Chemistry and Physics*, DOI 10.5194/acp-22-2955-2022; Yu et al. (2024), *Environmental Science & Technology*, DOI 10.1021/acs.est.3c08408.

Boeing 777 Sulfur Feedstock and Aerosol-Generation Energy

Prompt to AI: If the airplane makes H₂SO₄ aerosol, what fuel would be used to produce the aerosol (e.g. for the compressor)? How much would this fuel weigh and cost per Boeing 777 sortie? What would be the feedstock on the airplane (e.g. H₂SO₄, H₂O)? How much would that weight and cost per Boeing 777 sortie?

Response From AI:

The preferred C2 research architecture should not assume that a very large compressor mechanically manufactures the final submicron aerosol. The fund should compare mechanical atomization against a vapor route in which concentrated liquid H₂SO₄ is pumped into a hot exhaust or vaporizer region and then forms particles by nucleation and coagulation in the wake. Franke et al. (2026) finds that pump and power mass are negligible in its purpose-built-aircraft cases, while injection rate and cruise time dominate lofting economics.

A Boeing 777F is useful as a reference-scale mass, energy, and cost model because Boeing lists a 107-tonne gross structural payload. A conventional 777, however, has a maximum operating altitude of about 13.14 km, so it should not be assumed to be a 20-km tropical deployment aircraft. The fund should use the 777 as a scale reference and test platform where scientifically appropriate, while separately extrapolating to purpose-built high-altitude aircraft.

For an illustrative 95-tonne chemical payload—leaving roughly 12 tonnes of the 107-tonne structural payload for tanks and release equipment—95 tonnes of SO₂ contains about 47.5 tonnes of sulfur. Ninety-five tonnes of 98 wt% H₂SO₄ contains about 30.4 tonnes of sulfur. Using the large-scale material-cost assumptions in Franke et al. (2026), approximately \$150/t-S for SO₂ and \$300/t-S for H₂SO₄, the chemical cost is only about \$7,100 for the SO₂ payload and \$9,100 for the H₂SO₄ payload. These are planning values, not vendor quotations.

Architecture	Airborne feedstock	Primary energy source	Research question
C2A mechanical aerosol	Concentrated H ₂ SO ₄ and/or H ₂ SO ₄ /H ₂ O solution	Aircraft electrical/hydraulic power and possibly compressed air	Can atomizers produce sufficiently small droplets at multi-kg/s equivalent flow without excessive compressor mass or unevaporated droplets?
C2B exhaust-assisted vaporization	Concentrated liquid H ₂ SO ₄ ; normally little or no separately carried water	Waste heat in engine exhaust plus small pump/metering power	Can exhaust heat vaporize the acid at useful flow while producing a controlled

Architecture	Airborne feedstock	Primary energy source	Research question
			200–300 nm-class particle population after wake processing?
C2C dedicated vaporizer	Concentrated liquid H ₂ SO ₄	Dedicated heater/heat exchanger, potentially Jet A or electrical/thermal power	Does independent control of temperature and vaporization justify the added fuel, mass, complexity, and maintenance?

Franke et al. reports a latent heat of vaporization for H₂SO₄ of 67.6 kJ/mol, about 0.69 MJ/kg, and assumes H₂SO₄ can be carried as a liquid near ambient surface temperature. For approximately 95 tonnes of H₂SO₄, latent heat alone is roughly 65 GJ. If otherwise-wasted engine exhaust supplies that heat, incremental fuel for vaporization could be small, although engine integration, pressure loss, heat transfer, corrosion, and maintenance penalties still have to be measured. A dedicated burner would instead consume several tonnes of Jet A in a 95-tonne example once sensible heating and realistic heat-transfer efficiency are included.

For a fuel-cost sensitivity example, this chat used a 5-hour 777 mission burning about 40 tonnes of Jet A. At the U.S. Gulf Coast spot price of \$3.506/gal on August 3, 2026 and a standard reporting density of 0.8 kg/L, 40 tonnes corresponds to roughly \$46,000 of fuel. Actual 777 mission fuel must be obtained from a certified flight-performance model and should be varied with altitude, payload, route, loiter time, and release rate.

The feedstock should not automatically include enough water to make a finished ~75/25 H₂SO₄/H₂O droplet solution. Carrying water consumes payload, while H₂SO₄ aerosol can take up water from exhaust and ambient air. C2 experiments should therefore vary acid concentration and measure water uptake directly rather than assuming that the aircraft must lift all final aerosol water.

Key reference: Franke et al. (2026), “Trade-Offs between Sulfuric Acid Aerosol Precursors for Stratospheric Aerosol Geoengineering,” Environmental Science & Technology, DOI 10.1021/acs.est.5c15172.

Reflectivity per Dollar and per Kilogram of Sulfur

Prompt to AI: Let us assume we are working with a Boeing 777 aircraft and we have several options. These include things like: (a) release SO₂ gas, (b) release H₂SO₄ aerosol via compressor, and (c) release H₂SO₄ liquid and have aircraft exhaust cause it to become a vapor, etc. With each of these options, how much reflectivity do we get per dollar cost? Include operations cost, injected material cost, jet fuel cost, etc. Also, how much reflectivity do we get per kg of sulfur?

Response From AI:

The fund should report “reflectivity” primarily as global-mean top-of-atmosphere radiative forcing, because that is the quantity used by the SAI model literature. Two primary figures of merit should be mandatory: radiative forcing per unit sulfur injected and radiative forcing per total lifecycle dollar. The first captures atmospheric/material efficiency; the second captures whether a more sophisticated aerosol is actually worth the extra aircraft time and hardware.

Metric	SO ₂	Accumulation-mode H ₂ SO ₄ in 2022 global models	Interpretation
Net forcing at 5 Tg-S/year	−0.9 to −2.5 W/m ²	−1.6 to −3.8 W/m ²	Model spread is large; direct H ₂ SO ₄ is systematically more efficacious within each model
Radiative efficacy	0.18–0.50 W/m ² per Tg-S/year	0.32–0.76 W/m ² per Tg-S/year	A midpoint comparison is about 1.6× more forcing per kg-S for idealized AM-H ₂ SO ₄
Effective wet radius at 5 Tg-S/year	0.40–0.52 μm	0.27–0.39 μm	Smaller H ₂ SO ₄ -derived particles are closer to efficient scattering sizes in these simulations

A simple 777 payload calculation reveals why the per-kilogram advantage does not automatically become a per-sortie advantage. SO₂ is about 50% sulfur by mass, while 98% H₂SO₄ is about 32% sulfur by payload mass. In the illustrative 95-tonne chemical-load case, SO₂ carries about 47.5 tonnes S and 98% H₂SO₄ about 30.4 tonnes S—a 1.56× sulfur-payload advantage for SO₂. That is close to the ~1.6× midpoint radiative-efficacy advantage of idealized accumulation-mode H₂SO₄, so the two effects can roughly cancel at equal chemical payload.

Illustrative 777F reference case	SO ₂	Idealized direct H ₂ SO ₄
Chemical payload	95 t	95 t of 98 wt% H ₂ SO ₄
Sulfur payload	47.5 t-S	30.4 t-S
Approx. material cost using Franke et al. assumptions	\$7.1k	\$9.1k
Illustrative fuel cost used in this chat	~\$46k for a 40-t Jet A mission	~\$46k before any dedicated-heater penalty
Illustrative nonfuel aircraft/O&M allowance used in this chat	\$45k/sortie planning placeholder	\$45k/sortie planning placeholder
Approx. total before detailed release-system O&M	~\$100k/sortie	~\$104k/sortie
Modeled forcing per kg-S	Lower	Higher if the desired accumulation-mode distribution is actually achieved
Primary unknown	Long-term conversion, size growth, reacquisition	Whether practical generation rate preserves the modeled particle-size advantage

Using those deliberately simplified chat assumptions, \$1 billion/year would deliver about 0.475 Tg-S/year as SO₂ or about 0.293 Tg-S/year as H₂SO₄. Combining those sulfur delivery rates with the 2022 model ranges gives approximately 0.086–0.238 W/m² for SO₂ and 0.094–0.222 W/m² for idealized accumulation-mode H₂SO₄. Their midpoints are therefore almost identical. This is an illustrative sensitivity calculation, not a deployment estimate: the 777 altitude, actual flight fuel, release-system O&M, and especially the delivered H₂SO₄ particle distribution are unresolved.

A second benchmark comes from Franke et al. (2026), which uses a purpose-built 20-km aircraft rather than a 777. Their central total costs are about \$2,280/t-S for SO₂, \$5,150/t-S for H₂SO₄ at a 4.8 kg/s material injection rate, and about \$70,300/t-S for H₂SO₄ at 0.1 kg/s. The 0.1 kg/s case is intended to produce roughly 0.1 μm-diameter aerosols in the Yu et al. engine/plume configuration; Franke et al. estimates it is about 14× more expensive per kg-S and about 4× more expensive per watt of forcing than the high-rate H₂SO₄ case. They also note that a larger engine could permit higher useful injection rates.

The conclusion for the research fund is not that one option already wins. It is that forcing per kg-S and forcing per dollar can rank technologies differently. The optimum may be an intermediate H₂SO₄ vaporization rate rather than either the slowest “perfect particle” rate or the maximum throughput rate. The experiment should therefore map H₂SO₄ flow over a wide range—roughly 0.1, 0.25, 0.5, 1, 2, and several kg/s equivalent rates—against plume dilution, particle size, optical properties, aircraft time, energy use, and total cost.

The standardized final comparison should include three scenarios: equal sulfur mass, equal aircraft-hour, and equal dollar. A fourth systems analysis should ask what fleet, sulfur mass, energy, annual

cost, ozone-relevant surface area, and stratospheric heating would be required to produce a specified forcing such as -1 W/m^2 .

Key references: Weisenstein et al. (2022), DOI 10.5194/acp-22-2955-2022; Yu et al. (2024), DOI 10.1021/acs.est.3c08408; Franke et al. (2026), DOI 10.1021/acs.est.5c15172.

Project Phases (i.e. what happens first)

Prompt to AI: With the C1/C2/C3 experimental program, what are the phases of each experiment, what happens first, and how much might a representative campaign cost? Include an observability/value-of-information phase before outdoor releases.

Response From AI:

Each major field campaign should be treated as a complete decision cycle rather than merely a release flight. The program should pay for design, modeling, environmental review, engineering, monitoring, analysis, and independent evaluation as parts of one experiment.

Universal Phase 0: observability and value of information

Phase 0 occurs before expensive engineering or flight commitments. It produces an independent go/redesign/stop recommendation based on detection limits, predicted plume evolution, scientific decision value, alternatives, and environmental consequence. A campaign that cannot show a credible path to a useful measurement should stop at Phase 0.

Phase 0 work package	Deliverable	Typical planning cost
Current-knowledge review	Best estimate, uncertainty range, and decision that depends on the result	\$0.2M–\$0.5M
Forward plume/chemistry/optics ensemble	Predicted signal versus release mass, geometry, altitude, and time	\$0.3M–\$0.8M
Instrument detectability analysis	Detection limits, signal-to-background, sampling statistics and reacquisition probability	\$0.2M–\$0.5M
Alternative-method comparison	Why field work is superior to lab, chamber, natural analogues or existing data	\$0.1M–\$0.3M
Independent review	Published go/redesign/stop recommendation and release-size range	\$0.2M–\$0.5M
Engineering/economic decision model	How the experiment changes RF/kg-S, RF/aircraft-hour and RF/\$; whether a ground test is sufficient	\$0.2M–\$0.6M

Representative C2 campaign — competing direct H₂SO₄ architectures

A repeat C2 closure campaign is assumed to use one dedicated release aircraft and up to three monitoring roles, with multiple controlled release segments over approximately two to three weeks of

field operations. The release quantity might range from process-scale kilograms to closure-scale tonnes depending on the Phase 0 observability result.

Phase	Central estimate	Principal work
0. Observability/value-of-information	\$1.2M	Signal prediction, architecture comparison, mass/geometry selection, detectability, and decision analysis
1. Scientific and engineering design	\$1.3M	Hypotheses, generation-method matrix, uncertainty budget, pre-registered particle and cost predictions
2. Ground/plume architecture down-selection	\$4.5M	Mechanical atomization, exhaust vaporization, dedicated vaporizer, corrosion, heat transfer, and near-field microphysics
3. Governance/environmental review	\$1.2M	Permits, environmental assessment, public disclosure, independent approval
4. Release-system preparation	\$4.5M	Corrosion-resistant storage, metering, atomizers/vaporizers, exhaust interfaces, controls, and acceptance tests
5. Instrument preparation	\$3.0M	Particle sizing, H ₂ SO ₄ vapor, composition, optics, energy/mass flow, radiometry, lidar, and calibration
6. Aircraft integration/certification	\$3.0M	Release and monitoring payload integration, airworthiness, hazardous-material and engine-interface safety
7. Baseline/rehearsal	\$1.5M	Background measurement, plume-search rehearsal, timing, communications, and formation practice
8. Release operations	\$3.5M	Multiple architectures/flow segments and coordinated above/through/below measurements
9. Hours-to-days follow-up	\$1.5M	Revisits for aging, transport, coagulation, persistence, and water equilibration
10. Radiative, microphysical and economic closure	\$3.5M	Particle evolution, forcing/kg-S, forcing/\$, heating, chemistry, energy use, and uncertainty analysis
11. Independent evaluation/open archive	\$1.0M	Quality audit, blinded model comparison, technoeconomic audit, and complete public data package
12. Contingency	\$1.3M	Weather losses, repairs, replacement sorties and schedule extension

Representative C2 repeat campaign total: approximately \$31 million. The first year should also make a shared investment in near-field plume and exhaust testing so that each field campaign does not pay to rediscover basic generation physics. A first-of-a-kind architecture can cost substantially more; once the capability is mature, a new release segment can be much cheaper than starting a new campaign from scratch.

Representative C1 campaign — SO₂ gas

A major C1 campaign is assumed to create a sufficiently detectable sulfur perturbation for repeated sampling over approximately two to six weeks. One illustrative design is a multi-ton total release distributed across flight legs and multiple sorties. The expensive part is not the SO₂ itself; it is repeatedly finding and measuring an evolving three-dimensional plume.

Phase	Central estimate	Principal work
0. Observability/value-of-information	\$1.5M	Release mass/geometry, oxidation and dispersion ensemble, detectability and reacquisition probabilities
1. Scientific design	\$1.5M	Sampling schedule, tracer plan, chemistry targets, uncertainty budget
2. Laboratory/chamber support	\$2.0M	SO ₂ oxidation/particle-formation tests, instrument validation and chemistry checks
3. Governance/environmental review	\$2.0M	Permits, transboundary planning, environmental assessment and independent approval
4. SO ₂ release system	\$3.0M	Storage, metering, controls, tracer, safety systems and ground validation
5. Chemical instrument preparation	\$5.0M	SO ₂ , H ₂ SO ₄ , nanoparticles, sulfate, ozone/reactive chemistry and aerosol optics
6. Aircraft/balloon/lidar integration	\$3.5M	Release platform, sampling aircraft, remote sensing and communications
7. Background/rehearsal	\$2.0M	Baseline flights, lidar background, trajectory rehearsal and search practice
8. Release-day experiment	\$3.0M	Release and intensive near-field sampling
9. Days 1–7 sampling	\$7.0M	Repeated plume-interception sorties, balloons, lidar and twice-daily forecasts
10. Weeks 2–6 sampling	\$10.0M	Long-range reacquisition, repositioning, distributed observations and contingency opportunities
11. Plume reconstruction/data assimilation	\$3.0M	Three-dimensional sulfur mass balance, source inversion and trajectory evaluation
12. Chemistry/microphysics/radiative analysis	\$4.0M	Oxidation, growth, coagulation, optics, ozone and heating assessment
13. Independent evaluation/open archive	\$1.0M	Quality audit, blinded forecast challenge and public repository

Phase	Central estimate	Principal work
14. Contingency	\$4.5M	Weather, missed plume, instrument repair, repositioning and schedule extension

Representative major C1 campaign total: approximately \$53 million. Smaller process-focused C1 experiments could be much cheaper. A polar-vortex or other coherent-air-mass campaign with extensive reacquisition may be more expensive, especially if several high-altitude aircraft are required.

Representative C3 candidate-particle pathway

C3 is best budgeted by candidate and gate rather than as a fixed-size flight campaign. Most candidates should be eliminated before they ever reach an aircraft.

Stage	Typical cost per candidate	Purpose
1. Laboratory optical/chemical/toxicology screen	\$0.5M–\$2M	Determine whether the material deserves further work
2. Chamber and generation engineering	\$1M–\$4M	Measure aggregation, aging, surface chemistry and achievable size distribution
3. Small atmospheric test	\$3M–\$8M	Kilogram-scale dispersion, detectability, composition and microphysics
4. Closure-scale follow-on	\$8M–\$20M	Only for top candidates; radiative/chemical closure at a signal large enough to be useful

The \$60 million C3 portfolio is therefore enough to screen many materials in the laboratory but should support only a small number of atmospheric tests. If no candidate beats sulfate on the predefined scorecard, unused C3 funds should be redirected or returned rather than spent to complete a quota.

Shared capability should be built once and reused

Shared capability	Five-year planning investment	Use
High-altitude aerosol/trace-gas/radiation instrument suite	\$25M	Reusable C1/C2/C3 payloads, calibration and spares
Modular sulfur-generation and near-field test capability	\$15M	SO ₂ release plus mechanical H ₂ SO ₄ aerosol, exhaust-assisted H ₂ SO ₄ vaporization, dedicated vaporization, controls, corrosion testing, and ground validation

Shared capability	Five-year planning investment	Use
Plume location, imaging and navigation system	\$10M	Real-time 3-D plume location and geometry using lidar, optical/polarimetric imaging, UV-visible SO ₂ information, trajectory forecasts and in-situ plume hits
Forecasting, trajectory and data-assimilation center	\$10M	Observability studies, mission targeting, plume uncertainty envelopes, flight guidance and post-flight reconstruction
Balloon and distributed observation capability	\$8M	Profiles and extended spatial coverage
Common mission communications and data system	\$4M	Time synchronization, common GPS/INS, telemetry, live aircraft/plume displays and quick-look products
Common calibration/QA reserve	\$3M	Cross-campaign calibration and independent checks

These values total the \$75 million shared-infrastructure line in the portfolio. Laboratory and governance capabilities are budgeted separately in their own program lines.

Stage-gate rule

A phase is not automatically entitled to the money assigned to the next phase. Each major experiment should pass a gate after Phase 0, laboratory/chamber work, environmental review, engineering readiness, flight readiness, and post-experiment analysis. If the plume is not observable, the material cannot be generated reproducibly, a cheaper analogue answers the question, or the expected information no longer changes a decision, the activity should stop.

Purchase Orders

Prompt to AI: For the C1/C2/C3 program, what purchase orders, task orders, grants, or subawards might the Managing Organization issue? Keep the contracting structure practical and tie payments to measurable scientific deliverables.

Response From AI:

The program should not issue hundreds of unrelated purchase orders. It should establish a limited number of master agreements and then issue experiment-specific task orders. Universities and laboratories performing research would often receive subawards or sponsored-research agreements; commercial engineering, aircraft services, instruments, calibration, and consulting are more conventional procurements.

Recommended master agreements

Master agreement	Typical task orders covered	Illustrative five-year ceiling
Aircraft and flight operations	Baseline, rehearsal, release, plume-interception, long-range search and standby sorties	\$140M
Payload integration and airworthiness	Mechanical/electrical integration, inlets, racks, certification, hazardous-material review	\$45M
Release-system engineering	SO ₂ systems; mechanical H ₂ SO ₄ aerosol; exhaust-assisted and dedicated H ₂ SO ₄ vaporization; controls; near-field plume rigs; C3 dispersal prototypes	\$45M
Atmospheric instrumentation and calibration	Gas, aerosol, optics, radiation, lidar, spares, calibration and field operation	\$80M
Modeling, observability and mission forecasting	Phase 0 studies, plume forecasts, data assimilation, inversion and model challenges	\$75M
Balloon, lidar and distributed observing	Balloon launches, mobile lidar, remote sites and rapid-response deployments	\$35M
Scientific analysis consortium	University/laboratory subawards for chemistry, microphysics, radiation and downstream effects	\$105M
Governance, environmental review and independent evaluation	Environmental assessment, engagement, audit, safety review and open-data oversight	\$70M

The ceilings above overlap the portfolio categories and should not be summed as a second budget. They illustrate how money can be packaged contractually while remaining tied to the approved annual operating plan.

Common Phase 0 and cross-program orders

Code	Procurement package	Deliverable	Illustrative amount
G0-01	Observability ensemble	Release-mass/geometry versus signal and reacquisition probability for each proposed field campaign	\$0.3M–\$0.8M/campaign
G0-02	Independent value-of-information review	Written decision analysis and go/redesign/stop recommendation	\$0.2M–\$0.5M/campaign
G0-03	Instrument detectability audit	Detection limits, background variability, uncertainty budget and required sampling density	\$0.2M–\$0.5M/campaign
G0-04	Environmental consequence model	Expected transport, exposure, deposition and cumulative-release assessment	\$0.2M–\$0.6M/campaign
G0-05	Open data and reproducibility service	Repository, metadata schema, code archive, DOIs and release-deadline monitoring	\$2M–\$4M/year

C2 — direct H₂SO₄ delivery-architecture task orders

Code	Package	Required deliverable	Illustrative amount
C2-01	Science-design subaward	Architecture matrix, hypotheses, closure metrics, equal-S/equal-time/equal-cost comparisons, and pre-registered predictions	\$1.2M
C2-02A	Mechanical aerosol-generator task order	Acid-compatible atomizer/compressor system; measured droplet distribution, flow, power, and unevaporated fraction	\$3.0M
C2-02B	Exhaust-vaporization task order	Metered H ₂ SO ₄ injection into representative hot exhaust; vaporization fraction, corrosion, heat-transfer and engine-interface report	\$4.0M
C2-02C	Dedicated-vaporizer task order	Independent heater/heat-exchanger architecture; energy, mass, temperature-control and aerosol-generation performance	\$3.0M
C2-03	Near-field plume facility	Milliseconds-to-minutes measurements of nucleation, coagulation, dilution and particle-size evolution across flow rates	\$4.0M

Code	Package	Required deliverable	Illustrative amount
C2-04	777/reference-scale mass-energy-cost model	Open model of payload, tank/system mass, sulfur density, fuel, aircraft time, feedstock, maintenance and forcing/\$	\$1.0M
C2-05	Aerosol microphysics instrument package	Nanometre-to-micrometre size, number, H ₂ SO ₄ vapor, composition and inlet-loss characterization	\$1.4M
C2-06	Optics/radiation package	Scattering, absorption, extinction, lidar, polarimetry and above/below flux measurements	\$1.6M
C2-07	Release-aircraft mission order	Approved release sorties with certified material mass, flow, energy, time and location records	\$1.3M
C2-08	In-plume sampling-aircraft order	Immediate and repeated plume crossings with synchronized chemistry/microphysics and plume-age data	\$1.2M
C2-09	Above/below closure aircraft orders	Radiative legs and mapped plume geometry	\$1.5M
C2-10	Plume-location, forecasting and real-time guidance	Fused trajectory, lidar/imaging and in-situ information; live plume map, uncertainty envelope, quick-look data and updated aircraft targeting	\$0.5M
C2-11	Hours-to-days follow-up option	Additional sorties, lidar or balloons when plume aging remains measurable	\$1.2M
C2-12	Radiative/microphysical/economic closure	Measured-versus-predicted forcing/kg-S, forcing/\$, particle evolution, heating, chemistry and lifecycle cost	\$2.5M
C2-13	Independent data/model challenge	Withheld-observation predictions, calibration/QA audit, cost-model audit and final public report	\$0.8M

These line items show the major scientific procurements; campaign logistics, governance, integration, field personnel, and contingency are funded through the associated master agreements and annual plan. The key acceptance test is not that a flight occurred. It is that the specified measurement was delivered with the required uncertainty and documentation.

C1 — SO₂ pathway campaign task orders

Code	Package	Required deliverable	Illustrative amount
C1-01	Integrated science and observability design	Release mass/geometry, chemistry targets, sampling calendar and detectability envelope	\$1.5M
C1-02	SO ₂ release-system task order	Storage, metering, tracer, controls, leak detection, safety and ground validation	\$3.0M

Code	Package	Required deliverable	Illustrative amount
C1-03	SO ₂ /H ₂ SO ₄ gas measurement package	Fast-response calibrated gas measurements with spares and field support	\$1.3M
C1-04	Nanoparticle/sulfate package	New-particle formation, size growth, composition and sulfur mass measurements	\$1.6M
C1-05	Ozone/reactive chemistry and optics package	Ozone-relevant gases, aerosol surface area, scattering/absorption and heating diagnostics	\$1.4M
C1-06	Release-day aircraft block	Release aircraft plus immediate chase sampling and remote sensing	\$2.5M
C1-07	Days 1–7 sampling block	Approximately four to eight primary plume-interception sorties plus search support	\$5M–\$7M
C1-08	Weeks 2–6 high-altitude follow-up	Periodic direct sampling, long-range search and repositioning	\$7M–\$10M
C1-09	Distributed balloons/lidar	Profiles and aerosol-layer searches along predicted transport corridors	\$1M–\$2M
C1-10	Plume-location/forecast/data-assimilation center	Twice-daily and on-demand 3-D plume estimate, uncertainty envelope, lidar/imaging assimilation and aircraft flight recommendations	\$1.5M–\$2.5M
C1-11	Sulfur mass-balance reconstruction	3-D SO ₂ , gaseous products and sulfate mass through time	\$0.8M
C1-12	Chemistry/microphysics/radiative closure	Oxidation, nucleation, condensation, coagulation, optics and ozone assessment	\$2M–\$3M
C1-13	Independent trajectory/model challenge	Firewalled predictions and comparison against withheld observations	\$0.6M
C1-14	Open archive/final technical report	Calibrated dataset, code, uncertainty package, negative results and synthesis	\$0.4M

C3 — alternative-particle procurement packages

Code	Package	Required deliverable	Illustrative amount
C3-01	Candidate optical characterization	Spectral refractive index, scattering/absorption and particle-size sensitivity	\$0.3M–\$0.8M/candidate
C3-02	Stratospheric chemistry screen	Surface reactions, aging and ozone-relevant chemistry under controlled conditions	\$0.5M–\$1.5M/candidate
C3-03	Toxicology/environmental screen	Hazard, degradation products, deposition and exposure assessment	\$0.3M–\$1.0M/candidate

Code	Package	Required deliverable	Illustrative amount
C3-04	Particle-generation engineering	Generate and de-agglomerate a known size distribution; publish test data	\$0.5M–\$2M/candidate
C3-05	Small field-test package	Kilogram-scale controlled release with composition confirmation, plume mapping and microphysics	\$3M–\$8M/test
C3-06	Optional closure-scale follow-on	Only after board approval; optics, chemistry and radiative closure at a decision-useful signal	\$8M–\$20M/test

Other portfolio-level purchase orders and subawards

Portfolio activity	Example contracting package	Five-year program budget
Laboratory/chamber research	Multi-institution cooperative awards for ozone kinetics, aerosol aging, microphysics and toxicology	\$70M
Natural-analogue rapid response	Standing aircraft/balloon/instrument options activated by qualifying volcanic or wildfire events	\$60M
Targeted baseline/detection	Measurement campaigns only where model interpretation is limited; detection tests for meaningful perturbations	\$25M
Downstream effects	Deposition modeling/measurement, cirrus interaction, ecological and solar-resource studies	\$25M
Contingency	Unobligated program reserve plus pre-priced replacement-flight and rapid-response options	\$50M

Contract acceptance and payment rules

- Every major task order must state the uncertainty being reduced and the decision the result will inform.
- The final payment should depend on accepted data, calibration, metadata, source files, and documentation—not merely completion of field operations.
- Release quantity, material, major subcontractors, data rights, or acceptance criteria may not be changed by informal email; a formal change order is required.
- Negative or inconclusive results are deliverables and must be published.

- Fixed-price milestones are appropriate for defined engineering work; uncertain aircraft/research operations may require cost-reimbursement or time-and-materials structures with ceilings.
- Replacement-flight options should be pre-priced but exercised only when the scientific value remains high.
- The organization that performs the experiment should not be the sole organization auditing its calibration, environmental compliance, or model performance.

Recipients of Purchase Orders

Prompt to AI: For the C1/C2/C3 purchase orders and subawards, who might receive them? Identify plausible organizations by function, while preserving independent review and avoiding concentration of authority.

Response From AI:

The program should use a consortium model. No single institution has all of the required capabilities in high-altitude aircraft, stratospheric chemistry, aerosol microphysics, plume modeling, environmental review, data systems, and independent evaluation. The organizations listed below are plausible candidates based on the capabilities identified in the prior program analysis; actual participation would require competition, willingness, legal authority, and conflict review.

Program-wide functional leads

Function	Plausible lead or recipient	Possible partners
Managing Organization / fiduciary prime	RTI International; UCAR or Battelle as alternatives	Independent Agreement Oversight Board
Atmospheric-science integrator	UCAR/NSF NCAR	NOAA CSL/CIRES, Cornell, universities, international partners
Stratospheric measurements and chemistry	NOAA Chemical Sciences Laboratory / CIRES	Harvard, NASA instrument teams, university laboratories
High-altitude aircraft operations	NASA JSC WB-57 Program; NASA Armstrong ER-2 Program	NSF NCAR Gulfstream V or other research aircraft where altitude permits
Field-campaign logistics	NASA Earth Science Project Office or UCAR Earth Observing Laboratory	Aircraft operators, local logistics providers
Release-system engineering	Southwest Research Institute, Creare, or Battelle	Specialist valves, pumps, nozzle, gas and materials suppliers
Modeling and mission forecasting	NSF NCAR/UCAR and NOAA CSL	Cornell, NASA Ames, independent university teams
Environmental review	AECOM, Ramboll, or RTI unit independent of technical release leadership	Aviation/environmental law counsel
Public and jurisdictional engagement	Consensus Building Institute or Keystone Policy Center	Regional and international partners
Independent scientific review	National Academies or a separately appointed international panel	External aerosol, ozone, climate, ecology and aviation experts

C1 — SO₂ pathway recipients

C1 package	Likely recipient type / example	Key reason
C1-01 science/observability design	CIRES/CU Boulder + NOAA CSL; Cornell/NCAR as independent model partners	Combines stratospheric measurement, chemistry and plume modeling
C1-02 SO ₂ release system	SwRI or Creare; industrial gas suppliers such as Airgas/Linde for certified gases	Aircraft systems engineering plus qualified gas handling
C1-03 SO ₂ /H ₂ SO ₄ instruments	NOAA CSL instrument teams and atmospheric-chemistry collaborators	Direct relevance to sulfur conversion measurements
C1-04 nanoparticle/sulfate package	NOAA CSL, NASA/academic aerosol teams; commercial aerosol instrument OEMs	Nanometre-to-micrometre sizing and composition
C1-05 ozone/optics	NOAA CSL, NASA Langley, Harvard atmospheric chemistry	Reactive chemistry, aerosol optics and radiation
C1-06 through C1-08 aircraft	NASA WB-57/ER-2; NCAR GV for complementary long-range work where altitude allows	High-altitude in-situ sampling, overflight and search
C1-09 balloons/lidar	NOAA balloon teams, Aerostar, NASA/NOAA/university lidar groups	Distributed plume profiles and tracking
C1-10 forecast/assimilation center	NOAA CSL + NCAR	Operational trajectories, chemistry and data assimilation
C1-11/C1-12 analysis	NOAA CSL, NCAR, Cornell, Harvard, NASA Langley and selected international groups	Independent expertise across mass balance, microphysics, chemistry and radiation
C1-13 independent model challenge	A firewalled Cornell/NCAR/NASA team separate from the primary prediction team	Prevents the operating team from grading its own forecast

C2 — direct H₂SO₄ architecture recipients

C2 package	Likely recipient type / example	Key reason
C2-01 science design	CIRES/CU Boulder, NOAA CSL, Cornell, UChicago and aerosol specialists	Links particle-delivery settings and engineering alternatives to measurable closure questions
C2-02A mechanical aerosol	SwRI, Creare, Battelle or other aerosol-generation engineering groups; acid-compatible component suppliers	Atomization, compressor/pump design, corrosion-resistant storage and metering

C2 package	Likely recipient type / example	Key reason
C2-02B exhaust vaporization	SwRI/Creare/Battelle with a qualified engine OEM or propulsion research partner where required	Heat transfer, engine/exhaust integration, vaporization, materials and airworthiness
C2-02C dedicated vaporizer	Thermal-systems engineering firms, SwRI/Creare/Battelle, specialist heat-exchanger suppliers	Independent temperature and vaporization control with measured mass and energy penalty
C2-03 near-field plume	SUNY Albany/Yu collaborators, Harvard/Keutsch, UChicago, NOAA and plume/CFD laboratories	Nucleation and coagulation in the first milliseconds to minutes
C2-04 scale/economic model	UChicago climate-systems engineering, Cornell, independent aircraft-cost analysts	Translates measured particle performance into payload, aircraft-hour, fleet and forcing/\$ metrics
C2-05/C2-06 instruments	NOAA aerosol teams, NASA Langley, commercial aerosol instrument OEMs and academic groups	Fast particle sizing, vapor/composition, optics and radiative closure
C2-07 release aircraft	Certified operator selected for the approved release system	Dedicated emitter kept separate from measurement aircraft
C2-08/C2-09 monitoring aircraft	NASA WB-57/ER-2 and other qualified research aircraft	In-plume sampling and above/below radiometry
C2-10 forecasting	NOAA CSL and NCAR	Real-time plume targeting and post-flight reconstruction
C2-12 closure analysis	NOAA CSL, NASA Langley, NCAR, Cornell/UChicago and independent universities	Compare measured particles/optics/flux/cost to preregistered predictions
C2-13 independent evaluation	RTI or independent metrology/statistics/economics team plus a separate modeling group	Data quality, blind prediction testing and independent cost audit

C3 — alternative-particle recipients

C3 activity	Plausible recipients	Selection rule
Optical characterization	University/national-lab aerosol optics groups; NIST-traceable laboratories where appropriate	Capability to measure spectral optical properties of realistic particle distributions
Stratospheric chemistry	Harvard, NOAA CSL, other atmospheric-chemistry laboratories and international ozone experts	Demonstrated heterogeneous-chemistry capability
Toxicology/environment	Independent toxicology laboratories, environmental-health researchers and risk-assessment contractors	Must be organizationally separate from teams advocating the candidate

C3 activity	Plausible recipients	Selection rule
Particle-generation engineering	SwRI, Creare, Battelle or other qualified aerosol-generation engineering groups	Demonstrate controlled size and de-agglomeration before flight
Small field tests	Selected release-system contractor + NASA/NOAA/NCAR measurement stack as appropriate	Only candidates passing lab gates reach atmosphere
Closure follow-on	Same measurement stack plus independent external review	Requires evidence of material advantage and explicit board approval

Natural-analogue and shared-infrastructure recipients

Activity	Plausible recipients
Volcanic/wildfire rapid response	NOAA CSL, NASA airborne science, NCAR/UCAR, university and international volcanic-aerosol groups
Shared high-altitude instrument suite	NOAA/NASA instrument teams plus commercial OEMs and university laboratories
Mobile/airborne lidar	NASA Langley, NOAA/university lidar groups, qualified commercial lidar integrators
Balloon observations	NOAA stratospheric balloon teams, Aerostar and international meteorological partners
Open data and campaign archive	UCAR EOL, NOAA or NASA data systems; independent reproducibility support
Downstream/ecological studies	Rutgers, atmospheric/environmental universities, agriculture/ecology groups and selected Global South partners

Independence and conflict rules

The most important procurement rule is separation of roles. The organization designing or operating a release should not control the environmental approval, data-quality audit, public-engagement record, or final determination that a larger experiment should proceed.

- At least two qualified bidders should normally be sought for commercial engineering work.
- A scientific panel member must recuse from an award benefiting that member's institution or close collaborators.
- Major model predictions should be challenged by at least one separate team that does not see withheld observations.

- Environmental and toxicology contractors should report to independent oversight, not to the candidate-material developer.
- Aircraft operators retain flight-safety and airworthiness authority regardless of program schedule pressure.
- The Managing Organization should retain final payment until required open data, calibration, code, and documentation are accepted.

Project Management

Prompt to AI: In theory, a Project Manager (PM) would need to issue a Purchase Order (PO). The PO would specify the recipient, amount of money, and deliverable. The PM would be employed by a “Managing Organization” (MO) that receives money from the Funding Source (FS). Who is the MO? Is the MO a foundation (e.g. gates foundation), lab, university, non-profit, or company? What do you suggest and why? The MO would need to be trusted by the FS since the MO would be spending their money.

Response From AI:

The **Managing Organization should be an established, independent nonprofit research institute**, rather than a foundation, university, government laboratory, or conventional company.


My first choice would be:


RTI International as the Managing Organization, with UCAR/NCAR serving as the principal atmospheric-science integrator.

I would approach RTI first, although I do not know whether RTI would be willing to accept such a politically sensitive assignment.

RTI is an independent nonprofit research institute that carries out multidisciplinary projects for governments, foundations, universities, businesses, and other organizations. That makes it structurally suited to receiving restricted funding, employing a program-management team, issuing subcontracts, auditing costs, and coordinating multiple scientific institutions.

Proposed structure





RTI would **manage the money and contracts**, but it should not unilaterally determine the science or approve atmospheric releases.

Why RTI rather than the other choices?

Possible MO	Assessment
Independent nonprofit research institute	Best overall choice
UCAR	Excellent alternative; strongest atmospheric-science connection
Battelle	Excellent for very large technical programs, but less neutral scientifically
Foundation	Better as the Funding Source than as day-to-day MO
University	Good research recipient, but not ideal as sole \$1 billion manager
National laboratory	Good technical recipient, but generally not an independent program fiduciary
For-profit company	Strong execution capability, but weaker public-interest credibility
Newly created nonprofit	Potentially appropriate later, but too inexperienced initially

1. Foundation

A foundation such as the Gates Foundation could be the **Funding Source**, but I would not make it the operational MO unless it already had a substantial engineering and field-research organization.

A foundation should normally:

- Establish objectives and restrictions.
- Approve the five-year budget.
- Appoint some governing-board members.
- Receive financial and scientific reports.
- Decide whether to release the next funding tranche.

It should not ordinarily be choosing aircraft nozzles, negotiating university subawards, checking instrument calibration, or managing flight schedules. Those functions require a specialized operating organization.

2. University

A major university could legally manage the program, but I would not recommend a single university as the MO.

Potential problems include:

- The university's own professors would probably compete for research money.
- Decisions could become entangled in faculty and departmental politics.
- University procurement can be slow and designed primarily for conventional academic grants.
- The institution could be perceived as both judge and beneficiary.
- A controversial field release could expose the entire university to reputational risk.

Universities should receive competitively selected **scientific subawards**, not control the entire program.

3. Government or national laboratory

NOAA, NASA, NCAR, and national laboratories should be major performers. However, placing the whole privately financed program inside a government laboratory could make contracting, international participation, and philanthropic control more difficult.

A laboratory should answer questions such as:

What instrument should be flown, and what does the measurement mean?

The MO should answer:

Was the procurement competitive, was the money spent properly, were the milestones met, and can the next phase proceed?

Those are different functions.

4. For-profit company

A large engineering company might manage the program efficiently, but an ordinary commercial prime contractor would create public concerns about:

- Profit from expanding the program.

- Incentives to proceed to larger experiments.
- Proprietary intellectual property.
- Control of data.
- Potential future participation in deployment.

Commercial companies should build hardware and provide services under POs, but should not control the overall research agenda.

Why not make UCAR the MO?

UCAR would be a credible alternative. It is a nonprofit consortium of 129 North American universities focused on Earth-system science, manages NSF NCAR, and already has systems for procurement, grants, financial oversight, scientific projects, and outside funding.

UCAR has also demonstrated that it can manage work at substantial scale. An earlier five-year agreement to manage NCAR was valued at up to \$630 million, and NSF renewed UCAR's NCAR-management agreement in 2023.

Nevertheless, I would probably make UCAR the **Scientific Integrator**, rather than the financial MO, because:

1. UCAR and NCAR would likely conduct a significant portion of the modeling and atmospheric research.
2. UCAR represents universities whose investigators could compete for funding.
3. Separating financial control from scientific performance would reduce apparent conflicts of interest.
4. RTI could provide a somewhat more neutral interface among atmospheric scientists, engineers, environmental groups, governments, and the Funding Source.

A reasonable alternative would be to make **UCAR the MO** while appointing a completely independent procurement, safety, and environmental-review organization. That might be preferable if the Funding Source places the highest weight on acceptance by the atmospheric-science community.

Why not Battelle?

Battelle has the strongest record of managing extremely large technical research institutions. It currently states that it manages eight U.S. national laboratories and oversees more than 40,000 laboratory employees. It also operates the NSF-funded National Ecological Observatory Network. Battelle is a 501(c)(3) nonprofit rather than an ordinary shareholder-owned contractor.

Battelle could be the MO, particularly if the program became dominated by:

- Aircraft-system development.
- Specialized instrumentation.
- Large test facilities.
- Engineering safety.
- National-laboratory participation.
- Complicated international field logistics.

However, I would rank Battelle behind RTI and UCAR for the initial five-year research program because the program needs unusually strong perceptions of openness, environmental independence, and scientific neutrality—not merely large-project execution.

The strongest arrangement

I recommend a **two-organization model**:

RTI: fiduciary and program manager

RTI would:

- Receive funds.
- Employ the PMs.
- Maintain the approved budget.
- Issue POs and subawards.
- Conduct cost and responsibility reviews.
- Monitor milestones.
- Maintain insurance and risk controls.
- Commission audits.
- Report to the Funding Source.
- Suspend payments when deliverables are not accepted.

UCAR: atmospheric-science integrator

UCAR would:

- Coordinate the scientific plan.
- Organize modeling and measurement teams.
- Develop experiment requirements.
- Coordinate NCAR and university participation.
- Manage scientific data standards.
- Organize peer review.
- Help determine whether observations answer the stated questions.

RTI should be able to reject an invoice or procurement action. UCAR should be able to say that an experiment is scientifically inadequate. An independent safety board should be able to prevent a release.

How the Funding Source should protect its money

Trust should not rest entirely on the reputation of the MO. It should be built into the funding agreement.

Funding in tranches

The FS should commit up to \$1 billion but not transfer the entire amount immediately.

An illustrative schedule might be:

Funding stage	Amount
Initial organization and design period	\$25M
Year 1 authorization	\$100M
Year 2 authorization	\$175M
Year 3 authorization	\$225M
Year 4 authorization	\$250M
Year 5 authorization	\$225M
Maximum commitment	\$1.000B

Each tranche would depend on:

- An approved annual operating plan.
- Completion of previous deliverables.
- An independent audit.
- Safety and environmental approval.
- Publication of required data.
- Updated risk analysis.

Ring-fenced program

The MO should establish a legally and financially identifiable program account. Money should not disappear into the organization’s general research budget.

The agreement should require:

- Separate program accounting.
- Quarterly reports by PO and recipient.
- Annual independent financial audits.
- Disclosure of related-party transactions.
- Competitive procurement above specified thresholds.
- Written justification for sole-source awards.
- Return of unused funds.
- Funding Source audit rights.
- Clear rules governing interest earned on unspent funds.

Limits on self-awards

The MO should not be allowed to steer a large share of the scientific work to itself.

For example:

- MO program administration: approximately **5%–7%** of total funding.
- Direct research performed by the MO: no more than an additional **5%–10%** without FS approval.
- Remaining funding: competitively awarded to laboratories, universities, aircraft operators, and specialized companies.

At \$1 billion, a 6% program-management allowance would provide **\$60 million over five years**, or approximately \$12 million annually. That should be sufficient to operate a substantial PM, procurement, finance, legal, safety, data, and reporting organization without consuming a large portion of the research budget.

Governance that I would require

The program should have four distinct authorities:

Authority	Principal responsibility
Funding Source	Sets purpose, funding ceiling, and major restrictions
Managing Organization	Controls money, personnel, contracts, and execution
Scientific Steering Committee	Determines scientific merit and experiment design
Independent Safety and Governance Board	Approves, modifies, postpones, or rejects outdoor releases

The board should include expertise in:

- Atmospheric chemistry and aerosol physics.

- Aircraft and flight safety.
- Environmental assessment.
- Public health.
- International law and governance.
- Research ethics.
- Financial and procurement oversight.
- Representatives from potentially affected countries and communities.

Neither the PM nor the scientific principal investigator should be able to authorize a release alone.

Final selection

My ranking would be:

1. **RTI as MO, UCAR as scientific integrator — recommended.**
2. **UCAR as MO, with independent fiduciary and environmental oversight.**
3. **Battelle as MO, particularly for a more engineering-intensive program.**
4. **A special-purpose nonprofit managed initially by RTI or UCAR.**
5. A university, foundation, government laboratory, or for-profit company acting alone.

A new “SAI Research Foundation” could eventually become the MO, but it should not receive \$1 billion at inception. A safer approach would be to operate it for two or three years as a ring-fenced program inside RTI. After demonstrating sound procurement, audits, staffing, safety review, and public reporting, the FS could decide whether to spin it out as an independent nonprofit.

The key organizational principle is:

The MO should be trusted to manage money, but it should not be financially rewarded for concluding that progressively larger atmospheric experiments are necessary.

Funding Source Requirements

Prompt to AI: What requirements might be demanded by the Funding Source (FS) to maximize climate benefit per dollar spent (i.e. reduce risk of wasted time and money)? Might they require all produced materials and measurement data be placed on the internet and marked open-source (i.e. available to others for free), for maximum benefit?

Response From AI:

Yes — “open by default” should be a central funding requirement

For a \$1 billion SAI research program, the Funding Source should require that the results be treated as a **global public research asset**, not as proprietary property of the Managing Organization or individual contractors.

However, the agreement should not merely say “everything must be open-source.” Different outputs need different legal treatment:

Output	Recommended requirement
Papers, reports, presentations and protocols	Publicly posted under CC BY 4.0
Measurement data and metadata	Deposited in a permanent repository under CC0 or equivalent public-use terms
Models, analysis code and software	Source code released under an OSI-approved license , such as Apache-2.0, BSD-3-Clause or MIT
Hardware designs	Editable CAD files, drawings, bill of materials and test procedures under an open-hardware license
Laboratory and flight procedures	Publicly posted, subject to narrowly defined safety restrictions
Patents and inventions	Normally defensive publication or worldwide, royalty-free, nonexclusive licensing

CC BY permits reuse and modification, including commercial reuse, with attribution. OSI-approved software licenses permit software to be used, modified and redistributed. NASA’s current scientific-information policy similarly covers publications, data and software, while the Gates Foundation’s policy requires openly licensed preprints and underlying data for funded research.

The more important requirement: value of information

Research should not be judged principally by papers produced, flights completed or money spent. Before authorizing each major PO, the MO should answer:

1. **What important uncertainty will this work reduce?**
2. **What future decision will depend on the result?**
3. **What result would cause the program to stop, change direction or proceed?**
4. **Could the question be answered more cheaply through existing data, modeling, laboratory work or a natural volcanic plume?**
5. **Why is this contractor and price the best available approach?**

For example, a C2 direct-sulfate experiment should not merely promise “measurements of aerosol.” It should specify a generation architecture, mass-flow range, particle-size target after defined plume ages, uncertainty threshold, energy and aircraft-time measurements, and the future down-selection decision the result will inform. The required final outputs should include forcing/kg-S and forcing/\$, not merely particle size.

Determine whether an aircraft release system can produce particles with a median radius between specified limits four hours after release, with uncertainty below a stated threshold. The result will determine whether direct aerosol injection remains in the research portfolio.

This converts a scientific activity into a **decision-enabling deliverable**.

Requirements I would put in the FS–MO agreement

1. Stage-gated funding

The FS should commit up to \$1 billion but release money in annual or semiannual tranches. Large field experiments should pass separate gates:

- Scientific-design review.
- Modeling and laboratory review.
- Environmental and governance review.
- Engineering-readiness review.
- Flight-readiness review.
- Post-experiment data and audit review.

Failure at one gate should terminate or redesign the activity without creating an expectation that the remaining money will be spent.

2. Mandatory stop criteria

Every experiment should state in advance what findings would cause it to stop. Examples include:

- Instruments cannot meet their required detection limits.
- The plume cannot be located with adequate probability.
- Release-system output cannot be characterized.
- A cheaper natural analogue becomes available.
- Environmental analysis identifies an unacceptable risk.
- Earlier measurements make the experiment scientifically unnecessary.

A contractor should not be rewarded simply for completing additional flights after the scientific value has disappeared.

3. Preregistration

Before data collection, the project should publicly register:

- Hypotheses.
- Primary and secondary measurements.
- Experimental conditions.
- Statistical and modeling methods.
- Exclusion criteria.
- Success and failure thresholds.
- Planned decision resulting from each possible outcome.

Exploratory analysis would remain permissible, but it must be labeled exploratory. Preregistration creates a time-stamped distinction between planned tests and conclusions developed after examining the data.

4. Open-data deadlines

I would require:

Material	Publication deadline
Experiment description and release record	Within 7 days
Preliminary flight and instrument metadata	Within 30 days
Level-0 raw instrument data	Within 60–90 days
Calibrated and quality-controlled data	Within 6 months
Final analysis dataset and code	Within 12 months
Preprint or complete technical report	Within 12 months

A short embargo may be reasonable for calibration and quality control. A long embargo allowing investigators to reserve the data for several years would substantially reduce the public benefit.

NSF already requires funded proposals to include data-management and sharing plans, and NASA requires planning for sharing scientific information. The proposed SAI program could impose stronger and more specific deadlines.

5. Complete reproducibility package

Posting a spreadsheet of final results is not enough. Each experiment should publish:

- Raw and calibrated measurements.
- Calibration records.
- Instrument configuration and firmware versions.
- Sampling times and locations.
- Meteorological inputs.
- Data-quality flags.
- Processing algorithms.
- Model source code.
- Input files and configuration files.
- Analysis notebooks.
- Uncertainty calculations.
- Explanations of excluded observations.
- A machine-readable data dictionary.

An independent researcher should be able to reproduce every major figure and numerical conclusion.

6. Publication of failures and negative results

The MO must publish results even when:

- The plume was missed.
- Instruments failed.
- No expected particle growth was detected.
- Models performed poorly.
- The release system did not produce the intended size distribution.
- The experiment was terminated early.

These results may prevent another organization from wasting millions of dollars repeating the same mistake.

7. Open foreground intellectual property

The contract should distinguish:

- **Background IP:** technology a contractor already owned before the project.
- **Foreground IP:** designs, software, methods and inventions created using FS money.

Contractors may retain their background IP. Foreground IP should ordinarily become openly available.

A commercial vendor should not be allowed to accept \$5 million to design a nozzle and then prevent other researchers from manufacturing or improving that nozzle. The PO should require editable drawings, source files, manufacturing tolerances, test data and a license sufficient for others to use and modify the design.

8. Narrow exceptions process

Not literally everything should automatically be posted without review. Temporary restrictions might be justified for:

- Personal information.
- Aircraft-security information.
- Export-controlled technology.
- Proprietary background IP.
- Exact details that create a credible safety or unauthorized-deployment risk.
- Data whose release would violate an agreement with a government aircraft operator.

Each exception should require:

1. A written justification.
2. Approval by an independent information-review officer.
3. Publication of nonsensitive metadata and a summary.
4. The smallest possible redaction.
5. A date for reconsideration.

“Commercially sensitive” or “controversial” should not be sufficient reasons for secrecy.

Procurement controls

Every substantial PO should state:

- Recipient and subcontractors.
- Maximum authorized amount.

- Detailed deliverables.
- Acceptance tests.
- Schedule and milestones.
- Required data formats.
- Open licenses.
- Background and foreground IP.
- Data-release deadlines.
- Conflict-of-interest disclosures.
- Audit rights.
- Termination rights.
- Payment holdbacks.

I would withhold approximately **10%–15% of each PO** until the final data, documentation, source files and financial records have been accepted. A contractor that completes the flight but fails to deliver usable data has not completed the PO.

Competitive bidding should normally be required above perhaps **\$250,000**. Sole-source awards should require a public justification explaining why another organization cannot reasonably provide the deliverable.

Independent validation

Approximately **8%–12% of the research budget** should be reserved for independent replication, audits and model challenges.

The original team should not be the only group evaluating its own performance. For major experiments:

- One team predicts plume behavior before seeing the data.
- Another team reconstructs it afterward.
- An independent group audits calibration and uncertainty.
- Important findings are reproduced using separate code.
- Selected measurements are duplicated using different instruments.

The National Academies has emphasized that solar-geoengineering research should combine modeling, observations, laboratory work and potentially carefully governed outdoor experiments, rather than allowing field experiments to operate independently of the broader scientific program.

Program-wide financial transparency

The MO should maintain a public online dashboard showing:

- Every PO above a stated threshold.
- Recipient.
- Original and revised amount.
- Deliverables.
- Milestone status.
- Payments made.
- Administrative and indirect costs.
- Sole-source justification.
- Data and reports produced.
- Links to the open repository.

Sensitive employee information and protected pricing details could be omitted, but the public should be able to determine where the money went and what was received.

The FS should receive more detailed quarterly reports, including invoices and commitments by PO.

Suggested performance scorecard

The MO's success should be evaluated using metrics such as:

Metric	Illustrative target
Calibrated datasets released on time	At least 95%
Major results independently reproducible	At least 90%
POs completed within approved budget	At least 85%
Field experiments with preregistered decision criteria	100%
Sole-source commercial spending	Below 20%, absent justification
Experiments reusing common instruments and infrastructure	At least 75%
Reports including negative and inconclusive results	100%
Major POs with measurable acceptance tests	100%
Unresolved audit findings older than six months	Zero
Major sulfur experiments reporting forcing/kg-S, forcing/aircraft-hour and forcing/\$	100%
C2 architectures completing ground/plume gates before major outdoor release	100%

The central contractual principle

The FS should tell the MO:

The purpose of the program is not to spend \$1 billion or to demonstrate that SAI should be deployed. The purpose is to produce reliable, reusable knowledge that helps the world determine whether SAI could reduce climate risk, under what conditions, and with what limitations.

Open publication should therefore be mandatory, but it should be combined with **stage gates, preregistration, independent replication, transparent procurement and explicit stop rules**. Openness alone prevents private capture of the results; the other requirements prevent the program from efficiently publishing large quantities of low-value research.

Top SAI Papers

Prompt to AI: Summarize the top SAI papers, including recent work most relevant to the revised C1/C2/C3 program.

[Response From AI:](#)

Top 25 Stratospheric Aerosol Injection papers

“Top” is necessarily subjective. I ranked these **roughly by historical influence, scientific importance, and usefulness for designing an SAI research program**, rather than by citation count alone. The list emphasizes sulfur-based SAI, including SO₂ and H₂SO₄ injection, climate effects, ozone, field experiments, costs, and termination risk.

Foundations and early risk identification

1. Crutzen, 2006 — “Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?”

The modern starting point for SAI research. Crutzen argued that persistent failure to reduce greenhouse-gas emissions justified serious study of stratospheric sulfur injection as a temporary emergency measure. The paper emphasized that SAI would not solve the underlying CO₂ problem and could damage ozone, but it made the subject scientifically discussable again.

2. Rasch et al., 2008 — “An Overview of Geoengineering of Climate Using Stratospheric Sulphate Aerosols”

This early comprehensive review explains volcanic analogues, sulfate chemistry, particle microphysics, radiative forcing, climatic effects, and possible delivery systems. It concludes that sulfate could offset much global-mean warming, but not reproduce the original climate region by region. It remains an excellent introduction to the entire physical chain.

3. Robock, Oman and Stenchikov, 2008 — “Regional Climate Responses to Geoengineering with Tropical and Arctic SO₂ Injections”

One of the first influential global-model studies of explicit SO₂ injection. It found that both tropical and Arctic injections produced global effects and could weaken Asian and African summer monsoons. The paper helped shift attention from “Can SAI cool Earth?” to “Where would the residual climate damages occur?”

4. Tilmes, Müller and Salawitch, 2008 — “The Sensitivity of Polar Ozone Depletion to Proposed Geoengineering Schemes”

This paper showed that sulfate aerosol provides surfaces on which ozone-destroying chemistry can occur. It projected increased Arctic ozone loss during cold winters and delayed recovery of the Antarctic ozone hole under large sulfate interventions. It established ozone as a primary constraint on sulfur-based SAI.

5. Robock et al., 2009 — “Benefits, Risks, and Costs of Stratospheric Geoengineering”

A concise synthesis of potential benefits and adverse effects. The paper covers cooling, sea ice, glaciers, precipitation, ozone, diffuse light, astronomy, solar electricity, ocean acidification, termination, and governance. Although some quantitative assumptions are now dated, its checklist of risks remains influential.

Aerosol formation, chemistry, and efficiency

6. Heckendorn et al., 2009 — “The Impact of Geoengineering Aerosols on Stratospheric Temperature and Ozone”

This study integrated sulfate aerosol microphysics with atmospheric chemistry and radiation. It showed that particle growth, stratospheric heating, and ozone response are coupled, so one cannot estimate SAI effects from injected sulfur mass alone. Particle size and aerosol surface area are critical state variables.

7. Pierce et al., 2010 — “Efficient Formation of Stratospheric Aerosol for Climate Engineering by Emission of Condensible Vapor from Aircraft”

Especially relevant to the proposed H₂SO₄ experiments. The authors modeled releasing H₂SO₄ vapor from an aircraft so that nucleation and coagulation in the expanding plume create accumulation-mode particles. They concluded this could provide better particle-size control and greater radiative efficiency than releasing SO₂, although the near-field plume assumptions require experimental validation.

8. Kravitz et al., 2011 — “The Geoengineering Model Intercomparison Project (GeoMIP)”

GeoMIP established standardized experiments that different climate-model groups could run and compare. Before GeoMIP, differences between studies often reflected both differing models and differing scenarios. The framework made it possible to identify results that were robust across models and those that were not.

9. English, Toon and Mills, 2012 — “Microphysical Simulations of Sulfur Burdens from Stratospheric Sulfur Geoengineering”

This paper compared SO₂ gas, gaseous H₂SO₄, and particle injection in a three-dimensional model with sectional aerosol microphysics. It showed that increasing sulfur injection can produce disproportionately larger particles, which scatter less efficiently per unit mass and fall faster. Injection material, particle size, altitude, and spatial distribution all strongly influence performance.

10. Schmidt et al., 2012 — “Solar Irradiance Reduction to Counteract Radiative Forcing from a Quadrupling of CO₂”

The first major analysis of GeoMIP’s idealized G1 experiment. Reducing solar irradiance could largely restore global-mean temperature after quadrupling CO₂, but it could not restore all regional temperature and precipitation patterns. The paper clearly demonstrates why equal and opposite global radiative forcings do not create an unchanged climate.

Hydrology, termination, and field experiments

11. Tilmes et al., 2013 — “The Hydrological Impact of Geoengineering in the Geoengineering Model Intercomparison Project”

Using 12 Earth-system models, this paper found that solar reduction generally suppresses the global hydrological cycle relative to a climate with the same global temperature but lower CO₂. Precipitation and evaporation therefore cannot be simultaneously restored by selecting a global cooling target. It is a central paper on SAI’s precipitation tradeoffs.

12. Jones et al., 2013 — “The Impact of Abrupt Suspension of Solar Radiation Management—Termination Effect—in GeoMIP Experiment G2”

Eleven models simulated stopping solar geoengineering after decades of offsetting increasing CO₂. All produced rapid global warming, increased precipitation, and declining sea ice after termination. This paper established the physical basis of “termination shock” and the importance of preventing abrupt interruption.

13. Dykema et al., 2014 — “Stratospheric Controlled Perturbation Experiment: A Small-Scale Experiment to Improve Understanding of the Risks of Solar Geoengineering”

This is the most important conceptual paper on an actual small stratospheric perturbation experiment. It outlines a balloon platform, plume creation, in-situ sampling, aerosol and ozone chemistry, instrumentation, and external oversight. Its central argument is that laboratory work and global models cannot completely replace small, controlled atmospheric measurements.

14. Niemeier and Timmreck, 2015 — “What Is the Limit of Climate Engineering by Stratospheric Injection of SO₂?”

The authors modeled injection rates up to 100 teragrams of sulfur per year. Radiative forcing increased much less than proportionally because higher aerosol loading produced larger particles, faster sedimentation, and altered circulation. This established that SO₂ injection has strong diminishing returns and cannot be scaled linearly.

15. Keith et al., 2016 — “Stratospheric Solar Geoengineering Without Ozone Loss”

This paper proposed alkaline solid particles, particularly calcite, that might neutralize acidic chlorine and nitrogen compounds and potentially reduce ozone loss. It broadened SAI research beyond sulfate, but the favorable result depends on uncertain particle-surface chemistry, aggregation, aging, toxicity, and manufacturability. It should be read as a research hypothesis, not proof that calcite is safe.

Injection design and climate control

16. Visioni et al., 2017 — “Sulfate Geoengineering: A Review of the Factors Controlling the Needed Injection of Sulfur Dioxide”

A valuable technical review of how injection rate, altitude, latitude, season, transport, microphysics, particle size, and stratospheric heating determine the quantity of SO₂ required. It explains why estimates from different models vary substantially. For planning a research program, it provides a good map of the largest physical uncertainties.

17. MacMartin et al., 2017 — “The Climate Response to Stratospheric Aerosol Geoengineering Can Be Tailored Using Multiple Injection Locations”

The authors showed that injecting at several latitudes provides more control than injecting only at the equator. Different injection locations influence global mean temperature, interhemispheric temperature balance, and the equator-to-pole temperature gradient differently. This paper introduced the idea of SAI as a multivariable control problem rather than a single global thermostat.

18. Kravitz et al., 2017 — “First Simulations of Designing Stratospheric Sulfate Aerosol Geoengineering to Meet Multiple Simultaneous Climate Objectives”

This study used feedback control to adjust injections at four latitudes while pursuing three temperature objectives. It demonstrated that a climate model could approximately maintain global mean temperature and two broad temperature gradients despite rising greenhouse gases. It also showed that meeting temperature objectives does not automatically eliminate precipitation, ozone, circulation, or regional impacts.

19. Mills et al., 2017 — “Radiative and Chemical Response to Interactive Stratospheric Sulfate Aerosols in Fully Coupled CESM1(WACCM)”

This paper describes and evaluates a model that interactively calculates sulfate formation, particle evolution, radiation, atmospheric dynamics, and chemistry. Validation against volcanic aerosol observations made it an important foundation for subsequent SAI simulations. It helped replace prescribed aerosol layers with more physically complete sulfur-injection modeling.

20. Tilmes et al., 2018 — “CESM1(WACCM) Stratospheric Aerosol Geoengineering Large Ensemble Project”

The GLENS project produced a large ensemble of century-scale simulations using four injection locations and feedback control. Multiple realizations allow researchers to distinguish forced SAI responses from ordinary climate variability. GLENS became one of the principal datasets for studying regional climate, extremes, circulation, agriculture, water, and ecosystems under SAI.

Deployment, societal impacts, and newer model comparisons

21. Smith and Wagner, 2018 — “Stratospheric Aerosol Injection Tactics and Costs in the First 15 Years of Deployment”

This engineering-economic study examined aircraft requirements, payload, altitude, fleet size, flights, and operating costs. It estimated average direct deployment costs of roughly \$2.25 billion per year over an illustrative first 15 years, while requiring thousands and eventually tens of thousands of annual flights. Its key message is that SAI might be financially inexpensive relative to climate damages but operationally large and difficult to conceal.

22. Proctor et al., 2018 — “Estimating Global Agricultural Effects of Geoengineering Using Volcanic Eruptions”

Using volcanic eruptions as natural experiments, the paper estimated how cooling and altered sunlight affect major crops. The modeled benefits of lower temperature were approximately offset by yield losses associated with reduced direct sunlight. Volcanic eruptions are imperfect analogues for sustained SAI, but this paper demonstrates why diffuse and direct radiation must be included in impact studies.

23. Vattioni et al., 2019 — “Exploring Accumulation-Mode H₂SO₄ Versus SO₂ Stratospheric Sulfate Geoengineering”

This directly compares the two pathways discussed in your proposed experiments. Spatially distributed injection of accumulation-mode H₂SO₄ generally produced greater radiative forcing per unit sulfur than SO₂, but concentrated point or pulse injections caused enhanced coagulation and could lose much of that advantage. The paper shows that release geometry is almost as important as injected chemical form.

24. Vioni et al., 2021 — “Identifying the Sources of Uncertainty in Climate Model Simulations ... with the G6sulfur and G6solar GeoMIP Simulations”

This study separates uncertainty caused by the underlying climate models from uncertainty caused by how models represent sulfate aerosol and its transport. Models broadly agree that SAI could reduce temperature increases, but disagree more on aerosol distribution, required sulfur, stratospheric effects, and precipitation. It cautions against presenting one model’s injection requirement as a reliable engineering specification.

25. Weisenstein et al., 2022 — “An Interactive Stratospheric Aerosol Model Intercomparison ... by Injection of SO₂ or Accumulation-Mode Sulfuric Acid Aerosols”

Three interactive aerosol models ran coordinated SO₂ and accumulation-mode H₂SO₄ scenarios. All found greater radiative efficacy for direct accumulation-mode aerosol, but they differed in particle distributions, transport, and quantitative performance. Crucially, the study acknowledges that it prescribes the particles after near-field plume processing rather than modeling the first seconds and minutes—precisely the uncertainty that controlled C2 field experiments could address.

26. Yu et al., 2024 — “On Nucleation Pathways and Particle Size Distribution Evolutions in Stratospheric Aircraft Exhaust Plumes with H₂SO₄ Enhancement”

This study models the first moments of H₂SO₄-enhanced aircraft exhaust. Under modeled conditions that produce desirable roughly 200–300 nm diameter particles, nucleation occurs in the nascent hot, dry plume and subsequent coagulation governs evolution toward a relatively narrow size distribution. It gives the revised C2 program a concrete near-field mechanism and makes flow rate, exhaust conditions, dilution and plume age explicit experimental variables.

27. Franke et al., 2026 — “Trade-Offs between Sulfuric Acid Aerosol Precursors for Stratospheric Aerosol Geoengineering”

This engineering-economic study compares sulfur precursors, liquid handling, vaporization, safety, injection rate, material cost and aircraft lofting cost for a purpose-built 20-km aircraft. For H₂SO₄ it analyzes both a high 4.8 kg/s material rate and a 0.1 kg/s case intended to produce smaller particles, finding a strong particle-size-versus-aircraft-cost tradeoff and noting that larger engines could shift the optimum. It directly motivates C2A/C2B/C2C down-selection and mandatory forcing-per-dollar metrics.

[Recommended Reading](#)

The twelve papers I would have experiment-design teams read first are: *Pierce 2010; English 2012; Dykema 2014; Niemeier and Timmreck 2015; Vioni 2017; Mills 2017; Smith and Wagner 2018; Vattioni 2019; Vioni 2021; Weisenstein 2022; Yu 2024; and Franke 2026.*

Taken together, they define important experimental questions: can a real aircraft-generation system produce a known and reproducible sulfate particle-size distribution at economically useful mass flow, and can models predict its evolution from the first milliseconds through global dispersion?

Top Researchers

Prompt to AI: Who are the top researchers in this field? What are the top organizations working in this field?

Response From AI:

There is no authoritative ranking. My assessment below emphasizes **influence on SAI science, current activity, leadership of major datasets or programs, and relevance to the field experiments you are considering.**

Leading SAI researchers

Central technical leaders

Researcher	Current base	Principal contribution
David Keith	University of Chicago	Probably the field’s best-known interdisciplinary leader. Works across SAI physics, engineering, risk analysis, policy and governance; founded UChicago’s Climate Systems Engineering initiative.
Simone Tilmes	NSF NCAR	One of the strongest chemistry–climate modelers in SAI. Led GLENS-related work and recent efforts to establish scientific criteria for assessing SAI scenarios.
Ben Kravitz	Indiana University	Co-founder and coordinator of the Geoengineering Model Intercomparison Project, or GeoMIP; especially important for model comparison and climate-control design.
Douglas MacMartin	Cornell University	Leading expert on SAI system design, feedback control, injection-location optimization, uncertainty and research strategy.
Daniele Visioni	Cornell University	Major current researcher on sulfate-aerosol microphysics, injection strategies, climate effects and multi-model uncertainty.
Karen Rosenlof	NOAA Chemical Sciences Laboratory	Particularly important for stratospheric observations, aerosol budgets, intervention detection, field measurements and connecting observations to models.
Frank Keutsch	Harvard University	Atmospheric chemist and instrument developer; led the proposed SCoPEX field experiment. SCoPEX was discontinued in 2024, but Keutsch remains highly relevant to plume chemistry and small-field-experiment design.
Alan Robock	Rutgers University	Foundational researcher on volcanic analogues, regional climate effects, monsoons, risks and termination; also an influential critical voice within the field.

Researcher	Current base	Principal contribution
Claudia Timmreck	Max Planck Institute for Meteorology	Leads work on stratospheric forcing and climate, especially volcanic aerosols, aerosol transport and climate-model evaluation.
Jim Haywood	University of Exeter and UK Met Office community	Leading aerosol–radiation and climate-modeling expert; prominent in the expanding UK SRM research program.

Important specialists

Researcher	Particularly valuable for
Jadwiga Richter, NSF NCAR	Stratospheric dynamics, climate control and leadership of the ARISE-SAI simulation program.
Michael Mills, NSF NCAR	Interactive sulfate microphysics, chemistry, ozone and WACCM modeling.
Lili Xia, Rutgers	Agriculture, vegetation, food-system and regional-impact analysis under SAI.
Govindasamy Bala, Indian Institute of Science	Hydrological cycles, Indian monsoons, regional impacts and Global South scientific leadership.
Peter Irvine, University of Chicago	Earth-system impacts, risk comparisons, tipping elements and decision-oriented analysis.
Wake Smith, Yale	Aircraft, logistics, deployment architecture, costs and institutional organization.
Gernot Wagner, Columbia	Economics, risk, public policy and SAI deployment costs.
Jesse Reynolds, Degrees Initiative	International law, research governance and institutional arrangements.
Inés Camilloni, University of Buenos Aires/CONICET	South American hydroclimate impacts, equity and Global South participation in SRM decisions.
James Franke, University of Chicago	Sulfur-precursor engineering tradeoffs, aircraft lofting economics, feedstock and injection-rate cost comparisons
Fangqun Yu, University at Albany/SUNY	Near-field H ₂ SO ₄ nucleation, aircraft-exhaust plume microphysics and particle-size evolution

For the revised C2 engineering competition and C1 SO₂ pathway experiments, the most directly relevant initial group would probably include Rosenlof, Keutsch, Tilmes, Vioni, MacMartin, Mills, Richter, Keith, Yu and Franke. Wake Smith would be useful for aircraft and operational design, while Robock and Timmreck would provide strong independent scientific challenge.

Top organizations

1. NOAA Chemical Sciences Laboratory

NOAA CSL is probably the most important U.S. organization for the **actual measurement side** of SAI. Its strengths include:

- Stratospheric SO₂, aerosol, ozone, water-vapor and trace-gas measurements.
- High-altitude aircraft instrumentation.
- The SABRE stratospheric aircraft campaigns.
- Earth's Radiation Budget research.
- Aerosol detection, attribution and model validation.

For a controlled SO₂-release experiment, NOAA CSL would be difficult to replace.

2. NSF NCAR and UCAR

NCAR is probably the strongest single center for **integrated chemistry–climate modeling** of SAI. It operates CESM/WACCM and led major datasets including GLENS and ARISE-SAI. It also has field-campaign, aircraft, data-management and supercomputing infrastructure through the wider UCAR system.

3. Cornell Climate Engineering group

Cornell has one of the most concentrated academic teams devoted specifically to SAI. Its distinctive strength is integrating:

- Aerosol and climate modeling.
- Injection strategy.
- Engineering control.
- Scenario design.
- Uncertainty analysis.
- Regional and extreme-event impacts.

MacMartin and Vioni make Cornell particularly important for converting scientific questions into a structured experimental program.

4. University of Chicago Climate Systems Engineering initiative

UChicago is building a broad program around sunlight reflection, carbon removal and related climate interventions under David Keith. It is particularly promising for combining atmospheric science with

engineering, risk analysis, economics, governance and public policy. It is newer than NCAR or NOAA but may become one of the field's dominant centers.

5. Harvard Solar Geoengineering Research Program and Keutsch laboratory

Harvard has produced influential work on atmospheric chemistry, materials, governance and public policy. The proposed SCoPEX experiment gave Harvard substantial experience with the scientific and governance challenges surrounding outdoor research, even though the experiment was ultimately discontinued.

6. Geoengineering Model Intercomparison Project

GeoMIP is a research network rather than a single institution, but it is one of the field's most important scientific structures. It specifies common experiments for multiple climate models, allowing researchers to distinguish robust results from behavior specific to one model. It is essential for independent replication.

7. Max Planck Institute for Meteorology

MPI-M is especially strong in volcanic forcing, stratospheric aerosol evolution, circulation and Earth-system modeling. Claudia Timmreck's group would be a valuable independent European modeling and natural-analogue partner.

8. University of Exeter and the UK Met Office

Together, these organizations provide substantial expertise in aerosol–radiation interactions, SAI climate modeling, regional impacts and atmospheric observations. Exeter is also expanding interdisciplinary work on responsible climate intervention.

9. Rutgers University

Rutgers is particularly important for independent assessment of adverse effects, regional climate, agriculture, ecosystems and termination risk. Its team includes Alan Robock and Lili Xia and has long-standing involvement in GeoMIP.

10. Indian Institute of Science

IISc is a leading technical center for examining how SAI could affect India and its monsoon-dependent population. It would be an important partner for avoiding a research program dominated solely by U.S. and European perspectives.

Important funders and network builders

The **Simons Foundation** operates an international SRM collaboration focused on fundamental scientific gaps. The UK's **ARIA Exploring Climate Cooling program** is funding laboratory, modeling and carefully governed real-world work, while **NERC** funds complementary impact modeling. The **Degrees Initiative** is the leading organization building the capacity of Global South research teams to independently evaluate SRM.

Best consortium for the proposed \$1 billion program

I would build the scientific core around:

Function	Suggested lead
Stratospheric observations and instruments	NOAA CSL
Integrated global modeling	NSF NCAR/UCAR
Experiment and injection-strategy design	Cornell
Systems engineering and cross-disciplinary assessment	University of Chicago
Laboratory aerosol chemistry	Harvard/Keutsch group
Independent European modeling	Max Planck and Exeter/Met Office
Regional and agricultural impacts	Rutgers, IISc and Global South university teams
Multi-model replication	GeoMIP
Global participation and capacity building	Degrees Initiative

No single laboratory presently has all the capabilities required. The strongest program would deliberately distribute authority among **measurement teams, modelers, experiment designers, regional-impact researchers and independent critics**, rather than allowing one group to design, perform and evaluate its own experiments.

Scientific Advisory Panel

Prompt to AI: If a scientific panel was set up to oversee the research, who might you suggest be on that panel?

[Response From AI:](#)

Suggested scientific oversight panel

I would create a **13-member Scientific Advisory Panel**. It should contain leading SAI specialists, but no institution should hold more than two seats. At least one-third of members should be scientists who are not primarily identified with promoting SAI research.

Proposed membership

Role	Suggested person	Why
Chair	Susan Solomon, MIT	Outstanding independent authority on stratospheric chemistry, heterogeneous reactions, ozone depletion, and the interaction between atmospheric science and environmental policy. She would give the panel exceptional scientific credibility.
Vice-chair	Govindasamy Bala, Indian Institute of Science	Climate-modeling expert with extensive work on solar geoengineering, hydrology, monsoons, and regional effects. His appointment would ensure that countries highly dependent on monsoon rainfall are represented near the top of the panel.
Stratospheric observations	Karen Rosenlof, NOAA Chemical Sciences Laboratory	Strong choice for aircraft measurements, stratospheric transport, water vapor, aerosols, ozone, and connecting field observations with models. She has led NOAA's stratospheric research activities and participated in major ozone assessments.
Aerosol chemistry and field experimentation	Frank Keutsch, Harvard University	Atmospheric chemist and instrument developer with direct experience designing the proposed SCoPEX controlled-perturbation experiment. Particularly relevant to H ₂ SO ₄ particles, plume sampling, and chemical measurements.
Chemistry–climate modeling	Simone Tilmes, NSF NCAR	Leading expert on sulfate aerosols, ozone, atmospheric dynamics, and Earth-system consequences of SAI. She also helps lead CESM chemistry–climate modeling.
Injection strategy and control	Douglas MacMartin, Cornell University	Connects engineering control theory with climate dynamics. Particularly valuable for determining injection locations, measurement requirements, control objectives, and whether experiments answer decision-relevant questions.

Role	Suggested person	Why
Aerosol microphysics	Daniele Visioni, Cornell University	Specialist in stratospheric aerosols, chemistry, radiation, atmospheric dynamics, and climate-intervention strategies. Especially useful for comparing SO ₂ with directly injected accumulation-mode H ₂ SO ₄ .
Multi-model evaluation	Ben Kravitz, Indiana University	Co-founder and coordinator of GeoMIP and an expert in SAI climate simulations, uncertainty, control methods, and cross-model comparisons. He could make sure the program does not rely excessively on a single climate model.
Independent critical assessment	Alan Robock, Rutgers University	Foundational SAI and volcanic-climate researcher who has emphasized regional risks, precipitation changes, termination, and unanswered questions. He would be an important internal critic rather than an automatic supporter of larger experiments.
Volcanic analogues	Claudia Timmreck, Max Planck Institute for Meteorology	Leads research on stratospheric forcing and climate, with particular strength in volcanic aerosols, model evaluation, atmospheric transport, and natural analogues to SAI.
South American impacts	Inés Camilloni, University of Buenos Aires/CONICET	Climate and hydrological scientist studying SAI effects in South America and the La Plata Basin. She is currently a vice-chair of IPCC Working Group I, adding substantial international assessment experience.
Southern Hemisphere and ozone	Laura Revell, University of Canterbury	Atmospheric-chemistry and Earth-system-modeling expert working on aerosols, ozone, and stratospheric pollutants. She also brings an important New Zealand and Southern Hemisphere perspective.
Air quality and public health	Tracey Holloway, University of Wisconsin–Madison	Works at the intersection of atmospheric modeling, air quality, satellites, climate, and public health, and has led NASA's Health and Air Quality Applied Sciences Team. She would challenge the program to quantify surface exposure and health implications.

I do not know whether any of these individuals would accept an invitation.

Non-voting technical advisers

Several people should advise the panel without necessarily becoming permanent voting members.

David Keith should be a senior technical adviser. He has exceptionally broad knowledge spanning SAI science, engineering, public policy, and research-program design and currently leads the University of Chicago Climate Systems Engineering initiative. However, because his institution would likely seek substantial research funding, making him an adviser rather than chair would reduce concentration of influence.

Other rotating advisers should include:

- A propulsion/engine-integration expert independent of any selected C2 release-system vendor.
- A near-field aircraft-plume nucleation specialist, such as Fangqun Yu or an equivalent expert.
 - A NASA WB-57 or ER-2 aircraft-program representative.
 - A Federal Aviation Administration airspace and flight-safety specialist.
 - The principal scientist for each proposed experiment.
 - An independent statistician or uncertainty-quantification expert.
 - A research-data and open-source specialist.
 - An ecologist or agricultural scientist when downstream impacts are considered.
 - Representatives from countries over which a plume might travel.

The Funding Source and Managing Organization could each have **one non-voting observer**, but neither should control scientific conclusions.

Critical conflict-of-interest rule

Many of the best-qualified panelists belong to organizations that may apply for funding. Therefore:

1. The **chair and vice-chair should agree not to receive program research funding** during their terms.
2. A member must recuse from discussing or voting on a PO benefiting their institution, close collaborators, or current projects.
3. A recused member should not see confidential competing proposals.
4. At least **seven unconflicted members** should be required for a funding recommendation.
5. Awards above perhaps \$25 million should receive an additional external review.
6. The FS should publish recusals and conflict disclosures.

For example, MacMartin should not vote on a Cornell modeling award; Tilmes should not vote on an NCAR award; Rosenlof should not evaluate a NOAA instrument proposal; and Keutsch should not approve a Harvard field-experiment contract.

Do not give this panel sole release authority

The scientific panel should determine:

- Whether the question is important.
- Whether the proposed experiment can answer it.
- Whether the instrumentation is adequate.
- Whether the cost is justified.

- Whether the results would change a future decision.

A separate **Independent Safety, Ethics and Governance Board** should determine whether an outdoor release may legally and ethically proceed. That board should include international law, environmental health, research ethics, community engagement, aviation safety, and representatives from potentially affected regions.

An experiment should require approval from both bodies:

Scientific Panel: The experiment is worth doing.

Safety and Governance Board: The experiment is acceptable to conduct.

Suggested voting standard

Routine modeling and laboratory awards could require a simple majority. Controlled atmospheric releases should require:

- At least a **two-thirds majority** of the full scientific panel.
- Written findings explaining why laboratory work, natural analogues, or existing observations cannot answer the question.
- Confirmation that the expected scientific value justifies the cost.
- Publication of dissenting opinions.
- Separate approval from the safety and governance board.

This panel would contain strong SAI expertise, independent critics, field scientists, ozone specialists, regional-impact researchers, and Global South representation—without allowing a single university, laboratory, or intellectual faction to control the program.

Management Advisory Panel

Prompt to AI: If a panel was set up to oversee the agreement with the Managing Organization (MO) (e.g. RTI, UCHAR, Battelle) who might you suggest be on that panel?

Response From AI:

Recommended structure: a 9-member Agreement Oversight Board

This panel should be distinct from the **Scientific Advisory Panel**. Its job would be to protect the Funding Source’s money, enforce the MO agreement, and verify that the program is producing useful results.

I am treating “UCHAR” as **UCAR**.

The board should include:

- Seven independent members.
- Two representatives appointed by the Funding Source.
- No voting representatives from the selected MO.
- No voting representatives from organizations expecting major research awards.

The MO’s chief executive and the SAI program director should attend meetings as **non-voting participants**.

Illustrative membership

Seat	Suggested candidate	Reason for selection
Independent chair	France Córdova	Former NSF director, former NASA chief scientist, former university president and current Caltech trustee. She has experience overseeing large, multidisciplinary research enterprises without being closely identified with SAI.
Audit and financial accountability	Gene Dodaro	Served as U.S. Comptroller General and head of GAO for 15 years, through December 2025. He would bring unusually strong experience in financial controls, program evaluation, waste prevention and independent reporting.
Procurement and contracting	Lesley Field	Former deputy administrator for federal procurement policy at OMB and a National Academy of Public Administration fellow. She has direct experience with acquisition policy, competition, contract structures and procurement workforces.

Seat	Suggested candidate	Reason for selection
Large technical-program execution	Christine Fox	Former acting deputy secretary of Defense and a senior fellow at Johns Hopkins Applied Physics Laboratory. She has experience evaluating and overseeing large, technically complex, multi-institution programs.
Philanthropic effectiveness and evaluation	Fay Twersky	Led the Hewlett Foundation's Effective Philanthropy Group and later served as president of the Arthur M. Blank Family Foundation. Her expertise includes strategy, evaluation, organizational learning and determining whether grants actually produce intended outcomes.
Nonprofit governance	Anne Wallestad	Former president and CEO of BoardSource, with extensive experience in board accountability, nonprofit leadership and the division of responsibility between boards and executives.
International science governance	Peter Gluckman	President of the International Science Council through October 2026, former chief science adviser to New Zealand's prime ministers, and a specialist in science advice, risk and science diplomacy. He could help prevent the agreement from being governed solely as a U.S. private research project.
Funding Source finance representative	Selected by FS	Someone with authority over grants, investments, accounting or risk—not merely a climate-program officer.
Funding Source mission representative	Selected by FS	Someone responsible for ensuring that the program remains aligned with the donor's climate-benefit objectives.

These are examples of the type and caliber of people I would seek. I do not know whether any would be available or willing to serve.

My preferred leadership

I would select:

- **France Córdova as chair**
- **Gene Dodaro as chair of the Audit and Risk Committee**
- **Lesley Field as chair of the Procurement and Compensation Committee**

Córdova understands large research institutions; Dodaro understands whether money and programs are being properly controlled; Field understands how to write and enforce contracts.

The board should not be dominated by atmospheric scientists. Scientific experts would already serve on the Scientific Advisory Panel, and many would work for organizations receiving grants.

What authority should the board have?

The oversight board should be expressly authorized to:

1. Approve the MO's annual work plan and operating budget.
2. Approve appointment or removal of the program director.
3. Approve large POs and subawards.
4. Review performance against contractual milestones.
5. Commission independent financial and performance audits.
6. Suspend payments or experiments.
7. Require corrective-action plans.
8. reject excessive overhead, sole-source awards or unexplained budget growth.
9. Enforce open-data and intellectual-property requirements.
10. Recommend terminating or replacing the MO.

The MO should manage day-to-day activity, but it should not be able to substantially alter the program without board approval.

Recommended approval thresholds

The agreement could reserve the following decisions for the board:

Decision	Approval needed
Annual program plan and budget	Majority of full board
Individual PO above \$5 million	Majority, including Audit or Procurement chair
PO above \$25 million	Two-thirds of full board
Budget increase above 10%	Board approval
Sole-source award above \$1 million	Written public justification and board approval
Outdoor atmospheric release	Scientific and safety approvals plus two-thirds board authorization
Change in program director	Board approval
Change in MO indirect-cost rate	Board approval
Transfer between major research categories above \$10 million	Board approval
Suspension or replacement of MO	Two-thirds vote plus FS concurrence

These values can be adjusted, but the principle is that the board should focus on **material decisions**, not approve every \$50,000 instrument purchase.

Committee structure

Audit and Risk Committee

Led by the financial-accountability member, this committee would review:

- Financial statements.
- PO commitments and payments.
- Indirect costs.
- Fraud and whistleblower reports.
- Insurance and liability.
- Unused cash and investment earnings.
- Cost overruns.
- Annual independent audits.

The external auditing firm should report directly to this committee—not to the MO's CFO.

Procurement and Compensation Committee

This committee would review:

- Competition plans.
- Sole-source justifications.
- Related-party transactions.
- Executive and PM compensation.
- Contractor selection procedures.
- Conflicts of interest.
- Major contract amendments.
- Whether deliverables and acceptance tests are sufficiently specific.

Program Performance Committee

This committee would examine whether the work is producing value:

- Which uncertainties have been reduced?
- Which experiments were cancelled because they were no longer worthwhile?
- Are milestones meaningful or merely activity counts?
- Are datasets being published?
- Are independent teams reproducing the results?
- Is the MO terminating low-value work quickly?

One member of the Scientific Advisory Panel could attend this committee as a **non-voting liaison**.

Conflict-of-interest rules

The following should normally be excluded from voting membership:

- Current executives, employees or trustees of RTI, UCAR, Battelle or whichever institutions are bidding to become the MO.
- Current employees of organizations expecting major SAI research awards.
- Researchers whose laboratories receive substantial program funding.
- Individuals with investments in possible SAI contractors.
- Paid consultants to the MO or FS on the selection process.

A panel member whose former institution competes for a contract should disclose that relationship and recuse when necessary.

The MO should have **no authority to appoint or remove board members**. Members should be appointed by the FS, with terms that overlap so the FS cannot casually replace the entire independent board after a disagreement.

Compensation

These positions should be paid enough to demand serious attention, but not enough to create dependence on continuation of the program.

An illustrative structure would be:

- Chair: **\$150,000–\$250,000 per year**
- Committee chairs: **\$75,000–\$125,000 per year**
- Other independent members: **\$40,000–\$80,000 per year**
- Travel and professional expenses paid separately

At perhaps **\$1 million–\$1.5 million annually**, a strong oversight board would cost less than 1% of a \$1 billion program and could prevent tens or hundreds of millions of dollars in poor procurements or ineffective research.

A crucial structural distinction

The three governing bodies would have different questions:

Body	Central question
Scientific Advisory Panel	Is this scientifically valuable and properly designed?

Body	Central question
Safety and Governance Board	Is it acceptable, lawful and responsible to perform?
Agreement Oversight Board	Is the MO complying with the agreement and spending the FS's money effectively?

An experiment would proceed only when all necessary bodies approve it.

My recommended arrangement is therefore a small, powerful board dominated by **independent research-management, audit, procurement and philanthropic-effectiveness leaders**—not SAI proponents and not representatives of the organizations receiving the money.

Agreements

Prompt to AI: In theory, there would be an agreement between Funding Source (FS) and Managing Organization (MO), and an agreement between MO and recipient of Purchase Order (PO). What might those agreements look like?

Response From AI:

The two agreements should form a **contractual chain**:

Funding Source → Managing Organization → Research recipients and vendors

The first agreement governs the entire \$1 billion program. The second converts the program rules into specific work, deliverables, payments, data rights, and stop conditions.

A crucial distinction: not every downstream award should technically be called a “PO.” A university or laboratory making scientific decisions and carrying out part of the research program would normally receive a **research subaward or sponsored-research agreement**. A company supplying aircraft time, instruments, engineering, chemicals, or calibration services would receive a **procurement contract or PO**. The substance of the relationship matters more than the document’s name.

This is an illustrative structure, not a substitute for legal, tax, aviation, environmental, or insurance counsel.

1. FS–MO agreement

I would call this the:

Master Program Funding and Stewardship Agreement

It might be 50–80 pages, plus detailed schedules. It would resemble a restricted grant agreement combined with some features of a cooperative agreement.

Principal terms

Article	Illustrative provision
Parties	Funding Source and selected MO
Purpose	Produce reliable, openly available knowledge needed to evaluate the feasibility, risks, benefits, and controllability of SAI

Article	Illustrative provision
Maximum commitment	Up to \$1 billion over five years
Initial obligation	Only the first approved tranche—perhaps \$25M–\$100M—is legally obligated initially
Future tranches	Released after annual plan approval, audit, milestone completion, and board authorization
Restricted use	Funds may be used only for approved SAI research and program-management activities
MO status	Independent contractor and grantee; not an agent or partner of FS
Program governance	Agreement Oversight Board, Scientific Advisory Panel, and separate Safety and Governance Board
MO fee/overhead	Fixed or capped administrative rate, perhaps 5%–7% , separately identified
Procurement	Competition, documentation, conflicts rules, price reasonableness, and approval thresholds
Open science	Mandatory publication of data, reports, code, models, and funded designs
Safety	No outdoor release without all required scientific, safety, environmental, regulatory, and board approvals
Reporting	Monthly financial dashboard, quarterly program reports, annual audited statements
Audit	FS and independent auditors receive access to financial and program records
Remedies	Payment suspension, corrective-action plan, recovery of misused funds, termination, or replacement of MO
Transition	MO must transfer records, contracts, equipment, data, and remaining funds if replaced

Federal grant rules do not automatically govern a private philanthropic program, but they provide a useful model: documented internal controls, source-supported financial records, budget comparisons, written cost-allowability procedures, procurement documentation, contractor oversight, and conflict-of-interest controls.

Funding and payment structure

The agreement should distinguish between:

- **Maximum authorized program size:** \$1 billion.
- **Currently obligated amount:** money FS is presently committed to provide.
- **Approved annual budget:** money MO may commit through contracts.
- **Cash advanced:** money actually transferred to MO.
- **Expenditures:** money spent and accepted.
- **Uncommitted balance:** money that remains under FS control.

For example:

Event	Funding released
Agreement signing and organizational phase	\$25M
Approval of Year 1 operating plan	\$100M
Completion of design and governance gates	\$150M
Approval of initial field campaigns	\$200M
Later years	Released through additional tranches
Total possible	Up to \$1B

The MO should maintain separate program accounting and provide records identifying the source, authorization, commitment, expenditure, remaining balance, assets, and supporting documentation. That mirrors established federal financial-management principles.

Reserved decisions

The MO should manage daily operations, but the agreement should reserve major decisions to the FS or Agreement Oversight Board.

Decision	Required approval
Annual operating plan	Oversight Board and FS
PO or subaward above \$5M	Oversight Board
Award above \$25M	Two-thirds board vote and FS approval
Outdoor atmospheric release	Scientific, safety, board, and legal approval
Sole-source award above \$1M	Written justification and board approval
Major movement between budget categories	Board approval
Change in MO administrative rate	FS approval
Appointment/removal of program director	Board approval
Material alteration of open-access rules	FS approval
Use of program funds for deployment	Prohibited unless agreement formally amended

Performance requirements

The agreement should include a **Program Results Framework**, not merely a list of activities.

For every major research activity, the MO would document:

1. The uncertainty being addressed.

2. The decision that the result will inform.
3. Existing evidence and why additional work is needed.
4. Expected information value.
5. Estimated cost.
6. Measurable deliverables.
7. Success, failure, and stop criteria.
8. Data-release deadline.
9. Independent validation plan.

10. For sulfur pathways, standardized forcing/kg-S, forcing/aircraft-hour, forcing/\$, energy use, feedstock mass, and lifecycle-cost outputs.

Performance requirements should identify goals, indicators, targets, timelines, baseline information, and reporting methods. That is consistent with the performance-measurement structure used in major government research awards.

Open-science clause

An illustrative provision might say:

All research outputs produced using Program Funds—including reports, manuscripts, data, metadata, software, source code, model configurations, laboratory protocols, drawings, CAD files, bills of materials, calibration records, and analysis methods—shall be made publicly available under the licenses and timelines specified in the Open Science Schedule.

The schedule could require:

- Publications and reports: **CC BY 4.0**
- Data: **CC0**, where legally possible
- Software: **Apache-2.0, BSD-3-Clause, or MIT**
- Hardware: an approved open-hardware license
- Raw data: within 90 days
- Calibrated data: within six months
- Final code and reproducibility package: within 12 months

The Gates Foundation provides a relevant precedent: its current policy requires immediate open terms for funded manuscripts and availability statements covering underlying data, metadata, software, and materials needed to assess and reproduce findings. Its broader “Global Access” approach also requires access commitments to flow down to subgrantees and subcontractors.

FS remedies

The FS should be able to:

- Withhold an upcoming tranche.
- Stop new commitments.
- Suspend a particular experiment.
- Require replacement of a PM.
- Order an independent audit.
- Require a corrective-action plan.
- Recover funds spent outside the approved purpose.
- Terminate the agreement for cause.
- Terminate for convenience with appropriate wind-down funding.
- Transfer the program to another MO.

When a private foundation must exercise “expenditure responsibility,” IRS guidance specifically contemplates written commitments concerning permitted use, repayment of misused funds, annual reporting, expenditure records, and access to books. Whether those rules legally apply depends on the FS and MO’s tax status, but they provide a useful contractual minimum.

2. MO–recipient agreement

The correct form depends on what the recipient does.

A. Research subaward

Use this when a university, laboratory, or nonprofit:

- Helps determine scientific methods.
- Employs principal investigators.
- Interprets results.
- Carries out a portion of the program.
- Has meaningful programmatic discretion.

A research subaward might contain:

Section	Contents
Award identification	Recipient, project title, amount, period, PI, MO manager
Scientific objective	Question being answered and decision it informs
Statement of work	Tasks, experimental design, instruments, models, analysis
Milestones	Dates and measurable completion criteria
Budget	Personnel, travel, equipment, indirect costs, subawards
Payment	Advance or reimbursement, with milestone holdbacks
Key personnel	Named PI and approval needed for replacement

Section	Contents
Reporting	Monthly operational, quarterly technical, annual financial
Data management	Raw data, metadata, calibration, repository and release dates
Publications	Preprints, acknowledgment and open license
IP	Background IP, funded developments and public licenses
Equipment	Title, inventory, maintenance and disposition
Safety	Permits, reviews, incident reporting and stop-work authority
Audit	Records access, site visits and corrective actions
Subrecipients	Prior approval and mandatory flow-down terms
Termination	Cause, convenience, transition and final deliverables

Pass-through entities managing subawards are generally expected to identify the award, amount, period, requirements, indirect-cost treatment, records access, closeout rules, risk level, reporting, monitoring, corrective action, and audit obligations.

B. Procurement contract or PO

Use this for defined goods or services such as:

- Aircraft flights.
- Nozzle fabrication.
- Instrument purchases.
- Laboratory analysis.
- Calibration.
- Balloon launches.
- Lidar services.
- Cloud computing.
- Environmental consulting.

A PO should be shorter and more acceptance-oriented:

Recipient: NASA aircraft program or commercial aircraft operator

Maximum amount: \$3,500,000

Deliverable: Six instrumented high-altitude sorties, release and sampling support, flight records, telemetry, and accepted mission dataset

Period: March 1–September 30, 2028

Payment: Milestone-based, not simply hours billed

Example milestone schedule

Milestone	Deliverable	Payment
Contract signing	Bond, insurance, project plan	5%
Preliminary design review	Approved interface and safety design	10%
Critical design review	Final drawings and hazard analysis	15%
Ground acceptance test	Demonstrated release and measurement performance	20%
Flight-readiness review	Aircraft and payload authorized	10%
Field operations	Completed approved sorties	20%
Preliminary dataset	Raw data, logs and metadata	10%
Final acceptance	Calibrated data and complete technical package	10%

The final 10%–15% should remain unpaid until the recipient delivers the required data, documentation, source files, and open licenses.

Deliverable acceptance

A PO should avoid vague deliverables such as:

“Support aerosol research.”

Instead:

“Deliver a calibrated dataset measuring particle number concentration and size distribution from 10 nm to 2 μ m, with stated uncertainty, timestamps synchronized within one second, documented inlet losses, complete calibration records, and data delivered in the required repository format.”

The agreement should state:

- Who accepts the deliverable.
- Objective acceptance criteria.
- Review period—perhaps 30 days.
- Recipient’s right to correct deficiencies.
- Consequences if correction fails.
- Whether partial acceptance is possible.

Change-control clause

Neither party should be able to change scope through informal emails.

A formal change order should be required for changes to:

- Amount.
- Schedule.
- Experimental conditions.
- Release quantity or material.
- Principal investigator.
- Major subcontractors.
- Data rights.
- Safety requirements.
- Acceptance criteria.

Emergency operational decisions may be permitted, but they must be documented promptly and cannot increase the authorized amount without written approval.

Flow-down clause

Every downstream agreement should include language such as:

Recipient shall incorporate the applicable Open Science, Data Management, Intellectual Property, Safety, Conflict-of-Interest, Audit, Records-Retention, Incident-Reporting, and Prohibited-Use requirements of this Agreement into every lower-tier subaward and subcontract supported by Program Funds.

This is essential. Otherwise, the MO might promise the FS open data while a university's subcontractor claims ownership or confidentiality. Global-access frameworks similarly call for access commitments to be incorporated into agreements with project collaborators, subgrantees, and subcontractors.

Contract hierarchy

The documents should have a clear order of precedence:

1. Applicable law and regulatory approvals.
2. FS–MO Master Agreement.
3. Program-wide policy schedules.
4. Annual operating plan.
5. MO–recipient agreement.
6. Statement of work.
7. Approved budget.
8. Recipient proposal.
9. Informal communications.

This prevents a recipient’s proposal language from quietly overriding the FS’s data, safety, or IP requirements.

Recommended document package

FS–MO agreement attachments

- Schedule A: Program purpose and prohibited activities
- Schedule B: Five-year budget and tranche plan
- Schedule C: Governance and approval authorities
- Schedule D: Procurement manual
- Schedule E: Open science and IP policy
- Schedule F: Safety and outdoor-release approval process
- Schedule G: Performance metrics
- Schedule H: Financial reporting and audits
- Schedule I: Conflicts, ethics, and whistleblower policy
- Schedule J: Termination and transition plan

MO–recipient agreement attachments

- Exhibit 1: Statement of work
- Exhibit 2: Deliverable and acceptance matrix
- Exhibit 3: Budget and payment schedule
- Exhibit 4: Data-management and publication plan
- Exhibit 5: Background and foreground IP
- Exhibit 6: Safety and environmental requirements
- Exhibit 7: Required flow-down provisions
- Exhibit 8: Reporting templates

The most important principle is that the two agreements must be drafted together. Every material promise made by the MO to the FS—open data, competitive procurement, audit access, safety approval, IP access, reporting, and termination rights—must be converted into an enforceable obligation in the downstream subawards and POs.

Termination Shock

Prompt to AI: If we reflect sunlight for 10 years to offset the 0.3°C/decade global warming rate to 0.0°C/decade, and then we stop reflecting after 10 years, what is our new warming rate? Alternatively, if we reflect for 200 years and then stop reflecting when the concentration of CO₂ in atmosphere is 1200 ppm, for example, instead of the current 430 ppm, then what would be our new warming rate?

Response From AI:

There is **not one permanent “new warming rate.”** After SAI stops, warming would initially be very rapid and then slow as Earth approaches the temperature associated with the accumulated greenhouse-gas concentration. Climate models consistently show this termination effect, and its magnitude is approximately proportional to how much warming SAI was masking.

I will use your earlier assumptions:

- Present warming rate without SAI: **0.3°C per decade**
- SAI holds temperature change at **0.0°C per decade**
- ECS: **4°C per CO₂ doubling**
- Abrupt termination rather than a gradual phaseout
- CO₂ and other greenhouse forcing continue increasing

Results

SAI duration	Warming masked at termination	Approximate first-decade warming after stopping
10 years	About 0.3°C	About 0.45–0.6°C per decade
200 years, stopping at 1,200 ppm	About 6–7°C	Roughly 3–4°C per decade

Case 1: stop after 10 years

Without SAI, temperature would have risen approximately:

$$0.3^{\circ}\text{C/decade} \times 1 \text{ decade} = 0.3^{\circ}\text{C}$$

SAI has therefore created approximately a **0.3°C temperature gap** between the geoengineered world and the world without SAI.

After termination, two effects occur simultaneously:

1. The normal underlying warming continues at approximately **0.3°C per decade**.
2. Earth begins recovering the approximately **0.3°C of warming that SAI had suppressed**.

Climate-model experiments indicate that approximately **38%–61%—roughly half—of the equilibrium temperature response to an abrupt forcing change occurs during the first decade**.

A simple estimate is therefore:

$$0.30 + 0.5(0.30) \approx 0.45^\circ\text{C/decade}$$

Because the SAI cooling forcing at termination could represent somewhat more than the already realized 0.3°C temperature difference, a reasonable upper estimate is about **0.6°C per decade**.

Best estimate

Approximately 0.5°C per decade during the first decade after stopping.

After this initial catch-up, the rate would move back toward approximately **0.3°C per decade**, assuming the underlying greenhouse forcing continued increasing at the same rate.

This would be noticeable termination warming, but it would be small compared with terminating a centuries-long program.

Case 2: stop after 200 years at 1,200 ppm

Using ECS = 4°C and a preindustrial concentration of 280 ppm:

$$T_{eq}(C) = 4\log_2\left(\frac{C}{280}\right)$$

At 1,200 ppm:

$$T_{eq}(1200) = 4\log_2\left(\frac{1200}{280}\right) \approx 8.4^\circ\text{C}$$

That is the approximate equilibrium warming above preindustrial conditions associated with 1,200 ppm under the assumed ECS.

If the program had held global temperature near today's level—approximately 1.4°C above preindustrial—SAI would be masking approximately:

$$8.4 - 1.4 = 7.0^\circ\text{C}$$

A slightly narrower calculation that counts only the additional CO₂ increase from 430 to 1,200 ppm gives:

$$4\log_2\left(\frac{1200}{430}\right) \approx 5.9^{\circ}\text{C}$$

The difference is whether one includes the warming already committed but not yet realized at 430 ppm. Since your program is assumed to hold temperature completely flat for 200 years, the more appropriate masked amount is probably approximately **7°C**.

If roughly half emerged during the first decade:

$$0.5 \times 7.0 = 3.5^{\circ}\text{C}$$

Continuing greenhouse-gas growth might add roughly another **0.2–0.3°C per decade**, producing:

Approximately 3–4°C of global warming during the first decade after abrupt termination.

Model results support the general magnitude and behavior. One GeoMIP study involving a much shorter, roughly 70-year intervention found approximately **1.5°C warming in the first decade and 0.5°C in the second decade** after termination. Another study found a roughly **1.0–1.16°C-per-decade 20-year warming trend** after terminating a program masking a forcing jump of about 4–5 W/m². Your 1,200-ppm scenario would mask substantially more warming, so the termination rate would likely be considerably higher.

An approximate trajectory might be:

Period after termination	Approximate warming
First decade	3–4°C
Second decade	1–2°C
Following decades	Progressively slower
Long-run rate after catch-up	Returns toward the underlying CO ₂ -driven rate

The sulfate aerosol would not disappear instantaneously; its cooling influence would probably decay with an approximate one-year e-folding timescale. The strongest radiative change would therefore occur over the first several years rather than literally overnight.

Central Conclusion

Stopping after **10 years** probably raises the first-decade warming rate from **0.3°C to around 0.5°C per decade**.

Stopping abruptly after **200 years at 1,200 ppm** could produce approximately **3–4°C of warming in the first decade**, potentially followed by another **1–2°C in the second decade**.

That second scenario would be an extreme termination shock. Avoiding it would require gradually reducing SAI while simultaneously reducing CO₂ through emissions cuts and carbon removal. For example, limiting additional warming from SAI phaseout to **0.2°C per decade** would require approximately **50 years of phaseout for every 1°C of cooling being supplied**—potentially several centuries for a program masking 6–7°C.

SAI Aircraft, Instrumentation and Plume Tracking

Prompt to AI: The “sai_experiments_research_fund_ai, v1o.docx” file refers to four types of airplanes: (a) inject material, (b) fly above plume, (c) fly through plume, (d) fly below plume. Tell me more about these airplanes. What parameters does each measure? What instrumentation is needed to measure each parameter? Do we need to develop aircraft or modify existing aircraft? Do we need to develop instrumentation or equipment that is added to aircraft? What do we need to develop? Notice we are looking at a plume that moves about in the wind and disperses. Do we need an image (e.g. 1000 x 1000 pixels) or is one pixel sufficient? Do we need the imager (IR, UV, visible) on a gimbal to keep it steady as the airplane rocks back and forth? How do we see the plume? Are existing atmospheric research airplanes suited for this kind of work? They tend to look at other things. Might a phased-array microwave receiver be helpful?

Response From AI:

The monitoring problem should be treated as an instrument-and-mission-system problem more than an aircraft-development problem. Existing research aircraft are generally suitable as sensor buses. The largest new engineering efforts are the SAI release package, real-time plume finding and navigation, acid-compatible sampling/integration, and common radiative-closure payloads. A purpose-built high-altitude aircraft becomes relevant only if later work requires a combination of altitude and payload that the research fleet cannot provide.

Four aircraft roles

Aircraft role	Primary job and parameters	Representative instrumentation	Development need
Release/injector	Release a precisely known perturbation. Measure total mass and mass flow, feedstock pressure/temperature, aerosol-generator settings, exhaust or vaporizer conditions, position/attitude, ambient temperature/pressure/water vapor and release timing.	Precision mass-flow and tank metrology; pressure/temperature sensors; aerosol-generation diagnostics; GPS/INS; meteorological sensors; leak detection and safety controls.	High for the release module. Modify an existing aircraft initially; develop a modular, certifiable SO ₂ /H ₂ SO ₄ /C ₃ release system rather than a new airframe.
Above plume	Measure reflected and absorbed radiation plus plume geometry: up/down	Broadband SW/LW radiometers; spectral radiometer; multi-angle	Low-to-moderate. Existing aircraft and most sensors are suitable; improve

Aircraft role	Primary job and parameters	Representative instrumentation	Development need
	SW and LW flux, spectral reflectance, polarization, layer height, width/length and optical-depth heterogeneity.	polarimeter/imager; HSRL/CPL-class lidar; precision GPS/INS.	stabilization, common calibration and real-time processing.
Through plume	Measure what the released material becomes: SO ₂ /H ₂ SO ₄ , particle number and size, sulfate mass/composition/acidity, surface area, scattering/absorption/extinction, ozone-relevant chemistry, H ₂ O, winds and turbulence.	SO ₂ detector; CIMS for H ₂ SO ₄ /trace gases; nanoparticle-to-micrometre sizing; aerosol composition instruments; nephelometer/absorption optics; ozone/NO _y /H ₂ O; gust/turbulence package.	Low-to-moderate. Use existing stratospheric chemistry aircraft; emphasize low-loss/acid-compatible inlets, calibration, synchronized racks and plume-targeting displays.
Below plume	Measure radiation transmitted through the plume and the lower boundary for closure: SW/LW flux, spectral attenuation, column AOD, SO ₂ column early in C1, plume-base altitude and background aerosol/cloud state.	Up/down radiometers; spectral radiometer; UV-visible spectrometer/DOAS; upward-looking lidar; optional polarimeter; GPS/INS and meteorology.	Low-to-moderate. Existing aircraft are suitable; integrate matched radiation sensors and attitude correction.

Do we need to develop new aircraft?

For the five-year experimental fund, usually no. Existing high-altitude atmospheric research aircraft can perform the above-plume and through-plume roles, while Gulfstream-, C-130-, King-Air- or Twin-Otter-class research aircraft can perform many below-plume and lower-altitude roles. The program should develop removable payloads and certified interfaces so the same instrument packages can migrate among qualified aircraft.

The injector is the exception. For C1, the new equipment is a corrosion-resistant SO₂ storage, metering and safety package. For C2, the aircraft also becomes part of the aerosol-generation experiment: mechanical atomization, exhaust-assisted H₂SO₄ vaporization and dedicated-vaporizer architectures require tanks, pumps, heat exchangers or vaporizers, exhaust interfaces, controls, leak detection, acceptance testing and airworthiness approval.

Finding a moving and dispersing plume

For multi-day C1 work, plume reacquisition may be as important as individual instrument accuracy. The program should develop a Plume Location, Imaging and Navigation System that fuses trajectory forecasts, measured winds, lidar curtains, multispectral/polarimetric imagery, UV-visible SO₂ information and confirmed in-situ plume hits into a continuously updated three-dimensional plume estimate with an uncertainty envelope.

The operations display should function conceptually like weather radar for the science mission: predicted plume volume, latest confirmed detection, estimated plume age, uncertainty, and recommended cross-plume or along-plume flight legs. The same data system should guide the above, through and below aircraft and feed post-flight data assimilation.

One pixel versus an image

Measurement need	One-pixel / integrated measurement	Imaging requirement
Radiative forcing / flux closure	Often appropriate. A broadband or spectral flux radiometer deliberately integrates radiation over a defined field of view or hemisphere; this directly measures the quantity needed for the radiation budget.	An image is not required to obtain the flux number, but imaging helps establish whether the radiometer sampled plume center, edge, gaps or clean background.
Plume location and geometry	Insufficient by itself. A single signal cannot determine the two-dimensional plume boundary or heterogeneous optical depth.	Use a multi-pixel imager/polarimeter plus lidar. Specify ground sampling distance, swath, spectral bands, signal-to-noise and viewing geometry rather than an arbitrary 1000 x 1000 pixel count.
Practical spatial-resolution rule	N/A	Aim for roughly 20-50 independent samples across the narrowest plume width that must be resolved, then bin pixels when higher signal-to-noise is more valuable than native spatial resolution.

Stabilization and gimbals

Precision radiometry should use either a stabilized platform or sufficiently accurate, high-rate inertial measurements to correct for aircraft pitch and roll. A gimbaled imager is useful when the instrument

must hold the plume in view or acquire multiple viewing angles. A lidar can often remain fixed nadir- or zenith-looking if its beam direction is accurately referenced to the aircraft INS; scanning lidar can be added when cross-track geometry is scientifically valuable.

How the plume is seen

C1 and C2 are different. Immediately after a C1 SO₂ release, much of the perturbation is an optically weak gas plume. Early plume finding therefore relies heavily on the source trajectory, winds, trajectory models, UV-visible/SO₂ detection and direct SO₂ measurements when an aircraft crosses the plume. As sulfate forms, lidar, polarimetric imaging and aerosol optical measurements become progressively more useful.

For C2 direct H₂SO₄ aerosol, particles exist immediately, so lidar and optical imaging can map the layer almost at once. The combined observing logic is: lidar provides the vertical skeleton, imaging/polarimetry maps horizontal structure, the through-plume aircraft supplies physical and chemical truth, and above/below radiometers measure the energy-budget effect.

Are existing atmospheric research aircraft suited?

Yes. They already carry many of the required instrument families, even though current campaigns study volcanoes, ozone, air pollution, clouds, convection or satellite validation rather than deliberate SAI plumes. The SAI-specific work is mainly mission integration: acid-compatible and low-loss inlets, source-record synchronization, standardized instrument racks, real-time plume targeting, common calibration and coordinated multi-aircraft operations.

Illustrative platform assignment

Role	Illustrative existing platform class	Why it fits
Above plume	ER-2-class high-altitude remote-sensing aircraft	Very high altitude with optical ports and heritage carrying radiometers, lidar and imaging instruments.
Through plume	WB-57-class high-altitude in-situ aircraft	High altitude plus substantial payload for aerosol, gas chemistry and meteorological instruments.
Below plume	NSF/NCAR Gulfstream V or comparable high-performance research jet; smaller aircraft where altitude permits	Good payload/endurance for radiation, lidar, cloud/aerosol context and stacked flight legs.
Boundary-layer/context	C-130, King Air or Twin Otter class	Useful for lower-altitude radiation, aerosol/cloud context and remote

Role	Illustrative existing platform class	Why it fits
		sensing; not a substitute for stratospheric sampling.
Release	Modified existing jet initially	Keeps new development concentrated on the release module rather than the entire airframe.

Would a phased-array microwave receiver help?

It is not a priority instrument for directly detecting submicron sulfate aerosol. The particle size is far smaller than microwave wavelengths, so direct microwave backscatter is weak. Microwave radar/radiometry can still provide useful cloud, precipitation, temperature/humidity or dynamical context. If additional plume-detection money is available, the SAI program should prioritize lidar, UV-visible SO₂ detection, multi-angle optical/polarimetric imaging and real-time data fusion before a phased-array microwave receiver.

What the SAI fund should develop

- Modular SAI release-aircraft payloads for C1, C2A, C2B, C2C and later C3, including precision release metrology and corrosion/safety systems.
- A real-time three-dimensional plume finder and flight-guidance system combining forecasts, winds, lidar, imagery and in-situ detections.
- A stabilized radiative-closure payload with matched above/below SW and LW sensors, spectral radiometry, high-rate INS and common calibration.
- A standardized through-plume chemistry/microphysics payload assembled largely from existing instrument families but optimized for acid compatibility, inlet losses, rapid changes and synchronized data.
- A multispectral/polarimetric plume-mapping package adapting existing imaging-polarimetry and lidar technology rather than creating a new camera solely around an arbitrary pixel count.

The main conclusion is that the airplanes should be treated primarily as reusable buses. The unique new technology is the release hardware and the integrated measurement, plume-tracking and mission-navigation system.

Further Reading

Is it possible to solve the climate problem with a surge of R&D in key areas; and if so, what are those areas and how much might this cost? The following documents hope to answer this question.

Stratospheric Aerosol Injection (SAI)

SAI Reflectivity Plan (7-page PDF)

<https://drive.google.com/drive/folders/1Jf8jmP38b2P-vj9XTWRLGGMuahuLwxB6>

Proposed \$1B SAI Experimental Research Fund

https://drive.google.com/drive/folders/1KRtdfFEjKmzV5XaxiyzRL4JckPzrF_sM

Proposed \$1B SAI Equipment Development Fund

https://drive.google.com/drive/folders/1B0Sko2M2cl_V-_uZcPozD-s0cF8bjV-

Marine Cloud Brightening (MCB)

MCB Reflectivity Plan (7-page PDF)

<https://drive.google.com/drive/folders/1Aoz9vLxKLF4qBw1xnjlCfYGXERhTY3rI>

Proposed \$1B MCB Experimental Research Fund

<https://drive.google.com/drive/folders/1ItaJh7OQMeZMInUkArcShi-SGVFmzvD0>

Proposed \$1B MCB Equipment Development Fund

https://drive.google.com/drive/folders/1Leivbra4GXj5VlPg_jFZfPN1UQyHgR-S

Global Climate Models

Proposed \$1B GCM Experimental Research Fund

<https://drive.google.com/drive/folders/1i6uLaOv3Yfx9RzXsDUEnDvDjLlE0Dl-6>

Nuclear Fission Power

Proposed \$1B Nuclear Power Automated Construction Development Fund

https://drive.google.com/drive/folders/1HWNJZ2X3jgx4dXn_7VPq251jJJuIGlaS

Proposed \$1B Nuclear-Integrated Chemical Processing Platform Development

<https://drive.google.com/drive/folders/1-G2IHTerH0ETdfpCX-i8Nin0kh0lEBru>

Climate Laboratory

Climate Moonshot Business Plan (90-page PDF)

<https://drive.google.com/drive/folders/1EMb-Tqw49ewj16UeGhR7UOc9HY6POpKC>

Atmospheric Research

Proposed Experiments Measure Sunlight Reflecting off Pollution (47-pg PDF)

<https://drive.google.com/drive/folders/1X6PBoXXu-VhMOtQ5JOzlt29oW3qjzGhY>

Proposed \$250M Cloud Experimental Research Fund (13-page PDF)

https://drive.google.com/drive/folders/1apa-sQ4fyz-k6e5cnyvvy01QKQ8r_ICB

Proposed \$2B Atmospheric Experimental Research Fund (400-page PDF)

https://drive.google.com/drive/folders/1e5jvAmJUfBLCpkQVKwR0dS_Ur2UtFLkG

Ocean Iron Fertilization

Proposed \$25M Ocean Iron Fertilization Research Fund

https://drive.google.com/drive/folders/12Rn4oc_1sUcttW1_UjKNRLJxFrV0sMj0

Climate Moonshot Conference

Climate Moonshot Conference (3-page PDF)

https://drive.google.com/drive/folders/114b_1LPTVhBuBQ4ZFpTQsJlh29_rKaB7

Climate Moonshot Film

The Climate Kids (docudrama film)

https://drive.google.com/drive/folders/1UyEyyGock43PQpnMMcEmq_94gy4m3GHI