Mission Space
Space-environment intelligence for ground, Earth orbit, cislunar space and the lunar surface

SYSTEMIC RISK SERIES

Solar Storms and the Space Economy

Where the exposure sits in 2026, where it moves by 2035, and what it costs — in orbit, on the Moon, at Mars, and on the ground

Built on the Solar Storm Economic Exposure Model (SSEEM) v2.0

Nine loss channels, six of them off-planet. Five severity tiers. Forty-five tier-channel cells, each with a low, central and high value and a confidence grade. Published in full in Appendix A so that any parameter can be challenged and the result regenerated.

Prepared by Mission Space US Corp.
66 West Flagler Street, Suite 900, Miami, FL 33130 · mission.space · info@mission.space

August 2026

Audiences: corporate leadership and investors; government and policy; insurance and reinsurance; space and aerospace industry

Model: Solar Storm Economic Exposure Model (SSEEM) v2.0. Figures in 2026 US dollars unless stated.

Information cut-off: 24 August 2026. Programme status, contract values and orbital counts are stated as at that date and several are moving quickly; Appendix D dates every source.

Contents

Executive Summary **2**

Part I — The Hazard **6**

1. Physical Basis 6
2. The Historical Record 9
3. Frequency and Severity 13

Part II — The Exposure Model **15**

4. Model Architecture 15
5. Results 19

Part III — The Orbital Economy **25**

6. Space Infrastructure and Connectivity 25
7. Orbital Compute and In-Space Manufacturing 29
8. Human Spaceflight and Commercial Stations 32
9. Space Defence and National Security 35
10. Orbital Traffic and the Conjunction Problem 38

Part IV — Beyond Earth Orbit **40**

11. The Moon 40
12. Mars and Deep Space 45
13. Cislunar Competition 48

Part V — Earth Infrastructure **50**

14. Power Grids 50
15. Aviation 53
16. The GNSS-Dependent Ground Economy 55

Part VI — Implications and Actions **57**

17. Risk Transfer 57
18. Policy and Regulation 59
19. Recommendations by Audience 61
20. What Would Change These Estimates 63

Appendices **65**

A. SSEEM v2.0: Full Model Specification 65
B. Indices, Scales and Terms 72
C. Event Register 74
D. Sources 75

Executive Summary

Space weather has been analysed for two decades as a threat to terrestrial infrastructure. A storm that reaches down a transmission line and destroys a transformer. That framing was correct in 2010. It now describes the smaller and slower-growing half of the exposure.

There are 16,393 active satellites in orbit as of August 2026, against 1,980 in 2018. One operator holds two-thirds of them. Starlink runs at a $17 billion annualised revenue pace with 12 million subscribers and is now enterprise and government infrastructure as much as a consumer product. Nvidia H100s and Google TPUs are being flown as orbital data centres, with a European study proposing gigawatt-scale orbital compute. The ISS is authorised only to 2030 and four commercial stations are competing to succeed it. Golden Dome is a stated $185 billion programme placing missile-defence interceptors in proliferated low Earth orbit. Artemis II flew in April; Gateway was stood down in March; a $20 billion lunar surface base is being contracted; a nuclear-electric spacecraft is scheduled to leave for Mars in December 2028.

All of that sits either inside the layer of atmosphere that a geomagnetic storm heats and expands, or outside the magnetosphere entirely. This paper models what that costs.

Summary of findings

Annualised expected loss runs at **$20.6bn globally**, within a range of $6.5bn to $69.2bn. Seventy-seven per cent of it comes from Carrington-class events, which have no precedent in the electrified era.

The exposure is moving off the planet. Orbital and beyond accounts for 17 per cent of expected loss in 2026 and **42 per cent by 2035**, as the total rises to $44.3bn. Orbital compute grows twenty-five-fold from a near-zero base; lunar and cislunar twelve-fold.

A Carrington-class event would cost about **$2.4 trillion today and $4.7 trillion in 2035** on identical hazard assumptions. The increase is exposure growth, not a change in the Sun.

Across roughly thirty-five company announcements, agency programme pages and contract notices reviewed for this paper, the commercial space sector almost never names radiation, solar storms or orbital drag as a risk. Chapter 6 itemises the exceptions.

The sector's most consequential space-weather decision runs opposite to intuition. SpaceX is lowering roughly 4,400 Starlink satellites from 550 km to about 480 km through 2026, because the approaching solar *minimum* around 2030 would otherwise leave a failed satellite in orbit for four years instead of a few months.

A Nature paper published in July 2026 removes the assumed ceiling on storm severity. Sivadas and colleagues show the apparent saturation of Earth's response to strong solar wind is an artefact of measurement error. Doubling upper-tier severity takes expected annual loss to $37.8bn.

Published estimates of superstorm frequency still differ by a factor of twenty-six, from 0.46 to 12 per cent per decade. That single parameter moves expected annual loss between $5.9bn and $35.3bn.

Where the exposure sits, and where it moves

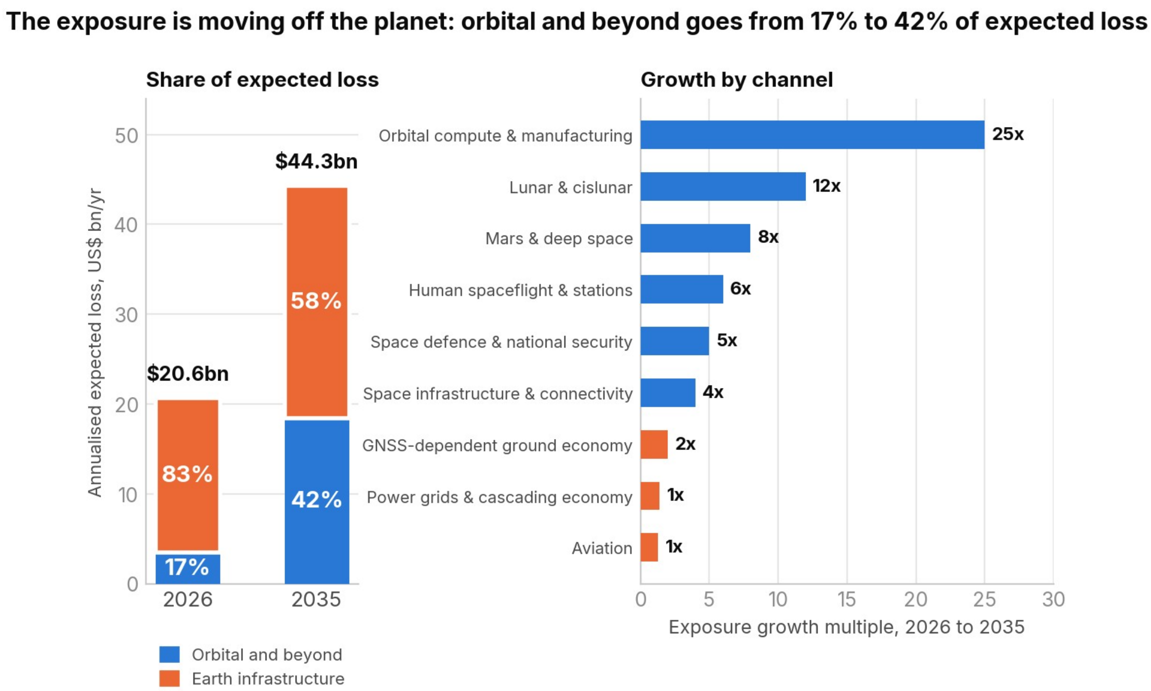

Exhibit 1. Where the exposure sits in 2026 and where it moves by 2035 Source: SSEEM v2.0. Growth multiples applied to 2026 severity per channel; basis for each in Appendix A.4. Space-economy growth follows the WEF/McKinsey projection of $1.8trn by 2035.

Power grids remain the single largest channel in 2026 at $12.9bn a year and hold the largest tail. That does not change by 2035 — grid exposure grows about 40 per cent with electrification, net of hardening. What changes is everything around it. Six channels that together account for $3.5bn a year today reach $18.4bn by 2035, and the composition of a Carrington-class loss shifts from 9 per cent orbital to 32 per cent.

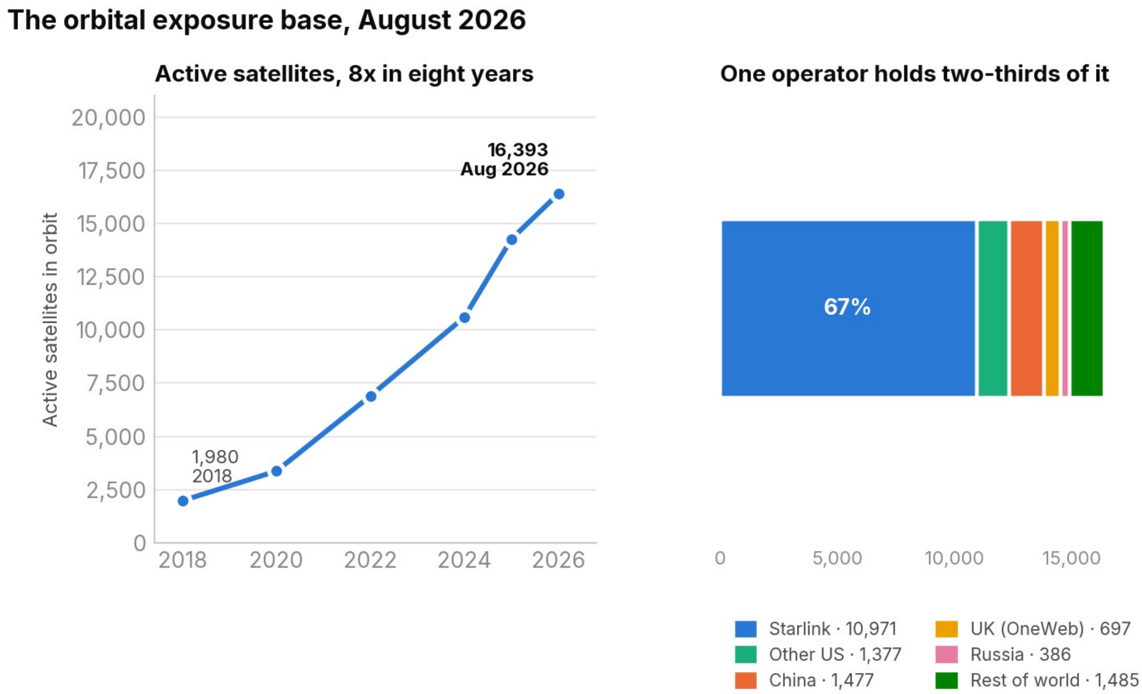

Exhibit 2. The orbital exposure base, 2018 to 2026 Sources: CelesTrak active catalogue via azmth.space (24 August 2026); Satellite Industry Association (2026) for the 2025 figure; historical counts from published annual summaries.

What the model says

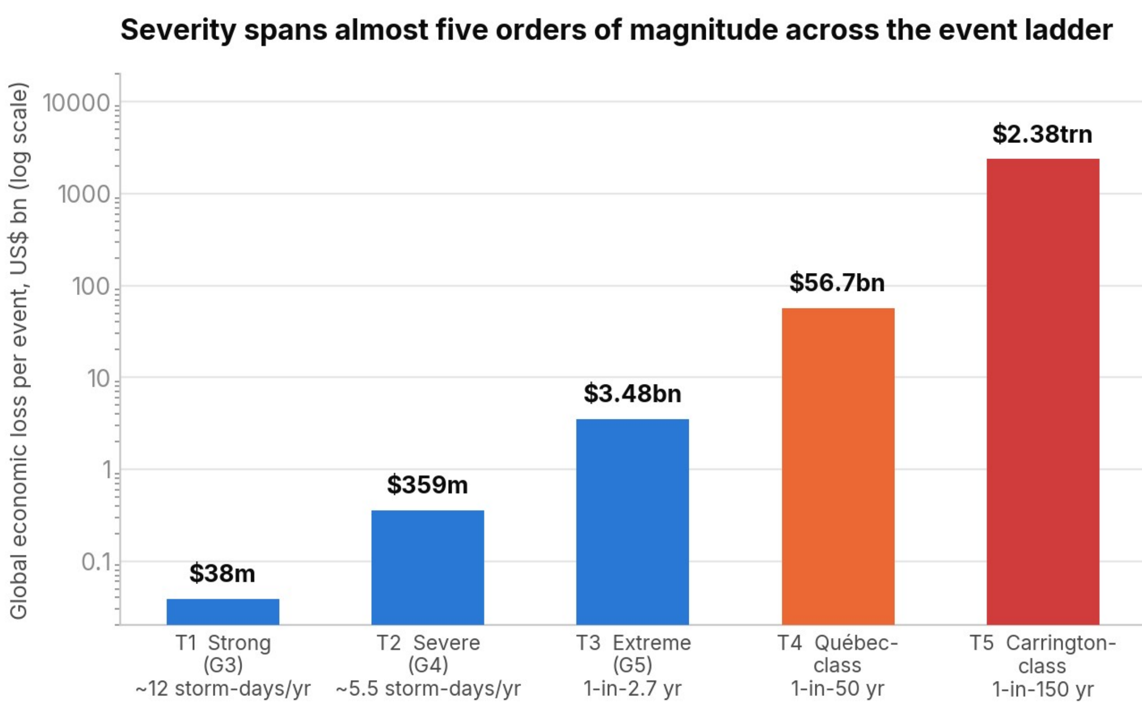

Exhibit 3. The severity ladder: frequency against loss, by tier Source: SSEEM v2.0. Frequencies from NOAA Space Weather Scales (T1–T3, storm-days per 11-year cycle) and Lloyd's/AER (2013) return periods (T4–T5). Severity anchors in Appendix A.

Almost five orders of magnitude separate a storm the grid manages roughly weekly at solar maximum from the storm that defines the risk. The failure modes differ in kind. A G3 storm is a reactive-power problem for a utility and a station-keeping problem for a constellation operator. A Carrington-class storm is a transformer replacement problem against a 36-to-60-month manufacturing queue, and simultaneously a conjunction-assessment failure across an orbital population of 16,000 objects with no ability to predict where any of them will be.

Recommendations

Audience	Priority action	Basis
Space and aerospace industry	Put space weather into the design basis. Publish drag, radiation and dose assumptions the way structural and thermal margins are published, and close the storm-time conjunction-assessment gap.	The sector is silent on the risk in public documentation, and in May 2024 collision screening failed for days.
Corporate leadership and investors	Run a 72-hour combined connectivity, positioning and timing loss scenario. Assign a named owner and a rehearsed manual fallback to every process that assumes any of the three.	The 2024 Gannon loss fell entirely on signal-dependent processes with no manual fallback.
Government and policy	Fund the frequency question, and extend planning benchmarks past 1-in-100 severity. Assign an owner for orbital-regime resilience.	Standards, stockpiles and capital are priced off a recurrence rate the literature cannot agree on within a factor of 26.
Insurance and reinsurance	Aggregate space weather as one correlated peril across property, energy, aviation, marine, agriculture and space. Build products for the orbital economy, where under 1 per cent of LEO satellites carry in-orbit cover.	The Lloyd's 2025 scenario models simultaneous loss across all lines; most books aggregate them as independent.

PART I · THE HAZARD

1. Physical Basis

Space weather is four distinct phenomena travelling at different speeds, arriving with different warning times, and affecting different assets. Treating them as one hazard produces resilience programmes aimed at the wrong failure mode, and for an orbital operator at the wrong scale entirely.

1.1 Four hazards

Hazard	Transit from Sun	Warning available	Assets affected
Electromagnetic radiation (solar flare, X-ray)	8 minutes	None — the flash and the effect arrive together	HF radio across the sunlit hemisphere; GNSS signal quality; radar returns
Solar energetic particles (SEP / proton event)	20 minutes to several hours	Minutes to an hour	Crew and aircrew dose; satellite electronics and memory; orbital compute; polar HF
Coronal mass ejection (CME) → geomagnetic storm	14.6 hours (record) to 3 days	15–60 minutes, from L1 sensors	Power transformers via GIC; pipelines; rail signalling; satellite drag, charging and conjunction assessment
High-speed solar wind stream (corotating)	Recurrent, ~27-day period	Days — the source region is predictable	Chronic satellite charging and drag; sustained moderate geomagnetic activity

The split matters for who is exposed. Terrestrial infrastructure is damaged almost exclusively by the third row, the geomagnetic storm and the currents it induces in long conductors. Orbital and beyond-orbit assets are exposed to all four, and their dominant hazard is the second row, which the terrestrial literature treats as a minor branch.

Warning times set what any operator can do. A flare's radio blackout cannot be forecast: the X-rays that cause it and the light that reveals it arrive together. A CME's geomagnetic storm can be forecast, but only from the moment the plasma cloud crosses the L1 Lagrange point 1.5 million kilometres upstream, which gives fifteen to sixty minutes depending on speed. NOAA's SWFO-L1 observatory, which reached its operational orbit in January 2026, improves the quality of that warning without extending it, because L1 sits at a fixed one per cent of the Sun-Earth distance. Longer lead times require an observatory at L5, trailing Earth in its orbit and viewing Earth-directed structures from the side. ESA's Vigil mission is scheduled for 2031.

Forecast accuracy during large events is poor. Through the growth phase of the May 2024 Gannon storm, geomagnetic index forecasts were underpredicted by 100 to 300 units, then overpredicted through recovery. The magnetic orientation of an arriving CME determines how strongly it couples to the magnetosphere, and it cannot be measured until the cloud reaches L1.

1.2 Two coupling mechanisms

Ground: geomagnetically induced currents

A CME's arrival disturbs the Earth's magnetosphere. The resulting rapid change in the geomagnetic field, measured as dB/dt in nanotesla per minute, induces an electric field in the conducting crust below. That geoelectric field (typically a few volts per kilometre, reaching an estimated 20 to 25 V/km in the largest reconstructed events) drives quasi-direct current through long grounded conductors. In a power system the current flows along transmission lines and down through transformer neutrals to earth, offsets the alternating magnetic flux, and drives the core into half-cycle saturation. The transformer then draws large distorted reactive currents, generates harmonics that trip protective relays on healthy circuits, and heats structural steel and winding insulation. Chapter 14 develops this.

Orbit: density, charge and particle flux

Above the atmosphere the mechanism is different in every respect. A geomagnetic storm deposits energy into the thermosphere, which heats and expands. Density at satellite altitudes rises (six-fold at 400 km within twelve hours during the May 2024 storm) and drag rises linearly with it. Orbits decay faster, orbit predictions fail, and the entire tracked catalogue moves in ways nobody modelled.

Separately, hot plasma injection charges spacecraft surfaces to thousands of volts, and energetic electrons deposit charge inside insulating materials over the days following a storm, producing internal discharge with no external signature. Solar energetic particles flip memory bits, latch up devices and degrade solar arrays permanently.

None of these mechanisms involves a conductor touching the ground, and none of them is measured by the index that terrestrial risk models use.

1.3 Measurement, and why the wrong scale is usually quoted

Four index families recur throughout the literature. They measure different quantities and are not interchangeable.

- Dst / SYM-H. Disturbance storm time, in nanotesla, measuring equatorial ring-current strength. The standard scalar for geomagnetic storm intensity. Carrington is reconstructed at about −900 nT; March 1989 measured −589 nT; May 2024 reached −412 nT.
- Kp / NOAA G-scale. A quasi-logarithmic three-hour planetary index from 0 to 9, mapped onto NOAA's G1 to G5 storm scale. Alerts are issued against it, but it saturates: a G5 storm at Kp 9 and a storm several times stronger both read as G5.
- dB/dt. Rate of change of the local magnetic field, in nT/min. This drives GIC, not Dst. The UK reasonable-worst-case places a 100-year event at 4,000–5,000 nT/min. A 2024 reconstruction of Carrington found mid-latitude rates above 700 nT/min, against a 1-in-100-year benchmark of 350–400 nT/min.
- Proton flux / NOAA S-scale. Integral flux of >10 MeV protons in particle flux units, mapped to S1 through S5. This governs crew and aircrew dose, satellite single-event effects, orbital compute reliability, and every shelter or EVA decision beyond LEO.

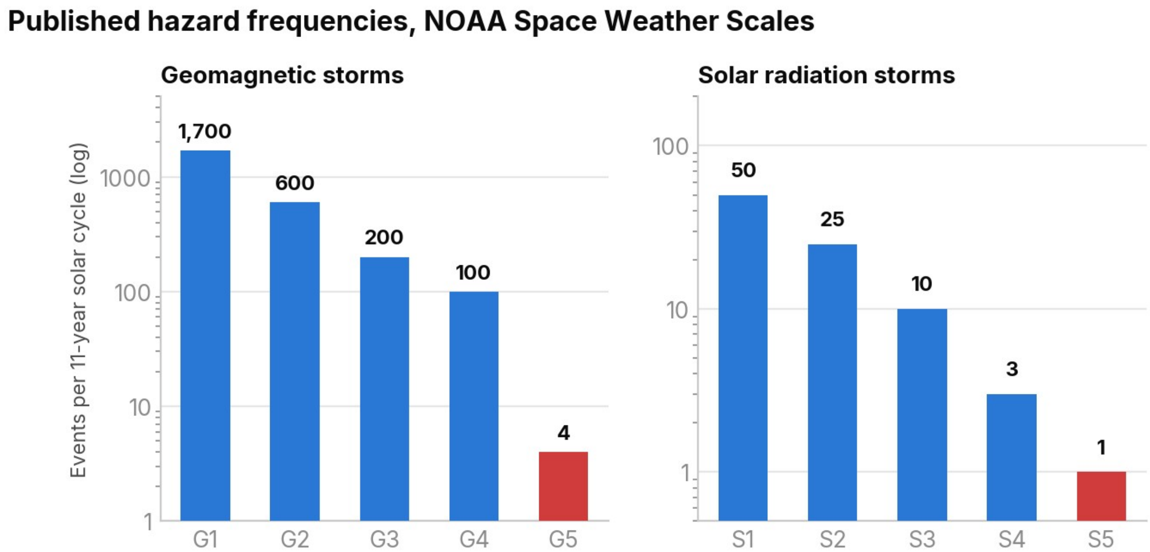

Exhibit 4. Geomagnetic, radiation and radio-blackout events per solar cycle Source: NOAA Space Weather Prediction Center, Space Weather Scales. Event counts per 11-year solar cycle. Both axes logarithmic.

> **The S-scale is the space economy's hazard index, and almost nobody quotes it**
>
> Public and boardroom attention tracks the G-scale, because that is what produces aurora and headlines. For an operator of satellites, orbital compute, or anything with crew aboard, the S-scale matters more: it drives bit-flips, latch-ups, array degradation and dose.
>
> The two scales are only loosely coupled. The **January 2026 event** produced an **S4 severe radiation storm** — among the most intense in the GOES record — from an X1.9 flare, alongside a G4 geomagnetic storm. The May 2024 Gannon storm was **G5 but only S1**.
>
> NOAA's published frequencies give roughly **three S4 events and about one S5 per solar cycle**, against four G5 storm-days. Any risk process keyed to the G-scale alone is blind to the hazard that dominates the orbital channels.

Kp saturation compounds the problem. The May 2024 Gannon storm and a Carrington-class event would both be announced as G5, while differing in economic consequence by a factor of about seven hundred. The Institute for Defense Analyses reviewed the NOAA scales for this and related reasons in 2025, and revisions remain under consideration.

2. The Historical Record

Eight events form the empirical base of every published loss model. Two are reconstructions from magnetograms and contemporary accounts rather than instrumental records. Together they set the severity ceiling and define the failure taxonomy. And the two most recent ones, in 2024 and 2026, define it for the orbital economy specifically.

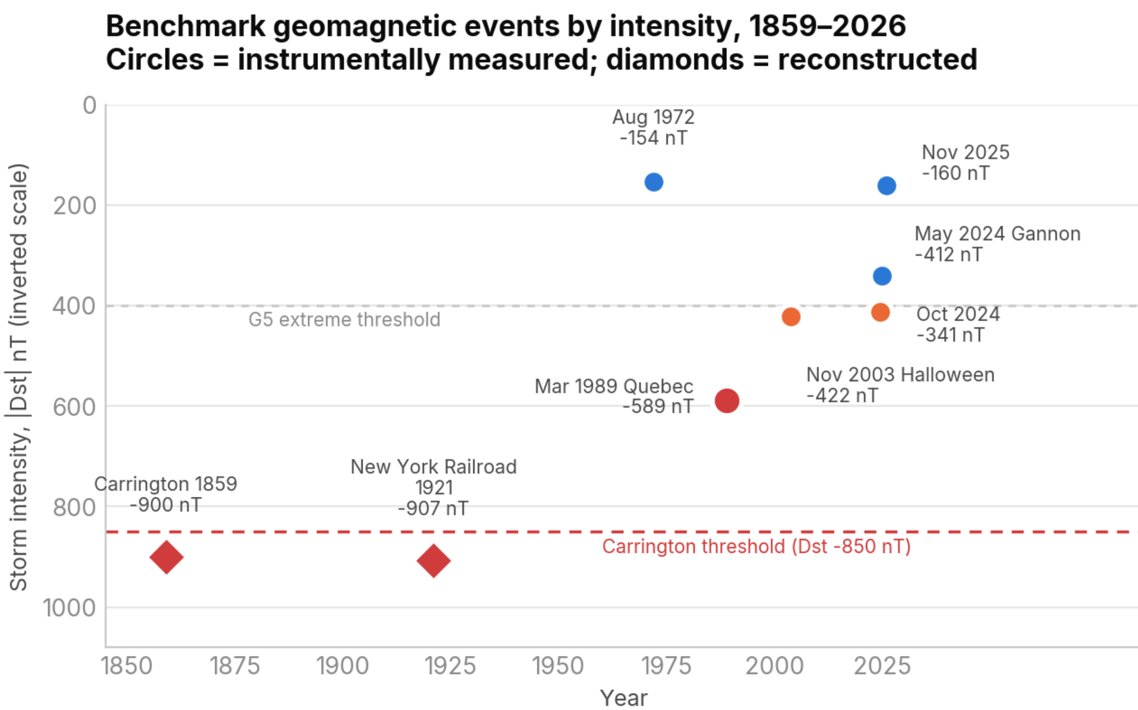

Exhibit 5. The documented record, 1859 to 2026 Source: compiled from USGS Geomagnetism Program, NOAA SWPC, and Beggan et al. (2024) for the Carrington reconstruction. August 1972 is plotted at its measured Dst of −154 nT; its significance is as a radiation event — see 2.4. The January 2026 event is not plotted: its significance is on the S-scale, not the Dst scale.

2.1 September 1859: Carrington

A white-light flare observed independently by Richard Carrington and Richard Hodgson on 1 September 1859 was followed by a CME that reached Earth in an estimated 17.6 hours. Aurorae were reported over Cuba and Hawaii. Telegraph systems across Europe and North America failed, operators received electric shocks, and pylons threw sparks. Two operators on the Boston-Portland line disconnected their batteries and transmitted for two hours on induced current alone.

The 1859 telegraph network was economically negligible. The event matters because it establishes what the Sun is capable of producing. Lloyd's and AER place Carrington-class recurrence at 1-in-150 years, with a range of 100 to 250. Tree-ring radiocarbon evidence indicates that the events of 774–775 CE and 993–994 CE were larger, the former perhaps ten or more times Carrington's magnitude, so the Carrington reference is a modelling convention rather than a physical ceiling.

2.2 May 1921: New York Railroad

Reconstructed at about −907 nT, marginally exceeding Carrington. Fires were reported at telegraph and railway signal installations in New York and Sweden. The 1921 event is the only near-Carrington storm to occur in an era with any electrical infrastructure, and its spacing (62 years after 1859) sits awkwardly with the lower published recurrence estimates.

2.3 May 1967: BMEWS

On 23 May 1967 a solar radio burst jammed all three Ballistic Missile Early Warning System radars in the northern hemisphere at once. Because jamming of missile-warning radar was itself treated as an act of war, additional nuclear-armed alert aircraft were placed in ready-to-launch status. Colonel Arnold Snyder at NORAD's Solar Forecast Center reported the solar event to the command post, and leadership recognised that all three affected sites were in sunlight and receiving direct solar radio emission. Forces stood down as the emission subsided.

The case establishes that the dominant military risk from space weather is misattribution rather than degradation. Chapter 9 develops this against a defence architecture now moving into proliferated low Earth orbit.

2.4 August 1972

The August 1972 storms hold the record for fastest CME transit at 14.6 hours, and matter for three separate reasons.

The associated solar particle event fell between Apollo 16 and Apollo 17. Had astronauts been on the lunar surface, unshielded doses are estimated at more than 10 Gy to skin and 2 Gy to bone marrow, producing acute radiation syndrome and potentially lethal during an EVA. This is the primary justification for the radiation-shelter requirements now written into Artemis mission design.

The storm also detonated dozens of US Navy Destructor magnetic-influence sea mines off Hòn La, North Vietnam, within about thirty seconds; one account puts the figure at 4,000. The cause was classified for decades and reconstructed in the open literature by Knipp and colleagues in 2018. It remains the only confirmed case of space weather destroying deployed munitions.

Terrestrial effects were substantial despite a Dst of only −154 nT: a 64 per cent voltage collapse on the North Dakota-Manitoba interconnection, a 120 MW drop in Manitoba Hydro exports within minutes, an AT&T L4 coaxial cable outage in the Midwest with an estimated local 800 nT/min, and a 60-volt surge on a Chicago-Nebraska cable. Intelsat IV F-2 lost 5 per cent of its solar-panel output (roughly two years of normal degradation in a single event) and a DSCS II satellite suffered a mission-ending power failure. Modelling indicates that a different magnetic-field orientation could have produced Dst near −1,600 nT. August 1972 is the clearest demonstration in the record that Dst is a poor proxy for damage.

2.5 March 1989: Hydro-Québec

The only large-scale space-weather blackout in the historical record. An X15 flare on 6 March was followed by two CMEs; the storm arrived at 01:27 UT on 13 March. The James Bay network went offline in under ninety seconds, leaving six million people without power for nine hours in a Québec March. Direct equipment damage was approximately CAD 6.5 million; broader economic losses exceeded CAD 1 billion.

The peripheral damage is the more instructive part. A step-up transformer at the Salem nuclear plant in New Jersey suffered permanent winding damage. NASA's TDRS-1 logged over 250 anomalies. GOES weather-satellite communications were interrupted and imagery lost. Space Shuttle Discovery recorded a spurious sensor reading on a hydrogen fuel-cell tank. Australian Army peacekeepers in Namibia lost HF communications entirely. In a related August event the Toronto Stock Exchange halted trading when three redundant disc drives failed.

One spectacular failure surrounded by unrelated-looking incidents across continents and sectors is the signature of a systemic peril, and it is why Chapter 17 argues for aggregating space weather as a correlated exposure.

2.6 October–November 2003: Halloween

The largest storm of the satellite era, including the strongest X-ray flare ever recorded at X28+, which saturated the GOES sensor. Effects included a one-hour blackout for 50,000 customers in Malmö, Sweden, at a direct repair cost of roughly $0.5m; the FAA's Wide Area Augmentation System disabled for about thirty hours; loss of vertical navigation guidance for satellite-based approaches exceeding ten hours; the loss of the Japanese ADEOS-II satellite; and extensive polar flight rerouting, with New York-Hong Kong flights burning an extra 26,600 pounds of fuel and shedding more than 16,500 pounds of payload.

Halloween 2003 remains the standard reference case for aviation loss modelling. It is also the last event before the orbital population began its eight-fold expansion, so it tells us nothing useful about the constellation era.

2.7 February 2022 to November 2025: the constellation era begins

- February 2022 — Starlink. A storm barely registering on the G-scale raised thermospheric density enough to deorbit 38 of 49 newly launched Starlink satellites before they could raise orbit. SpaceX deploys at very low insertion altitude by design, so that failed units deorbit quickly; that choice creates a window in which a modest density enhancement exceeds the vehicle's ability to climb out. The vulnerability follows from an operational decision rather than from storm severity, and it recurs on every launch.
- May 2024. The Gannon storm. The largest geomagnetic storm in two decades, at Kp 9 and Dst −412 nT, but only S1 on the radiation scale. Thermospheric density at 400 km reached six times its pre-storm baseline. Roughly half of the approximately 10,000 active LEO payloads manoeuvred, with decay rates reaching 180 m/day, and conjunction assessment became, in the study authors' words, very difficult or impossible during the storm and for days afterwards. A wall of ionospheric plasma produced GPS positioning errors up to 70 metres across the central United States for about six hours, costing US

Midwest farmers more than $500 million during the planting window. ICAO issued 148 advisories.

- October 2024. G4, Dst −341 nT, producing 266 ICAO advisories — more than May, despite lower intensity, because the GNSS-disturbance footprint was larger. At least one Starlink satellite reentered prematurely.
- November 2025. G4, Dst −160 nT. Aurora to Utah and Pennsylvania; magnetic declination at Fairbanks varied by 1.8 degrees. Aeromagnetic survey, directional drilling, satellite operations, GPS and radio were affected.

2.8 January 2026: the radiation event

An X1.9 long-duration flare erupted at 18:09 UTC on 18 January 2026, launching an Earth-directed CME that arrived in about 25 hours, implying a transit speed near 1,700 km/s against an initial model estimate of 1,400. High-energy particle levels peaked at 19:15 UTC on 19 January at **S4, severe** — placing the event among the most intense radiation storms in the GOES record. NOAA issued a G4 severe geomagnetic storm watch for 20 January. Some coverage described it as the most powerful solar storm since 2003.

ESA's assessment noted that infrastructure operators who were informed could activate mitigation actions, and no major confirmed damage has been reported. That outcome should be read carefully. The event is the first S4 radiation storm to occur with roughly 16,000 satellites in orbit, commercial GPUs operating in LEO, and four commercial station programmes in development. It arrived seven months before this paper, and its principal documented consequence is that operators who had been warned took precautions.

For the purposes of this model, January 2026 is a T2-class geomagnetic event with an S4 radiation profile. It is the closest thing the orbital economy has to a live-fire exercise, and the absence of published loss data from it is itself a finding — discussed in Chapter 6.

2.9 Where the cycle stands

Solar maximum passed in October 2024. 2025 produced 19 X-class flares against more than 50 in 2024, and 2026 is a declining-phase year. The strongest individual events historically occur during the declining phase rather than at maximum, so the fall in event count does not translate proportionately into reduced tail risk over the next three to four years.

The approaching solar minimum around 2030 matters in the opposite direction for orbital operators, and Chapter 6 covers why: a quiet Sun means a cooler, thinner thermosphere, less drag, and derelict objects that persist for years instead of months.

3. Frequency and Severity

Expected loss, capital requirements, the cost-benefit case for hardening and insurance pricing all depend on two parameters. One is disputed by a factor of twenty-six. The other was revised upward by a factor of two in July 2026.

3.1 The frequency disagreement

Riley's 2012 analysis, fitting a power law to the tail of the Dst distribution, estimated a 12 per cent probability per decade of a Carrington-magnitude event. It remains the most-quoted figure in the field. Lloyd's and AER, working from a different reconstruction, adopted a 1-in-150-year recurrence — about 6.5 per cent per decade — with a stated range of 100 to 250 years.

Moriña and colleagues (2019) used a Weibull counting process rather than a homogeneous Poisson process, conditioning on the absence of any Carrington-level event in the 58,000 days since 1859. Their estimate is 0.92 per cent per decade, with a bootstrap 95 per cent interval of 0.46 to 1.88 per cent — roughly thirteen times below Riley. They attribute the divergence to the Weibull model's greater flexibility and to the Dst tail being better fitted by an exponential than by a power law.

Three structural reasons keep the question open. The sample is one observed event, reconstructed indirectly, which cannot distinguish between tail forms that diverge by an order of magnitude in the region of interest. Dst is a proxy for the wrong quantity — damage follows dB/dt and the geoelectric field, and August 1972 caused a 64 per cent voltage collapse at Dst −154 nT. And Riley (2025) finds material inter-solar-cycle variability in extreme-storm occurrence, meaning the hazard rate is not constant, so both the Poisson and the simple Weibull models approximate a process with cycle-dependent intensity.

3.2 The severity revision

Sivadas and colleagues, publishing in Nature on 15 July 2026, address a separate question with direct consequences for the tail. Geomagnetic response has long appeared to saturate at high solar-wind driving, and ten competing physical theories have been proposed to explain why. Using over a million solar-wind measurements from 1995 to 2019 and a Monte Carlo error model, the authors show that the saturation is an artefact of regression to the mean: solar-wind measurement uncertainty is at least 30 per cent and varies with magnitude, and correcting for it leaves the relationship linear to the top of the observed range, up to 15 mV/m.

What the no-saturation result changes

If Earth's response does not saturate, the magnetosphere does not damp extreme storms, and the severity of the largest events has been systematically underestimated. The authors state that extreme geomagnetic storm impacts may be **twice as large as previously thought**.

Applying a 2× multiplier to T4 and T5 severity, holding frequency constant, takes a Carrington-class event from $2.4trn to **$4.8trn** and total annualised expected loss from $20.6bn to **$37.8bn**.

> The result is recent and has not been absorbed into commercial catastrophe models, national benchmarks or the NERC planning standard. It is carried here as a sensitivity case rather than in the base model; Appendix A.2 gives the joint frequency-severity grid.

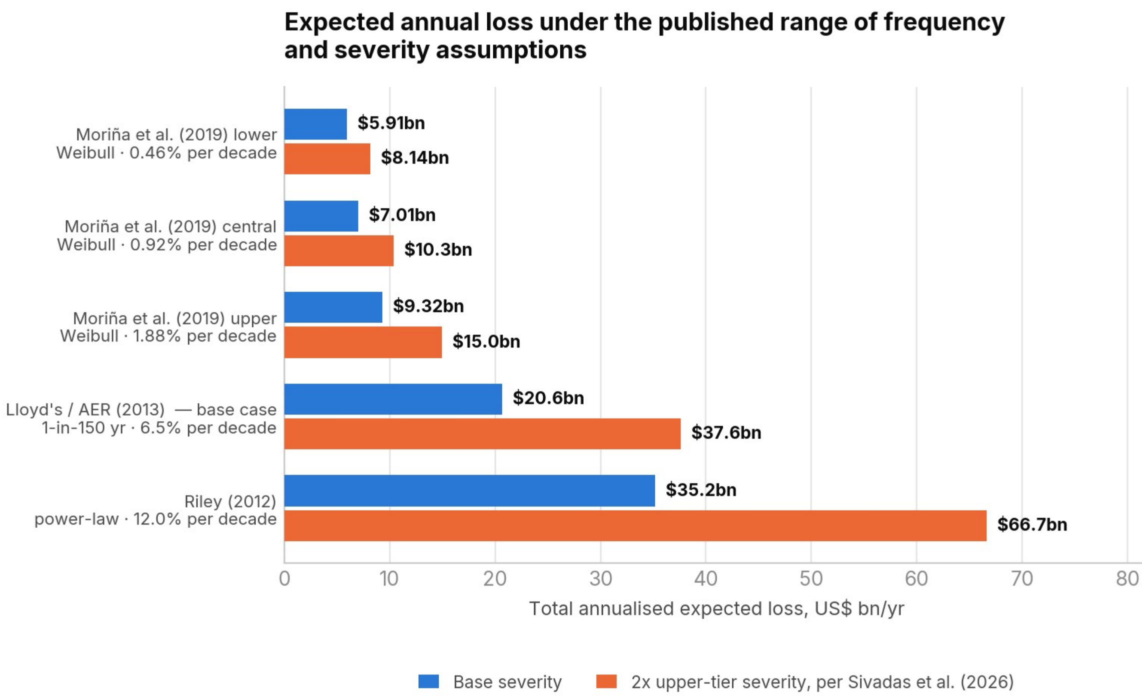

Exhibit 6. Joint sensitivity to superstorm frequency and severity Source: SSEEM v2.0 sensitivity run. Blue bars hold severity at the base case and vary only the Carrington-class recurrence rate. Orange bars additionally apply a 2× multiplier to T4 and T5 severity, following Sivadas et al. (2026). Appendix A.2 gives the underlying table.

The base case uses the Lloyd's/AER 1-in-150 figure, which sits near the geometric centre of the published range and is the assumption most widely used in commercial risk modelling. On the Weibull estimate, total annualised expected loss falls from $20.6bn to about $7.0bn. On Riley, it rises to $35.3bn. Combining Riley frequency with doubled severity gives $67.0bn.

3.3 The research priority

Resolving recurrence is worth more than any single hardening programme. The disagreement moves global annualised expected loss by roughly $29bn a year at base severity. A programme combining extended radiocarbon and ice-core surveys of historical events, systematic reanalysis of nineteenth-century magnetograms, expanded GIC monitoring, and physics-based rather than statistical extreme-event modelling would cost a small fraction of one year's disputed loss. Chapter 18 sets this out as the first recommendation for government readers.

A second, newer priority now sits alongside it: independent replication of the Sivadas result. A single paper overturning ten prior explanations is a strong claim, and the difference between it being right and wrong is $17bn a year.

PART II · THE EXPOSURE MODEL

4. Model Architecture

SSEEM v2.0 uses the standard catastrophe-modelling structure (hazard, vulnerability, exposure, financial) adapted to a peril with almost no loss history and an exposure base changing faster than the hazard. Nine channels, five severity tiers, forty-five cells, each with a low, central and high value and a confidence grade.

This chapter describes the architecture. Appendix A publishes the model itself: every one of the forty-five severity cells with its source family and grade, the frequency basis for each tier, the aggregation arithmetic in closed form, the nine growth multiples behind the 2035 projection, a joint frequency-severity sensitivity grid, and a section setting out in order which parameters a critic should attack and how far each moves the answer. The implementation is a single file of roughly 300 lines with no hidden state, published alongside this paper, and every table in Appendix A is generated from it.

The reason for that level of disclosure is specific to this peril. Thirty-one of the forty-five cells rest on structured judgement rather than published loss data, because operators do not publish what a storm costs them. A model built on that basis has no claim on the reader unless its judgements are visible and substitutable.

4.1 What changed from v1.1

1. Nine loss channels, six of them off-planet. The single satellite channel used in earlier versions of this work has been split into space infrastructure and connectivity; orbital compute and manufacturing; human spaceflight and commercial stations; space defence and national security; lunar and cislunar; and Mars and deep space. These have different owners, different failure modes, different insurance treatment and radically different growth rates. Aggregating them hides the story.
2. A 2035 projection. A static 2026 exposure base understates a sector growing at roughly twice global GDP. Each channel carries a growth multiple with a stated basis.
3. The Sivadas severity case. Carried as a sensitivity, not in the base.

4.2 Hazard layer

Tier	Intensity band	Analogue	Frequency	Basis
T1 Strong (G3)	Dst −100 to −200 nT	Routine cycle-maximum storm	~11.8 storm-days/yr	NOAA scales: 130 days/cycle
T2 Severe (G4)	Dst −200 to −350 nT	Oct 2024; Jan 2026 (S4)	~5.5 storm-days/yr	NOAA scales: 60 days/cycle
T3 Extreme (G5)	Dst −350 to −500 nT	May 2024 Gannon (−412 nT)	0.36/yr · 1-in-2.7 yr	NOAA scales: 4 days/cycle
T4 Québec-class	Dst −500 to −750	March 1989 (−589 nT)	0.020/yr · 1-in-50 yr	Lloyd's/AER (range 35–

Tier	Intensity band	Analogue	Frequency	Basis
	nT			70 yr)
T5 Carrington-class	Dst below −850 nT	September 1859 (~−900 nT)	0.0067/yr · 1-in-150 yr	Lloyd's/AER (range 100–250 yr)

The unit for T1 to T3 is the storm-day rather than the storm event. NOAA's scales publish both an event count and a days-per-cycle figure; the event count treats each three-hour Kp interval as an event, which overstates the number of distinct episodes an operator has to manage.

Tiers are keyed to Dst despite its limitations (3.1). Nearly all published frequency work is expressed in Dst, and adopting a different intensity metric would sever the model from its own sources. For the orbital channels this is a material simplification, because their dominant hazard is the S-scale rather than Dst. The January 2026 event sits at T2 on this ladder while being a top-tier radiation event. The tier severities for those channels are calibrated to include that radiation exposure, but the tier labels do not show it. This is the model's largest structural compromise and it is discussed again in 5.6.

4.3 Severity layer

Each of forty-five tier-channel cells carries a low, central and high global loss per event, a confidence grade, and a named source.

- H. a published figure applied at the scope for which it was estimated. Example: the $500m US Midwest agricultural loss from May 2024.
- M. a published figure scaled to a different geographic footprint, exposure base or year. Example: the Hydro-Québec 1989 economic loss scaled to a modern global exposure base.
- L. structured judgement, calibrated against adjacent published work, where no directly applicable published figure exists. Every orbital channel except the two anchored on documented Starlink events falls here, because the sector does not publish loss data.

One cell is graded H, thirteen M and thirty-one L. The L-graded cells contribute 23 per cent of central annualised expected loss in 2026 and 41 per cent in the 2035 projection. The channel ordering in this paper is robust; the absolute magnitudes of the orbital channels are indicative, and they are indicative because the underlying industry does not disclose.

4.4 The nine channels

Channel	Contents	Dominant mechanism
Space infrastructure & connectivity	Starlink, Kuiper, OneWeb, Guowang, Qianfan; GEO comms; Earth observation; the GNSS space segment	Thermospheric drag; charging; conjunction-assessment failure; service degradation
Orbital compute & manufacturing	Orbital data centres; in-space manufacturing; pharmaceutical and semiconductor production in microgravity	Single-event upsets and latch-up; total ionising dose; drag on large low-altitude platforms

Channel	Contents	Dominant mechanism
Human spaceflight & commercial stations	ISS to 2030; commercial LEO destinations; private astronaut missions; crew transport; tourism	SEP dose; shelter-driven stand-down; launch scrub; station reboost demand
Space defence & national security	Golden Dome and the pLEO interceptor layer; SDA tracking; missile warning; military SATCOM, HF and PNT	Radio blackout; radar false returns; navigation denial; sensor saturation; attribution ambiguity
Lunar & cislunar	Artemis; the lunar surface base; CLPS landers and rovers; LunaNet and Moonlight; ILRS	Unattenuated SEP and GCR; surface charging and dust; relay and PNT degradation
Mars & deep space	SR-1 Freedom; Mars Telecommunications Orbiter; Tianwen-3; Rosalind Franklin; cislunar SDA; asteroid missions	GCR accumulation over long transits; SEP events with no shelter and long comms latency
Power grids & cascading macroeconomy	Transmission and distribution assets; the business-interruption and supply-chain loss following an outage	GIC → transformer saturation, heating, relay misoperation, voltage collapse
GNSS-dependent ground economy	Precision agriculture, financial timing, telecoms sync, logistics, surveying, offshore and directional drilling	Ionospheric scintillation and TEC gradients degrading positioning and timing
Aviation	Airlines, air navigation service providers, passengers and freight	HF blackout on polar routes; GNSS approach unavailability; aircrew and passenger dose

4.5 Exposure growth to 2035

Each channel carries a growth multiple applied to 2026 severity. The multiples are the most contestable numbers in the model and are stated explicitly so they can be replaced.

Channel	2035 multiple	Basis
Orbital compute & manufacturing	25×	Near-zero 2026 base; market projected $1.77bn (2029) to $39.1bn (2035)
Lunar & cislunar	12×	Surface base operational; six-month landing cadence; ILRS construction from 2031
Mars & deep space	8×	SR-1, MTO, Tianwen-3 return, crewed-Mars precursor infrastructure
Human spaceflight & stations	6×	Multiple commercial stations reach continuous crew as the ISS winds down
Space defence & national security	5×	Golden Dome complete architecture mid-2030s; pLEO interceptor layer
Space infrastructure & connectivity	4×	Constellation population and revenue tracking the space economy to ~$1.8trn
GNSS-dependent ground economy	2×	Deeper penetration, net of multi-constellation receiver adoption
Power grids & cascading	1.4×	Electrification of final energy, net of TPL-007-era hardening

Channel	2035 multiple	Basis
macroeconomy		
Aviation	1.3×	Traffic growth, net of improved forecasting and routing

4.6 Exclusions

- Cascading infrastructure interdependency. Water, telecoms, healthcare and fuel distribution all fail when power fails. The models that resolve this describe their own estimates as lower bounds.
- **Mortality and morbidity**, terrestrial or in space. A crew casualty is discussed qualitatively in Chapters 8 and 11 and not monetised.
- Kessler-type debris cascade. A collision cascade triggered by storm-time conjunction-assessment failure would remove an orbital regime for decades. No defensible probability can be assigned; Chapter 10 treats it qualitatively.
- **Deliberate exploitation** of a degraded sensing and communications environment by an adversary.
- Events larger than Carrington. The 774 CE radiocarbon event is estimated at ten or more times Carrington's magnitude. Each of these exclusions biases the results downward.

5. Results

5.1 Annualised expected loss, 2026

Central annualised expected loss across all nine channels is $20.6 billion a year, within a range of $6.5bn to $69.2bn generated by taking the low and high severity anchor in every cell at once. That range is a scenario band rather than a confidence interval, and it excludes both the frequency disagreement in 3.1 and the severity question in 3.2.

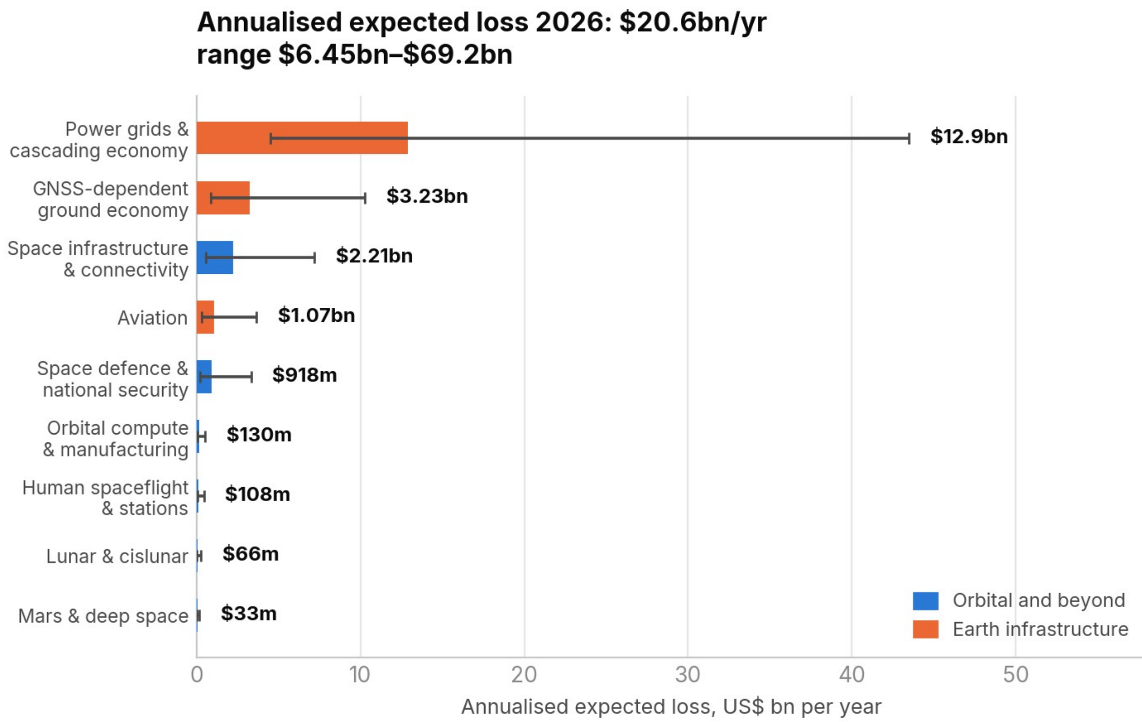

Exhibit 7. Annualised expected loss by channel, 2026 Source: SSEEM v2.0. Bars show central annualised expected loss; whiskers show the low-to-high range across all severity anchors.

Channel	Low	Central	High	Share	2035
Power grids & cascading macroeconomy	$4.52bn	$12.88bn	$43.49bn	62.4%	$18.03bn
GNSS-dependent ground economy	$0.85bn	$3.23bn	$10.26bn	15.7%	$6.46bn
Space infrastructure & connectivity	$0.53bn	$2.21bn	$7.20bn	10.7%	$8.83bn
Aviation	$0.28bn	$1.07bn	$3.63bn	5.2%	$1.40bn
Space defence & national security	$0.22bn	$0.92bn	$3.32bn	4.4%	$4.59bn
Orbital compute & manufacturing	$0.03bn	$0.13bn	$0.50bn	0.6%	$3.26bn
Human spaceflight & commercial stations	$0.02bn	$0.11bn	$0.43bn	0.5%	$0.65bn
Lunar & cislunar	$0.01bn	$0.07bn	$0.27bn	0.3%	$0.79bn

Channel	Low	Central	High	Share	2035
Mars & deep space	$0.01bn	$0.03bn	$0.14bn	0.2%	$0.26bn
Total	**$6.45bn**	**$20.65bn**	**$69.23bn**	**100%**	**$44.27bn**

Power grids remain the largest single channel at 62 per cent, and their tail still dominates the extreme scenarios. GNSS-dependent ground loss at $3.23bn exceeds space infrastructure at $2.21bn, which is the finding that most contradicts how resilience spending is currently allocated: the economy loses more from six hours of degraded positioning than from damage to the constellations providing it.

The four newest channels (orbital compute, human spaceflight, lunar and Mars) together account for $0.34bn a year, 1.6 per cent of the total. In 2026 they are genuinely small. That is the last year in which the sentence will be true.

5.2 The 2035 projection

Exhibit 1, in the executive summary, plots this shift: the left panel gives total expected loss and its split between orbital and terrestrial channels in each year, the right panel the growth multiple applied to each channel.

On the growth assumptions in 4.5, total annualised expected loss reaches $44.3bn by 2035, and the orbital and beyond-orbit channels move from 17 per cent of the total to 42 per cent. Space infrastructure at $8.8bn becomes the second-largest channel, ahead of GNSS-dependent ground loss. Orbital compute at $3.3bn overtakes aviation and space defence.

None of this requires the Sun to change. The hazard frequencies are held constant across both years. The entire movement is exposure — more satellites, more compute in orbit, more crew outside the magnetosphere, more capital committed to the lunar surface.

Two caveats apply symmetrically. The growth multiples are judgement, and a slower space economy compresses the shift. And the model holds vulnerability constant, when in practice a sector that is currently silent on space weather is likely to start hardening once it takes a loss. Chapter 20 sets out what would move these numbers in each direction.

5.3 Where the expected loss originates

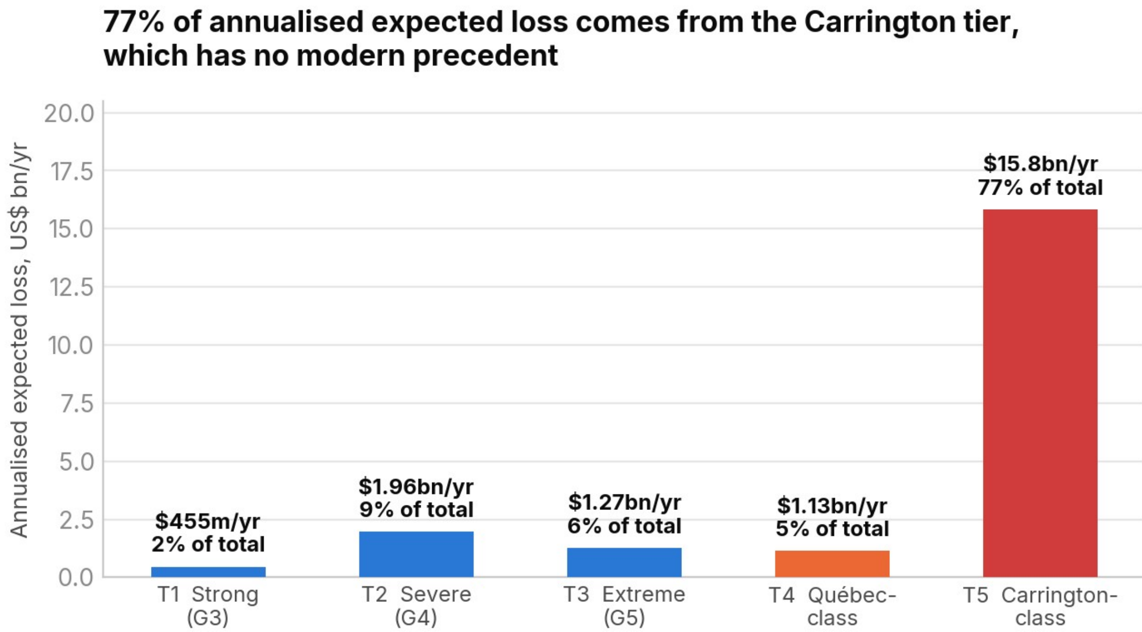

Exhibit 8. Annualised expected loss by severity tier Source: SSEEM v2.0. Percentages are shares of central annualised expected loss.

Seventy-seven per cent of expected loss comes from T5, the tier with no modern precedent. Seventeen per cent comes from the T1-to-T3 range, which constitutes the entire operational experience of everyone currently working in grid operations, satellite operations, air traffic management or space-weather forecasting.

Institutional learning is proportional to event frequency; economic exposure is inversely proportional to it. An organisation calibrated on what it has experienced is calibrated to 17 per cent of its exposure. For the orbital channels the problem is sharper, because most of the operators now flying have existed for less than one solar cycle and have never seen a G5 storm, let alone anything above it.

5.4 The scenario ladder

Channel	T1	T2	T3	T4	T5	T5 in 2035
Space infrastructure & connectivity	$7m	$63m	$590m	$8.2bn	$210bn	$840bn
Orbital compute & manufacturing	$0.5m	$4m	$35m	$500m	$12bn	$300bn
Human spaceflight & stations	$0.5m	$4m	$35m	$400m	$9bn	$54bn
Space defence & national security	$5m	$35m	$280m	$5.0bn	$70bn	$350bn
Lunar & cislunar	$0.2m	$2m	$20m	$250m	$6bn	$72bn
Mars & deep space	$0.1m	$1m	$10m	$120m	$3bn	$24bn
Power grids & cascading macroeconomy	$15m	$150m	$1.8bn	$28bn	$1,600bn	$2,240bn
GNSS-dependent ground economy	$6m	$80m	$600m	$12bn	$340bn	$680bn

Channel	T1	T2	T3	T4	T5	T5 in 2035
Aviation	$4m	$20m	$110m	$2.2bn	$125bn	$163bn
Total per event	**$38m**	**$359m**	**$3.5bn**	**$56.7bn**	**$2,375bn**	**$4,723bn**
Probability in any 10-year period	**certain**	**certain**	**97%**	**18%**	**6.4%**	**6.4%**

A Québec-class event carries an 18 per cent probability in any ten-year window and a $56.7 billion global cost. Six million people have already lived through one. It belongs in planning cases rather than in tail scenarios.

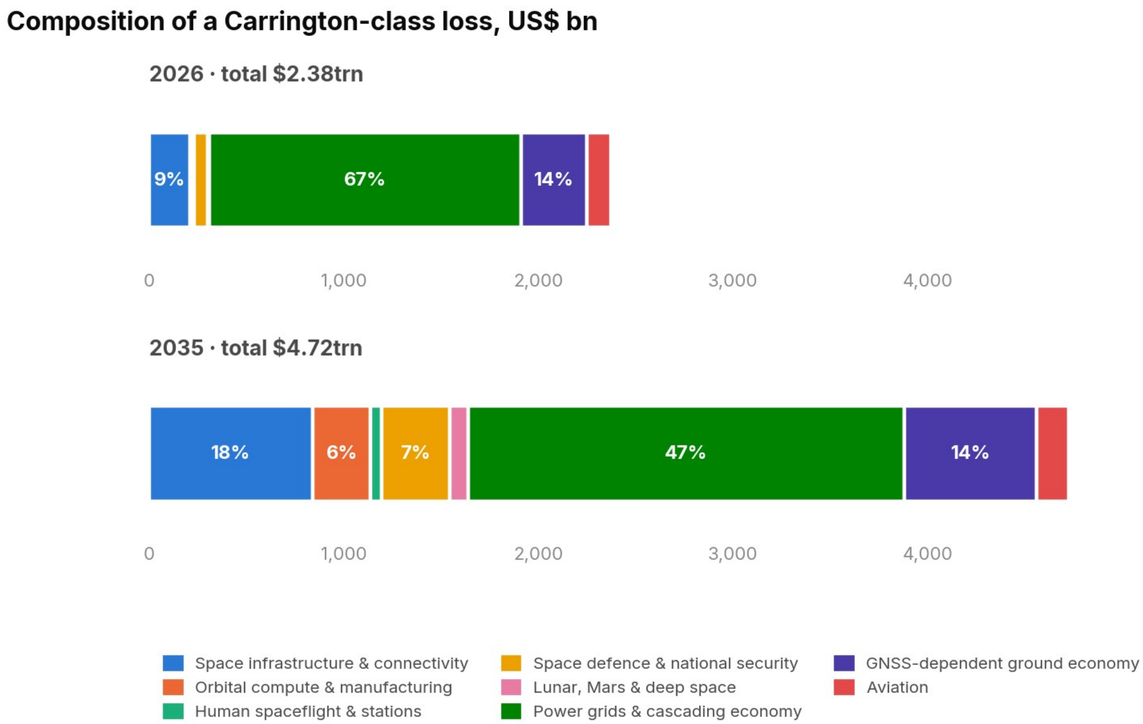

Exhibit 9. Composition of a Carrington-class loss, 2026 against 2035 Source: SSEEM v2.0. Lunar, Mars and deep space are combined into one slice to stay within eight distinguishable categories. The orbital share of a Carrington-class loss rises from 9 per cent in 2026 to 32 per cent in 2035.

5.5 Exceedance and reconciliation

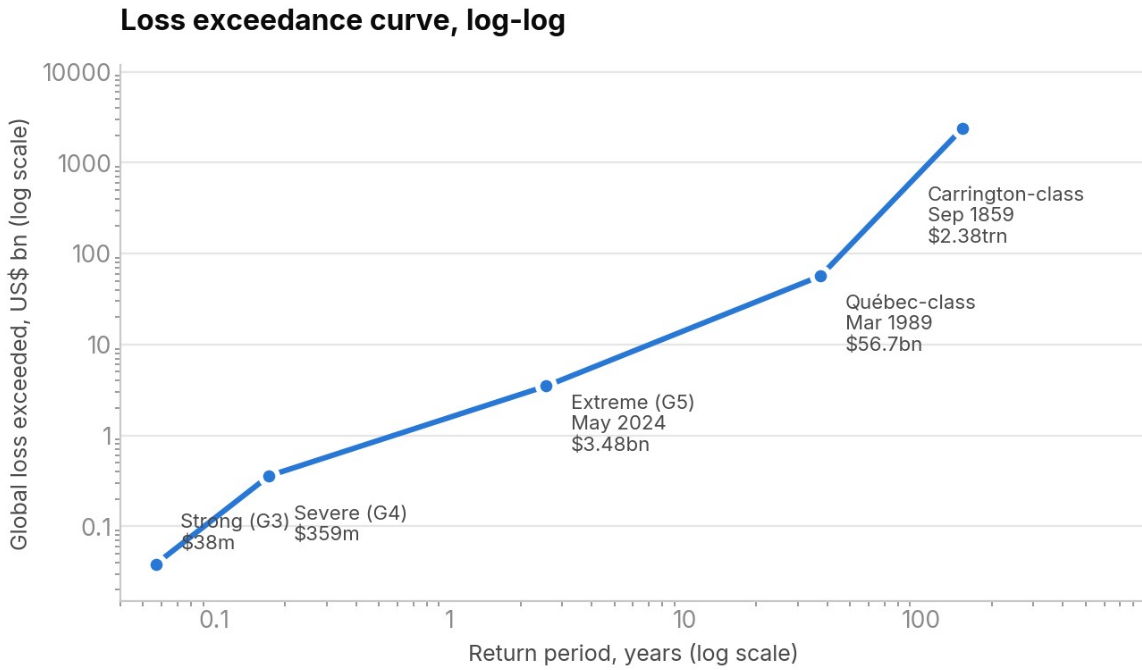

Exhibit 10. Occurrence exceedance probability Source: SSEEM v2.0. Occurrence exceedance probability, log-log. Return periods derived from cumulative tier frequencies.

Plotted log-log, the loss curve is close to linear across four decades of return period, which is the signature of a heavy-tailed exposure in which no single return period dominates. There is no elbow at which retention should stop and transfer should begin, which is one reason space-weather cover remains thinly traded.

SSEEM's T5 central estimate of $2.375 trillion sits close to Lloyd's severe-variant figure of $2.4 trillion. That is calibration, not independent confirmation: the T5 power-grid anchor was set against the Lloyd's/Cambridge scenario range. The independent comparisons hold. The North America share sits in the lower part of the Lloyd's/AER $0.6–2.6trn band, and the implied daily loss over an illustrative recovery period is consistent with Oughton's mid-range scenarios.

One divergence is worth flagging. Lloyd's probability-weights its $2.4 trillion severe scenario down to a $17 billion expected loss over five years, implying an annual rate near 1-in-700 — far below the 1-in-150 recurrence established in the same organisation's grid study. Readers using both numbers should know they do not come from the same model.

5.6 Limitations

1. The orbital channels rest on judgement, because the sector does not publish losses. Thirty-one of forty-five cells are graded L. No commercial constellation operator publishes storm-attributable propellant consumption, anomaly rates or revenue impact. The February 2022 Starlink loss is documented only because 38 satellites reentered visibly.
2. Dst tiering understates the radiation hazard. The orbital and beyond-orbit channels are driven by the S-scale, which correlates only loosely with Dst. January 2026 was a T2

geomagnetic event and a near-record radiation event. A future version should tier those channels on proton fluence separately.

3. Channel correlation is unmodelled. SSEEM sums channel losses at each tier as if additive. In a real T5 event, an aviation network disrupted by HF blackout is also disrupted by GNSS loss, at airports without power, while the satellites providing both are manoeuvring.
4. Growth multiples are judgement. They are stated per channel in 4.5 precisely so a reader can substitute their own.
5. Adaptation is not projected. NERC compliance, neutral-blocking deployment, multi-constellation receivers, radiation-tolerant orbital compute and improved forecasting all reduce vulnerability over time.
6. Regional resolution is coarse. Geoelectric hazard is intensely geographic; orbital hazard is intensely altitude- and inclination-dependent. Neither is resolved here.

PART III · THE ORBITAL ECONOMY

6. Space Infrastructure and Connectivity

16,393 active satellites, two-thirds of them Starlink, against 1,980 in 2018. Expected loss is $2.21bn a year, rising to $8.83bn by 2035. The dominant mechanism is not damage but drag: a storm heats and expands the thermosphere, orbits decay, orbit predictions fail, and the fleet manoeuvres en masse. The sector's most significant space-weather decision (SpaceX lowering roughly 4,400 satellites by 70 km) is being taken in anticipation of the coming solar minimum, not a storm.

6.1 The exposure base

The global space economy reached $429 billion in 2025, of which the commercial satellite industry was $303 billion. Ground equipment at $165.2bn and satellite services at $105.0bn dwarf manufacturing at $20.4bn and launch at $12.4bn. Most of the value sits in the signal and the terminals using it, which is why the GNSS-dependent ground economy is modelled as a separate and larger channel in Chapter 16.

Measure	Value	Date
Active satellites in orbit	16,393	24 August 2026
Of which in LEO	15,580	
Of which Starlink	10,971 (67%)	19 August 2026
US-registered total	12,348	
China	1,477	
UK (largely OneWeb)	697	
Satellites launched during 2025	4,434 across 296 launches	SIA (2026)
Starlink subscribers	~12 million	2026
Starlink revenue run-rate	~$17bn annualised	2026

Two features of this table drive the risk profile. The population has grown roughly eight-fold in eight years, so nearly all of it was designed, launched and operated by organisations with no institutional memory of a severe storm. And the concentration is extreme: a single operator holds two-thirds of everything in orbit, which means one company's design choices determine most of the sector's aggregate exposure.

6.2 Drag: the dominant mechanism

Geomagnetic storms deposit energy into the thermosphere, which heats and expands, raising density at satellite altitudes. During the May 2024 Gannon storm, density at 400 km reached six times its pre-storm baseline within twelve hours. Drag scales linearly with density, so orbital decay rates can quadruple or worse: the KANOPUS-V 3 spacecraft went from about 38 m/day to 180 m/day.

The February 2022 Starlink loss remains the clearest case, and its lesson is counterintuitive. The storm was minor. It destroyed 38 of 49 satellites because the deployment profile places them at very low insertion altitude by design, so that failed units deorbit quickly. A modest density enhancement exceeded their ability to climb out. The vulnerability follows from a business-model choice rather than from storm severity, and it recurs on every launch into a storm window.

The consequence for an operator is continuous rather than episodic: propellant consumed on unplanned station-keeping, service degradation during manoeuvres, shortened mission life, and replacement launches brought forward. None of this is published, which is why the severity anchors for this channel are graded M and L.

6.3 The solar-minimum problem

SpaceX is lowering 4,400 satellites because the Sun is going quiet

On 1 January 2026, Michael Nicolls, Starlink's VP of engineering, announced that SpaceX would lower roughly **4,400 satellites from about 550 km to about 480 km over the course of 2026 —** around half the operational constellation.

The stated driver is the solar minimum expected around 2030. At solar maximum a hot, expanded thermosphere drags derelict objects down within months. At solar minimum the atmosphere contracts, drag falls, and a failed satellite persists far longer. Nicolls quantifies the gain as **a greater than 80 per cent reduction in ballistic decay time at solar minimum — from four-plus years down to a few months**.

The secondary reason is that debris density and planned constellation traffic are both markedly lower below 500 km.

This is a deliberate, fleet-scale, multi-year capital decision taken on space-weather grounds, and it is the clearest evidence in the sector that at least one operator models the environment properly. It also runs opposite to the intuition that a quiet Sun means low risk: for a constellation operator, the storm years and the quiet years each carry their own failure mode.

The general point generalises beyond one company. Any operator whose disposal plan relies on atmospheric drag has a disposal timeline that varies by a factor of several across the solar cycle, and most published post-mission disposal compliance assessments do not state which part of the cycle they assume.

6.4 Charging, single-event effects and array degradation

Surface charging from hot plasma injection produces differential potentials of thousands of volts, and the resulting electrostatic discharge damages surfaces, degrades solar arrays and injects noise into avionics. Deep dielectric charging from energetic electrons (the population that builds in the outer radiation belt over days following a storm) deposits charge inside insulating materials, causing internal discharge with no external signature. Single-event upsets from solar energetic particles flip memory bits and can latch up devices.

Documented instances: TDRS-1 logged over 250 anomalies in March 1989; Intelsat IV F-2 lost 5 per cent of its array output in August 1972, two years of normal degradation in one event; a DSCS II satellite suffered a mission-ending power failure; ADEOS-II was lost in 2003.

Commercial small satellites are more exposed to this than the traditional GEO fleet. Cost pressure has pushed operators toward commercial-grade rather than radiation-hardened parts, with redundancy at the constellation level substituting for redundancy at the spacecraft level. That is a rational trade against random failures and a poor one against a correlated event.

6.5 What the sector says about space weather

For this paper we reviewed roughly thirty-five public sources across the commercial orbital sector — company announcements, product pages, funding releases, agency contract notices and programme descriptions covering commercial stations, orbital data centres, in-space manufacturing, satellite servicing, and non-US constellations. The purpose was to establish how the industry describes this risk to its own investors and customers.

The disclosure gap

Radiation, solar storms, orbital drag and space weather are **almost never named as risks** in the sector's public material. Design documentation emphasises thermal management, power, launch cadence, propulsion and regulatory approval.

The consistent exceptions are: **SpaceX**, whose 2022 loss is documented and whose constellation-lowering decision is explicitly space-weather-driven; **Google's Project Suncatcher**, which published TPU radiation-tolerance testing showing the Trillium part withstood nearly three times its expected five-year mission dose before irregularities appeared; and **NASA and ESA**, whose crewed and lunar programmes treat it as a design driver throughout.

The absence is not evidence that operators ignore the hazard internally. It is evidence that the hazard is not being priced, disclosed or underwritten, which is the condition under which a correlated loss arrives unreserved.

The January 2026 S4 event tests this directly. It was among the most intense radiation storms in the GOES record, it occurred with 16,000 satellites in orbit, and seven months later there is no public accounting of what it cost anyone. ESA's summary is that operators who were informed could activate mitigations. That may be the whole story. It may also be that the losses were absorbed quietly as anomaly statistics and propellant budgets, in which case the sector has no basis for pricing the next one.

6.6 Insurance

Metric	Value	Implication
Launch and in-orbit premium volume	~$1.35bn (2024)	Against a $429bn space economy
Total underwriting capacity	$550–700m/yr	Below the insured value of a single large GEO cluster
2023 loss ratio	179%	$557m premium against $995m claims —

Metric	Value	Implication
		worst in over two decades
Active LEO satellites with in-orbit cover	<1%	~300 of ~10,000, overwhelmingly GEO
In-orbit annual rate	~1.2%	Roughly doubled from 0.6%

Fewer than one per cent of LEO satellites carry in-orbit cover, so a severe storm's satellite losses fall almost entirely on operators' balance sheets. The 2023 loss year, driven by Viasat-3 Americas, Inmarsat 6-F2 and SES O3b mPower rather than by space weather, pushed Brit, Allianz, AIG and Swiss Re to cut or exit space underwriting: capacity contracted while the exposure base grew. Trade reporting in January 2026 described a Munich Re satellite product addressing solar-flare risk; the report has not been confirmed by a Munich Re publication and is treated here as an indication of underwriter interest rather than an established market instrument. Either way, market capacity remains an order of magnitude below a severe-storm loss.

7. Orbital Compute and In-Space Manufacturing

The newest channel and the fastest-growing: $130m a year in 2026, projected 25-fold to $3.26bn by 2035. Nvidia H100s and Google TPUs are already operating in orbit. The concepts being funded run to gigawatt scale. This is the first industry to put the most radiation-sensitive commercial hardware ever built into the environment least suited to it, and it is doing so for a reason (orbital solar arrays are five to eight times more productive than terrestrial ones) that has nothing to do with radiation tolerance.

7.1 What is already flying

Programme	Status	Detail
Starcloud-1	Launched November 2025	Carried an Nvidia H100 GPU. Starcloud's claim that this is 100× more compute than anything previously operated in space is the company's own and has not been independently verified. Starcloud states it trained an LLM in orbit and ran a version of Gemini on-orbit.
Axiom orbital data centre nodes	Deployed 11 January 2026	Two nodes in LEO on Kepler Communications' optical relay network; 2.5 Gbps inter-satellite links, 10 Gbps maximum throughput.
Starcloud-2	Slipped from October 2026 to 2027	Multiple H100s with Blackwell-generation integration.
Google Project Suncatcher	Learning mission early 2027	TPUs in orbit with Planet Labs. Target architecture is an 81-satellite compute cluster within a 1 km radius.
ASCEND (EU)	Feasibility study complete; demonstration proposed for 2028	The 1 GW-before-2050 target is Thales Alenia Space's study conclusion, not a funded programme objective. Figures circulating for committed European funding are not traceable to an EU or ESA document and are omitted here.
Starcloud long-term	2030s	5 GW stated vision.
Starlink V3	Announced; no operational unit in orbit	SpaceX specifies 1 Tbps per satellite, roughly 10× V2 mini, at ~60 per Starship launch. These are design figures for a satellite that depends on Starship reaching operational cadence; none is yet providing service.

The economic case rests on power and cooling. Space solar arrays generate five to eight times more energy than terrestrial installations, and panels in dawn-dusk sun-synchronous orbit can be up to eight times more productive without batteries. Against that, launch costs of $1,500–2,500/kg must fall below roughly $200/kg for parity with terrestrial data centres. The $200/kg threshold is Google's own published break-even estimate for Project Suncatcher, and the mid-2030s date by which the industry expects to reach it is an industry expectation rather than a demonstrated trajectory. The market projections quoted in this sector, running from $1.77bn in 2029 to $39.1bn by 2035, are commercial forecasts from BIS Research and comparable vendors; they are cited here to show the direction of expectation, not as evidence of realised revenue. SSEEM's exposure base uses in-orbit assets and contracted programmes, not these forecasts.

7.2 Why this hardware is the wrong hardware for this environment

A modern AI accelerator is optimised for density and power efficiency. Feature sizes are small, operating voltages are low, and on-die memory is vast. Every one of those properties increases susceptibility to single-event effects. A high-energy proton passing through a dense logic die deposits enough charge to flip a bit or, in the worst case, trigger a latch-up that destroys the device unless power is cycled quickly.

Radiation-hardened processors exist, and they are typically one to two decades behind the commercial state of the art in performance per watt. The entire premise of orbital compute is to fly current-generation commercial silicon. The industry is therefore committed to accepting single-event effects and managing them architecturally — through error-correcting memory, checkpointing, redundancy across nodes, and power-cycle recovery.

Google's published Trillium TPU testing is the useful data point: the part withstood nearly three times its expected five-year mission dose before irregularities emerged. That is a total-ionising-dose result, and it is encouraging. It is not a single-event-effect result, and it is not a result from an S4 storm.

> **The failure mode nobody has priced**
>
> In an S4 or S5 radiation storm, proton flux at LEO rises by three to five orders of magnitude for hours to days. Upset rates rise with it.
>
> For a conventional communications satellite this produces anomalies that ground control works through. For a compute cluster running paid workloads it produces something different: **correlated, simultaneous corruption across every node in the constellation**, precisely because the nodes are co-located by design — Suncatcher's target architecture puts 81 satellites within one kilometre.
>
> Terrestrial data centres survive on the assumption that failures are independent and localised. That assumption does not hold in orbit during a radiation storm, and the architectural mitigations that work for random upsets (cross-node redundancy, retry, failover) work least well when the whole cluster is hit at once.
>
> SSEEM places this channel at $12bn in a Carrington-class event in 2026 and **$300bn by 2035**. The 2035 figure carries the widest relative uncertainty of any cell in the model.

7.3 Drag, at the altitude and scale this needs

Orbital compute wants low altitude for launch cost and communications latency, and large deployed area for solar collection and thermal radiation. Those two preferences multiply the drag problem: drag force scales with area, and density at low altitude is exactly where storm enhancement bites hardest. A 1 km-scale array cluster at low altitude is, from an orbital-mechanics standpoint, a very large sail in a variable atmosphere.

The February 2022 Starlink loss involved compact, mass-produced spacecraft with substantial propulsion margin. A power-and-radiator-dominated compute platform has a far worse area-to-mass ratio, and the sector has not published how it intends to hold altitude through a Gannon-class density enhancement.

7.4 In-space manufacturing

A smaller but related exposure. Varda Space Industries has flown a series of W-series reentry capsules producing pharmaceutical crystals in microgravity; Space Forge has flown ForgeStar-1 targeting semiconductor substrate production. The value proposition is materials that cannot be made in gravity, and the products are high-value-per-kilogram by construction.

The space-weather exposure here is mostly schedule and reentry rather than radiation: a storm-driven change in atmospheric density shifts the reentry corridor and the landing dispersion for a capsule returning from orbit, and the recovery operation is licensed against a predicted footprint. This is a manageable problem and a small one today. It becomes less small as the sector scales.

8. Human Spaceflight and Commercial Stations

$108m a year in 2026, projected six-fold to $651m by 2035 as commercial stations take over from the ISS. The loss function is not asset damage: it is stood-down operations, scrubbed launches, slipped missions and consumed crew career dose. The tail case is a fatality, whose programme consequences exceed any modelled figure. The commercial station sector has published almost nothing about how it will handle a solar particle event.

8.1 The transition

ISS operations are authorised through 2030, and NASA holds a contract with SpaceX for a US Deorbit Vehicle at $843 million. That end date is no longer settled. NASA's March 2026 Ignition realignment restructured the Commercial LEO Destinations programme, and GAO has documented the transition decisions still pending, including whether to extend station operations past 2030 while commercial replacements mature. The report treats the retirement date as a range rather than a fixed year, because the exposure question (how many crewed platforms are flying, and under whose radiation procedures) is unchanged by which year is chosen.

Commercial LEO Destinations is funding replacements with roughly $550m of public money against more than $3bn of private capital. Every date in the table below is a target announced by the operator or by NASA, not a demonstrated capability; commercial station schedules have slipped repeatedly and should be read as intent.

Programme	Status	Detail
ISS	Authorised to 2030; extension under review	SpaceX US Deorbit Vehicle contracted at $843m. Retirement date reopened by the March 2026 Ignition realignment; GAO records the decision as pending.
Vast Haven-1	NET Q1 2027	Over $1bn raised. Single-module commercial station; crewed visits on Crew Dragon.
Axiom Station	In development	$350m raise February 2026. Modules initially berth to ISS, then detach as a free flyer.
Starlab	CCDR complete Feb 2026	Voyager Technologies-led; single large module launched on Starship.
Orbital Reef	In development	Blue Origin and Sierra Space.
Private astronaut missions	PAM-5 and PAM-6 targeted for 2027	Commercial crewed flights to the ISS.

The ISS carries three decades of accumulated operational practice for radiation events: dose monitoring, defined shelter locations in the most shielded modules, and crew trained on the procedure. None of that transfers automatically. A single-module commercial station with a smaller pressurised volume has fewer places to put mass between crew and particles, and the tourism-oriented business models involve crew with no professional radiation training on short-duration visits.

8.2 The dose environment

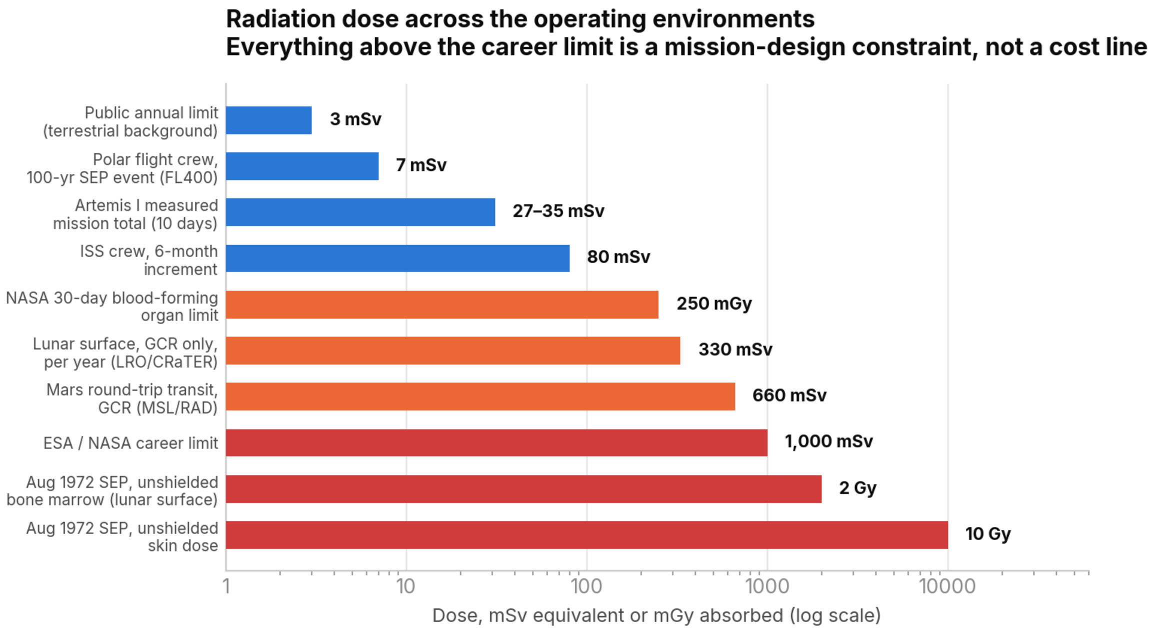

Exhibit 11. Radiation dose against exposure limits, by environment Sources: LRO/CRaTER lunar surface measurements; NASA spaceflight radiation standards; Hapgood et al. (2021) for aviation; the Artemis I measured totals from Nature (2024); MSL/RAD for Mars transit; Physics Today for the August 1972 reconstruction. Logarithmic scale. The first four bars are effective dose in mSv; the NASA 30-day limit and the 1972 figures are absorbed dose in mGy.

NASA's 30-day limit for blood-forming organs is 250 mGy; the 30-day skin limit is 1,500 mGy. The career limit is derived from a mission-specific 3 per cent risk of exposure-induced death, a formulation the National Academies has criticised as difficult to communicate and inconsistent with terrestrial occupational practice. ESA uses a 1,000 mSv career limit.

LEO is comparatively benign: the magnetosphere deflects most solar particles, and an ISS six-month increment delivers roughly 80 mSv. The exposure is not zero (the South Atlantic Anomaly drives most of it, and a severe SEP event raises polar-passage dose materially) but LEO crew are shielded in a way that lunar and Mars crew are not.

The Artemis I measurements, published in Nature in 2024, give the best available deep-space calibration. Total mission dose equivalent was 26.7 to 35.4 mSv, of which radiation-belt passages contributed 1.8 to 3.9 mSv, against deep-space GCR of 0.96 to 1.24 mSv/day. Position inside the vehicle mattered by a factor of four: 69 µGy/min at the most shielded location during the proton-belt transit against 287 µGy/min at the least shielded. A deliberate reorientation during one burn cut measured levels nearly in half.

8.3 The loss function

1. Launch scrub and slip. Space-weather constraints are a launch commit criterion. A scrub costs propellant, range time and schedule.

2. Operational stand-down. Sheltering crew suspends the activity the mission exists to perform. For a commercial station selling research time or tourist days, this is directly revenue-linked in a way it never was for the ISS.
3. EVA cancellation. NOAA's S3 threshold recommends radiation avoidance for astronauts on EVA, and S3 events occur about ten times per solar cycle.
4. Reboost demand. A station in LEO loses altitude faster during a storm, and the propellant to recover it is a recurring operating cost. The ISS has absorbed this for decades with government resupply. A commercial operator carries it on a P&L.
5. Career dose consumption. For a small and expensively trained professional population, dose is a capital asset. For paying visitors it is a liability question nobody has litigated.

SSEEM places this channel at $35m for a T3 event, $400m for a T4 and $9bn for a T5, all graded L. The T5 figure depends heavily on whether crew are aboard and where, which is a coincidence probability rather than a severity estimate.

8.4 What the sector has not addressed

None of the four commercial station programmes has published a radiation shelter specification, a dose budget for visiting non-professional crew, or a stated policy on suspending operations during an S3 or higher event. Nor has any published a storm-time reboost reserve.

These are answerable questions with known engineering content (the ISS answers all of them) and the absence of public answers at this stage of development is the same disclosure gap documented in 6.5, in a domain where the consequence is a person rather than a spacecraft.

9. Space Defence and National Security

$918m a year, rising five-fold to $4.59bn by 2035 as Golden Dome builds out. The programme places missile-defence interceptors in proliferated low Earth orbit (the orbital regime most affected by storm-driven drag) at a stated cost of $185 billion. The dominant risk remains what it was in 1967: misattribution. A degraded sensor that appears to have been degraded deliberately is an escalation pathway, and the modern architecture multiplies the number of sensors that can degrade at once.

9.1 The architecture being built

Element	Detail
Golden Dome, total programme	$185bn is the administration's stated programme cost, not an appropriation or an independent estimate; CBO's published range for a space-based interceptor layer alone is far higher. Initial capability is targeted for 2028 and the complete architecture for the mid-2030s. FY2027 request $17.5bn.
Space-based interceptors	A planned constellation in proliferated LEO carrying kinetic interceptors to engage missiles in boost, midcourse and glide phases. Prototype demonstration targeted for 2028. No interceptor has been flown.
SBI contractors	Twelve companies hold other-transaction agreements under a $3.2bn shared ceiling. The ceiling is the maximum value of the vehicle, not money awarded or obligated; individual task orders are a fraction of it. Anduril, Booz Allen Hamilton, General Dynamics Mission Systems, GITAI USA, Lockheed Martin, Northrop Grumman, Quindar, Raytheon, Sci-Tec, SpaceX, True Anomaly, Turion Space.
SDA tracking layer	36 accelerated missile-warning and tracking satellites for Tranche 3, launching by end-2028.
Space Force budget	The FY2026 base request of about $26.3bn was a reduction of roughly 9 per cent on FY2025 enacted. The headline growth reported at the time came from reconciliation funds added outside the base request, which lift the FY2026 total above $40bn. Quoting either number without the other misstates the direction of travel.
Cislunar Coordination Office	Announced 15 April 2026 at Space Symposium. Mandate covers roadmaps and acquisition for cislunar operations supporting a lunar base by 2030.
Department of War	Operating under the name since a September 2025 executive order. On 22 July 2026 the House passed its FY2027 NDAA including language formally renaming the department, 216–212, the first time either chamber has legislated the change. The Senate has not yet acted and it is not yet law. Roughly $50m spent on implementation; CBO estimates $125m or more.

The strategic logic of proliferated LEO is resilience through numbers: an adversary cannot attrit a constellation of hundreds the way it can a handful of exquisite satellites. That logic is sound against kinetic and directed-energy attack. It offers no protection against space weather, which is the one threat that reaches every satellite in the constellation simultaneously and by the same mechanism.

Proliferation trades one correlation for another

A proliferated LEO architecture is designed so that no single loss matters. Its failure modes are therefore assumed to be independent.

Storm-driven drag is not independent. It acts on every satellite in the shell at once, in proportion to area-to-mass ratio, which for a constellation of identical buses is identical. So is the radiation environment during an SEP event.

For an interceptor constellation the operational consequence is specific: during and after a severe storm, orbit knowledge degrades across the whole layer at once, exactly as it did for the commercial fleet in May 2024. An interceptor whose position is uncertain by kilometres cannot be committed against a boost-phase target.

SSEEM does not attempt to price a degraded-capability window in a missile-defence system. The figure in this chapter covers communications, sensing, PNT and asset costs only.

9.2 Four operational effects

1. HF and SATCOM denial. Military HF remains essential for beyond-line-of-sight communication where SATCOM is unavailable, contested or under emission control — submarine and maritime operations, special forces, and the high Arctic where geostationary coverage fails. Solar flares deny HF across the sunlit hemisphere with no warning; polar cap absorption can deny it at high latitude for days. The March 1989 storm cut HF for Australian peacekeepers in Namibia entirely.
2. Radar degradation and false returns. In 1967 solar radio emission entered BMEWS receivers directly and presented as jamming. Over-the-horizon radar, which depends on ionospheric refraction, degrades more fundamentally: the propagation medium becomes unpredictable. The November 2015 European air traffic control incident shows the same physics affecting civil secondary surveillance radar.
3. Navigation denial. GPS-guided munitions, platform navigation and timing-dependent networked systems all degrade under ionospheric scintillation. The May 2024 error magnitudes, up to 70 metres, exceed the accuracy requirements of most precision-guided weapons. The Defense Science Board and GAO have both documented DoD's PNT dependency and slow progress on complementary systems.
4. Direct effects on ordnance and platforms. The August 1972 detonation of dozens of magnetic-influence sea mines off Hòn La within thirty seconds remains the only confirmed case of space weather destroying deployed munitions. Any system with a magnetometer in its trigger or sensing chain inherits the failure mode.

9.3 Attribution

Military systems are designed to degrade and commanders are trained to operate through degradation. The risk is that degradation with a natural cause is indistinguishable in real time from degradation with a hostile one.

In May 1967, jamming of missile-warning radar was itself treated as an act of war, and nuclear-armed alert aircraft were placed in ready-to-launch status before a solar forecaster's report reached the NORAD command post. The resolution depended on a specific individual with domain knowledge being on duty, and on someone thinking to ask whether the affected sites were in sunlight.

The modern environment is worse in three respects. Decision timelines are shorter. Far more systems can degrade simultaneously. A proliferated architecture means hundreds of sensors reporting anomalies at once, which destroys the coincidence pattern that once suggested a natural cause. And a plausible adversary capability now exists to produce similar effects deliberately, so "it is probably the Sun" is no longer the natural first hypothesis it was in 1967.

The 2023 Space Weather Implementation Plan assigns DOC, DOD and NASA a joint action to implement a satellite-anomaly attribution information system. That is the direct institutional answer to the 1967 problem. It remains an action item rather than a capability in service.

9.4 The competitive dimension

Space weather now sits inside a strategic competition rather than alongside it. China operates 1,477 satellites, is building the Guowang and Qianfan constellations, maintains the Queqiao lunar relay architecture, and targets a crewed lunar landing before 2030. The International Lunar Research Station has partners across Asia, Africa and Latin America and enters its construction phase from 2031.

Two consequences follow. First, a severe storm degrades both sides at once, and the side that understands what is happening faster has a temporary advantage. Space-weather situational awareness is therefore a competitive capability, not a utility function. Second, the pace of the competition creates schedule pressure that works directly against the conservative launch-commit and shelter criteria that space weather demands. A programme racing a rival to a lunar landing has an incentive to fly through a marginal forecast.

Allied capability is uneven. The Canadian RCAF strategy identifies space weather as a key infrastructure vulnerability. NATO has no dedicated space-weather command function. The effects are hemispheric (a flare denies HF to every alliance member on the sunlit side at once) and a shared understanding of what is happening is exactly what the attribution problem requires.

10. Orbital Traffic and the Conjunction Problem

The risk in this chapter is not modelled as a dollar figure because its consequence is not a dollar figure. A severe storm degrades orbit prediction and forces mass manoeuvring simultaneously, which is the specific combination that turns a crowded orbital regime into a cascading one. Peer-reviewed analysis puts the interval between losing manoeuvre capability and a likely catastrophic collision at 5.5 days on 2025 orbital population, against 164 days under 2018 conditions.

10.1 What happened in May 2024

Roughly half of the approximately 10,000 active LEO payloads manoeuvred within a few days, most autonomously in response to drag. Thousands of debris objects lost several kilometres of altitude. The study authors concluded that the combination of unpredictable drag and bulk manoeuvring made it very difficult or impossible to identify potential conjunctions during the storm and for days afterwards.

Atmospheric-density forecasts fail at the same moment they are most needed. Geomagnetic index forecasts during Gannon were underpredicted by 100 to 300 units through the growth phase, and orbit propagation depends on those forecasts. The industry loses its ability to predict where objects will be exactly when everything moves unpredictably and simultaneously.

Operator requirements are well ahead of capability. The Space Weather Advisory Group's 2024 user-needs survey found that space traffic management operators need 3-day forecasts for collision avoidance, 7-day forecasts for manoeuvring decisions, and 5-year solar cycle predictions for mission planning. Sixteen participants contributed to that assessment.

10.2 The margin

5.5 days

Thiele, Heiland, Boley and Lawler, published in Acta Astronautica in 2026, estimate that if operators lost the ability to command avoidance manoeuvres, a catastrophic collision would become likely within approximately **5.5 days** on the 2025 orbital population. Under 2018 conditions the same calculation gave **164 days**. The authors call this the CRASH clock.

Losing control for a single day carries roughly a **17 per cent chance** of a major collision.

Within the Starlink constellation alone, close approaches under one kilometre occur roughly **every 22 minutes**, and each satellite performs an average of **41 course corrections a year**.

A severe geomagnetic storm is one of the few plausible mechanisms for simultaneously degrading command capability, corrupting orbit prediction and forcing mass manoeuvring. That is the pathway by which a space-weather event becomes a debris-cascade event.

The 5.5-day figure assumes total loss of manoeuvre command, which a storm does not produce. What a storm produces is worse in one respect and better in another: manoeuvre capability is retained, but the information needed to manoeuvre correctly is degraded, and every operator is manoeuvring at once on bad information. Whether that is more or less dangerous than doing nothing is genuinely unclear, and no published analysis answers it.

10.3 Why this is not in the loss number

A Kessler-type cascade in the 500–600 km shell would remove the orbital regime that carries Starlink, most Earth observation, the proposed orbital compute clusters and much of the planned defence architecture. The loss would be measured in the multi-decade unavailability of the regime and the sunk cost of everything in it, not in replacement value.

SSEEM does not price it, because assigning a probability requires knowing the conditional likelihood that a specific storm produces a specific collision that produces a specific cascade, and none of those three conditionals has a defensible published estimate. It is recorded here because a loss model that omits it should say so, and because it is the single scenario in this paper with the potential to exceed every figure in it.

10.4 What reduces this exposure

- Better thermospheric density forecasting. The binding constraint is not sensing but modelling. Operators need 3-day and 7-day drag forecasts and currently have neither at usable accuracy through a storm.
- Shared, machine-readable ephemerides during storms. Deductible waivers for contributing to the Space Force's Unified Data Library already exist. The collective benefit during a storm (when everyone's predictions degrade at once) exceeds the individual competitive cost by a wide margin.
- Disposal plans that state a solar-cycle assumption. A post-mission disposal timeline computed at solar maximum and executed at solar minimum can be wrong by years.
- Altitude discipline. SpaceX's decision to lower its primary shell (6.3) is the sector's clearest example of taking the orbital-lifetime problem seriously at fleet scale.
- Pre-agreed storm-time coordination protocols. There is currently no agreed procedure for how operators sequence manoeuvres when everyone needs to move and nobody can predict where anyone will be.

PART IV · BEYOND EARTH ORBIT

11. The Moon

The Moon has no magnetosphere and effectively no atmosphere, so it receives the full solar particle flux unattenuated. Expected loss is $66m a year today, rising twelve-fold to $787m by 2035 as a $20 billion surface programme builds out. The financial figure is not the operative risk: an unshielded August-1972-class event on the lunar surface delivers roughly eight times NASA's 30-day marrow limit, and since Gateway was stood down in March 2026 the architecture has no orbital safe haven.

11.1 A different risk regime

Protection Earth provides	On the Moon	Consequence
Magnetosphere deflects charged particles	Absent. Only weak, patchy crustal magnetic anomalies; the Moon spends most of each month in the solar wind	Full SEP and GCR flux reaches the surface unattenuated
Atmosphere absorbs particle radiation (~10 t/m² of shielding)	Absent. Surface pressure is ~3×10^{-15} bar	Only spacecraft structure, regolith or purpose-built mass shields crew
Ground provides somewhere to shelter	Regolith works, but must be moved, and the equipment to move it is itself exposed	Shelter must be built before it is needed, not improvised
Terrestrial comms and navigation as backup	None. Everything routes through a thin, new relay and PNT layer	A single degraded link severs both command and position

A fourth difference is less obvious. On Earth and in LEO a solar particle event is primarily a hazard to equipment, because crew are shielded by the magnetosphere. Beyond the magnetosphere the crew are the most radiation-sensitive component of the system, so the shelter drives the architecture and the mission plan is subordinate to the dose budget.

11.2 The radiation environment

Galactic cosmic rays

LRO's CRaTER instrument measures approximately 130 mGy per year on the lunar surface, almost entirely from GCR — around 140 times Earth's surface background, corresponding to roughly 330 mSv/yr in equivalent dose. GCR is high-energy heavy nuclei. Shielding against it is counterproductive at intermediate thicknesses, because primary particles fragment in the shield and produce secondary showers that can raise dose rather than lower it. Only metres of regolith reliably help.

GCR flux is highest at solar minimum, when the weakened heliospheric magnetic field admits more galactic particles. The quiet years safest for solar particle events are the worst for chronic dose, so no window in the solar cycle is optimal for both.

Solar particle events

The August 1972 reconstruction is the design case: more than 10 Gy skin dose and approximately 2 Gy to bone marrow, unshielded on the lunar surface. Against NASA's 250 mGy 30-day marrow limit that is a factor of eight exceedance, within acute radiation syndrome. ESA's assessment is that inside a command module the dose would likely have caused acute radiation poisoning, and that for an astronaut on EVA it could have been lethal.

Shielding works here and the requirements are known. Two centimetres of aluminium reduces organ dose significantly but not below limits. Approximately 7.4 cm of aluminium, or its areal-density equivalent in regolith (roughly 20 to 30 cm depending on packing) provides adequate protection for almost all SEP events.

NOAA's S-scale gives roughly three S4 events and about one S5 per solar cycle, with S3 events at which EVA avoidance is recommended occurring about ten times per cycle. A lunar surface programme should plan for around ten EVA-suspension events and three serious shelter events per eleven-year cycle. Against a target cadence of crewed landings every six months, these are schedule inputs rather than contingencies.

Surface charging and dust

The lunar surface charges under solar wind and UV exposure, strongly positive on the dayside and, in the terminator region and permanently shadowed polar craters, strongly negative. The resulting electrostatic fields levitate and transport fine regolith. Apollo experience established lunar dust as a first-order engineering problem: abrasive, electrostatically adhesive, and degrading to seals, optics, radiators and spacesuit joints. Solar events intensify the charging environment, and the south-polar sites Artemis targets sit in exactly the terminator-like illumination geometry where charging gradients are strongest.

11.3 The programme as of August 2026

NASA restructured its exploration architecture on 24 March 2026, in a direction that increases space-weather exposure.

Element	Status
Artemis II	Flown 1–11 April 2026. Four crew; maximum distance 406,771 km, beyond the Apollo 13 record. First crewed flight beyond LEO since 1972. Carried six HERA sensors, crew dosimeters, the German M-42 EXT heavy-ion detector and organ dosimetry.
Lunar Gateway	Stood down on 24 March 2026. NASA's own language is that the programme is paused pending an international review rather than terminated, and Congress has not ratified the decision; the practical effect on the 2026–2030 architecture is the same, which is how this report treats it. Sunk cost is widely reported as roughly $4.4bn for 2019–2025, a figure NASA has not confirmed and which different accountings put lower. The Exploration Upper Stage and the SLS upgrades Gateway depended on were stood down alongside it. ESA, JAXA, CSA and the UAE had already invested in flight hardware; the Power and Propulsion Element became the SR-1 Mars bus.

Element	Status
Artemis III	Restructured to a LEO demonstration at ~460 km, NET late 2027. Orion and four crew rendezvous with Starship HLS Pathfinder and Blue Moon MK2. No landing.
Artemis IV	2028. First crewed lunar landing since 1972, south pole.
Artemis V	Late 2028. Base build-out begins.
Lunar surface base	~$20bn over seven years in three phases, Shackleton Connecting Ridge. Phase 1 rovers and payloads; Phase 2 habitats; Phase 3 permanent infrastructure including nuclear and solar power, comms network, lunar GPS and relay satellites. Target cadence: crewed landings every six months with at least two providers.
Fission surface power	100 kWe closed-Brayton reactor, targeted to be ready to launch by Q4 2029 under a NASA–DOE memorandum of January 2026. TRISO fuel enriched below 19.5%; requirement to run autonomously for at least ten years. In July 2026 NASA named 37 companies under an unfunded Announcement for Collaborative Opportunity. That is an in-kind collaboration instrument: no reactor has been procured, no money has changed hands, and selection under it is not a contract award.

What standing Gateway down did to the risk profile

Gateway would have provided a permanently crewed platform in near-rectilinear halo orbit with a hardened radiation shelter and an independent communications relay. Under the surface-first architecture that has replaced it for planning purposes, a crew on the lunar surface during a major SEP event has no orbital safe haven within reach and no relay asset outside the storm.

Direct Earth-to-surface mission profiles are faster and cheaper. They are also less resilient: abort options are longer, the shelter must be pre-positioned on the surface, and there is no second node in the communications architecture.

The decision has strategic and scientific logic that Gateway did not. The narrower point is that the space-weather risk profile of the lunar programme changed in March 2026, and the mitigation architecture has not been re-baselined against it.

11.4 Landers, rovers and the surface build

The Commercial Lunar Payload Services programme is the delivery mechanism, and its record is mixed: Astrobotic's Peregrine failed in January 2024; Intuitive Machines' IM-1 landed but tipped over in February 2024; Firefly's Blue Ghost M1 succeeded fully in March 2025 with fourteen days of surface operations; IM-2 landed on its side in March 2025 and ended after a few hours. Roughly half the attempts to date have failed or partially failed.

Programme	Value	Status and radiation content
CLPS Moon Base awards	~$600m total, June 2026	Four deliveries in late 2028. Astrobotic $297.9m (two Peregrine flights), Intuitive Machines $148.3m, Firefly $144.2m. Every one of the four carries LETS, a radiation spectrometer explicitly for astronaut safety planning.
Astrobotic Griffin-1	$199.5m	Slipped to NET November 2026. Carries the Astrolab FLIP rover.
Intuitive Machines IM-3	$77.5m	H2 2026, Reiner Gamma. Carries Lunar Vertex (magnetometer and plasma spectrometer, studying a mini-magnetosphere shielding the surface from

Programme	Value	Status and radiation content
		solar wind) and KASI's LUSEM lunar space environment monitor.
Firefly Blue Ghost M2	$112m	Slipped to NET 2027, far side. Carries ESA's Lunar Pathfinder comms relay and LuSEE-Night.
Lunar Terrain Vehicle	$219m Astrolab; $220m Lunar Outpost	Demonstration task orders May 2026, delivery November 2027 ahead of Artemis IV. The stated survival driver is thermal (surviving 150 days of lunar darkness at −400 °F) not ionising radiation.
VIPER	$190m ceiling, not an award	Revived September 2025 on a second Blue Moon MK1. The $190m is the maximum potential value of the task order rather than money obligated. NASA's stated landing date was late 2027; with New Glenn out of service the realistic date is 2028 or later. 100-day south-pole water-prospecting mission.
New Glenn / Blue Moon MK1	Return to flight targeted before end-2026	New Glenn's third flight, NG-3, launched successfully on 19 April 2026 and landed its booster. A subsequent vehicle was destroyed on 28 May 2026 during a static-fire test at Launch Complex 36, damaging the pad. The distinction matters: the vehicle has flown successfully and the loss was a ground test, not a launch failure. Blue Origin has completed two of five recovery phases and is moving to hybrid horizontal-vertical integration. Blue Moon MK1 has not flown. Both the VIPER landing and the November 2027 rover delivery depend on this vehicle returning to service.
JAXA/Toyota pressurised rover	NET 2032	Fuel-cell electric, supporting 30–45 day crewed excursions. The canonical long-duration element where shielding should be a design driver; no published shielding specification found.

Three observations from that table. The instrumentation story is good: LETS flies on all four Moon Base landers, and IM-3 carries two dedicated space-environment payloads. NASA is measuring the environment systematically before committing crew to it.

The schedule story is fragile, and it has become more so since these dates were set. Griffin-1 has moved from July to November 2026 and Blue Ghost M2 from December 2026 into 2027. Two further dates, the VIPER landing and the November 2027 rover delivery ahead of Artemis IV, both ride on New Glenn, which has not flown since the 28 May 2026 static-fire loss damaged Launch Complex 36. Whatever margin the programme holds for space-weather stand-downs is being consumed by vehicle availability first. Readers should treat every date in the table as a target that has already moved once.

The engineering story is thinner. The Lunar Terrain Vehicle's published survival requirement is thermal rather than radiological. No public shielding-thickness requirement exists for the Moon Base habitats, no shadow-shield mass or exclusion radius has been published for the 100 kWe surface reactor, and no dose specification is public for the JAXA/Toyota pressurised rover. These may all exist internally. They are not in the public record, which means they are not being scrutinised.

11.5 Communications and navigation

Lunar operations depend on LunaNet, NASA's interoperability framework, together with commercial relays and ESA's Moonlight programme. A five-satellite constellation, one communications and four navigation, contracted to Telespazio for €123m with a further €176m committed in November 2025, with a maiden launch in 2028. Intuitive Machines holds a Near Space Network contract with a maximum potential value of $4.82 billion for lunar relay and navigation services. China operates the Queqiao-2 relay, a 1,200 kg spacecraft with an eight-to-ten-year design life supporting Chang'e-6, -7 and -8 and the ILRS.

The layer is new, thin, and exposed on two fronts. The relay spacecraft sit outside the magnetosphere with the same SEP exposure as crew, without the option of sheltering, and there is no on-orbit servicing at lunar distance. And the Earth-Moon link crosses a disturbed ionosphere, while the Sun itself raises the receiver noise floor when the Moon is near solar conjunction from a ground station's viewpoint.

The result is a correlated failure with no terrestrial analogue: the same event that requires the crew to shelter can degrade the link over which the shelter order is given, and the navigation signal a returning vehicle depends on.

11.6 Recommendations for lunar operators

1. Pre-position shelter before crew. The 7.4 cm aluminium-equivalent requirement is known and stable. Regolith-shielded shelter emplaced robotically before crewed arrival should lead the surface build. The three-phase plan currently sequences habitats into Phase 2.
2. Publish the shielding specification. A stated regolith depth and crew dose budget for the Moon Base habitat, and a shadow-shield mass and exclusion radius for the surface reactor. These exist as engineering answers; putting them in the public record allows them to be reviewed.
3. Re-baseline abort and safe-haven planning post-Gateway. Whatever replaces the orbital node's shelter function should be an explicit programme requirement with a stated dose-limit rationale.
4. Build relay redundancy. Two independent paths with different orbits and ground segments. The correlated-failure mode in 11.5 is what turns a manageable event into an unrecoverable one.
5. Institutionalise far-side observation. Perseverance currently provides up to two weeks' notice of active regions rotating toward Earth, as an opportunistic use of a Mars asset. A dedicated far-side solar monitor is a modest investment against a $20bn programme.
6. Schedule against the S-scale. Ten EVA-suspension events and three shelter events per solar cycle, with matching margin. A six-month landing cadence leaves little room to absorb them otherwise.

12. Mars and Deep Space

$33m a year today, rising eight-fold to $261m by 2035. The exposure is small in dollars and decisive in mission design: radiation is now the binding constraint on crewed Mars architecture, and the two levers that act on it (trajectory and propulsion) are both being pulled. A 2026 analysis concludes that current dose measurements restrict crewed Mars missions to under four years. ESA finds that launching at solar maximum is safer, which inverts the intuitive answer.

12.1 The dose numbers

The Mars Science Laboratory's RAD instrument measured 1.8 mSv/day in cruise, giving roughly 0.66 Sv for a round-trip transit on current propulsion. Only about 5 per cent of that came from solar energetic particles; the remainder is GCR, which is the hard problem because it cannot be practically shielded. Against a 1,000 mSv career limit, a single Mars round trip consumes most of a career.

Mission duration compounds it. A conjunction-class profile runs 1,005 days total — 198 out, 558 on the surface, 197 back. An opposition-class profile is shorter at 560 days but spends more of it in transit, where dose rates are highest, and only 40 days on the surface. Windows open roughly every 26 months. Zhang and colleagues, publishing in Space Weather in 2026, conclude that current radiation measurements restrict crewed Mars missions to under four years total.

> **The solar-maximum paradox**
>
> ESA analysis published in March 2026, examining sixty years of Mars transfer orbits with data from the Trace Gas Orbiter, finds that **launching at solar maximum is safer**.
>
> The mechanism is solar modulation: an active Sun's enhanced heliospheric magnetic field suppresses galactic cosmic rays, which supply roughly 95 per cent of transit dose. The increased solar particle risk is real but is acute, episodic and shieldable; GCR is chronic and is not.
>
> Fast trajectories cut GCR exposure by up to 55 per cent, fuel-efficient ones by up to 45 per cent. The Mars surface is roughly 60 per cent lower dose than transit.
>
> The conclusion is that a round trip is achievable within career limits **if** it launches at solar maximum on a targeted trajectory. Solar-cycle phase becomes a first-order input to mission scheduling rather than a secondary consideration.

12.2 What is actually flying

Mission	Date	Detail
NASA SR-1 Freedom / Skyfall	Launch December 2028	20+ kWe HALEU fission reactor with closed Brayton conversion driving Hall-effect thrusters. Nuclear electric propulsion, encased in a boron carbide

Mission	Date	Detail
		radiation shield. Repurposes the Power and Propulsion Element built for Gateway. NASA has not published a transit duration; nuclear-electric trajectories in this power class are spiral transfers measured in months to years rather than the fast transit the technology is often assumed to deliver. Carries three Ingenuity-class helicopters with ground-penetrating radar to map subsurface water and identify human landing sites.
Mars Telecommunications Orbiter	Delivery targeted by end-2028	~$700m fixed-price solicitation. Five bidders are publicly known, including Rocket Lab, Blue Origin and Lockheed Martin. No contract has been awarded and the delivery date is NASA's requirement rather than a committed schedule. The ~$4bn total mission cost against an $11bn legacy estimate is Rocket Lab's own bid-stage figure, not an independent assessment.
ESA Rosalind Franklin	NET late 2028	NASA approved implementation April 2026. Falcon Heavy from LC-39A; NASA supplies braking engines, radioisotope heater units and MOMA electronics. Drills to 2 m specifically because surface organics are destroyed by ionising radiation.
China Tianwen-3	Launch 2028, return 2031	Mars sample return. NASA's own Mars Sample Return was cancelled in the January 2026 appropriation, leaving 30+ cached Perseverance samples with no funded retrieval path.
SpaceX Starship to Mars	NET 2028	The 2026 window was abandoned. Cargo priced at $100m per tonne. Flight 12 on 22 May 2026 was the first Block 3 vehicle; the booster was lost and the FAA required a mishap investigation, which it closed on 13 July 2026, clearing Flight 13.
AFRL Oracle	In launch-readiness review	First US cislunar space situational awareness capability, operating near EML1 at ~326,000 km on a two-year campaign.

Two features of this table matter for this report. Nuclear electric propulsion is, in the crewed case, a radiation-mitigation technology: shorter transit means less accumulated GCR and fewer opportunities to be caught by a large SEP event without shelter. NASA intends to scale from 20 kWe to megawatt class in the 2030s, and the reason that matters for crew safety is not usually how it is described publicly.

And the Gateway hardware went to Mars. The Power and Propulsion Element that was to have provided the lunar programme's orbital safe haven became the bus for the first nuclear-electric interplanetary spacecraft. That is efficient reuse and it is also a compact illustration of how fast this architecture is being rearranged.

12.3 The communications problem

Everything about deep space compounds the SEP problem in one direction: latency. Light travels between Earth and Mars in four to twenty-four minutes each way. A crew or a spacecraft at Mars distance cannot be given a real-time shelter order, cannot be talked through an anomaly, and cannot receive an updated forecast in time to act on a fast event.

The operational consequence is that autonomy is not optional. Shelter decisions, safe-mode entry and anomaly response must be onboard functions with onboard thresholds. That is a solved problem for robotic missions and an unsolved one for crewed architecture, because the thresholds involve a judgement about mission value against dose that no programme has yet delegated to software or to a crew without ground concurrence.

12.4 What deep-space operators should do

- Treat solar-cycle phase as a scheduling constraint. The ESA result makes launch-window selection a dose decision, not only a propulsion one.
- Fund propulsion as radiation mitigation. The clearest path to a survivable crewed Mars mission is a shorter transit, and the argument for nuclear electric propulsion is stronger on dose grounds than on cost grounds.
- Design onboard autonomy for shelter decisions. Ground concurrence is not available at Mars distance.
- Publish the dose architecture. No public radiation shielding specification or dose budget exists for Starship's crew configuration. For a vehicle intended to carry people to Mars, that is the single most consequential unpublished number in the sector.

13. Cislunar Competition

Space weather is usually framed as a shared environmental hazard, which it is. It now also sits inside a competition between two lunar programmes running on schedules that leave little room for conservatism.

13.1 The two architectures

	United States and partners	China and ILRS partners
Crewed landing target	Artemis IV, 2028	Before 2030
Crew vehicle status	Artemis II flown April 2026; Artemis III restructured to a LEO demonstration NET late 2027	Long March 10 / Mengzhou max-dynamic-pressure abort test February 2026, including first sea-based crew recovery. Long March 10B flew its maiden mission on 10 July 2026 and recovered its first stage, China's first orbital-class booster recovery
Surface programme	~$20bn over seven years; 100 kWe reactor by Q4 2029	ILRS construction from 2031; Russia–China lunar reactor targeting 2033–2035
Robotic precursors	CLPS: mixed record, four Moon Base deliveries late 2028; VIPER late 2027	Chang'e-6 far-side sample return complete June 2024; Chang'e-7 postponed on 23 August 2026, the day before this paper's cut-off, after Tropical Depression Narra closed the launch window; CNSA said the probe did not meet launch conditions. Revised timing is uncertain, from September 2026 to 2027. Chang'e-8 ISRU demonstration ~2028–29
Relay architecture	LunaNet; ESA Moonlight from 2028; Intuitive Machines Near Space Network, a $4.82bn IDIQ ceiling across the full contract period rather than an award	Queqiao-2 operational since March 2024, plus Tiandu-1 and -2 constellation precursors
Partners	ESA, JAXA, CSA, UAE and Artemis Accords signatories	Roughly a dozen ILRS partners across Asia, Africa and Latin America

Both programmes carry substantial space-environment instrumentation. Chang'e-6 flew a French radon detector and a Swedish negative-ion instrument measuring solar-wind interaction with the surface, and Chang'e-7's objectives include magnetic field measurement and Moon-based observation of Earth's magnetotail. NASA flies LETS on every Moon Base lander.

13.2 Where competition and space weather interact

1. Schedule pressure works against conservative launch criteria. A programme racing a rival has an incentive to fly through a marginal forecast, which is not a hypothetical failure mode in spaceflight history.

2. A severe event degrades both sides at once. Space-weather situational awareness becomes a competitive capability rather than a shared utility, and the side that understands the environment faster recovers faster.
3. Crew safety standards may diverge. Dose limits are set nationally. If one programme accepts higher crew risk to meet a date, the comparison becomes a policy problem rather than an engineering one.
4. Cislunar space is becoming a monitored domain. AFRL's Oracle operates near EML1, and the Space Force stood up a Cislunar Coordination Office in April 2026. That infrastructure sits outside the magnetosphere with no published radiation specification.

Space weather is also one of the few areas where cooperation costs nothing strategically. Solar observations are not meaningfully dual-use, warning data has no offensive value, and both sides benefit from the same forecasts. It is a natural candidate for technical cooperation that survives political competition, on the model of nuclear test-detection data sharing.

PART V · EARTH INFRASTRUCTURE

14. Power Grids

Still the largest single channel at $12.9bn a year, 62 per cent of the 2026 total, and still effectively all of the tail. It is also the slowest-growing: 1.4× by 2035, against 25× for orbital compute. The binding constraint is not whether transformers fail but how long they take to replace — US lead times have gone from under twelve months before 2020 to 36–60 months, against a fleet averaging 38–40 years old with 82 per cent import dependency.

14.1 Exposure

About 90 per cent of electricity consumed in the United States passes through at least one large power transformer, and the Department of Energy estimates between 4,900 and 6,799 in service on the transmission system. Average age was 38 to 40 years at the most recent comprehensive survey, against a nominal design life of 40 years, and more than 70 per cent had passed 25 years. Age matters directly: insulation degradation reduces the thermal margin available to absorb stray-flux heating, so a storm that a new unit survives can destroy an old one.

Vulnerability concentrates geographically. Coupled physics-engineering-economic modelling across 10,464 US substations finds the largest geoelectric responses in the Great Lakes region and along the north-eastern seaboard, with additional exposure up the Appalachians into the Northeast Corridor. The intersection of high geomagnetic latitude, resistive igneous basement rock, and the densest concentration of load and long transmission in North America. Lloyd's placed 20 to 40 million Americans at risk of extended outage in a Carrington-class event.

The same physics applies to the Nordic grids, the UK's north-south transmission spine and the Canadian shield. Low-latitude systems were long assumed safe, but modelling of the Japanese, Chinese and South African grids under Carrington-class forcing shows GIC magnitudes sufficient to matter, driven by coastal conductivity contrasts rather than auroral electrojets.

14.2 Failure sequence

Stage	Mechanism	Timescale	Effective mitigation
1. Reactive power collapse	Saturated transformers draw large reactive currents; system voltage falls across a wide area	Seconds to minutes	Pre-storm de-rating, reserve dispatch, capacitor availability
2. Protective misoperation	Harmonics trip relays and capacitor banks on healthy circuits, removing the reactive support the system needs	Seconds	Harmonic-tolerant relay settings; this was the Québec 1989 mechanism
3. Voltage collapse / cascade	The system loses the ability to hold voltage and separates	Under 2 minutes	Pre-authorised islanding and network-switching plans

Stage	Mechanism	Timescale	Effective mitigation
4. Transformer thermal damage	Stray flux heats structural steel and windings; insulation degrades or fails	Minutes to hours of exposure; failure may present weeks later	Neutral blocking devices; GIC monitoring; post-event dissolved-gas analysis
5. Restoration constraint	Damaged units must be replaced from spares or new manufacture	16 days to 1–2 years	Spare strategy — the only mitigation at this stage

GIC damage at stage 4 is often latent. A transformer stressed during a storm may not fail during it; insulation degraded by a thermal excursion can fail weeks or months later under normal load, by which point the causal link is invisible in outage statistics. The historical loss record therefore almost certainly understates GIC damage, and post-storm dissolved-gas analysis is a high-value, low-cost practice.

14.3 The restoration constraint

Transformer replacement

Before 2020: under 12 months from order to delivery. **2024–2026:** 36 to 60 months, driven by global demand from data-centre buildout and electrification.

Import dependency: 82 per cent of large power transformers used in the US were imported — 617 of 754 units in the most recent survey year. Domestic capacity peaks at about 343 units a year and was running at roughly 40 per cent utilisation. 89 per cent of surveyed manufacturers reported difficulty finding qualified workers.

Lloyd's modelled restoration at 16 days to 1–2 years depending on spare availability. At current lead times the pessimistic end is the base case for any event requiring more than a handful of replacements.

Four spare-transformer programmes exist in the United States and none is designed for a correlated event. STEP triggers only on a presidentially declared emergency and has never been activated. SpareConnect is voluntary and informal. RESTORE requires members to maintain designated spares. Grid Assurance holds inventory for 31 subscribers, sized by Monte Carlo analysis to a 98 per cent confidence of meeting anticipated emergencies — calibrated against conventional contingencies rather than a continental event damaging units at many substations at once. Conventional spare strategy assumes uncorrelated failures. A severe geomagnetic disturbance violates the premise.

14.4 What a storm costs

Event or scenario	Loss	Note
Hydro-Québec, March 1989 — direct	CAD 6.5m	Equipment and substation damage only
Hydro-Québec, March 1989 — economic	> CAD 1bn	9-hour outage, 6 million people
Malmö, Sweden, October 2003	~$0.5m	1-hour outage, 50,000 customers
PJM market, June 2000–Dec 2001	~$500m	3.3% increase in real-time prices attributed to

Event or scenario	Loss	Note
		geomagnetic activity
US blackout scenario S1 (8% of population)	$4.0bn/day	Oughton et al. (2017)
US blackout scenario S4 (66% of population)	$48.5bn/day	$41.5bn domestic, $7.0bn international
UK 1-in-100-year storm	~£10.2bn	GDP loss, dynamic inoperability input-output model
Aotearoa New Zealand, extreme CME	NZ$8.36bn	Without mitigation
Carrington-class, United States	$0.6–2.6trn	Lloyd's/AER (2013)
250-year storm, US, physics-coupled model	$2.04bn/day	Lower bound: excludes voltage collapse and cascading failure

Indirect losses dominate. Oughton and colleagues find that direct disruption to electricity represents 49 per cent of total macroeconomic cost, with the remainder propagating through domestic and international supply chains. Manufacturing takes the largest hit, followed by government, finance and insurance, and property. China, Canada and Mexico are the economies most exposed to indirect cost from a US blackout. A global multi-regional input-output analysis of a Québec-scale event produced $2.4–3.4 trillion, or 3.9 to 5.6 per cent of global GDP.

14.5 Regulation and mitigation economics

The United States has the only mandatory enforceable geomagnetic disturbance standard for grids. NERC TPL-007, approved by FERC in 2016 and revised through TPL-007-4, requires transmission planners to assess their systems against a benchmark event and develop corrective action plans. Chapter 18 covers the wider policy architecture.

The cost-benefit case for mitigation is unusually clear. Operational measures (pre-storm de-rating, network switching, controlled islanding, harmonic-tolerant relay settings) were assessed in the New Zealand study at a benefit-cost ratio of approximately 700:1. GIC blocking devices reach up to 80:1. GIC monitoring is cheap and is the prerequisite for knowing whether anything is working; coverage remains sparse globally, which is itself a principal source of model uncertainty. Backup generation is modelled to reduce GDP loss by 11 to 14 per cent during long outages.

The Space Weather Advisory Group's 2024 user-needs survey collected input from 125 electric power participants. The sector remained stable through the Gannon storm, which the survey attributes to preparatory action, though isolated transformer monitoring exceeded thresholds and some transmission equipment tripped unexpectedly.

15. Aviation

$1.07bn a year, growing 1.3× to 2035. The exposure is dominated by precautionary action rather than damage: airlines lose money by rerouting away from a hazard because the regulatory framework requires it. ICAO issued 2,350 advisories between November 2019 and September 2025, of which 62.7 per cent concerned GNSS and 37.1 per cent HF communications. The tail case is multi-week disruption to polar and high-latitude operations.

15.1 Three hazards

1. HF communications blackout. Above about 82 degrees latitude there is no geostationary coverage, so polar operations depend on HF radio for the position reporting that separation standards require. A flare's X-rays absorb HF across the sunlit hemisphere within eight minutes and with no warning; polar cap absorption from a particle event can persist for days, with the July 1959 event lasting fifteen. When HF fails the flight cannot legally remain on the polar route.
2. GNSS degradation. Scintillation and steep total-electron-content gradients degrade positioning accuracy and, more importantly, integrity. During the October 2003 storms the FAA's WAAS was unavailable for about thirty hours and aircraft lost vertical guidance for satellite-based approaches for over ten hours. In May 2024 positioning errors exceeded the four-metre tolerance for precision approach.
3. Radiation dose. The UK reasonable-worst-case places aircrew dose at approximately 7 mSv at FL400 in polar regions during a 100-year event, against an EU occupational limit of 20 mSv per year and a public limit of 1 mSv. A single flight would consume a third of a year's occupational allowance.

15.2 The advisory record

Metric	Value	Detail
Total operational advisories	2,350	Nov 2019 – Sep 2025
GNSS advisories	1,472 (62.7%)	
HF communication advisories	872 (37.1%)	
Radiation advisories	6 (0.26%)	
May 2024 (Kp 9)	148 advisories	40 GNSS, 108 HF COM
October 2024 (Kp 8)	266 advisories	206 GNSS, 60 HF COM
April 2025 (multiple CMEs)	180 advisories	137 GNSS, 43 HF COM

GNSS advisories outnumber HF nearly two to one, which inverts the conventional framing of space weather as primarily a polar HF problem. GNSS impacts concentrate over the equatorial South Atlantic and South America, geography unrelated to the auroral oval and affecting a different set of carriers. October 2024 generated 80 per cent more advisories than May 2024 despite being a weaker storm, because advisory count tracks the geographic footprint of ionospheric disturbance rather than storm intensity.

Advisories show clear equinoctial seasonality consistent with the Russell-McPherron effect: GNSS peaks in October (304) and April (202); HF peaks in May (143) and September (114).

15.3 Cost and forecast gaps

Published event-scale costs for the US airline sector run $0.4m to $5m for moderate conditions and $1m to $30m for more extreme scenarios. In the 2003 Halloween storms, New York-Hong Kong flights diverted, consuming 26,600 extra pounds of fuel and shedding more than 16,500 pounds of payload. In January 2005, 26 United Airlines flights diverted in a single event.

Carrier behaviour follows from the asymmetry. An unnecessary reroute costs a known five-to-six-figure sum; an unmitigated radiation exposure or lost separation event is unbounded. Carriers therefore reroute conservatively, and most of aviation's space-weather cost is the price of caution.

The SWAG survey found that accurate space weather forecasts at 12 hours and beyond, actionable by the aviation industry, do not exist but are needed. Forecast uncertainty is itself a cost driver: in January 2012 cross-polar rerouting was required not because HF had failed but because uncertainty about whether it would fail was operationally unacceptable.

On 4 November 2015 a solar radio burst near 1 GHz caused false plots on secondary surveillance radar at multiple European airports. The civil analogue of the 1967 BMEWS incident, in which the system does not fail but reports something that is not there.

SSEEM places aviation loss in a Carrington-class event at $125bn. The mechanism is not aircraft damage: the UK reasonable-worst-case describes HF unavailable for days, GNSS errors exceeding 100 metres with periodic total loss over several days, and satellite communications disrupted for one to three days. The National Audit Office summarises it as widespread air travel disruption lasting multiple weeks.

16. The GNSS-Dependent Ground Economy

$3.23bn a year, the second-largest channel, doubling to $6.46bn by 2035. This is the channel where the largest documented modern loss occurred, and it involved no damage to anything: six hours of degraded RTK-GPS during the 2024 planting window cost US Midwest farmers over $500 million. Coverage follows breakage, and this loss did not break anything.

16.1 The dependency

GPS has generated an estimated $1.4 trillion in US economic value since 1983, with 90 per cent accruing since 2010. RTI's analysis for NIST estimated that a complete GPS outage would cost the US economy approximately $1 billion per day, and flagged that the figure would be higher during the agricultural planting season. Ten sectors were identified as materially GPS-dependent: telecommunications, precision agriculture, electricity grid management, financial services, location-based services, maritime, surface mining, oil and gas, professional surveying, and vehicle telematics.

Two of those close a loop with the rest of this paper. Electricity grid management depends on GPS timing for phasor measurement and protection coordination, and financial services depend on it for transaction timestamping under regulations such as MiFID II. A space-weather event that degrades GNSS is simultaneously degrading a control input to the grid it is stressing. SSEEM does not model that coupling, and we are not aware of a published study that quantifies it.

16.2 May 2024: anatomy of the loss

1. A wall of ionospheric plasma stretched across North America on 10 May 2024, producing positioning errors up to 70 metres in the central United States. Younas and colleagues documented the spatio-temporal evolution using roughly 100 high-accuracy GPS receivers from a seismic research network.
2. Peak disruption lasted approximately six hours, with residual errors persisting up to two days.
3. Approximately 70 per cent of planted US acreage depends on GPS-guided automated equipment. RTK guidance operates at centimetre accuracy, so a 70-metre error is total loss of function rather than degradation. Farmers reported tractors behaving erratically and stopped work. The SWAG survey recorded errors of 10 to 30 feet and significant cumulative loss of planting area.
4. Illinois corn planting-date yield curves place 10 May at roughly 95 per cent of maximum yield and 31 May at roughly 85 per cent; at $4.50 per bushel and 200 bushels per acre, a 10 per cent yield reduction is a $90 per acre loss.
5. Aggregate documented loss exceeded $500 million, with no equipment damaged and no satellite affected.

Four properties that make this the template

It was a precision failure, not an outage. The signal was present throughout. It was wrong, and the receiving systems had no way to know by how much.

Timing determined magnitude. Six hours in May cost what six hours in November would not have.

There was no manual fallback. A generation of operators has trained on guided equipment, and the capacity to work without it has atrophied, the same pattern documented in maritime celestial navigation and in aviation's automation-dependency literature.

Nobody insured it. Neither the farmers nor the equipment manufacturers nor the GNSS operator carried cover for this loss.

16.3 Where the same pattern sits

- Financial services. Regulatory timestamping has made GNSS timing a compliance dependency. Most large institutions hold over with local oscillators, but holdover accuracy degrades over hours and the depth of that resilience is rarely tested against a multi-day scenario.
- Telecommunications. Mobile network synchronisation, particularly for TDD and carrier aggregation, depends on GNSS timing. Holdover exists; multi-day holdover generally does not.
- Maritime and offshore. Dynamic positioning for drilling and offshore construction is a GNSS-dependent safety-critical function. Directional drilling additionally depends on magnetic surveying, which a storm corrupts directly; the November 2025 USGS advisory named directional drilling and aeromagnetic survey specifically.
- Surveying, construction and mining. RTK-dependent workflows halt entirely under scintillation, with the same absence of manual fallback as agriculture.

The SWAG survey's GNSS sector assessment was still in progress when the 2024 report was published, so the second-largest loss channel identified in this paper is also the one for which national user requirements have not been formally collected.

16.4 Mitigation

- Multi-constellation, multi-frequency receivers. GPS, Galileo, GLONASS and BeiDou on L1/L5 provide redundancy and the dual-frequency measurement that allows first-order ionospheric delay to be estimated and removed. The highest-leverage technical mitigation, and a procurement decision rather than a research problem.
- Terrestrial timing backup. eLoran and fibre-based time distribution for fixed infrastructure. Adoption is uneven; the UK, South Korea and parts of the US have programmes.
- Documented manual fallback. A named owner and a rehearsed degraded-mode procedure for every process that assumes centimetre positioning.
- Integrity monitoring rather than accuracy alone. The May 2024 failure was dangerous because systems did not know they were wrong. RAIM and SBAS integrity outputs should drive automatic degradation to a safe mode.

PART VI · IMPLICATIONS AND ACTIONS

17. Risk Transfer

Space weather has three properties that make a peril hard to insure, and it has all three at once.

1. It is systemic. A single event damages property, energy, aviation, marine, agriculture, space and political-risk books simultaneously, across every continent in daylight and every orbital shell at once. Diversification, the mechanism insurance rests on, does not operate.
2. **Its frequency is disputed by a factor of 26, and its severity was revised upward by a factor of two in July 2026.** Pricing requires a loss cost; a loss cost requires both parameters.
3. Its exceedance curve has no elbow. Loss scales smoothly with return period across four decades, so there is no natural attachment point at which retention should stop and transfer begin.

17.1 Current coverage

Line	Typical treatment	Gap
Space — launch and in-orbit	Mature market for GEO; ~$1.35bn premium; capacity $550–700m/yr	Under 1% of LEO satellites carry in-orbit cover; the fastest-growing exposure is essentially uninsured
Orbital compute & manufacturing	No established market	Data loss, workload interruption and correlated multi-node corruption have no product
Human spaceflight & tourism	Launch liability and some crew cover	Radiation dose to non-professional visiting crew is an unlitigated liability question
Property / energy	Physical damage to transformers is generally covered; the GIC mechanism is not usually excluded	Business interruption without physical damage (the dominant loss) often is not
Business interruption	Usually contingent on physical damage at the insured location	A grid outage caused by damage elsewhere frequently fails the trigger
Aviation	Hull and liability cover physical loss	Precautionary rerouting and cancellation (the actual loss mechanism) is uninsured operating cost
Agriculture	Multi-peril crop insurance covers yield loss by cause	The May 2024 loss was a planting-delay yield loss from signal degradation; coverage is untested
Marine / offshore	Physical damage and some operational cover	GNSS-driven dynamic-positioning loss and magnetic-survey corruption sit ambiguously

Coverage follows breakage; the loss does not. Almost every documented modern space-weather loss (the 2024 agricultural loss, aviation rerouting, satellite station-keeping propellant, operational stand-down) is economic loss without physical damage. Policy wordings built for perils that break things do not reach it. The orbital channels compound this, because the loss is often invisible even to the operator: propellant consumed, life shortened, upsets absorbed.

17.2 Parametric structures

Space weather suits parametric cover better than most perils for which parametric products already exist.

- The index is objective and public. Dst, Kp and >10 MeV proton flux are measured by government agencies, published in near real time, and archived. There is no loss-adjustment dispute available.
- The measurement is global. One index serves policyholders on every continent. The property that makes the peril hard to diversify and easy to index.
- Basis risk is manageable within a sector. An airline's rerouting cost correlates with Kp and proton flux. A farmer's planting-delay loss correlates with scintillation indices and calendar date. A constellation's propellant burn correlates with thermospheric density.
- It pays without a claim. For losses with no physical damage to inspect, parametric payment is the only practical mechanism.

The space market has already moved this way for constellation risk, with parametric triggers eliminating claims overhead that would otherwise exceed the value of an individual small satellite. Trade press reported a Munich Re satellite product addressing solar flares in January 2026, which the reinsurer has not itself published. The obvious extensions are a proton-fluence trigger for orbital compute downtime and a thermospheric-density trigger for constellation propellant.

17.3 Underwriting actions

1. Build a cross-line space-weather aggregate. Most carriers can produce an aggregate for windstorm, earthquake and cyber. Few can produce one for space weather, because the exposure sits in nine books that do not communicate. The Lloyd's 2025 scenario models exactly this simultaneity.
2. Test wordings against a no-physical-damage loss. Run the May 2024 agricultural loss and a 72-hour GNSS degradation through existing policy language. The answer is more often "nothing pays" than most underwriters expect.
3. Price the frequency and severity uncertainty explicitly. A loss cost derived from a single point estimate is one of ten defensible answers on the grid in Appendix A.2.
4. Underwrite the orbital economy before it is large, not after. Orbital compute grows twenty-five-fold by 2035 on this model's assumptions. Building the data, the wordings and the triggers now is far cheaper than doing it after a correlated loss.
5. Require disclosure as a condition of cover. The disclosure gap in 6.5 is a solvable problem if the sector's insurers ask for drag assumptions, radiation tolerance data and storm procedures the way they ask for structural margins.

18. Policy and Regulation

18.1 The United States

Instrument	Date	Content
National Space Weather Strategy and Action Plan	2015, revised 2019	First whole-of-government strategy. Establishes benchmarks, vulnerability assessment and response planning as national objectives.
PROSWIFT Act	2020	Formalises NASA and NOAA forecasting responsibilities; creates the Space Weather Advisory Group; requires user-needs assessment.
Space Weather Implementation Plan	December 2023	Five-year OSTP plan under the SWORM Subcommittee. Three objectives: protect assets and operations; produce accurate and timely forecasts; establish response and recovery procedures.
NERC TPL-007 (to -007-4)	2016 onward	The world's only mandatory enforceable GMD grid standard.
SWAG National Survey of User Needs	2024	First systematic user requirements assessment. 46 findings and 113 recommendations across seven of ten identified sectors.
SWFO-L1	Operational orbit January 2026	Sustains the operational L1 solar-wind measurement the 2023 plan requires.

The 2023 Implementation Plan assigns named agencies to each action. Under Objective I, refining space weather benchmarks falls to DHS, DOC, DOD, DOE, DOI, NASA and NSF; critical infrastructure vulnerability assessment to DHS and DOD; geoelectric hazard mapping from magnetotelluric survey data to DHS, DOC, DOE, DOI and FERC; and a pilot test of hardening strategies at military installations to DHS, DOD and DOE. Under Objective III, DHS updates the Federal Operating Concept and exercises response plans, and DOC, DOD and NASA jointly implement a satellite-anomaly attribution system.

The plan sets biennial and quadrennial review cadences rather than numerical targets. Phase 1 benchmarks cover five space weather phenomena, with agencies directed to address gaps rather than meet stated thresholds.

18.2 Other jurisdictions

Jurisdiction	Position
United Kingdom	On the National Risk Register since 2011; government estimates a 5–25% likelihood of a severe event within five years. Met Office forecasting is internationally respected. The National Audit Office found in 2025 that roles and accountabilities remain unclear, there is no stated position on acceptable residual risk, and exercising has been limited.
European Union	The EU Space Programme Regulation (2021) identifies space weather as a critical-infrastructure protection priority. ESA's Vigil mission to L5 is scheduled for 2031 and would materially extend warning times. ESA's Moonlight programme and the ASCEND orbital-compute study are both funded. No binding grid standard.
China	The National Medium- and Long-Term Development Plan (2024–2050) designates space-weather detection and forecasting as key research areas. Operates an independent lunar relay architecture

Jurisdiction	Position
	and flies dedicated space-environment instruments on Chang'e missions.
Australia / New Zealand	Exercise Aurora (2024) tested electricity and communications resilience. New Zealand has produced the most rigorous published national mitigation cost-benefit analysis.
Canada	RCAF strategy identifies space weather as a key infrastructure vulnerability. Hydro-Québec holds the deepest operational experience of any utility.

18.3 Five gaps

1. Benchmarks stop at 1-in-100. Standard practice sets the planning benchmark at roughly 100-year severity, excluding by construction the tier carrying 77 per cent of expected loss. Defensible for design standards; less so for national risk assessment.
2. Recovery is unregulated. TPL-007 and its analogues assess how a system responds to a storm. No jurisdiction mandates spare-transformer holdings against a correlated event. With lead times at 36 to 60 months, recovery capacity is now the binding constraint.
3. GNSS resilience has no owner. The second-largest loss channel falls between space agencies, transport regulators and sector regulators. No jurisdiction has an accountable owner for the economy-wide consequence of losing positioning and timing for a week.
4. The orbital regime has no resilience regulator at all. Licensing regimes address debris mitigation, spectrum and, increasingly, disposal. None requires an operator to state a space-weather design basis, a storm-time manoeuvre protocol, or a solar-cycle assumption for post-mission disposal. This is the fastest-growing exposure in the model and the least regulated.
5. Both key parameters are unfunded research questions. The frequency disagreement is a factor of 26; the Sivadas severity result is unreplicated. No government has a dedicated programme for either.

18.4 The first recommendation

A multilateral programme to resolve superstorm recurrence — extended radiocarbon and ice-core surveys, systematic reanalysis of nineteenth- and early-twentieth-century magnetograms, expanded GIC monitoring, and physics-based extreme-event modelling — would cost on the order of tens of millions of dollars a year. The disagreement it would resolve moves global annualised expected loss by approximately $29 billion a year. Independent replication of the Sivadas no-saturation result is a smaller programme with a comparable payoff: the difference between that result being right and wrong is $17bn a year.

No infrastructure investment in this paper has a comparable ratio. Every other recommendation — how many spares to hold, how many blocking devices to install, how much capital to reserve, what altitude to fly a constellation at — is currently sized against numbers that cannot be defended.

19. Recommendations by Audience

19.1 Space and aerospace industry

1. **Put space weather into the published design basis.** State the drag environment a constellation is designed to hold altitude through, the total-dose and single-event tolerance of orbital compute, the shielding specification for crewed volumes, and the solar-cycle phase assumed in disposal analysis. The sector publishes structural and thermal margins as a matter of course and is silent on this one (6.5).
2. **Close the storm-time conjunction assessment gap.** In May 2024 collision screening became impossible for days. Operators need 3-day and 7-day drag forecasts they do not have, and a pre-agreed protocol for sequencing manoeuvres when everyone needs to move at once (10.4).
3. **Design orbital compute for correlated upset, not random upset.** Cross-node redundancy fails when the whole cluster is in the same storm. Checkpointing cadence, power-cycle recovery and workload placement should be specified against an S4 event, not against a background rate.
4. **For commercial stations, publish the radiation architecture** before flying paying visitors: shelter location and shielding, dose budget for non-professional crew, and the policy on suspending operations at S3 and above (8.4).
5. **For lunar operators, follow 11.6.** Pre-position shelter, publish the shielding specification, re-baseline safe haven post-Gateway, build relay redundancy, and schedule against the S-scale.
6. **Reconsider low insertion altitudes, or budget the recurring loss.** The February 2022 Starlink loss followed from an operational choice and recurs on every launch into a storm window.
7. **Share tracking data during storms.** The collective benefit when everyone's predictions degrade at once exceeds the individual competitive cost.

19.2 Corporate leadership and investors

1. **Run a 72-hour combined connectivity, positioning and timing loss scenario.** All three together, because that is how the event arrives. Assign a named owner and a rehearsed degraded-mode procedure to every process that assumes any of them.
2. **Ask portfolio companies in the space sector for their space-weather design basis.** If the answer is not available, that is diligence information.
3. **Audit supplier concentration in the geoelectric hazard zones.** Roughly half of a US blackout's macroeconomic cost propagates through supply chains, with China, Canada and Mexico most exposed as suppliers of intermediate goods.
4. **Check whether business interruption cover actually triggers.** Most BI wordings require physical damage at the insured location.
5. **Establish transformer lead time before an event rather than after.** A 36-to-60-month replacement queue is a strategic fact about any business owning or depending on high-voltage equipment.

6. **Subscribe to forecasts and define an action on receipt.** NOAA SWPC, the UK Met Office and ESA all publish free real-time alerts.

19.3 Government and policy

1. **Fund the frequency question and replicate the severity result** (18.4).
2. **Extend national planning benchmarks past 1-in-100 severity**, and review them against the Sivadas no-saturation finding.
3. **Assign an owner for orbital-regime resilience.** Require a stated space-weather design basis as a licensing condition, in the same way debris mitigation became one.
4. **Assign an owner for economy-wide GNSS resilience.** Currently unassigned in every jurisdiction reviewed.
5. **Address the recovery constraint.** Strategic transformer reserves, standardised designs for interchangeability, and domestic manufacturing capacity. US capacity runs at roughly 40 per cent utilisation against an 82 per cent import dependency.
6. **Mandate or incentivise operational GMD procedures**, assessed at roughly 700:1 benefit-cost — procedural rather than capital, and the cheapest large-scale risk reduction available.
7. **Fund L5.** ESA's Vigil would extend CME warning from tens of minutes to days. It is scheduled for 2031; acceleration merits consideration, and the case is stronger now that the exposure includes crewed lunar operations.
8. **Treat space-weather data sharing as a candidate for cooperation across the strategic competition.** Solar observations have no offensive value and both sides benefit from the same forecasts (13.2).

19.4 Insurance and reinsurance

1. **Build a cross-line space-weather aggregate** covering all nine channels as one correlated event (17.3).
2. **Develop products for the orbital economy now**, while it is small and the data can still be built from a standing start.
3. **Develop parametric triggers for the no-physical-damage losses** — proton fluence for orbital compute downtime, thermospheric density for constellation propellant, scintillation indices for agriculture.
4. **Stress-test wordings against the 2024 Gannon loss.** A documented half-billion-dollar loss with no physical damage is the ideal test case.
5. **Require disclosure as a condition of cover**, which is the fastest available route to closing the gap in 6.5.
6. **Recognise that the space book understates the space exposure.** Under 1 per cent of LEO satellites are insured in orbit; the real exposure sits on operator balance sheets and on those of the companies depending on them.

20. What Would Change These Estimates

20.1 Findings that would raise them

- Confirmation of the no-saturation result. Sivadas et al. (2026) is a single paper on a question with ten competing prior explanations. If independent work confirms it, upper-tier severity roughly doubles: a Carrington-class event moves from $2.4trn to $4.8trn and annualised expected loss from $20.6bn to $37.8bn.
- A resolved frequency question favouring the power-law tail. Riley's estimate takes annualised expected loss to $35.3bn on base severity, or $67.0bn combined with doubled severity.
- Faster space-economy growth than assumed. The 2035 projection uses a roughly 4× space economy. A faster buildout of orbital compute or lunar infrastructure moves the orbital share above 42 per cent.
- A documented correlated loss in the orbital sector. The first published account of what an S4 event actually costs a constellation or a compute cluster would almost certainly move the orbital anchors upward, because they are currently set conservatively in the absence of data.
- Better constraint on pre-Carrington events. The 774 CE and 993 CE radiocarbon events sit above the model's ceiling.
- Demonstrated correlation between channels. SSEEM sums channels additively; in a real event they compound.

20.2 Findings that would lower them

- A resolved frequency question favouring the Weibull tail. Moriña's central estimate takes annualised expected loss to $7.0bn.
- Demonstrated effectiveness of TPL-007-era hardening. The standard has been in force since 2016 and has never been tested by a severe event.
- Radiation-tolerant orbital compute at commercial performance. If error-correction and architectural mitigation prove sufficient through an S4 event, the fastest-growing channel's vulnerability function falls sharply. Google's Trillium result is the first encouraging data point.
- Rapid multi-constellation, multi-frequency receiver penetration. Dual-frequency GNSS removes most first-order ionospheric error and could halve the ground channel.
- Materially better forecasting. ESA's Vigil at L5 in 2031 would convert much of the aviation and satellite loss (which is precautionary and driven by forecast uncertainty) into avoidable cost.
- A slower space economy. Launch costs that fail to fall, or an orbital-compute business case that does not close, compress the 2035 projection substantially.

20.3 Where the evidence stands

The terrestrial severity side is reasonably well understood at the level of order of magnitude. Several independent research groups, using different methods and different data, arrive at extreme-event losses in the trillions globally, with indirect effects roughly doubling the direct infrastructure cost. The 2026 Sivadas result introduces the first serious challenge to the upper bound in over a decade, and it points upward.

The frequency side is not understood. The field does not know whether a Carrington-class event is a 1-in-80-year or a 1-in-1,000-year proposition, and every probability in this paper inherits that uncertainty.

The orbital side is not measured at all. Thirty-one of forty-five severity cells rest on judgement because the industry does not publish loss data, and the one recent event that would have generated it (January 2026, an S4 radiation storm among the most intense in the GOES record) has produced no public accounting seven months later. The numbers in Part III are the most contestable in this paper and they are contestable for a reason that the sector could fix.

What is not in doubt is the direction. In 2018 there were 1,980 active satellites. There are now 16,393. The ISS is authorised only to 2030 and four commercial stations are competing to succeed it. Commercial GPUs are running workloads in orbit. A $185 billion missile-defence architecture is moving into proliferated LEO. A $20 billion lunar base is contracted, and a nuclear-electric spacecraft leaves for Mars in 2028. The hazard is a constant of the solar system. The exposure is a set of decisions, being made now, at a pace that the risk analysis has not kept up with.

Appendix A. SSEEM v2.0: Full Model Specification

This appendix is the model, not a summary of it. It contains every parameter that produces every figure in this paper, the arithmetic that combines them, the frequency basis, the growth assumptions, the sensitivity runs and the confidence grade attached to each input. It is published at this level of detail so that a reader who disagrees with a number can replace it, recompute, and say precisely how much the conclusion moves. That is the standard the model asks to be held to.

A.0 Structure and arithmetic

SSEEM v2.0 is a frequency-severity catastrophe model in the standard actuarial form, applied to a peril for which no loss portfolio exists. It has four layers.

1. **Hazard.** Five severity tiers indexed on Dst, from T1 (Strong, G3) to T5 (Carrington-class). Each tier carries an occurrence rate λ in events per year, derived from the NOAA scale frequencies for T1–T3 and from the published Carrington-class return-period literature for T4–T5.
2. **Vulnerability and exposure.** Nine loss channels, six orbital and beyond, three terrestrial. Each channel is crossed with each tier to give 45 cells. Each cell carries a low, central and high per-event global loss in 2026 US dollars, and a confidence grade.
3. **Financial.** Losses are combined across channels within a tier by simple summation, on the assumption that a single storm affects all channels simultaneously. This is the correct assumption for this peril and is what distinguishes it from most catastrophe perils: the hazard is global and arrives everywhere at once, so channel losses are perfectly correlated in occurrence, not independent.
4. **Projection.** A per-channel exposure growth multiple carries the 2026 severity matrix to 2035. Frequencies are held constant, because solar activity is not a function of how much hardware is exposed to it.

The annualised expected loss is the sum over all 45 cells of the cell's central severity multiplied by its tier's occurrence rate:

The aggregation

AEL = Σ over tiers t, channels c of (λ(t) × S(t,c)), where λ(t) is the occurrence rate of tier t in events per year and S(t,c) is the central per-event loss in channel c at tier t.

Scenario loss for tier t = Σ over channels c of S(t,c). This is the figure quoted for a single Carrington-class or Québec-class event. No frequency term enters it.

Occurrence exceedance probability. P(at least one event of tier t or worse in n years) = 1 − exp(−n × Σ over tiers t' ≥ t of λ(t')), the Poisson complement. The OEP curve in Exhibit 10 is this expression evaluated across tiers.

2035 projection. $S_{2035}(t,c) = S(t,c) \times g(c)$, where g(c) is the channel growth multiple in A.4. λ is unchanged.

Low and high totals are computed by substituting the low or high severity in every cell simultaneously. They are therefore comonotonic bounds, not confidence intervals: they describe

> the result if every judgement erred in the same direction at once, which is conservative at the high end and optimistic at the low end.

Three properties of this construction should be stated plainly because they bound what the output means. The model has no time dimension within an event, so it cannot represent a loss that compounds with outage duration. It has no dependency structure between channels beyond the shared occurrence, so it cannot represent a lunar mission failing because a relay satellite failed. And it prices no cascade: the debris scenario in Chapter 10 is deliberately excluded, and it is the one pathway that could exceed every figure the model produces.

A.1 The severity matrix: all 45 cells

Every severity anchor in SSEEM v2.0, with its confidence grade and source family. Losses are global, per event, in 2026 US dollars. Confidence grades: **H** = published figure applied at its estimated scope; **M** = published figure scaled to a different footprint, exposure base or year; **L** = structured judgement calibrated against adjacent published work. Channels above the rule are orbital and beyond; below it, terrestrial. Thirteen cells are graded M and thirty-one L: the grades are part of the output, not a disclaimer attached to it.

Channel	Tier	Low	Central	High	Conf.	Source family
Space infra.	T1 Strong	$2m	$7m	$22m	L	Constellation operating economics
	T2 Severe	$14m	$63m	$215m	M	Constellation operating economics
	T3 Extreme	$140m	$590m	$1.62bn	M	Starlink 2022 / Gannon 2024
	T4 Québec	$1.85bn	$8.20bn	$25.50bn	M	Starlink 2022 / Gannon 2024
	T5 Carrington	$52.00bn	$210bn	$700bn	L	SIA 2026 exposure base
Orbital compute	T1 Strong	$100k	$500k	$2m	L	Orbital data centre programmes
	T2 Severe	$800k	$4m	$16m	L	Orbital data centre programmes
	T3 Extreme	$7m	$35m	$130m	L	Orbital data centre programmes
	T4 Québec	$90m	$500m	$1.90bn	L	Orbital data centre programmes
	T5 Carrington	$2.20bn	$12.00bn	$46.00bn	L	Orbital data centre programmes
Human LEO	T1 Strong	$100k	$500k	$2m	L	Commercial LEO destinations
	T2 Severe	$900k	$4m	$16m	L	Commercial LEO destinations
	T3 Extreme	$8m	$35m	$135m	L	Commercial LEO destinations
	T4 Québec	$70m	$400m	$1.60bn	L	Commercial LEO destinations
	T5 Carrington	$1.80bn	$9.00bn	$36.00bn	L	Commercial LEO destinations
Space defence	T1 Strong	$1m	$5m	$18m	L	Golden Dome / SDA / historical precedent

Channel	Tier	Low	Central	High	Conf.	Source family
	T2 Severe	$8m	$35m	$125m	L	Golden Dome / SDA / historical precedent
	T3 Extreme	$60m	$280m	$900m	L	Golden Dome / SDA / historical precedent
	T4 Québec	$1.00bn	$5.00bn	$18.00bn	L	Golden Dome / SDA / historical precedent
	T5 Carrington	$18.00bn	$70.00bn	$260bn	L	Golden Dome / SDA / historical precedent
Lunar	T1 Strong	$50k	$200k	$1m	L	Artemis and lunar surface programme
	T2 Severe	$400k	$2m	$8m	L	Artemis and lunar surface programme
	T3 Extreme	$4m	$20m	$80m	L	Artemis and lunar surface programme
	T4 Québec	$40m	$250m	$1.00bn	L	Artemis and lunar surface programme
	T5 Carrington	$1.10bn	$6.00bn	$25.00bn	L	Artemis and lunar surface programme
Mars	T1 Strong	$20k	$100k	$500k	L	Mars and deep space programmes
	T2 Severe	$200k	$1m	$4m	L	Mars and deep space programmes
	T3 Extreme	$2m	$10m	$40m	L	Mars and deep space programmes
	T4 Québec	$20m	$120m	$500m	L	Mars and deep space programmes
	T5 Carrington	$500m	$3.00bn	$13.00bn	L	Mars and deep space programmes
Power grids	T1 Strong	$4m	$15m	$40m	L	Operating-cost derivation
	T2 Severe	$30m	$150m	$500m	L	Operating-cost derivation
	T3 Extreme	$400m	$1.80bn	$5.00bn	M	Gannon 2024 grid record
	T4 Québec	$8.00bn	$28.00bn	$90.00bn	M	Hydro-Québec 1989, scaled
	T5 Carrington	$600bn	$1.60trn	$5.50trn	M	Lloyd's/AER + Lloyd's-Cambridge
GNSS ground	T1 Strong	$1m	$6m	$22m	L	RTI/NIST GPS valuation
	T2 Severe	$15m	$80m	$280m	M	RTI/NIST GPS valuation
	T3 Extreme	$250m	$600m	$1.20bn	H	Younas et al. 2025 (JGR)
	T4 Québec	$3.00bn	$12.00bn	$35.00bn	M	RTI/NIST GPS valuation
	T5 Carrington	$90.00bn	$340bn	$1.10trn	M	RTI/NIST GPS valuation

Channel	Tier	Low	Central	High	Conf.	Source family
Aviation	T1 Strong	$500k	$4m	$15m	M	ICAO advisory record
	T2 Severe	$3m	$20m	$70m	M	Published US airline event cost
	T3 Extreme	$30m	$110m	$350m	M	Published US airline event cost
	T4 Québec	$500m	$2.20bn	$7.00bn	M	Xue et al. 2023
	T5 Carrington	$35.00bn	$125bn	$420bn	L	UK reasonable worst case

A.2 Hazard frequencies

Tier	λ (events/yr)	Return period	P(≥1 in 10 yrs)	Frequency basis
T1 Strong	11.8182	1 per 31 days	100.0%	NOAA scales (H)
T2 Severe	5.4545	1 per 67 days	100.0%	NOAA scales (H)
T3 Extreme	0.3636	1-in-2.8 yr	97.4%	NOAA scales (H)
T4 Québec	0.0200	1-in-50.0 yr	18.1%	Lloyd's/AER (M)
T5 Carrington	0.0067	1-in-150.0 yr	6.4%	Lloyd's/AER (M)

A.3 Joint sensitivity: frequency and severity

Rows vary the Carrington-class recurrence rate across the published range. Columns vary upper-tier (T4 and T5) severity between the base case and the 2× multiplier implied by Sivadas et al. (2026). The base case used throughout this paper is the Lloyd's/AER row, base severity column.

Frequency source	P(Carrington) per decade	AEL, base severity	AEL, 2× severity
Moriña et al. (2019), Weibull, low	0.46%	$5.91bn/yr	$8.13bn/yr
Moriña et al. (2019), Weibull, central	0.92%	$7.01bn/yr	$10.34bn/yr
Moriña et al. (2019), Weibull, high	1.88%	$9.32bn/yr	$14.96bn/yr
Lloyd's/AER (2013), 1-in-150	**6.45%**	**$20.65bn/yr**	**$37.62bn/yr**
Riley (2012), power-law	12.0%	$35.17bn/yr	$66.67bn/yr

Carrington-class event severity: $2.38trn base, $4.75trn under the 2× case. Total annualised expected loss: $20.64bn base, $37.61bn under the 2× case at Lloyd's/AER frequency.

A.4 Exposure growth assumptions to 2035

Channel	Multiple	Basis
Space infra.	4×	Constellation population and revenue; ~4x space economy
Orbital compute	25×	From a near-zero base; $1.77bn (2029) to $39bn (2035) market

Channel	Multiple	Basis
Human LEO	6×	ISS retires 2030; multiple commercial stations reach continuous crew
Space defence	5×	Golden Dome complete architecture mid-2030s; pLEO interceptor layer
Lunar	12×	Surface base operational; six-month landing cadence; ILRS construction from 2031
Mars	8×	SR-1, MTO, Tianwen-3 return, crewed-Mars precursor infrastructure
Power grids	1.4×	Electrification of final energy, net of TPL-007-era hardening
GNSS ground	2×	Deeper penetration, net of multi-constellation receiver adoption
Aviation	1.3×	Traffic growth, net of improved forecasting and routing

Resulting totals: $20.64bn/yr in 2026 and $44.27bn/yr in 2035. Orbital and beyond moves from 16.8% to 41.5% of expected loss.

A.5 Solar radiation storm frequencies

The S-scale governs crew and aircrew dose, satellite and orbital-compute single-event effects, and every shelter or EVA decision beyond LEO. It is the dominant hazard index for six of the nine channels.

Level	Threshold	Per cycle	Per year	Operational effect
S1 Minor	>=10 pfu	50	4.55	Minor polar HF impact
S2 Moderate	>=1e2 pfu	25	2.27	Elevated aircrew dose; possible SEUs
S3 Strong	>=1e3 pfu	10	0.91	EVA radiation-avoidance recommended
S4 Severe	>=1e4 pfu	3	0.27	Memory-device faults; polar HF blackout
S5 Extreme	>=1e5 pfu	1	0.09	Unavoidable high radiation hazard to crew

A.6 Source families in full

Constellation operating economics — Derived from constellation operating economics: 16,393 active satellites (Aug 2026), of which Starlink is 10,971; Starlink at ~$17bn annualised revenue and 12m subscribers; station-keeping propellant, service degradation and manoeuvre load during storm-days.

Starlink 2022 / Gannon 2024 — Feb 2022 Starlink loss (38 of 49 satellites) and the May 2024 Gannon fleet-wide manoeuvre event (~half of ~10,000 active LEO payloads; 180 m/day decay; 6x thermospheric density at 400 km). SpaceX is lowering ~4,400 satellites from 550 km to ~480 km ahead of the ~2030 solar minimum.

SIA 2026 exposure base — Satellite Industry Association (2026): $429bn global space economy, $303bn commercial satellite industry, 14,266 operational satellites at end-2025. Under 1% of the LEO fleet carries in-orbit cover.

Orbital data centre programmes — Orbital data centres and in-space manufacturing: Starcloud-1 (Nvidia H100, Nov 2025), Axiom orbital data centre nodes (Jan 2026), Google Project Suncatcher (81-satellite cluster, learning mission 2027), ASCEND (1 GW before 2050). Varda six W-series capsules; Space Forge ForgeStar-1. Market projected $1.77bn (2029) to $39.1bn (2035).

Commercial LEO destinations — ISS deorbit 2030 (SpaceX USDV, $843m); NASA Commercial LEO Destinations (~$550m public against >$3bn private); Vast >$1bn raised, Haven-1 NET Q1 2027; Axiom $350m (Feb 2026); Starlab CCDR complete (Feb 2026); private astronaut missions PAM-5 and PAM-6 in 2027.

Golden Dome / SDA / historical precedent — Golden Dome ($185bn programme; space-based interceptors in proliferated LEO; 12 contractors, $3.2bn OTA; FY2027 request $17.5bn); Space Force FY2026 $26.3bn; SDA tracking layer. Historical precedent: BMEWS 1967, Aug 1972 sea mines. PNT exposure per DSB (2024) and GAO-21-320SP.

Artemis and lunar surface programme — Artemis II flown Apr 2026; Gateway cancelled Mar 2026; ~$20bn seven-year lunar surface programme; CLPS Moon Base awards ~$600m (Jun 2026); LTV demonstration task orders $219m and $220m; 100 kWe fission surface power for Q4 2029; ESA Moonlight EUR123m + EUR176m; Intuitive Machines Near Space Network ceiling $4.82bn.

Mars and deep space programmes — SR-1 Freedom / Skyfall (Dec 2028, 20+ kWe HALEU reactor, boron carbide shield); Mars Telecommunications Orbiter (~$700m fixed price, delivery 2028); ESA Rosalind Franklin (2028); China Tianwen-3 (2028 launch, 2031 return); AFRL Oracle cislunar SDA.

Operating-cost derivation — Reactive-power procurement, load re-dispatch and manual operating-procedure costs at G3/G4; calibrated on the Forbes & St Cyr (2004) PJM price-effect study (~$500m over 19 months of moderate activity).

Gannon 2024 grid record — May 2024 Gannon operational record: no transformer losses, but widespread voltage-control action, HVDC curtailment and pre-emptive de-rating across NERC, ENTSO-E and NZ grids.

Hydro-Québec 1989, scaled — Hydro-Quebec 1989: CAD 6.5m direct equipment damage; >CAD 1bn broader economic loss. Scaled to a modern, larger, more interconnected and more electricity-intensive global economy.

Lloyd's/AER + Lloyd's-Cambridge — Lloyd's/AER (2013): US $0.6-2.6trn. Oughton et al. (2017): $4.0-48.5bn/day. Lloyd's/Cambridge CRS (2025): global $1.2trn / $2.4trn / $9.1trn.

RTI/NIST GPS valuation — RTI/NIST (2019): GPS has generated $1.4trn in US economic value; a complete GPS outage costs ~$1bn/day to the US economy alone.

Younas et al. 2025 (JGR) — Younas et al. (2025), JGR Space Physics, and associated reporting: >$500m US Midwest agricultural loss from the 10 May 2024 RTK-GPS outage; up to 70 m positioning error over ~6 hours.

ICAO advisory record — ICAO space-weather advisory record Nov 2019-Sep 2025: 2,350 advisories (1,472 GNSS; 872 HF COM); May 2024 produced 148 and Oct 2024 produced 266.

Published US airline event cost — Published event-scale US airline cost band: ~$0.4-5m (moderate) and $1-30m (extreme) per event, scaled to global traffic.

Xue et al. 2023 — Xue et al. (2023), Space Weather: Halloween-2003-class disruption priced against 2019 Northern-Hemisphere flight data; 2005 precedent of 26 United polar diversions in a single event.

UK reasonable worst case — UK Reasonable Worst Case (Hapgood et al., 2021): days-long HF/GNSS loss, ~7 mSv aircrew dose at FL400 polar; NAO (2025) 'weeks of air-travel disruption'.

A.7 Reproducibility

SSEEM v2.0 is implemented as a single Python module of roughly 300 lines, published alongside this paper. The hazard layer, severity matrix, growth projection, aggregation, exceedance curve and both sensitivity analyses are in one file with no hidden state and no external dependencies beyond the standard library. Every table in this appendix is generated from that file, so the paper and the model cannot drift apart. Changing any value in the tables above and re-running regenerates the full result set, including the exhibits.

A.8 How to challenge this model

The model is more useful contested than accepted. A reader who wants to test it should attack it in this order, because this is the order in which the parameters actually move the answer.

1. **The T5 severity anchors, which dominate the AEL.** The Carrington-class tier contributes the majority of annualised expected loss despite a 1-in-150-year rate, because its severity is three orders of magnitude above T3. On the base case it contributes about three-quarters of the total. If you believe the T5 column is wrong, the total is wrong by roughly the same proportion. The T5 grid and macroeconomic figure is anchored to the Lloyd's-Cambridge systemic scenario and is the most defensible cell in the column; the six orbital T5 cells are the least.
2. **The Carrington-class recurrence rate.** A.3 runs the model across the full published range, from Moriña's low estimate to the highest rate in the literature, and the AEL moves from $5.91bn to $35.17bn a year across that range, a factor of six. This is the single largest source of uncertainty in the paper and it is a question for the space-physics community, not for the loss modeller. Anyone who prefers a different frequency source can read their answer straight off the A.3 grid.
3. **The Sivadas correction.** Sivadas et al. (2026) find that solar-wind measurement uncertainty produces an apparent saturation in the geomagnetic response that is a statistical artefact, and that corrected, extreme storm impacts may be twice as large as previously assessed. The model applies a 2× multiplier to T4 and T5 severity only, as a sensitivity rather than a base case. Applying it to the base case roughly doubles the headline figure. A reader who accepts the paper should read the right-hand column of A.3 as the central estimate.
4. **The growth multiples.** The 2035 projection rests on nine numbers in A.4, of which the 25× for orbital compute is by far the most aggressive and the most contestable. It reflects a sector moving from single-GPU demonstrators to funded cluster architectures, which is a judgement about programme execution rather than about physics. Halving it, from 25× to 12.5×, moves the 2035 total from $44.27bn to $42.64bn and the orbital share from 41.5 to 39.3 per cent. The finding that orbital exposure overtakes terrestrial exposure in relative terms is robust to that substitution; a reader who wants to overturn it has to move several multiples at once, which A.4 makes it easy to try.
5. **The independence assumption inside a tier.** Channel losses are summed. If channels are correlated in magnitude as well as occurrence, the tail is fatter than modelled; if they substitute for each other under stress, it is thinner. Neither has been tested.
6. **The exclusions.** No debris cascade, no time dimension within an event, no reinsurance or capital-market feedback, no sovereign response cost, no fatality valuation for the crewed channels. Each is stated where it arises, and every one pushes the modelled figure down rather than up.

What would change this analysis most is not a better model but disclosure. Thirty-one of forty-five cells are graded L because operators do not publish what a storm costs them, and any one of them holds figures that would replace a judgement here with an observation.

Appendix B. Indices, Scales and Terms

Intensity indices

- **Dst / SYM-H** (nT) — Disturbance storm time, measuring equatorial ring-current strength. The standard scalar for geomagnetic storm intensity. Carrington ≈ −900; March 1989 = −589; May 2024 = −412.
- **Kp / NOAA G-scale** (0–9) — Quasi-logarithmic planetary three-hour index. Saturates at 9, so it cannot distinguish a G5 storm from one many times stronger.
- **dB/dt** (nT/min) — Rate of change of the local geomagnetic field; the quantity that drives GIC. UK 100-year benchmark 4,000–5,000 nT/min.
- **Geoelectric field** (V/km) — The induced crustal field driving GIC. Estimated 20–25 V/km in the largest reconstructed events; strongly dependent on ground conductivity.
- **pfu** (particle flux units) — Protons cm^{-2} s^{-1} sr^{-1} above a stated energy, conventionally >10 MeV. The basis of the NOAA S-scale.

NOAA Space Weather Scales

Scale	Range	Measures	Events per 11-yr cycle
G — Geomagnetic storm	G1–G5	Kp 5 to 9	1,700 / 600 / 200 / 100 / 4
S — Solar radiation storm	S1–S5	>10 MeV proton flux, 10 to 10^5 pfu	50 / 25 / 10 / 3 / <1
R — Radio blackout	R1–R5	X-ray flux, M1 to X20	2,000 / 350 / 175 / 8 / <1

Physical phenomena

- **CME** (coronal mass ejection) — An eruption of magnetised plasma from the corona. Transit to Earth takes 14.6 hours (record) to several days.
- **SEP** (solar energetic particle) event — A burst of high-energy protons and ions arriving 20 minutes to hours after a flare. The radiation hazard to crew, aircrew, electronics and orbital compute.
- **GIC** (geomagnetically induced current) — Quasi-DC current driven through grounded conductors by the geoelectric field; the mechanism of transformer damage.
- **Half-cycle saturation.** GIC offsetting a transformer's AC flux, driving the core into saturation for half of each cycle. Produces reactive-power draw, harmonics and stray-flux heating.
- **Thermospheric expansion.** Storm energy heating the upper atmosphere, raising density at satellite altitudes and drag with it. Six-fold at 400 km in May 2024.
- **SEU / latch-up.** Single-event upset: a particle flipping a memory bit. Latch-up: a particle triggering a parasitic conducting path that destroys the device unless power is cycled.
- **Deep dielectric charging.** Energetic electrons depositing charge inside insulating materials over days, causing internal discharge with no external signature.
- **Polar cap absorption.** Ionospheric absorption of HF radio in the polar caps from SEP precipitation. Can persist for days; the July 1959 event lasted fifteen.

- **Scintillation.** Rapid fluctuation in amplitude and phase of a trans-ionospheric radio signal from plasma irregularities. The principal mechanism of GNSS degradation.
- **GCR** (galactic cosmic rays) — High-energy particles of extra-solar origin. Chronic, effectively unshieldable, strongest at solar minimum.
- **Solar modulation.** The suppression of GCR flux by an active Sun's enhanced heliospheric magnetic field. The mechanism behind the solar-maximum paradox in 12.1.

Programme and market terms

- **pLEO.** Proliferated low Earth orbit: an architecture of many small satellites rather than few large ones, adopted for resilience against attack.
- **CLPS.** Commercial Lunar Payload Services, NASA's fixed-price lunar delivery programme.
- **CLD.** Commercial LEO Destinations, NASA's programme to replace the ISS.
- **LunaNet / Moonlight.** NASA's lunar comms and PNT interoperability framework, and ESA's five-satellite constellation implementing part of it.
- **ILRS.** International Lunar Research Station, the China-led lunar base programme.
- **AEL / OEP.** Annualised expected loss; occurrence exceedance probability.
- **LPT.** Large power transformer. 4,900–6,799 in US service; 36–60 month lead times; average age 38–40 years.
- **Parametric cover.** Insurance paying a fixed sum on a measured index threshold rather than on assessed loss.

Appendix C. Event Register

Date	Intensity	Documented consequences
Sep 1859 Carrington	Dst ≈ −900 nT (recon.) >700 nT/min	Telegraph failures across Europe and North America; operators shocked; pylons sparking; two-hour transmission on induced current alone; aurora over Cuba and Hawaii
May 1921 New York Railroad	Dst ≈ −907 nT (recon.)	Fires at telegraph and railway signal installations in New York and Sweden
May 1967	Major radio burst	All three BMEWS sites jammed simultaneously; nuclear-armed alert aircraft placed in ready-to-launch status; resolved by NORAD solar forecaster
Aug 1972	Dst = −154 nT CME transit 14.6 hr (record)	Dozens (one account: 4,000) of magnetic sea mines detonated off Hòn La; 64% voltage collapse on ND-Manitoba interconnection; AT&T L4 cable outage; Intelsat IV F-2 lost 5% array output; DSCS II mission-ending failure; est. >10 Gy skin / 2 Gy marrow unshielded on the lunar surface
Mar 1989 Hydro-Québec	Dst = −589 nT	James Bay network collapsed in <90 s; 6 million without power for 9 hours; CAD 6.5m direct, >CAD 1bn economic; Salem NJ transformer damaged; TDRS-1 logged 250+ anomalies; HF lost for Australian peacekeepers in Namibia
Oct–Nov 2003 Halloween	Dst = −422 nT X28+ flare (record)	Malmö 1-hr outage, 50,000 customers, ~$0.5m; FAA WAAS disabled ~30 hrs; vertical guidance lost >10 hrs; ADEOS-II lost; NY–Hong Kong diversions burning 26,600 lb extra fuel
Jan 2005	S-class SEP event	26 United Airlines flights diverted in a single event to avoid HF blackout
Jul 2012	Carrington-class	Missed Earth by approximately nine days of orbital position
Feb 2022 Starlink	Minor storm	38 of 49 newly launched Starlink satellites lost to enhanced drag before orbit raising
May 2024 Gannon	Dst = −412 nT, Kp 9, S1 Density 6× at 400 km	>$500m US Midwest agricultural loss; GPS errors to 70 m for ~6 hrs; ~half of ~10,000 active LEO payloads manoeuvred; decay to 180 m/day; conjunction assessment impossible; 148 ICAO advisories
Oct 2024	Dst = −341 nT, Kp 8	266 ICAO advisories (206 GNSS); at least one premature Starlink reentry
Nov 2025	Dst = −160 nT, G4	Aurora to Utah and Pennsylvania; 1.8° declination variation at Fairbanks; aeromagnetic survey, directional drilling, GPS and radio affected
Jan 2026	X1.9 flare; CME ~1,700 km/s **S4 severe**, G4	Among the most intense radiation storms in the GOES record. Occurred with ~16,000 satellites in orbit and commercial GPUs operating in LEO. Operators who were warned activated mitigations; no public loss accounting has been published
Apr 2026 Artemis II	Quiet	First crewed lunar flyby since 1972 completed without a significant space-weather event; cabin-reconfiguration shelter procedure tested; ~5% of a career dose limit accrued

Appendix D. Sources

All sources were accessed in August 2026. Where a paywalled primary paper is cited, the accessible summary actually used is listed alongside or in place of it, so that every figure in this paper can be traced to something the reader can open.

HAZARD, HISTORICAL EVENTS, FREQUENCY AND SEVERITY

Sivadas, Sibeck, Subramanyan, Walach, Ozturk, Ferdousi & Michotte de Welle (2026), Nature 655, 1143–1147 — Regression to the mean can explain saturation of geomagnetic storms. Solar-wind measurement uncertainty of at least 30% produces the apparent saturation; corrected, the response stays linear to 15 mV/m, and extreme storm impacts may be twice as large as previously thought. https://www.nature.com/articles/s41586-026-10757-4

Nature research briefing (2026) — Geomagnetic storms caused by solar wind might be more extreme than previously thought. https://www.nature.com/articles/d41586-026-02245-6

Lancaster University / NASA Goddard press release (2026) — Risks of solar storms may be underestimated. https://www.eurekalert.org/news-releases/1135865

ESA (2026) — Monitoring the January 2026 space weather event: X-class flare 18 January 18:09 UTC, CME ~1,700 km/s, S4 severe radiation storm peaking 19 January 19:15 UTC, among the most intense in the GOES record. https://www.esa.int/Space_Safety/Space_weather/ESA_monitoring_January_2026_space_weather_event

NOAA SWPC (19 January 2026) — S4 (Severe) Solar Radiation Storm in progress. https://www.swpc.noaa.gov/news/s4-severe-solar-radiation-storm-progress-january-19th-2026

NOAA SWPC (January 2026) — G4 (Severe) Geomagnetic Storm Watch for 20 January. https://www.spaceweather.gov/news/g4-severe-geomagnetic-storm-watch-20-january-utc-day

NOAA SWPC (19 January 2026) — G4 (Severe) geomagnetic storm levels reached. https://www.swpc.noaa.gov/news/g4-severe-geomagnetic-storm-levels-reached-19-jan-2026

CNN (19 January 2026) — The Sun released the largest solar radiation storm in over 20 years, corroborating the S4 ranking used here. https://www.cnn.com/2026/01/19/science/largest-solar-radiation-storm-auroras

Beggan et al. (2024), Space Weather — Digitized continuous magnetic recordings for the August/September 1859 storms from London: the source of the >700 nT/min reconstruction. https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023SW003807

British Geological Survey Geomagnetism — The largest magnetic storm on record: the Carrington event of 1859. http://www.geomag.bgs.ac.uk/education/carrington.html

Knipp et al. (2018), Space Weather — On the little-known consequences of the 4 August 1972 ultra-fast coronal mass ejecta: sea-mine detonations, grid and satellite effects, Apollo dose reconstruction. https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2018SW002024

Knipp et al. (2016), via AGU — The May 1967 great storm and radio disruption event: BMEWS jamming and the military response. https://news.agu.org/press-release/1967-solar-storm-nearly-took-us-to-brink-of-war/

USGS Geomagnetism Program — Severe magnetic storm, 11 November 2025. https://www.usgs.gov/programs/geomagnetism/science/severe-magnetic-storm-november-11-2025

Riley (2012), Space Weather — On the probability of occurrence of extreme space weather events (12% per decade). https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2011sw000734

Moriña et al. (2019), Scientific Reports — Probability estimation of a Carrington-like geomagnetic storm (0.92% per decade, CI 0.46–1.88%). https://www.nature.com/articles/s41598-019-38918-8

Riley (2025), Space Weather — Inter-solar-cycle variability of extreme geomagnetic storms. https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2025SW004552

Hapgood et al. (2021) — Development of space weather reasonable worst-case scenarios for the UK National Risk Assessment. https://eprints.lancs.ac.uk/id/eprint/151500/1/885472_2_merged_1611426408.pdf

NOAA SWPC — Space Weather Scales (G, S and R, with average frequency per solar cycle). https://www.spaceweather.gov/noaa-scales-explanation

Institute for Defense Analyses (June 2025) — Revisions to NOAA's Space Weather Scales. https://www.ida.org/research-and-publications/publications/all/r/re/revisions-to-noaas-space-weather-scales

NOAA SWPC — SWFO-L1 status update, January 2026. https://www.spaceweather.gov/news/noaa-swfo-l1-update-january-5th-2026

THE ORBITAL ECONOMY: CONSTELLATIONS, COMPUTE AND STATIONS

azmth.space, from the CelesTrak active catalogue — 16,393 active satellites as of 24 August 2026; orbital-regime and operator breakdown. https://azmth.space/stats

Space Daily (19 August 2026) — Starlink passes 11,000 satellites in orbit, roughly two-thirds of all active satellites. https://spacedaily.com/t-starlink-passes-11000-satellites-august-2026/

Satellite Industry Association (2026), 29th State of the Satellite Industry Report — $429bn global space economy, $303bn commercial satellite industry, segment breakdown, 14,266 operational satellites at end-2025, 4,434 satellites launched across 296 launches. https://sia.org/affordability-productivity-drive-historic-satellite-industry-growth-satellite-industry-association-releases-29th-annual-state-of-the-satellite-industry-report/

Baruah et al. (2024), Space Weather — The loss of Starlink satellites in February 2022. https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023SW003716

Space.com (January 2026), reporting Michael Nicolls of SpaceX — SpaceX lowering the orbits of 4,400 Starlink satellites from ~550 km to ~480 km over 2026: a greater than 80% reduction in ballistic decay time at solar minimum, from 4+ years to a few months. https://www.space.com/space-exploration/satellites/spacex-lowering-orbits-of-4-400-starlink-satellites-for-safetys-sake

SpaceNews (January 2026) — SpaceX to lower orbits of some Starlink satellites. https://spacenews.com/spacex-to-lower-orbits-of-some-starlink-satellites/

Parker & Linares (2024) — Satellite drag analysis during the May 2024 Gannon storm: 6× density at 400 km, 180 m/day decay, fleet-wide manoeuvring, conjunction-assessment failure, forecast error of 100–300 units. https://arxiv.org/html/2406.08617v1

Thiele, Heiland, Boley & Lawler (2026), Acta Astronautica — An orbital house of cards: frequent satellite close conjunctions. Peer-reviewed version of arXiv:2512.09643, whose figures it revises. CRASH clock 5.5 days on the 2025 population against 164 days in 2018; 17% chance of a major collision from 24 hours without control; close approaches under 1 km within Starlink roughly every 22 minutes; 41 course corrections per satellite per year. The earlier preprint figures (2.8 days, 121 days, 30%, 11 minutes) are superseded and should not be cited. https://www.sciencedirect.com/science/article/pii/S0094576526004091

Boley (2026), UN COPUOS Scientific and Technical Subcommittee — The CRASH clock: a key environmental indicator for orbital sustainability. Presentation of the peer-reviewed figures to the Committee on the Peaceful Uses of Outer Space. https://www.unoosa.org/documents/pdf/copuos/stsc/2026/List_of_TP/Wednesday11/AM/PDF_Slides_OSI_Aaron_Boley_Item_11.pdf

Introl (January 2026) — Orbital data centre nodes reach space: Axiom's two LEO nodes on Kepler's optical relay, Starcloud-1 with an Nvidia H100, Starcloud-2 October 2026, Google Project Suncatcher, ASCEND 1 GW, Starlink V3 at 1 Tbps, launch-cost parity thresholds, and Google's Trillium TPU radiation-tolerance result. https://introl.com/blog/orbital-data-center-nodes-launch-space-computing-infrastructure-january-2026

Data Center Frontier — Starcloud launches an orbital AI data centre with an Nvidia H100 GPU. https://www.datacenterfrontier.com/site-selection/article/55337494/starcloud-launches-orbital-ai-data-center-with-nvidia-h100-gpu

Data Center Dynamics — Project Suncatcher: Google to launch TPUs into orbit with Planet Labs; 81-satellite compute clusters within a 1 km radius. https://www.datacenterdynamics.com/en/news/project-suncatcher-google-to-launch-tpus-into-orbit-with-planet-labs-envisions-1km-arrays-of-81-satellite-compute-clusters/

Scientific American — Space-based data centres could power AI with solar energy, at a cost. https://www.scientificamerican.com/article/data-centers-in-space/

New Space Economy (2026) — The orbital insurance market: premium volume, capacity, 2023 loss ratio, rate movements, constellation structures, <1% LEO in-orbit cover, Munich Re's January 2026 solar-flare product. https://newspaceeconomy.ca/2026/03/29/the-orbital-insurance-market-how-underwriters-are-pricing-constellation-scale-risk/

DEFENCE, POLICY AND NATIONAL SECURITY SPACE

Aerospace Center for Space Policy and Strategy — FY 2026 defense space budget and the emergence of Golden Dome: $26.3bn Space Force request, a near-40% increase; $21.6bn through reconciliation. https://csps.aerospace.org/papers/fy-2026-defense-space-budget-emergence-golden-dome

DefenseScoop (24 April 2026) — Space Force names 12 companies to develop Golden Dome's space-based interceptors: $3.2bn combined OTA value, $185bn total programme, FY2027 request $17.5bn, prototype demonstration by 2028, complete architecture mid-2030s. https://defensescoop.com/2026/04/24/golden-dome-space-based-interceptor-missile-defense-contractors/

Space Development Agency — Awards to build 36 accelerated missile-defence tracking-layer satellites for Tranche 3 in support of Golden Dome. https://www.sda.mil/space-development-agency-issues-awards-to-build-36-accelerated-missile-defense-tracking-layer-satellites-for-tranche-3-in-support-of-golden-dome-for-america/

Stars and Stripes (April 2026) — Pentagon asks Congress to codify the Department of War name; ~$50m spent, ~7,600 conforming changes to federal law required. https://www.stripes.com/theaters/us/2026-04-28/pentagon-congress-codify-dow-name-21516668.html

Military.com (July 2026) — House passes the Department of War renaming in its FY2027 NDAA, 216–212 on 22 July 2026; the Senate has not yet acted and the change is not yet law. https://www.military.com/house-passes-plan-to-officially-rename-the-department-of-defense-to-department-of-war-heres-why-it-matters

Congressional Budget Office (January 2026) — The costs of using the name Department of War. https://www.cbo.gov/system/files/2026-01/61942-DoD-Name-Change.pdf

Breaking Defense (April 2026) — Space Force launches a Cislunar Coordination Office supporting a NASA moon base by 2030; AFRL Oracle operating near EML1. https://breakingdefense.com/2026/04/with-eyes-on-future-nasa-moon-base-space-force-launches-cislunar-acquisition-task-force/

White House OSTP (December 2023) — Space Weather Implementation Plan: three objectives, action-by-action agency assignments, review cadences, Phase 1 benchmarks, SWORM coordination. https://www.weather.gov/media/nws/Implementation-Plan-for-National-Space-Weather-Strategy-12072023%20(2)%20(1).pdf

White House OSTP (2019) — National Space Weather Strategy and Action Plan. https://www.weather.gov/media/nws/National-Space-Weather-Strategy-and-Action-Plan-2019.pdf

Space Weather Advisory Group (2024) — Results of the First National Survey of User Needs for Space Weather: 46 findings, 113 recommendations, seven of ten sectors, participant counts, the aviation 12-hour forecast gap, space traffic management 3-day and 7-day requirements, Gannon storm sector outcomes. https://www.weather.gov/media/nws/Results-of-the-First-National-Survey-of-User-Needs-for-Space-Weather-2024.pdf

FERC — Geomagnetic Disturbance Reliability Standard (TPL-007) rulemaking record. https://www.federalregister.gov/documents/2018/05/23/2018-11001/geomagnetic-disturbance-reliability-standard

UK National Audit Office (2025) — The UK's resilience to severe space weather. https://www.nao.org.uk/reports/the-uks-resilience-to-severe-space-weather/

National Academies (2024) — The Next Decade of Discovery in Solar and Space Physics: the 2024–2033 heliophysics decadal survey. https://www.nationalacademies.org/projects/DEPS-SSB-21-03/publication/27938

Defense Science Board (2024) — PNT control executive summary. https://dsb.cto.mil/wp-content/uploads/2024/07/DSB_PNT-Control_ExecutiveSummary_5-1-24.pdf

GAO-21-320SP — Defense navigation capabilities: DoD PNT technologies complementing GPS. https://www.gao.gov/assets/gao-21-320sp.pdf

GAO-26-107805 — Low Earth orbit: NASA faces impending decisions on the International Space Station transition and Commercial LEO Destinations. Source for treating the ISS end date as an open decision rather than a fixed 2030 retirement. https://www.gao.gov/assets/gao-26-107805.pdf

CSIS (2026) — NASA changes course on commercial space stations: analysis of the Commercial LEO Destinations restructuring. https://www.csis.org/analysis/nasa-changes-course-commercial-space-stations

Air & Space Forces Magazine — FY2026 budget request for the Air Force and Space Force. Source for the Space Force base request falling year-on-year before reconciliation funds are added. https://www.airandspaceforces.com/trumps-budget-boosts-usaf-cuts-space-force-spending/

SpacePolicyOnline — Adding reconciliation, the Space Force budget tops $40bn in FY2026. Source for the distinction between the base request and the reconciliation-inclusive total. https://spacepolicyonline.com/news/adding-reconciliation-space-force-budget-tops-40-billion-in-fy2026/

SpaceNews — Defense appropriations bill for 2026 funds the Space Force at $26bn and presses the Pentagon on Golden Dome. https://spacenews.com/defense-appropriations-bill-for-2026-funds-space-force-at-26-billion-presses-pentagon-on-golden-dome/

THE MOON, MARS AND DEEP SPACE

Spaceflight Now (25 March 2026) — NASA outlines its $20 billion moon base plan: three phases over seven years, crewed landings every six months, two launch providers, nuclear and solar power, rovers, site preparation, communications network, lunar GPS and relay satellites. https://spaceflightnow.com/2026/03/25/nasa-outlines-ambitious-20-billion-plan-for-moon-base/

CNBC (24 March 2026) — NASA to spend $20 billion on a moon base and cancel the orbiting lunar station. https://www.cnbc.com/2026/03/24/nasa-to-spend-20-billion-on-moon-base-cancel-orbiting-lunar-station.html

CSIS (2026) — What comes next for Artemis: Gateway cancellation, Artemis III restructured to a LEO demonstration, Artemis IV as first landing, $20bn surface base, HLS status, $4.4bn Gateway sunk cost. https://www.csis.org/analysis/what-comes-next-artemis

NASA — Artemis II mission page: 1–11 April 2026 crewed lunar flyby, crew, trajectory and distance records, onboard radiation payloads. https://www.nasa.gov/mission/artemis-ii/

NASA Science — To protect Artemis II astronauts, NASA experts keep eyes on the Sun: monitoring architecture, two-tier dose thresholds, cabin-reconfiguration shelter procedure, Perseverance far-side observation. https://science.nasa.gov/missions/artemis/artemis-2/to-protect-artemis-ii-astronauts-nasa-experts-keep-eyes-on-sun/

Nature (2024) — Artemis I radiation measurements: 26.7–35.4 mSv total mission dose equivalent, fourfold variation by location inside the vehicle, deep-space GCR 0.96–1.24 mSv/day. https://www.nature.com/articles/s41586-024-07927-7

NASA (June 2026) — Moon Base science awards and upcoming opportunities: ~$600m across Astrobotic, Firefly and Intuitive Machines for four late-2028 deliveries, each carrying the LETS radiation spectrometer for astronaut safety planning. https://www.nasa.gov/news-release/nasa-awards-more-moon-base-science-previews-new-opportunities/

NASA (May 2026) — Update on Moon Base rovers, landers and missions: Astrolab $219m and Lunar Outpost $220m Lunar Terrain Vehicle demonstration task orders; Blue Origin MK1 delivery contracts. https://www.nasa.gov/news-release/nasa-provides-update-on-moon-base-rovers-landers-missions/

Spaceflight Now (June 2026) — Next-generation astronaut moon rovers aim for deployment ahead of Artemis IV crew arrival; thermal survival requirements for lunar night. https://spaceflightnow.com/2026/06/26/next-gen-astronaut-moon-rovers-aim-for-deployment-ahead-of-artemis-4-crew-arrival/

NASA (2025) — Blue Origin selected to deliver the VIPER rover to the lunar south pole: $190m maximum potential value, landing late 2027, 100-day mission. https://www.nasa.gov/news-release/nasa-selects-blue-origin-to-deliver-viper-rover-to-moons-south-pole/

ANS Nuclear Newswire (January 2026) — NASA and DOE solidify collaboration on a lunar surface reactor: 100 kWe closed Brayton, ready to launch by Q4 2029, TRISO fuel below 19.5% enrichment, ten years autonomous operation. https://www.ans.org/news/article-7674/nasa-doe-solidify-collaboration-on-a-lunar-surface-reactor/

Physics Today — Space weather on the Moon: 130 mGy/yr GCR from LRO/CRaTER, the August 1972 dose reconstruction, the 7.4 cm aluminium shelter requirement, NASA 30-day limits. https://physicstoday.aip.org/quick-study/space-weather-on-the-moon

ESA — Protecting Artemis and lunar explorers from space radiation. https://www.esa.int/Space_Safety/Space_weather/Protecting_Artemis_and_lunar_explorers_from_space_radiation

ESA Moonlight — €123m Telespazio prime contract plus €176m committed November 2025; five satellites, maiden launch 2028. https://www.esa.int/Newsroom/Press_Releases/ESA_launches_Moonlight_to_establish_lunar_communications_and_navigation_infrastructure

Intuitive Machines — NASA Near Space Network contract, $4.82bn maximum potential value for lunar relay and navigation services. https://investors.intuitivemachines.com/news-releases/news-release-details/nasa-awards-intuitive-machines-near-space-network-contract

NASASpaceflight (March 2026) — SR-1 Freedom / Skyfall: 20+ kWe HALEU fission reactor, nuclear-electric propulsion with Hall-effect thrusters, closed Brayton cycle, boron carbide shield, December 2028 launch, three Ingenuity-class helicopters, built on repurposed Gateway PPE hardware. No transit duration is stated in this or any other NASA source located; figures of roughly one year circulating in secondary coverage are not attributable and are not used here. https://www.nasaspaceflight.com/2026/03/nasa-sr1-freedom-mars-2028/

Science / AAAS (2026) — NASA plans to send a nuclear-powered spacecraft to Mars in 2028. https://www.science.org/content/article/nasa-plans-send-nuclear-powered-spacecraft-mars-2028

ESA (March 2026) — The radiation paradox: why solar maximum is the safest time to travel to Mars. Sixty years of transfer orbits; fast trajectories cut GCR exposure by up to 55%; the Mars surface is ~60% lower dose than transit; ESA career limit 1,000 mSv. https://www.esa.int/Science_Exploration/Human_and_Robotic_Exploration/The_radiation_paradox_why_solar_maximum_is_the_safest_time_to_travel_to_Mars

Zeitlin et al. (2013), Science, via SwRI — MSL/RAD cruise measurements: 1.8 mSv/day, ~0.66 Sv round-trip transit, ~5% from solar energetic particles. https://www.swri.org/newsroom/press-releases/swri-led-team-calculates-the-radiation-exposure-associated-trip-mars

Zhang et al. (2026), Space Weather — The constraint of crewed Mars missions based on current radiation dose measurements. https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2025SW004724

NASA Science (April 2026) — NASA begins implementation for ESA's Rosalind Franklin mission to Mars: Falcon Heavy launch NET late 2028, NASA-supplied braking engines and radioisotope heater units, 2 m drill rationale. https://science.nasa.gov/blogs/mars-rosa/2026/04/16/nasa-begins-implementation-for-esas-rosalind-franklin-mission-to-mars/

Live Science (2026) — NASA's Mars Sample Return is dead, leaving China's Tianwen-3 (2028 launch, 2031 return) as the retrieval path. https://www.livescience.com/space/mars/nasas-mars-sample-return-is-dead-leaving-china-to-retrieve-signs-of-life-from-the-red-planet

Payload (2026) — The commercial race to Mars: Mars Telecommunications Orbiter ~$700m fixed price, delivery required by end-2028, five known bidders including Rocket Lab, Blue Origin and Lockheed Martin. No award had been announced at this paper's cut-off. The ~$4bn total mission cost is Rocket Lab's own bid-stage estimate. https://payloadspace.com/the-commercial-race-to-mars-heats-up/

Planetary Society — Chang'e-6: far-side sample return, June 2024, with the French DORN radon detector and the Swedish NILS negative-ion instrument measuring solar-wind interaction with the surface. https://www.planetary.org/space-missions/change-6

Global Times (February 2026) — Long March 10 / Mengzhou max-dynamic-pressure abort test and China's first sea-based crewed spacecraft recovery; crewed lunar landing targeted before 2030. https://www.globaltimes.cn/page/202602/1355194.shtml

Space.com (August 2026) — Chang'e-7 launch postponed from 23 August 2026 by Tropical Depression Narra; revised timing uncertain between September 2026 and 2027. https://www.space.com/space-exploration/launches-spacecraft/bad-weather-delays-launch-of-chinas-historic-change-7-mission-to-moons-south-pole

NASASpaceflight (July 2026) — Blue Origin accelerates LC-36 recovery via a hybrid integration plan; return to flight targeted before the end of 2026. New Glenn NG-3 launched successfully on 19 April 2026 with booster recovery; the vehicle lost on 28 May 2026 was destroyed during a static-fire test on the pad, not in flight. https://www.nasaspaceflight.com/2026/07/blue-origin-lc-36-recovery-hybrid-plan/

Contemporaneous reporting (July 2026) — Long March 10B maiden flight, 10 July 2026, with recovery of the first stage: China's first orbital-class booster recovery.

NASA (July 2026) — Announcement for Collaborative Opportunity on lunar fission surface power: 37 companies named under an unfunded in-kind collaboration instrument. No reactor procurement, contract award or obligation of funds is associated with the selection.

Spaceflight Now (27 May 2026) — FAA requires a SpaceX-led mishap investigation before Starship launches resume; the investigation closed on 13 July 2026. https://spaceflightnow.com/2026/05/27/faa-requires-spacex-led-mishap-investigation-before-resumption-of-starship-launches/

PwC lunar market assessment, via SpaceQ — Lunar economy projected at up to $127bn by 2050. https://spaceq.ca/pwc-report-lunar-economy-projected-to-generate-up-to-us127-billion-by-2050/

TERRESTRIAL INFRASTRUCTURE

Lloyd's / Cambridge Centre for Risk Studies (2025) — Extreme space weather systemic risk scenario: $1.2trn / $2.4trn / $9.1trn, regional breakdown, $17bn probability-weighted expected loss. https://www.lloyds.com/insights/media-centre/press-releases/extreme-space-weather-scenario

Lloyd's / AER (2013) — Solar storm risk to the North American electric grid: $0.6–2.6trn, 20–40 million at risk, 150-year recurrence. https://assets.lloyds.com/assets/pdf-solar-storm-risk-to-the-north-american-electric-grid/1/pdf-Solar-Storm-Risk-to-the-North-American-Electric-Grid.pdf

Oughton et al. (2017), Space Weather / AGU — Daily economic impact scenarios S1–S4 ($4.0–48.5bn/day); 51% of loss via indirect and international channels. https://news.agu.org/press-release/extreme-space-weather-induced-electricity-blackouts-could-cost-u-s-more-than-40-billion-daily/

Coupled physics-engineering-economic modelling (2024) — 10,464 US substations; 100/150/250-year daily losses; explicit lower-bound caveat. https://arxiv.org/html/2412.18032v3

Frontiers in Astronomy and Space Sciences (2026) — Review of economic consequences and national policy: Québec and Malmö costs, VoLL ranges, New Zealand benefit-cost ratios, international policy comparison. https://www.frontiersin.org/journals/astronomy-and-space-sciences/articles/10.3389/fspas.2026.1733091/full

US Department of Energy (2024) — Large Power Transformer Resilience Report to Congress: lead times, import dependency, fleet age, spare programmes. https://www.energy.gov/sites/default/files/2024-10/EXEC-2022-001242%20-

%20Large%20Power%20Transformer%20Resilience%20Report%20signed%20by%20Secretary%20Granholm%20on%207-10-24.pdf

npj Space Exploration (2025) — Space weather disrupts aviation: 2003 diversion fuel and payload figures, HF outage cost estimates, WAAS and vertical guidance loss. https://www.nature.com/articles/s44453-025-00011-y

IntechOpen — Space Weather and Civil Aviation: ICAO advisory statistics November 2019–September 2025, event-scale US airline costs, the 2015 solar radio burst radar incident. https://www.intechopen.com/online-first/1233638

Xue et al. (2023), Space Weather — Economic costs of 2003 Halloween storm effects on Northern-Hemisphere aviation using 2019 flight data. https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2022SW003381

RTI International / NIST (2019) — Economic benefits of GPS: $1.4trn cumulative value, ~$1bn/day outage cost, ten dependent sectors. https://www.rti.org/impact/gps-14-trillion-economic-engine

Younas et al. (2025), JGR Space Physics — Spatio-temporal evolution of mid-latitude GPS scintillation and position errors during the May 2024 solar storm. https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2025JA033839

Space.com reporting on Younas et al. — the >$500m US farm loss figure, six-hour peak, 70% of planted acreage GPS-dependent. https://www.space.com/astronomy/sun/may-2024-solar-storm-cost-usd500-million-in-damages-to-farmers-new-study-reveals

farmdoc daily, University of Illinois (2024) — Farm profitability analysis of the 10 May GPS outage: $90/acre planting-delay penalty, yield response curves. https://farmdocdaily.illinois.edu/2024/05/how-will-the-gps-outage-on-may-10-affect-us-farm-profitability.html

Patterson et al. (2024), Scientific Reports — Modelling electrified railway signalling misoperations during extreme space weather in the UK. https://www.nature.com/articles/s41598-024-51390-3

Schulte in den Bäumen et al. (2014), NHESS — How severe space weather can disrupt global supply chains. https://nhess.copernicus.org/articles/14/2749/2014/

About this analysis. SSEEM v2.0 was built specifically for this paper. It is not a commercial catastrophe model and has not been validated against a loss portfolio — no such portfolio exists for this peril, and for the orbital channels no loss data is published at all. It should be used to compare the relative scale of exposures and to structure a conversation about where resilience investment goes, not as a basis for capital adequacy.

Verification. Every current-programme statement in Parts III and IV was checked line by line against a primary or contemporaneous source on 24 August 2026, and the following were corrected as a result: the Thiele et al. conjunction figures, which the peer-reviewed Acta Astronautica version revises from the December 2025 preprint; the ISS end date, which the March 2026 Ignition realignment reopened; the direction of the FY2026 Space Force budget, whose base request fell rather than rose; the New Glenn sequence, where NG-3 flew successfully on 19 April 2026 and the vehicle lost on 28 May was destroyed in a static-fire test rather than in flight; the status of Lunar Gateway, which NASA describes as stood down pending review rather than terminated; the SR-1 thruster type; the Mars Telecommunications Orbiter competition, which has five known bidders and no award; the lunar fission-power selection, which is an unfunded collaboration announcement rather than a procurement; and slipped dates for Starcloud-2, Griffin-1, Blue Ghost M2 and VIPER. The Long March 10B maiden flight and booster recovery of 10 July 2026 were added.

What is a fact and what is not. Programme statements in this paper are labelled by type. Contract ceilings and maximum potential values are identified as ceilings, not as money awarded or obligated; this applies to Golden Dome's $3.2bn interceptor agreements, VIPER's $190m, and the $4.82bn Near Space Network vehicle. Programme cost totals attributed to an administration or an operator, including Golden Dome's $185bn, are stated as claims rather than as appropriations or independent estimates. Vendor performance claims, including Starcloud's compute comparison, Starlink V3's per-satellite throughput and Thales Alenia Space's gigawatt orbital-compute concept, are attributed to the party making them. Commercial market forecasts are identified as forecasts and are excluded from the model's exposure base, which uses in-orbit assets and contracted programmes only. Two items are flagged in the text as absent from the public record rather than verified: shielding specifications for the lunar surface habitat and reactor, and any radiation architecture for Starship's crew configuration. Where a figure could not be traced to a source that a reader can open, it has been attributed to the reporting that carries it or removed.

On the orbital estimates specifically. Thirty-one of forty-five severity cells are graded L. This reflects the state of public disclosure rather than the state of the underlying engineering: operators plainly model drag and radiation internally, and do not publish. The figures in Part III should be read as an argument about relative magnitude and direction of travel, and as an invitation to the sector to replace them with something better.

A note on precision. Figures are presented to two or three significant figures because that is the arithmetic the model produces, not because the underlying evidence supports that precision. Where a number is graded L in Appendix A, treat the order of magnitude as the claim and the digits as bookkeeping.

About Mission Space

Mission Space builds space-environment intelligence for operators: real-time radiation and space-weather monitoring for agencies, orbital data centres and defence, spanning ground, Earth orbit, cislunar space and the lunar surface. The company operates its own sensing hardware and the software layer that turns the measurements into operational decisions.

This paper sits in Mission Space's Systemic Risk Series. It was written because the exposure it describes is growing faster than the analysis of it, and because the sector that carries most of the new exposure publishes almost nothing about how a storm affects its economics. SSEEM v2.0 and its full parameter set are published in Appendix A so that operators, insurers and researchers can challenge the numbers with better information than we have.

Mission Space US Corp. 66 West Flagler Street, Suite 900, Miami, FL 33130. mission.space · info@mission.space

Correspondence on the model, including corrections and offers of loss data, is welcome and will be reflected in future revisions.