

SEE Mitigation for FPGAs in Low Earth Orbit

Sep 27, 2026 · @Seba

1. Summary

For SRAM FPGAs in LEO, the stack that works is: latchup-safe parts or per-rail latchup protection, continuous configuration scrubbing, distributed TMR with triplicated I/O on control logic, ECC with scrubbing on every memory, and an independent rad-hard supervisor that owns power and reconfiguration. Once scrubbing runs every few minutes or faster, TMR implementation quality and latchup handling decide reliability, not scrub speed.

- **The environment is proton-driven and bursty.** At 400–800 km essentially all >30 MeV trapped-proton fluence arrives in the SAA, 90% of it in 2.5–4.6% of orbit time, and flux rises nearly 60× from 400 to 1,200 km (§2).
- **Rates will rise through 2030.** Cycle 25 peaked in October 2024; low-altitude protons and GCR climb toward the next minimum, so size designs for solar minimum (§2).
- **Latchup is the mission-ending risk for commercial parts.** 7-series and UltraScale+ show 0.1–0.3 A steps on VCCAUX that only a power cycle clears; one UltraScale+ event was destructive without a current limit (§3, §5).
- **Configuration upsets are the most frequent effect.** Even the space-grade XQRKU060 takes about two CRAM upsets a day in GEO, so every SRAM part needs a scrubber with an external watchdog (§4, §6).
- **TMR done wrong can hurt.** Flip-flop-only TMR raised a cross-section 30–50%; shared I/O held the gain to 3× against 86× with triplicated I/O; placement-aware TMR reached 367× (§7).
- **Faster scrubbing stops paying early.** In the worked 800 km model, accumulation MTTF with a 60 ms scan is about 370,000 years, while implementation-limited ceilings run from 0.65 to 80 years (§10).
- **Memories and verification hide traps.** ECC needs write-back scrubbing, DDR with ECC still fails mainly through SEFIs (vendor spread above 100×), and fault injection misses the MCUs, SEFIs and latchup that matter (§8, §11).

Section 12 turns this into baselines by mission class.

2. The LEO radiation environment

In LEO, trapped protons in the South Atlantic Anomaly (SAA) set most single-event rates. Galactic cosmic rays (GCR) and solar particle events (SPEs) add a high-latitude component that grows with inclination. Altitude and solar-cycle phase each move rates severalfold.


Source	Particles	Where it reaches a LEO spacecraft	What it does to an FPGA
Trapped protons (inner belt)	Protons, ~10–400 MeV	The SAA, where the inner belt dips to ~300–1,200 km	Most CRAM/BRAM upsets and SEFIs; nuclear recoils trigger SEL in susceptible CMOS; main TID source
Galactic cosmic rays	Protons to iron and beyond, GeV energies	Only where geomagnetic cutoff is low: polar caps and high latitudes (< 2 GV)	High-LET ions drive MBUs, SEFIs and destructive SEL; cannot be shielded
Solar particle events	Protons plus heavy ions, hours to days	High-inclination orbits during events	Upset rates rise sharply for the event's duration; the most intense events are heavy-ion rich
Trapped electrons	Electrons up to several MeV	Outer-belt horns at high latitude, and the SAA	TID and charging, not a primary SEE driver

Sources: [Barth, Dyer & Stassinopoulos 2003](#); [NASA/TM-20220011775](#).

Altitude. Trapped-proton flux climbs steeply with altitude; >30 MeV flux peaks near 2,500 km at the equator ([Barth et al.](#)). A 1,200 km constellation orbit sits far deeper in the belt than a 500 km rideshare orbit, so rate estimates never transfer between them.

Trapped-proton flux rises nearly 60× from 400 km to 1,200 km

Orbit-averaged trapped-proton flux above 30 MeV (protons/cm²/s), log scale



computed here · AP8MIN/AP8MAX via IRBEM · 7 circular orbits, 3-day average · >30 MeV

Computed here with AP8 in [IRBEM](#) over three days of ground track. Solar-minimum flux is 1.8–1.9× solar-maximum at 400–500 km but only 1.2–1.3× at 800 km and above, the same altitude trend SAMPEX and POES observed. AP8 is a static 1960s-epoch model that does not follow the SAA's drift, so cross-check design margins with AP9.

The fluence arrives in bursts. In the same runs, essentially all >30 MeV fluence at 400–800 km arrived inside the SAA, and 90% of it in just 2.5–4.6% of orbit time (11.9% at 1,200 km). Section 10 shows why that clustering matters for TMR and scrub-rate sizing.

Inclination. Low-inclination orbits are SAA-dominated and well shielded from GCR; NASA modeling puts the equatorial GCR contribution about five times below an inclined orbit ([NASA TM](#)). ISS-like (51.6°), sun-synchronous and polar orbits add polar-cap GCR and SPE exposure, which brings the high-LET ions that cause SEL and SEFI.

Solar cycle. Cycle 25 peaked in October 2024 at a smoothed sunspot number of 161 ([SIDC](#)). Low-altitude trapped protons are anticorrelated with solar activity: SAMPEX saw ~300% modulation at 500 km versus ~21% at POES altitude, with the flux peak lagging the F10.7 minimum by ~500 days ([Bregou et al. 2022](#)). GCR also rises toward minimum, and Cycle 26 is expected to begin between 2029 and 2032 ([NOAA SWPC](#)).

A mission launched in 2026–2030 will therefore see SEE rates climb through its life. Size the design for solar-minimum protons and GCR, plus a worst-case SPE for high-inclination orbits.

The SAA is moving. ESA Swarm data show the anomaly grew by nearly half the area of continental Europe since 2014, with its eastern lobe weakening fastest since 2020 ([ScienceDaily, Feb 2026](#)). NASA also reports a northwesterly drift and a split into two lobes ([NASA](#)). Operations gated on an SAA map should use a current field model or on-board upset counts, not a fixed polygon.

TID is modest. NASA's modeling behind 5 mm Al gives 136, 430 and 333 rad(Si)/year for 500 km orbits at 0°, 51° and 89° ([NASA TM](#)). SmallSat avionics built on COTS FPGAs are typically specified to 20–30 krad(Si) ([NASA SoA 2026](#)). Models and LEO flight dosimetry differ by factors of 3–30, usually with the models high ([Electronics 2023](#)).

Scaled cells see more of the spectrum. Sub-100 nm cells can be upset by low-energy protons through direct ionization, raising SEU cross-sections by up to several orders of magnitude at those energies ([NASA proton SEE guideline](#)). For 28, 16 and 7 nm FPGAs this makes proton test energy selection part of the rate estimate.

Flight evidence. The Cibola Flight Experiment (560 km, 35.4°) logged 849 configuration upsets in nine Virtex XQVR1000s over 3,160 device-days: 0.27 per device-day, mostly in the SAA, below the 0.5 predicted ([Caffrey et al. 2009](#)).

3. Single-event effects in FPGAs

In an SRAM FPGA the most frequent effect is a configuration upset that silently rewires the circuit until it is scrubbed. Latchup is the effect that can end the mission. Everything else sits between those two.

Effect	What happens	Where it lands	Lasts until	Main mitigations
Configuration SEU	A configuration bit flips, changing a LUT, route or setting	SRAM-configured FPGAs only; flash, SONOS and antifuse cells are immune	Scrub or reconfiguration	Scrubbing (§6), TMR (§7)
User-state SEU	A flip-flop, BRAM/URAM, LUTRAM/SRL or DSP register flips	Every FPGA type	Overwritten or reset	ECC, TMR, safe FSMs (§7–8)

Effect	What happens	Where it lands	Lasts until	Main mitigations
MCU / MBU	One strike flips several adjacent cells	Dense CRAM and BRAM; grows with scaling	As an SEU	Interleaving plus ECC, floorplanned TMR (§6-7)
SET	A transient on a combinational, clock or reset net gets latched	Every type; the main residual in flash and antifuse parts	Next clock edge, or as an SEU once latched	SET filters, triplicated clocks and resets (§7)
SEFI	An upset in control logic stalls a block or the device: configuration controller, POR, JTAG/ICAP, CPU, DDR controller, transceiver, NoC	Every type, especially hard IP	Block reset, reconfiguration or power cycle	Watchdogs, supervisor, recovery ladder (§9)
Micro-SEFI	One control-cell upset corrupts many bits at once	Observed in 28, 20 and 16 nm SRAM FPGAs	As an SEU	System-level checks (§8)
SEL / micro-latchup	A parasitic thyristor turns on; current steps from ~0.1 A to amps	Bulk CMOS, often I/O and auxiliary rails	Power cycle; may be destructive or leave latent damage	Part selection, current sense plus power cycle (§5)
SEB / SEGR	Burnout or gate rupture in power MOSFETs	POL converters and load switches feeding the FPGA	Permanent	Derating, rad-hard power parts (§5)

Only essential bits matter. Designs typically use under 30% of configuration bits, and the share of upsets that affect operation is often under 10% and typically under 5% ([AMD WP395](#)). AMD's own example design has 0.49% essential bits ([PG187](#)), so the rate that matters is set by the design, not the die.

Per-bit sensitivity falls; bit counts rise. 16 nm UltraScale+ CRAM measured 2.5×10^{-16} cm²/bit under neutrons, about 28× below 28 nm devices, while BRAM, OCM and caches sat at $1.5\text{--}3 \times 10^{-15}$ cm²/bit ([Anderson et al. 2018](#)). Meanwhile XQRVC1902 carries 363 Mb of CRAM against 193 Mb in the XQRKU060 ([AMD, SEFUW 2025](#)). Device-level rates fall much more slowly than per-bit figures suggest.

MCUs are what break TMR. Interleaving turns most MCUs into single-bit errors per ECC word: a ZU9EG neutron campaign saw 93.8% single-bit upsets and no same-frame MBUs ([arXiv 2303.08098](#)). Yet in beam tests of TMR designs on 7-series parts, every observed TMR failure traced to MCUs hitting routing shared by two domains ([Cannon et al. 2018](#)).

Micro-SEFIs are rare but defeat ECC. BYU recorded single events corrupting 16 to 5,311 BRAM bits; one 1,024-bit event hit the same bit in 512 words, so SECDED flagged every word uncorrectable. They were at least 10× rarer than MCUs: 16 against 7,628 in a KU040 ([Pérez-Celis et al. 2021](#)).

Latchup in commercial FPGAs is usually a small current step. 28 nm 7-series parts showed 105–125 mA steps on VCCAUX that only a power cycle cleared ([Swift, XRTC 2014](#)). 16 nm UltraScale+ showed VCCAUX high-current events at every LET tested (3.2–20 MeV·cm²/mg); one destroyed the part with no current limit, and only a full power cycle cleared the rest ([Lee et al. 2018](#)). A step that small can hide inside normal load swings, which drives SEL detector design in §5.

The power chain fails too. SEB and SEGR strike the MOSFETs in point-of-load converters and load switches, not the FPGA, but they take the FPGA down just the same.

4. FPGA technology and device landscape

Configuration technology decides where the mitigation effort goes. Non-volatile cells remove configuration upsets but leave user flip-flops and memories exposed. SRAM parts need scrubbing plus TMR, and commercial SRAM parts add latchup protection on top.

- **Antifuse, flash and SONOS** (RTAX, RTG4, RT PolarFire): no configuration upsets and no scrubber. Effort moves to user flip-flops, memory EDAC and SETs.
- **Rad-hard-by-design SRAM** (Virtex-5QV, NanoXplore): hardened configuration cells cut CRAM rates by orders of magnitude, but scrubbing stays.

- **Rad-tolerant SRAM** (XQRKU060, XQR Versal, Frontgrade/Lattice RT): SEL-immune and characterized, but standard-density CRAM still upsets daily.
- **Commercial SRAM** (7-series, UltraScale+, Versal): best performance per watt and dollar, with SEL, lot variation and every mitigation layer on you.

Space-grade and rad-tolerant families

Family	Configuration	TID, krad(Si)	SEL-free to LET (MeV·cm ² /mg)	Configuration upsets	Built-in hardening	Qualification / status
Microchip RTAX-S/SL	Antifuse, one-time programmable	300 functional, 200 parametric	>117	Immune	Built-in TMR flip-flops (<1×10 ⁻¹⁰ errors/bit-day, GEO); optional SRAM EDAC	Long flight heritage
Microchip RTG4	Flash, 65 nm	>100	>103	Immune	SEU-hardened flip-flops (immune to LET >37), SECDED SRAM, optional SET filters	QML Class V (Synopsis)
Microchip RT PolarFire	SONOS, 28 nm	100	80 (RTPF500ZT); 37 at 3.3 V GPIO (RTPF500T)	Immune	SECDED LSRAM; flip-flops unhardened, ~1×10 ⁻⁷ → ~1×10 ⁻¹¹ errors/bit-day (GEO) with TMR	RTPF500ZT: MIL-STD-883 Class B and QML Class Q (2025); RT PolarFire SoC sampling
AMD Virtex-5QV	RHBD SRAM, 65 nm	1,000	>100	3.8×10 ⁻¹⁰ upsets/bit/day (GEO)	Dual-node CRAM latches (~1,000× harder), TMR configuration control, 800 ps SET filters	Space-grade since 2011

Family	Configuration	TID, krad(Si)	SEL-free to LET (MeV·cm ² /mg)	Configuration upsets	Built-in hardening	Qualification / status
AMD XQRKU060	SRAM, 20 nm	100–120	80 at max V and T _j	9.5×10^{-9} upsets/bit/day (GEO): ~1.8 per device-day	CRAM ECC and CRC for SEM IP, BRAM EDAC	Class B and Y, –55 to +125 °C
AMD XQR Versal (VC1902, VE2302)	SRAM, 7 nm	100 (passed 120)	80 at max V and T _j	EDAC plus interleaving, no uncorrectable events in LEO or GEO	XilSEM on hard-wired TMR MicroBlaze in the PMC; ECC caches and OCM	Class B; VE2302 shipping from Q4 2025 (AMD)
NanoXplore NG-Medium	RHBD SRAM, 65 nm	>100	>60	$<1.7 \times 10^{-4}$ per device-day (GEO)	RHBD configuration and registers, CMIC bitstream check, EDAC	Space-grade
NanoXplore NG-Ultra	RHBD SRAM, 28 nm FD-SOI	50	Immune (FD-SOI)	RHBD cells	RHBD configuration and registers, EDAC, TMR SPI boot, Reed-Solomon DDR; quad Cortex-R52	First ESCC 9030 part, Jan 2026 (EE Times Asia)
Frontgrade Certus-NX-RT / CertusPro-NX-RT	SRAM, 28 nm FD-SOI	100–300	≥80	Frame-based SED/SEC	EBR ECC; plastic package	PEM-INST-001 Level 1/2

RT PolarFire shows why the "SEL-free" column needs its test conditions: the RTPF500T's threshold was 63–68.5 at 2.5 V I/O and 100 °C ([EDN 2020](#)) but 37 on 3.3 V GPIO. Rad-tolerant is not upset-free either: an XQRKU060 still takes about two CRAM upsets a day in GEO.

Commercial families common in LEO

Family	Configuration	What testing showed	What you must add
Artix-7, Kintex-7, Zynq-7000	SRAM, 28 nm bulk	105–125 mA steps on VCCAUX, cleared only by power cycle (Swift 2014); flight products rated ~20–25 krad	SEL detect and power cycle, SEM scrubbing, TMR by criticality, supervisor
Kintex UltraScale	SRAM, 20 nm	No latchup-like events in NASA heavy-ion tests (Berg et al. 2017)	Scrubbing, distributed TMR
Kintex UltraScale+, Zynq UltraScale+ MPSoC	SRAM, 16 nm FinFET	VCCAUX high-current events from LET 3.2 up, one destructive without a current limit (Lee et al. 2018); in 15+ SmallSat products, typically rated 30 krad	Latchup protection is mandatory; SEM, PS ECC and lockstep, supervisor
PolarFire, PolarFire SoC	SONOS, 28 nm	Configuration immune; SmallSat products rated 30 krad with software EDAC	TMR for flip-flops, SECEDED, lot-specific SEL data
Lattice Nexus (CrossLink-NX, Certus-NX)	SRAM, 28 nm FD-SOI	CrossLink-NX LIFCL-40: no TID failures at 200 krad, first at 250 krad (NEPP 2023)	SED/SEC scrubbing, TMR by criticality

Flight-product ratings come from NASA's 2026 SmallSat avionics survey, which notes the usual pattern: COTS processor and FPGA first, surrounded by rad-hard ECC, watchdogs, scrubbing and redundancy ([NASA SoA 2026](#)).

5. Latchup (SEL) mitigation

Treat SEL as a part-selection problem first and a circuit problem second. Buy parts that are SEL-free to LET 60–80 MeV·cm²/mg at maximum voltage and temperature where you can. Where you fly commercial parts, detect current steps rail by rail and power-cycle the whole device fast.

Selecting and qualifying parts

- **Heavy-ion data is not optional for COTS.** In NASA's dataset about 50% of commercial CMOS parts were SEL-susceptible and about half of those failed destructively. Rates spanned seven orders of magnitude, and vendor or process gave no predictive power ([Ladbury et al.](#)).
- **Proton susceptibility follows the heavy-ion threshold.** NASA's proton guideline waives proton SEE testing when heavy-ion LETth exceeds 37 MeV·cm²/mg, and requires 100–200 MeV proton SEL testing at LETth ≤20 ([LaBel et al. 2009](#)). The UltraScale+ VCCAUX events started at LET 3.2, so they are reachable by SAA protons.
- **Test at worst case.** SEL cross-section roughly doubles from 25 to 100 °C, and thresholds fall as supply voltage rises ([Sexton 2003](#)). Test at maximum junction temperature and at the I/O voltages you will fly, as the RT PolarFire example in §4 shows.
- **Buy the tested lot.** Hardness of commercial parts can be lot-specific, and short product life cycles leave a narrow window to buy what you tested ([NASA RHA guidelines](#)).

Why the auxiliary rails latch. A latched thyristor needs roughly 1.0 V to hold, so sub-1 V core rails rarely sustain latchup ([Sexton 2003](#)). The 1.8 V VCCAUX and the 2.5–3.3 V I/O banks can, which matches where 7-series and UltraScale+ events appeared.

Designing the protection circuit

1. **Sense each susceptible rail on its own.** A 100–330 mA step vanishes inside a core budget that exceeds 30 A on an XQRKU060 VCCINT ([TI SLVAFM2](#)). Monitor VCCAUX, each VCCO bank group and transceiver auxiliary rails separately.
2. **Use two detectors.** A fast hardware trip catches classic latchup; published designs set it near 4× nominal ([Layton et al. 1997](#)). A slower step detector against a tracked baseline catches micro-latchups that never reach the hard limit.
3. **Cut every rail, in sequence.** Power down all FPGA rails in the vendor's order, not just the rail that tripped, then restore and reconfigure.
4. **Hold off long enough.** Latched states draw hundreds of mA but hold with under 10 mA ([Maurer & Kinnison](#)), so bulk capacitance must decay fully. Layton et al. recovered devices with off-times from 45 μs to 2.5 ms.
5. **Block back-powering.** Powered neighbors can keep a latch alive through input protection diodes. Layton et al. needed 511 Ω input series resistors where 91 Ω failed; isolation switches or tri-stated neighbors do the same job.

6. **Let thresholds move with dose.** TID raises supply current over the mission, so a fixed threshold drifts toward false trips ([Layton et al.](#)). Track the baseline in telemetry and re-set thresholds by command.
7. **Log and trend every event.** Latent damage appeared in about 30% of protected parts that latched, with aluminum cross-sections cut by up to two orders of magnitude ([Sexton 2003](#), citing Becker et al.). Record rail, magnitude, time and orbit position, and set a policy for parts that latch repeatedly.
8. **Calibrate against beam data.** Only testing your lot at temperature shows the real step size; ZU9EG VCCAUX events averaged 330 mA under neutrons ([Lee et al. 2018](#)).

Redundancy does not cure SEL-prone parts. Destructive-SEE failures grow nonlinearly with mission time, and rates above $\sim 10^{-5}$ per day erode a redundant system's resilience. In NASA's analysis, adding current limiting and power-cycle recovery raised confidence for a 50-part COTS system from 26% to 52% ([Ladbury et al.](#)). Protection helps, but part count and part choice dominate.

The power chain: SEB and SEGR

N-channel power MOSFETs have burned out at 22–90% of rated breakdown voltage, and SEB thresholds often sit near 50% of rating ([Sexton 2003](#)). Derate V_{DS} below the characterized SEE safe operating area and use regulators with published SEE data. TI's rad-tolerant (-SEP) and rad-hard (-SP) FPGA power families are characterized to 43 and 75 MeV·cm²/mg respectively ([TI SLVAFM2](#)).

6. Configuration memory mitigation

For SRAM FPGAs the scrubber is the most important single mitigation. TMR without scrubbing only delays failure, because configuration upsets accumulate until two domains are hit. The design questions are where the scrubber lives, how fast it runs, and what happens when it cannot repair.

Approach	How it works	Strengths	Weaknesses	Examples
Blind scrubbing	An external controller rewrites every frame from a golden image on a timer, with no readback	Simple; fixes errors it never detects	No telemetry; rewriting dynamic frames corrupts LUTRAM and SRLs unless masked	GRSCRUB blind mode
Internal readback	An on-chip controller reads frames through ICAP, checks frame ECC and device CRC, and repairs	No extra parts; fast; classifies upsets with essential-bit data; built-in error injection	The scrubber can itself be upset; cannot recover a configuration-controller SEFI	SEM IP , XiISEM, Lattice SED/SEC, NanoXplore CMIC
External readback	A rad-hard FPGA, ASIC or MCU reads frames over SelectMAP, compares CRC or full frames against the golden image with masks, and rewrites	Scrubber immune to the target's upsets; can detect SEFIs and reconfigure	Board area, golden image and mask storage, interface bandwidth	GRSCRUB (Kintex UltraScale, Virtex-5), RTG4 or RT PolarFire supervisors

GRSCRUB checks the configuration interface by reading a known frame before it writes, so a SEFI in the port cannot turn a repair into an overwrite. That check is worth copying in any external scrubber.

Sizing the scrub rate

- **Know your scan time.** Reading 193 Mb through a 32-bit, 100 MHz port takes about 60 ms before overhead. Bedi reports about 1.2 ms to mitigate a single-bit error on the XQRKU060 at 90 MHz ICAP ([Bedi, EDN 2020](#)). NASA GSFC scrubbed Kintex UltraScale at 10 Hz in beam tests ([Berg et al. 2017](#)).

- **Set the period from the TMR math.** The scrub period must be far shorter than the mean time between essential-bit upsets in one TMR domain; §10 works an example.
- **Scrub hardest in the SAA.** With 90% of the proton fluence arriving in 3–5% of orbit time (§2), an orbit-average rate hides bursts. Run the scrubber continuously, or at its maximum rate during SAA passes.
- **Make the period tunable.** NanoXplore's CMIC runs from 5.3 ms to 65 days per cycle at 50 MHz ([Bedi, EDN 2020](#)); a configurable period lets operations tune it in flight.

Golden images and the boot chain

1. **Store golden bitstreams in radiation-tolerant NVM**, such as rad-hard NOR flash or MRAM, with CRC-protected copies. A 25 nm NAND flash showed about 100× the SEU cross-section of DDR3 at low LET and much poorer TID tolerance ([Grürmann et al. 2014](#)).
2. **Use MultiBoot with a fallback image.** UltraScale+ falls back to the golden image on an IDCODE error, a CRC error or a configuration watchdog timeout (about 15 s at default CCLK). If the fallback also fails, configuration stops with INIT_B and DONE low ([AMD XAPP1296](#)), so an external supervisor must own the last resort.
3. **Version masks with bitstreams.** A readback scrubber must skip dynamic bits, and after partial reconfiguration the golden-frame ECC and CRC must be recomputed or the scrubber reports false errors ([Bedi, EDN 2020](#)).

What scrubbing cannot do

- **Repair contents that change in operation.** LUTRAM, SRLs and BRAM data are masked, so they need ECC or TMR (§8). Keep LUTRAM and SRLs out of critical paths on SRAM parts.
- **Recover the configuration controller.** The XQRKU060's CRAM SEFI rate is 4.5×10^{-4} per device-day in GEO ([DS882](#)), about one in six years. AMD puts XilSEM on Versal at one in 200 years, because it runs on the PMC's hard-wired TMR MicroBlaze ([AMD, SEFUW 2025](#)).
- **Protect itself.** A fabric scrubber like SEM IP has its own small failure rate ([PG187](#)). An external supervisor should watch its heartbeat and status, and treat an "uncorrectable" report or a micro-SEFI signature as a trigger for full reconfiguration.
- **Prevent errors between scans.** Scrubbing bounds how long an upset lives; only TMR or ECC keeps outputs correct during that window.

7. User-logic mitigation

On SRAM FPGAs, TMR pays off only when logic and registers are both triplicated, voters sit on feedback paths, and the three domains share no clock buffers, I/O or routing. Half-measures can make things worse: triplicating only flip-flops raised an Artix-7 design's neutron cross-section by 30–50% ([Keller & Wirthlin 2021](#)).

Variant	What is triplicated	Result on SRAM FPGAs	Where it fits
Local TMR (LTMR)	Flip-flops, with voters	NASA GSFC: "not acceptable in SRAM-based FPGAs for any design" (Berg et al. 2017)	Flash and antifuse parts, where configuration is immune; built into RTAX and RTG4 flip-flops
Distributed TMR (DTMR)	Flip-flops and combinational logic, voters after registers	The strongest scheme NASA GSFC tested on SRAM FPGAs	Default for SRAM parts
Global TMR (GTMR)	Everything, including clock sources and I/O	NASA GSFC advises against it: clock-skew races between domains	Triplicate clock buffers and I/O, not clock sources
Block TMR (BTMR)	Whole blocks voted at their boundaries	Protects IP you cannot edit; internal state needs a reset to resynchronize	Encrypted or vendor IP (Synplify)
Partial TMR	Selected clusters	Full TMR cut cross-section ~99%; one feedback cluster gave 21% at 12% of components; random picks hurt (Keller & Wirthlin)	Area-limited designs, protecting large related clusters first

Implementation rules that decide effectiveness

1. **Vote on every feedback path.** Feedback voters resynchronize state after the scrubber repairs a domain. With feedback TMR plus scrubbing, a LEON3 survived 14,455 upsets on average against 282 unmitigated ($51\times$), and its sensitive bits fell from 240,000 to 4,700 ([Wirthlin 2016](#)).
2. **Triplicate I/O.** On an Artix-7 under neutrons, TMR with triplicated I/O gained $86\times$; the same TMR with shared I/O gained $3\times$. Splitting nets early and placing reduction voters beside the pins recovered $26\times$ where I/O could not be triplicated ([Cannon et al. 2020](#)). At board level, vote outputs externally or use minority-voted tri-state outputs.
3. **Give each domain its own clock buffer and reset.** Routing failures, clock and non-clock, were the largest single points of failure in the same study. Drive each domain from its own buffer off one source, and release resets synchronously per domain.
4. **Floorplan against routing common modes.** Every TMR failure in 7-series beam tests traced to MCUs in routing shared by two domains. SEU-aware incremental placement reached $367\times$ over unmitigated and $5\times$ over baseline TMR ([Cannon et al. 2018](#)). On RT PolarFire, Libero's spatial separation puts triplets on different cluster clocks ([Microchip](#)).
5. **Stop the tools from undoing it.** Synthesis merges identical logic. Use a TMR-aware flow, such as Synplify (LTMR, DTMR, BTMR, DWC, Hamming-3 and safe FSMs), [Precision Hi-Rel](#), AMD's ITAR-controlled [TMRTTool](#) for Virtex-4QV/5QV, or BYU's open-source [SpyDrNet-TMR](#). Then check the routed netlist and fault-inject it.
6. **Budget the cost.** Across 15 designs, TMR averaged $3.6\times$ area and added 16% to the critical path, ranging from 4.9% to 41.1% ([Wirthlin 2016](#)). Power scales roughly with area.

Cheaper protection where detection is enough

- **Duplicate with compare (DWC).** Twice the area; flags a mismatch so the frame can be dropped, recomputed or retried. Fits streaming image and SDR pipelines.
- **Algorithm-level checks.** Checksums across matrix operations, end-to-end CRCs on packets, and range checks on outputs catch datapath errors at a fraction of TMR's cost. AMD reports fault-aware training cut SEU-induced datapath errors about $4\times$ on Versal AI Engines ([AMD, SEFUW 2025](#)).
- **Safe state machines.** Safe-case or Hamming-3 encodings guarantee recovery from illegal states; pair them with a progress watchdog.

- **SET filtering.** In flash and antifuse parts, SETs are the main residual. RTG4 offers optional SET filters and Virtex-5QV filters up to 800 ps; elsewhere, the chance of latching a transient grows with clock rate.

Where to spend partial TMR

Protect in this order: control and feedback logic (state machines, counters, sequencers, configuration registers), then anything that can hang the system (bus arbiters, DMA engines, command decoders), then datapaths only if a wrong output is worse than a dropped one. Triplicate whole clusters of related logic rather than scattered cells, which is the pattern that performed best in the Artix-7 study.

8. Memories, processors and hard IP

ECC corrects on read but does not clean the memory, so single-bit errors pile up into uncorrectable doubles unless something writes corrected data back. Once ECC handles bit flips, external DDR fails mainly through SEFIs in its control logic. Hard processors bring their own ECC and lockstep, but they also bring new SEFI modes.

On-chip memories

Memory	Built-in protection	What you add
AMD block RAM	One 64-bit SECDED encoder per 36 Kb block; single-bit errors are corrected in the output data and flagged on SBITERR/DBITERR (UG473)	Enable ECC; add a background read-correct-write-back scrubber; treat bursts of DBITERR as a reload trigger
LUTRAM and SRLs	None, and masked from configuration scrubbing	TMR, or move the storage into ECC block RAM
RT PolarFire LSRAM	Optional SECDED: $\sim 5 \times 10^{-8}$ falls to $\sim 5 \times 10^{-15}$ errors/bit-day in GEO (Microchip)	Enable it; scrub long-lived tables
RTG4 LSRAM and uSRAM	Optional SECDED	Enable it

Block RAM ECC needs a scrubber because nothing else rewrites the stored word. BYU's micro-SEFIs make a second point: a single event that corrupts a whole BRAM column makes SECEDED report every word uncorrectable ([Pérez-Celis et al. 2021](#)). Reload or recompute the whole memory when that pattern appears.

External memory

- **Pick DDR parts on SEFI data.** In one DDR3 study, row-SEFI cross-sections ranged from 1.7×10^{-5} to 4.4×10^{-3} cm² per device across vendors, a spread of more than 100× ([Grürmann et al. 2014](#)).
- **Refresh the mode registers.** Periodically rewriting DDR3 control registers brought device-SEFI cross-sections from as high as 1.25×10^{-4} down to about 5×10^{-6} cm² per device. Power cycling clears what remains but erases the contents.
- **Use symbol codes that survive a lost chip.** Reed-Solomon ECC handles sparse SEUs well, and a code sized to tolerate a whole device's failure turns a DDR SEFI into a correctable erasure. NG-Ultra's controller uses Reed-Solomon that survives losing two 16-bit devices ([ESCIES](#)).
- **Check the ECC width.** Zynq-7000's DDR controller supports ECC only in 16-bit mode ([DS190](#)); a 32-bit design needs a PL-side ECC controller.

Hard processor subsystems

Device	Built-in protection	Gaps to close
Zynq-7000	Byte parity on L2 and OCM; three watchdog timers (DS190)	Parity only detects: trap the error, invalidate and reload; route DDR through ECC
Zynq UltraScale+ MPSoC	Cortex-R5F lockstep, 64-bit ECC on OCM, L2 ECC, DDR ECC, system watchdogs (UG1085 , UG1137)	Run FDIR on the lockstep R5F; treat the A53 cluster as payload that may crash
XQR Versal	ECC caches and OCM; hard-wired TMR MicroBlaze in the PMC running XilSEM	AMD estimates ~1.3 PS SEFIs per device-year at 500 km, 20°; top signatures are core hangs (20%) and NoC NPS errors (26%) (AMD, SEFUW 2025)

Workload matters as much as silicon. Simple ZU9EG benchmarks showed no crashes or silent errors under neutrons ([Anderson et al. 2018](#)). Heavier multi-threaded, vision and DPU workloads on the same part produced crashes and silent errors at 10^{-11} to 10^{-9} cm² per device ([arXiv 2303.08098](#)). The DPU's tolerable output errors ran higher still, at 5×10^{-8} cm². Characterize the software you will fly.

Soft processors

- **AMD TMR MicroBlaze subsystem.** Three MicroBlaze cores with a TMR Manager, comparators and voters, a lockstep fail-safe option, fault injection (TMR Inject) and integrated SEM scrubbing ([PG268](#)).
- **Fault-tolerant cores on rad-hard fabric.** A LEON3FT on RTG4 is a common pairing in current SmallSat computers ([NASA SoA 2026](#)). On SRAM parts, feedback TMR plus scrubbing gave a LEON3 51× more upsets to failure (§7).


Software-level mitigation

- Walk DDR and OCM with a background scrub task so single-bit errors are fixed before a second hit lands.
- Kick watchdogs only after health checks pass, not from a timer loop that survives a hung application.
- Keep critical variables in triplicate or with checksums, and verify code segments against a golden CRC.
- Report every corrected error, parity trap and reset cause to telemetry; these counters are your in-flight radiation monitor.

9. SEFI handling and system-level FDIR

Every mitigation above eventually meets a fault it cannot see: a hung configuration port, a stalled CPU, a latched rail. An independent supervisor that is itself SEL-free and immune to configuration upsets, with authority over power and configuration, is the backbone of an FPGA payload in LEO.

An independent rad-hard supervisor owns power, configuration and recovery



supervisor architecture · 6 elements · who powers, scrubs and watches whom

The supervisor is the only element that can power-cycle or reconfigure the FPGA, so it is built from a flash FPGA such as RTG4 or RT PolarFire, or a rad-hard microcontroller. Everything it watches reports into it, and it reports to the spacecraft.

The recovery ladder

Escalate one rung at a time, and log every rung taken:

1. **Correct in place:** ECC, SEM or XiISEM repair, with no interruption.
2. **Reset the function:** restart the IP block, DMA engine or software task that stopped making progress.
3. **Reboot the processor or reload a partition:** warm reset of the PS, or partial reconfiguration of one region.
4. **Reconfigure the device:** full reload through PROGRAM_B or IPROG, falling back to the golden image.
5. **Power-cycle in sequence:** the only cure for latchup and for SEFIs no reset clears, as the UltraScale+ VCCAUX events showed.
6. **Switch to the spare or safe mode:** after repeated failures, or any sign of permanent damage.

Watchdogs that actually watch

- Use a challenge-response or windowed heartbeat, not a toggling pin; a free-running counter keeps toggling while the application is hung.
- Nest them: a PS watchdog inside the SoC, the supervisor watching the SoC, and the spacecraft watching the supervisor.
- Watch the scrubber too. A silent SEM or XilSEM is a SEFI in its own right.

SAA-aware operations

With 90% of the proton fluence landing in 3–5% of orbit time (§2), gating a few operations during SAA passes removes most of the exposure for very little lost time. Operators commonly switch off non-essential components through the SAA, and ISS instruments such as GEDI see reset events that cost a couple of hours of data about once a month ([NASA](#)).

- Flag SAA passes from the propagated orbit and confirm with live counters, such as the SEM correction rate.
- During a pass: scrub at maximum rate, defer flash writes, reconfiguration and firmware updates, and checkpoint state before entry.
- Update the SAA boundary as the anomaly drifts rather than freezing a pre-launch polygon.

Telemetry that closes the loop

Time-stamp and geolocate every corrected upset, ECC event, SEFI, SEL trip and reset. The Cibola Flight Experiment's upset map is what showed its rates were SAA-dominated and below prediction ([Caffrey et al. 2009](#)). Trend supply currents too: a slow rise is TID, and a step that stays is a latent latchup warning.

10. Quantitative analysis: rates and a worked TMR model

At realistic scrub rates, upset accumulation almost never limits a TMR design in LEO; the measured quality of the TMR implementation does. The worked example below shows where that crossover sits, and why SAA bursts move it by 10–44×

Rate-prediction workflow

1. **Environment.** Trapped protons from AP8 or AP9 at solar minimum; GCR and SPEs from CREME96 with the orbit's geomagnetic transmission; worst-day and peak-5-minute SPE cases for high inclinations. SPENVIS and OMERE wrap these models ([NASA TM](#)).
2. **Device response, per resource.** Heavy-ion Weibull fits (LETth, W, S, σ_{sat}) and proton Weibull or Bendel fits for CRAM, BRAM, flip-flops, SEFI and SEL, at worst-case voltage and temperature. For 16 nm UltraScale+ CRAM: LETth 1.5, σ_{sat} 0.01 $\mu\text{m}^2/\text{bit}$, W 8.0, S 2.2 ([Lee et al. 2018](#)).
3. **Rates.** Fold heavy-ion curves with LET spectra (RPP or IRPP) and proton curves with the orbit's proton spectrum, then add them: SEE rates are additive ([Ladbury 2017](#)).
4. **Design scaling.** Multiply the CRAM rate by essential bits, or better, fault-injection-measured critical bits. Add BRAM without ECC. Keep SEFI and SEL separate, since their consequences differ.
5. **Margins.** NASA's proton guideline suggests a 10 $\times$ margin on SEU rates built from higher-energy data, and 200–400 $\times$ on SEL rates for parts with LETth of 20–37 ([LaBel et al. 2009](#)).
6. **Flight check.** Compare predictions with on-orbit counters (§9).

For protons the fold reduces to a simple product when the cross-section saturates above an energy threshold:

$$R_p = \int \sigma(E) \phi(E) dE \approx \sigma_{\text{sat}} \Phi(E > E_{\text{th}})$$

A TMR plus scrubbing model

A TMR design fails through two paths: two domains holding unrepaired upsets at once (accumulation), or one upset that defeats TMR on its own through shared routing, I/O or clocks (the residual). With per-domain essential-bit upset rate λ_d , scrub period T_s and measured TMR gain G over the unmitigated design:


$$\lambda_{\text{fail}} \approx \frac{\lambda_{\text{unmit}}}{G} + 3 \lambda_d^2 T_s P(T_s), \quad \lambda_d = n_e \sigma \Phi$$

The burst penalty $P(T_s)$ corrects for protons arriving in SAA bursts rather than uniformly. Computed from the AP8 time series, P is 44 $\times$ at 525 km SSO, 22 $\times$ at 800 km SSO and 10 $\times$ at 1,200 km for scrub periods under a minute, falling to 1 by about 12 hours.

Worked example assumptions: 800 km, 98.6° SSO; AP8MIN >30 MeV flux of 116.6 protons/cm²/s over 30 days; a CRAM cross-section of 2.5×10^{-16} cm²/bit, the 16 nm neutron value ([Anderson et al. 2018](#)) used as a proxy for high-energy protons; 5 Mb of essential bits per domain, every essential upset counted as a domain failure; protons only. That gives 0.0126 essential upsets per domain-day, so the unmitigated design's MTTF is about 79 days.

Once scrubbing runs every few minutes, TMR quality sets the MTTF

Mean time to TMR failure from upset accumulation (years) vs scrub period; 800 km SSO example, log-log



computed here · AP8MIN 30-day ground track at 800 km, 98.6° · assumptions listed above · ceilings from Cannon et al. 2018 and 2020

With a 60 ms full-device scan, accumulation MTTF is about 370,000 years even with bursts. It drops below the 367× ceiling (80 years) once scrub periods pass 5–10 minutes, and below the 86× ceiling (19 years) past 3–6 hours. Below a few minutes, only TMR implementation quality moves the result; the shared-I/O case caps MTTF at 0.65 years whatever the scrub rate.

Scaling rules. The residual term scales with flux, essential bits and cross-section; the accumulation term scales with their square times P. Moving the same design from 800 to 1,200 km raises flux 4.9× and shortens accumulation MTTF about 10× (P falls from 22 to 10); moving to 525 km SSO lengthens it about 15×.

11. Verification and qualification

Fault injection proves the mitigation logic works; only beam testing proves it works against the real mechanisms. The failures that break TMR (MCUs, SEFIs, routing common modes) and the one that ends missions (SEL) are invisible to single-bit injection.

Beam testing

Test	Particles	What it establishes	Conditions that matter
SEL screening	Heavy ions to LET 60–80; 100–200 MeV protons if LET _{th} ≤ 20	Latchup threshold, step size per rail, destructive or not	Maximum temperature and voltage; per-rail current logging; current-limited supplies; flight lot
Cross-sections	Heavy ions for Weibull fits; protons at three or more energies, such as 60, 120 and 190 MeV (LaBel et al. 2009)	CRAM, BRAM, flip-flop and MCU inputs to rate prediction	Readback after each run; MCU pattern analysis
SEFI modes	Heavy ions and protons	Which recovery rung each SEFI needs	Realistic workloads: AMD generated over 10 million processor tests under beam for XQR Versal (AMD, SEFUW 2025)
Mitigated design	Neutrons, protons or heavy ions	The measured TMR gain G and the residual single points of failure	Scrubber and FDIR running as in flight

Modern FPGAs are flip-chip, so ions must reach the active layer through the substrate: thin the backside or use high-energy beams. The JESD57 revision defines a worst-case penetration range and notes that tilt and roll angles can both reveal SEEs ([NASA, JESD57 update](#)). Proton-only campaigns are cheaper and closer to LEO, but 10¹⁰ protons/cm² at 200 MeV would miss SEL in most susceptible devices ([Ladbury 2017](#)).

Fault injection

- **Use it to iterate.** SEM IP, XilSEM and AMD's TMR Inject provide built-in injection. BYU-style campaigns flip millions of CRAM bits: 2.4 million in one TMR study ([Cannon et al. 2018](#)) and about 11 million in the partial-TMR study ([Keller & Wirthlin 2021](#)).
- **Know what it misses.** Single-bit injection cannot produce MCUs, SEFIs, SETs or SEL. In the 2018 study, every TMR failure under beam came from MCUs, and none from single-bit upsets.
- **Inject into the FDIR too.** In hardware-in-the-loop tests, force SEL trips, scrubber "uncorrectable" reports, watchdog timeouts, corrupted golden images and SAA flags, and check that each lands on the right recovery rung.


Standards and guidance

- [ECSS-Q-ST-60-15C Rev.1](#) (20 March 2025): radiation hardness assurance for EEE components, covering TID, displacement damage and SEE.
- [ECSS-E-HB-20-40A](#) (11 October 2023): the mitigation-techniques handbook for ASICs and FPGAs, superseding ECSS-Q-HB-60-02A.
- [ESCC 25100](#): SEE test method and guidelines; JESD57: heavy-ion SEE test procedures.
- NASA's [Avionics RHA Guidelines](#) and [Proton SEE Guideline](#).

12. Recommended baselines by mission class

Match the mitigation stack to what an outage costs. A demo can reboot through its upsets, an operational constellation must budget availability, and a high-reliability mission should keep SRAM FPGAs out of any function that cannot stop. Apply the questions below per function, not per spacecraft.

Three questions set the mitigation stack



decision flow · 3 questions · 4 outcomes, applied per function

Every path ends with memory ECC, and datapaths get duplicate-and-compare or end-to-end CRCs wherever a dropped output is acceptable.

Layer	Demo or CubeSat (≤ 2 years, ≤ 600 km)	Operational constellation (3–7 years)	High reliability, or above 1,000 km
Part choice	COTS SRAM SoC with published heavy-ion data	COTS for payload processing, lot-tested; rad-tolerant parts for control	Rad-hard or rad-tolerant only, SEL-free to LET 80
Latchup	Per-rail detect and sequenced power cycle	Same, with thresholds from your own heavy-ion test at maximum temperature	Designed out by part choice; keep current limits as a backstop

Layer	Demo or CubeSat (≤ 2 years, ≤ 600 km)	Operational constellation (3–7 years)	High reliability, or above 1,000 km
Configuration	SEM or XiSEM running continuously; reload on uncorrectable	Internal scrubber watched by an external supervisor; golden images in rad-hard NVM	External readback scrubber, or non-volatile configuration
User logic	TMR on control and state machines; safe FSMs; CRCs on data	DTMR with triplicated I/O on control; partial TMR ranked by fault injection	Full DTMR with floorplanning; DWC or TMR on datapaths
Memories	BRAM ECC with scrubbing; DDR ECC	Reed-Solomon DDR with mode-register refresh	Rad-hard memories with chip-kill ECC
Processor	Lockstep R5F for housekeeping where available	Lockstep core plus nested watchdogs	Rad-hard processor for control; SoC cores as payload
System	Spacecraft watchdog and power cycle	Dedicated rad-hard supervisor; SAA-aware operations; full telemetry	Adds cold or hot spares
Verification	Fault injection plus literature beam data	Proton test of the mitigated design; heavy-ion SEL test of the flight lot	Full RHA to ECSS-Q-ST-60-15C or NASA guidelines

The altitude split follows §2: at 1,200 km the proton flux is about $26\times$ that of a 525 km SSO, so a design that is comfortable in one is marginal in the other.

13. Common pitfalls

Most in-flight surprises come from a mitigation that exists on paper but is defeated by one shared resource, one untested condition or one averaged number.

Pitfall	Why it bites	Fix
TMR without scrubbing	Configuration upsets accumulate until two domains fail	Continuous scrubbing (§6)
Triplicating only flip-flops on SRAM parts	Raised an Artix-7 cross-section 30–50%	DTMR (§7)
Shared I/O or clock buffers across domains	Held TMR gain to 3× instead of 86×	Triplicated I/O, per-domain buffers (§7)
Synthesis merging the redundancy	Silent loss of a domain	TMR-aware flow, netlist check, fault injection (§7)
LUTRAM or SRLs in critical logic	Masked from scrubbing	TMR, or ECC block RAM (§8)
Block RAM ECC with no write-back	Single errors pile up into doubles	Background memory scrubber (§8)
A fabric scrubber recovering itself	A scrubber SEFI goes unnoticed	External heartbeat check and reload (§9)
SEL detection on the aggregate supply	0.1–0.3 A steps hide in a 30 A budget	Per-rail step detection (§5)
Power-cycling one rail, or back-powered I/Os	The latch survives the off-time	Sequence every rail off; isolate neighbors (§5)
SEL tests at room temperature, other I/O voltages or with protons only	Thresholds look better than flight	Heavy ions at maximum temperature and flight voltages (§5, §11)
Sizing scrub rates on orbit-average flux	Accumulation risk understated 10–44×	Burst-aware model; scrub hardest in the SAA (§10)
Designing to solar maximum or an old SAA map	LEO rates climb toward the next minimum, and the SAA moves	Solar-minimum environment; updated SAA boundary (§2)
Fault injection as the final verdict	Misses MCUs, SEFIs and SEL	Beam test the mitigated design (§11)

Pitfall	Why it bites	Fix
DDR chosen without SEFI data	Vendor spread above 100×	SEFI-based selection, Reed-Solomon ECC (§8)
Flying a different lot than you tested	Hardness can be lot-specific	Buy and test the flight lot (§5)
Watchdog kicked from a timer interrupt	A hung application keeps kicking	Kick after health checks (§9)

14. Sources

All pages below were opened for this analysis. The proton-flux, burst-penalty and TMR-model figures were computed here with AP8 through [IRBEM](#).

Environment

- [Space, Atmospheric, and Terrestrial Radiation Environments](#), Barth, Dyer & Stassinopoulos, 2003
- [Mission Radiation Environment Modeling and Analysis](#), NASA/TM-20220011775
- [Solar Cycle 25 reached its maximum in October 2024](#), SIDC
- [Solar Cycle Progression](#), NOAA SWPC
- [Gleissberg Cycle Dependence of Inner Zone Proton Flux](#), Bregou et al., 2022
- [SAA growth from Swarm data](#), ScienceDaily, February 2026
- [NASA researchers track slowly splitting dent in Earth's magnetic field](#), NASA
- [Toward the Use of COTS in Space: the True LEO Environment](#), Electronics, 2023
- [On-orbit results of the Cibola Flight Experiment](#), Caffrey et al., 2009
- [Small Spacecraft Avionics, State of the Art 2026](#), NASA

Device data

- [DS882 XQR Kintex UltraScale data sheet, v1.5](#), AMD, April 2026
- [XQR Versal adaptive SoCs, SEFUW 2025](#) and [availability announcement](#), AMD, 2025; [space portfolio](#)
- [DS192 Virtex-5QV overview](#), Xilinx
- [Microchip space FPGAs, ESA 2023](#); [RT PolarFire qualification news](#), July 2025; [RT PolarFire introduction](#), EDN, 2020
- [RTG4 introduction](#); [RTAX-S/SL](#)

- NanoXplore [NG-Medium](#), [NG-Ultra](#), [NG-Ultra presentation](#), [ESCC 9030 qualification](#), January 2026
- [Certus-NX-RT and CertusPro-NX-RT brief](#), Frontgrade
- [Single-Event Characterization of 16 nm FinFET UltraScale+ Devices](#), Lee et al., 2018
- [Neutron Radiation Beam Results for the UltraScale+ MPSoC](#), Anderson et al., 2018; [SEE Assessment of UltraScale+ MPSoC](#), 2023
- [XRTC 7-series single-event results](#), Swift, 2014; [Kintex UltraScale heavy-ion report](#), Berg et al., NASA, 2017
- [Identifying Radiation-Induced Micro-SEFIs in SRAM FPGAs](#), Pérez-Cellis et al., 2021
- [NAND flash and DDR3 SDRAM test results](#), Grürmann et al., 2014; [NEPP compendium](#), 2023
- AMD [UG473](#), [DS190](#), [UG1085](#), [UG1137](#)

Mitigation techniques

- [WP395 Mitigating Single-Event Upsets](#), [PG187 SEM IP](#), [SEM IP](#), [XAPP1296 MultiBoot](#), [PG268 TMR subsystem](#), [TMRTool](#), AMD
- [Scrubbing solutions for SRAM-based space-grade FPGAs](#), Bedi, EDN, 2020; [GRSCRUB configuration supervisor](#)
- [Improving TMR with SEU-aware incremental placement](#), Cannon et al., 2018; [TMR with non-triplicated I/O](#), Cannon et al., 2020; [Partial TMR](#), Keller & Wirthlin, 2021; [Feedback TMR](#), Wirthlin, 2016
- [Three flavors of TMR in Synplify](#); [RT PolarFire TMR and spatial separation](#); [Precision Hi-Rel](#); [SpyDrNet-TMR](#)
- [Single Event Latchup Protection of Integrated Circuits](#), Layton et al., 1997; [SEU and Latchup Sensitive Devices in Satellite Systems](#), Maurer & Kinnison
- [Destructive single-event effects in semiconductor devices and ICs](#), Sexton, 2003; [Redundant systems and destructive SEE](#), Ladbury et al.
- [Using space-grade power components to power AMD FPGAs](#), TI SLVAFM2

Standards and guidance

- [ECSS-Q-ST-60-15C Rev.1](#), March 2025; [ECSS-E-HB-20-40A](#), October 2023; [ESCC radiation standards](#)
- [Proton SEE Guideline](#), LaBel et al., NASA, 2009; [Strategies for SEE Hardness Assurance](#), Ladbury, 2017; [Avionics RHA Guidelines](#), NASA; [JESD57 revision update](#), NASA