☑ Fact-Checked Science

EVIDENCE REVIEW · S0 NEWS RESEARCH NOTE

S0's storm-explosion note: the 23% adds up, but the excess is borderline and unconfirmed

S0 News's research note counts 1,677 utility explosions in S0 News's own outage log since 1999 and finds 77 on the three days after 127 strong solar storms, where about 62.5 were expected: 23% more [1 Claims 1]. Every number reproduces from its published tables and GFZ's Kp, and the excess survives placebo storm calendars, at a one-sided p of 0.02 to 0.05 (our analysis) [2][4]. But it is borderline: it rests on a window that is the best of 24 and shrinks when the window moves a day or the threshold is NOAA's strict G3, the note doesn't say how its rows were found, and no independent register has been tested (our analysis). Studies of grid and insurance records find effects of this kind, measured in other ways and some of them delayed [7 Abstract][8 Abstract], so the idea is plausible; this log alone can't confirm it.

By Fact-Checked Science · Published November 7, 2026

OUR ANALYSIS marks our own reasoning and calculations: our recomputation of the note's numbers from its tables and GFZ's Kp, the placebo calendars and the other controls, the spot check of 35 rows, the splits by equipment and region, and our reading of what the note does and doesn't test. "Post hoc" marks checks we added after seeing the results, some of them after a fresh review, and "exploratory" those we fixed in advance as descriptive only. The verdicts are our assessment. Everything else comes from the numbered sources. For the note [1], locators are its own section names; for papers, Abstract or § with a section number is the part we read. Kp is the planetary geomagnetic index, in thirds of a step, so 7− is 6⅔; observed ÷ expected is the ratio of explosions logged to those expected.

The note's arithmetic holds; its excess after storms is borderline and unconfirmed

OUR ASSESSMENT

77 vs 62.5

logged explosions on the three days after 127 strong storms, against the number expected. It reproduces exactly, and placebo storm calendars give a one-sided p of 0.02 to 0.05 (our analysis) [1 Claims 1][2]

1st of 24

days 1–3 is the best of the 24 three-day windows from 10 days before a storm's peak to 15 after. Fixed in advance, 4.1% of shifted storm calendars match it; picked after looking, half of them would (post hoc) [2]

1 in 43

the odds of a one-day spike as rare as Pipeline B's day 3 among the 28 tests the note's tables show for all its rows, against its "about 1 in 70" (post hoc) [1 Claims 2][2]

THE NOTE'S CLAIMS

UNCONFIRMED

"+23% Explosions in the 3 Days Following Kp7 (G3) Solar Storms"

The note's subtitle, with its title and first claim [1 Subtitle][1 Claims 1]

The 77 explosions where 62.5 were expected reproduce exactly, as does the adjusted 1.28 (1.01–1.62), and the excess passes placebo calendars that keep each storm's year and month (p = 0.019) and shifts of the whole storm calendar (p = 0.041) (our analysis) [2]. But it is one-sided and borderline, and it rests on the note's own choices. Days 1–3 is the best of 24 possible windows: fixed in advance, as the note says it was, 4.1% of shifted calendars match it, but picked after looking, half of them would (post hoc), and the note's "fixed in advance" has no dated record we could check (Unverified). A day earlier or later gives +10% to +14%, and a strict Kp 7 gives +18% (p = 0.13). How the rows were found isn't stated, and no independent register has been tested (our analysis). See section 4.

OVERSTATED

"day 3 was +44% (73 vs 51 expected; chance ≈ 1 in 600, still about 1 in 70 after allowing for the extra day×threshold comparisons)"

The note's second claim, on Pipeline B [1 Claims 2]

The 73 against 50.7 reproduces, and taken alone the day is rare: only 0.13% of shifted storm calendars match it (our analysis) [2]. But "1 in 70" allows only for days 1–3 at three thresholds. Among the 28 tests the note's tables show for all its rows, a spike this rare turns up about 1 in 43 times, and anywhere in its four lag tables about 1 in 24; its rarest number, day 3 among rows with no outside trigger, holds up better, at about 1 in 90 to 110 (post hoc). The window the note set, days 1–3 pooled, gives +12% (p = 0.07), and its "29% jump from before to after" divides by a pre-storm dip that is itself within chance (two-sided p = 0.11). In Pipeline A, day 3 isn't the loudest day, day −4 is, and the best-dated rows peak on day 2 (our analysis) [2]. See section 5.

MISLEADING

"Explosions with no outside cause show it more strongly."

The note's third claim, which adds that "the 125 with a named outside cause show no bump" [1 Claims 3]

The numbers reproduce [2]. But the rows kept are 1,630 of the 1,677 in Pipeline A and 1,552 in Pipeline B, 97% and 93%, so "more strongly" compares nearly the same data: 1.28 against 1.23 in A, 1.18 against 1.12 in B. And the 125 rows with a named cause are too few to show anything either way: on Pipeline B's days 1–3 they give 9 where 13.3 were expected, with a 95% interval of 0.31 to 1.29 that takes in the main result (our analysis). "No bump" there says nothing about outside causes. See section 5.

OVERSTATED

"Bigger storms, bigger effect."

The note's fourth claim [1 Claims 4]

There is some basis: minor and moderate storms give about 1.06 and strong ones 1.23 [1 Claims 4]. But every tier's 95% interval overlaps the others': 0.95–1.18 at Kp 5, 0.90–1.28 at Kp 6, 0.97–1.54 at the note's Kp 7− and 0.83–1.77 at Kp 8−. The severe storms are no higher than the strong ones, and a strict Kp 7 gives 1.18; counted by their peak alone, the G2 storms come out at 0.89 (our analysis) [2]. See section 6.

MISLEADING

"Two independent storm definitions, same pattern."

The note's fifth claim [1 Claims 5]

There are two definitions, and their patterns are alike [1 Claims 5]. But they are one index at two thresholds, so one sits inside the other: all 173 of Pipeline A's storm days are among Pipeline B's 341, and all 76 explosions behind A's headline are among the 138 on B's days 1–3 (our analysis) [2]. Their agreement on days 1–3 is built in: outside A's window days, B's days 1–3 give 62 against 64.1. Their shared day-3 peak comes from different storms, and at Kp 7 B has none (post hoc). It isn't a replication. See section 5.

TRUE

"Each row keeps its source link and a date-certainty grade"

The note's methods, on its inventory [1 Methods, Inventory]

They do, and the rows we checked are real. Of 35 rows drawn at random, ten of them from the days that carry the headline, 31 sources answered us, and all 31 report an explosion of utility equipment at the stated place, one of them as a loud bang and fire. The dates check out too: the four whose local and UTC dates differ had all been put on the UTC day, and one row has dates in conflict that it flags itself (our analysis) [2]. Four sources refused our requests. See section 3.

---

1 · WHAT THE NOTE DID

S0 News's own log of 1,677 explosions, set against 28 years of storms

S0 News's research note takes every "Electrical Explosions" row of S0 News's own outage database, 1,677 of them from 1999 to 3 November 2026, and asks whether they come more often in the three days after strong geomagnetic storms [1 Methods, Inventory]. It credits an "internal GrokBot Team 003" [1 Credit], and S0 News published it as a Google Doc, with its derived tables in a public folder [1][2]. The title of S0 News's episode of 7 November 2026 names its subject [3].

STEP 1

The log

Every row of the "SWN Tech Trouble & Infrastructure Outage Database" in the category Electrical Explosions, "Power-only", each with its source links and a grade for how sure its date is; 677 are "Confirmed" [1 Methods, Inventory].

STEP 2

Two kinds of storm

Each UTC day's highest three-hourly Kp from GFZ [4]. Pipeline A: the 127 runs of days at Kp 7− or more, timed from each run's peak day. Pipeline B: all 341 days at Kp 6 or more [1 Methods, Storms]. NOAA calls Kp 6 a G2 storm and Kp 7 a G3 [5].

STEP 3

What to expect

For A, a centred 365-day running mean of the daily counts; for B, the mean of the same month of the same year, to take out the log's growth and, for B, the seasons [1 Methods, Expected counts].

STEP 4

Chance

Observed against expected on each day from before the storms to after, with random storm calendars and a rate ratio adjusted for year, month and weekday for A, and shifts of the whole storm calendar for B [1 Methods, Chance].

In the note's words A and B are its two pipelines.

> *"The inventory is a floor – regional gap-filling is still open."*
>
> The note [1 Methods, Inventory]

> *"Lag window days 1–3 was fixed in advance; day-by-day values are shown."*
>
> The note [1 Methods, Chance]

> *"An independent utility incident register finds no lag."*
>
> The note, among the results that would prove it wrong [1 What would prove this wrong]

What the note reports

On days 1–3 after the 127 strong storms, 77 explosions where 62.5 were expected, 1.23 times the expected rate (permutation p = 0.045), or 1.28 (95% interval 1.01–1.62) adjusted for year, month and weekday [1 Headline values (A)]. After the 341 storm days, day 3 alone at 73 against 50.7, +44% [1 Headline values (B)]; a stronger excess where no outside cause is named [1 Claims 3]; and bigger excesses after bigger storms [1 Claims 4].

Open about its work

The note publishes its derived tables, with a manifest and a recipe for every headline number [1 How to recreate][2]. It keeps its two pipelines apart, sets aside a control group of explosions with a named outside cause, calls its inventory a floor, and lists five results that would prove it wrong [1 What would prove this wrong]. It hasn't been peer reviewed. Its companion, note 005, which leans on this note's result, is checked on our page on that note.

2 · THE NUMBERS REPRODUCE

Every number reproduces, from the note's tables and from GFZ's Kp

We recomputed the note's numbers from its event table with our own code, and rebuilt both of its storm lists from GFZ's file. They come out as the note gives them (our analysis).

The note's numbers beside our recomputation

NUMBER	THE NOTE	OURS
Rows; with no outside cause named; with a named cause	1,677; 1,552; 125	the same
Pipeline A storms, Kp 7− or more	127	127, identical in every date and peak
Pipeline B storm days, Kp 6 or more	341	341, the same days
A: days 1–3, observed and expected	77 vs 62.5	77 vs 62.50
A: random calendars by year	p = 0.045	p = 0.047
A: adjusted rate ratio	1.28 (1.01–1.62)	1.276 (1.009–1.615)
A: negative-binomial check	1.27 (0.98–1.65)	the same with a dispersion of 0.61, a moment estimate from the raw daily counts; 1.28 (1.01–1.62) with the dispersion estimated in the model, 0.08
B: day 3, observed and expected	73 vs 50.7	73 vs 50.69
B: day 3 against shifted calendars	p = 0.0016	p = 0.0013, over all shifts
B: days 1–3	138 vs 122.9, +12%	the same; p = 0.068
B: ladder at Kp 5, 6 and 7	+5%, +12%, +18%	+5.2%, +12.3%, +17.6%

The note's values are from its headline values [1 Headline values (A)][1 Headline values (B)] and tables [2]; ours are our analysis of its event table [2] and of GFZ's Kp [4]. p values from random calendars differ a little from run to run.

From GFZ's file of 7 November, each UTC day's highest Kp matches the note's on all 10,169 days, so both storm lists rebuild to the day; the note calls values after 30 June 2026 provisional, and GFZ's own flag marks Kp as a nowcast only from 1 October [1 Methods, Storms][4]. The note's statement that "No row cites space weather" also holds: the only two rows that mention a word like it name a town called Aurora and a solar-inverter battery [1 Methods, Cause sets][2].

One number needs a word: the negative-binomial check, a model that lets daily counts vary more than chance alone would. The note's version, 1.27 with an interval of 0.98 to 1.65 [2], comes out with a dispersion of 0.61, which is what a moment estimate from the raw daily counts gives before year, month and weekday are allowed for. Estimated within the model, the dispersion is 0.08, and the model then gives the same answer as the simpler one, 1.28 (1.01–1.62) (our analysis). The raw estimate is inflated by the log's growth, from 13 rows a year to 155; once year, month and weekday are allowed for, the daily counts vary almost exactly as chance predicts.

OUR ANALYSIS

The note's arithmetic is right, and its method is what it says it is. A reproduction can't show what the log measures, which is section 3, or how firm the excess is, which is section 4.

3 · THE LOG

Real explosions, led by street-level equipment, gathered by a method the note doesn't describe

The rows are news reports from 78 countries, and those we checked are real. What the note doesn't say is how they were found, and for a test like this one that matters (our analysis).

Where and when

The log grows fast: 13 rows in 1999, 75 in 2020 and 155 in 2026 to 3 November, a third of them since 2023. A third are from North America (591), a quarter from Latin America and the Caribbean (422) and 18% from Europe (307), then Africa (93), South Asia (79), the Middle East and Turkey (70), East and Southeast Asia (55) and Oceania (33). The leading countries are the US (562), Brazil (155), the UK (140), Mexico (84), Venezuela (82) and South Africa (69). The links go to 1,309 different websites, none of them S0 News's own, led by energyincidentregister.com (106 rows), cbsnews.com (45) and ARIA, France's database of industrial accidents (37). By grade, 677 dates are "Confirmed" and 577 "Estimated" (our analysis) [2].

What exploded

By keyword rules we fixed before counting, 41% of the rows describe distribution equipment, such as pole transformers, manholes, vaults and underground cables; 24% a substation of unstated size; 21% a transformer of unstated kind; and 13% transmission or generating plant, at 100 kV or more or named as such. Read by eye, 16 of 20 rows drawn at random were classed as the rules intend: a pumped-storage hydro plant fell into "other", a substation was classed as transmission on the words "high voltage", and two substations of transmission companies were classed as of unstated size (our analysis) [2].

A spot check of 35 rows

We drew 25 rows at random and 10 more from the 76 on the days that carry Pipeline A's headline, with fixed seeds, and read each row's first source, fetched under a plain client name; where a source refused us we didn't try again (our analysis) [2].

What the sources say

OF 35 ROWS	ROWS
Sources that answered us (4 refused, with HTTP 403)	31
An explosion of utility equipment at the stated place, one of them a loud bang and fire	31 of 31
Date the same in local time and in UTC	21
Local and UTC dates differ, and the row is on the UTC day: three evening events in the Americas and an early-morning blast in Australia	4 of 4
Dates in conflict between sources, flagged by the row itself	1
No time given, or only "afternoon"; the date is the source's local date	5
Equipment: distribution; substation; transmission or generation; transformer of unclear kind	13; 10; 3; 5

Our analysis. The four rows on the UTC day record no time themselves; their sources give it. The rows are listed by their sheet number with our findings in spotcheck_findings.csv, which gives a date only where the check concerned it; no source page is hosted [2].

How the rows were found

The note doesn't say [1 Methods, Inventory]. That matters because a log kept by people who watch for storm effects can find more events after storms, by searching harder then, without any effect being real; only a register kept for other reasons can rule that out, which is why the note's own list of what would prove it wrong includes one [1 What would prove this wrong]. Its event table has no column for when a row was entered. Two traces remain: 207 rows sit in a block appended to the end of the sheet, with dates from 1999 to 2026, and 38 rows carry a "Catch-up" note: 34 read "Catch-up 2026-11-03", the day the data were frozen, and 4 are cut short (our analysis) [2]. Five of those 38 fall on Pipeline A's days 1–3, against 2.35 expected from their own running mean, a ratio of 2.1, but so few rows can't settle anything: its interval runs from 0.7 to 5.0 (our analysis).

SOURCES THAT ANSWERED

OUR ANALYSIS

31 of 31

report an explosion of utility equipment at the stated place, one of them as a loud bang and fire [2]

ON THE UTC DAY

OUR ANALYSIS

4 of 4

the sampled events whose local and UTC dates differ were already on the UTC day [2]

MARKED "CATCH-UP 2026-11-03"

OUR ANALYSIS

34

rows carry that note, the date on which the note says its data were frozen [1 Methods, Inventory][2]

4 · HOW FIRM THE 23% IS

The 23% survives placebo calendars, but it moves with the window, the threshold and the dating

We put Pipeline A through the controls the note didn't run, each with its rules written down and hashed first (our analysis). It passes those that ask whether storms' timing alone could produce it, and gives ground to those that ask how much it rests on the note's own choices.

Logged explosions, observed ÷ expected, on each day around 127 strong storms

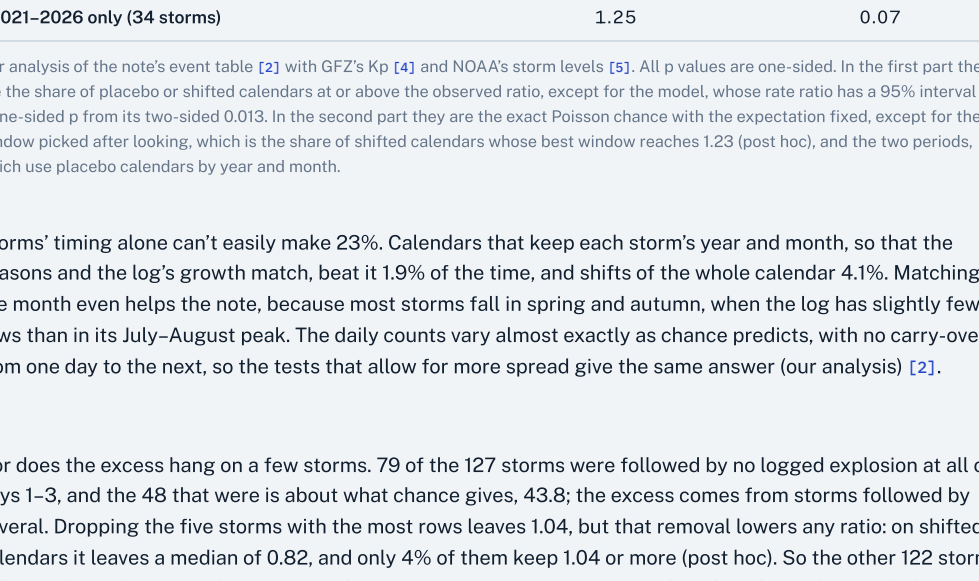

Our analysis, from the note's event table and its 365-day expectation [2], with storms from GFZ's Kp [4]. The band is where 95% of the 10,050 circular shifts of the storm calendar put any single day. Pooled, days 1–3 give 1.23, where 95% of shifted calendars give 0.76 to 1.26.

Pipeline A's days 1–3 under each control

CONTROL	OBSERVED ÷ EXPECTED	ONE-SIDED P
WHAT IT PASSES		
Random calendars by year, the note's test	1.23	0.047
Random calendars by year and month	**1.23**	**0.019**
Shifts of the whole storm calendar	**1.23**	**0.041**
Expected from the same month and year	1.29	0.010
Model with year-and-month strata and weekday	1.37 (1.07–1.76)	0.006
Rows from one place on one day counted once	1.23	0.021
Days with any explosion, not counts	1.19	0.023
WHAT MOVES IT		
Days 1–3 picked after looking, as the best of 24 windows	**1.23**	**0.50**
Window a day earlier, days 0–2	1.10	0.22
Window a day later, days 2–4	1.14	0.16
Days 1–4	1.15	0.09
Strict Kp 7, NOAA's G3 line (96 storms)	**1.18**	**0.13**
Kp 8− and above (50 storms)	1.23	0.15
1999–2020 only (93 storms)	1.21	0.05
2021–2026 only (34 storms)	1.25	0.07

Our analysis of the note's event table [2] with GFZ's Kp [4] and NOAA's storm levels [5]. All p values are one-sided. In the first part they are the share of placebo or shifted calendars at or above the observed ratio, except for the model, whose rate ratio has a 95% interval and a one-sided p from its two-sided 0.013. In the second part they are the exact Poisson chance with the expectation fixed, except for the window picked after looking, which is the share of shifted calendars whose best window reaches 1.23 (post hoc), and the two periods, which use placebo calendars by year and month.

Storms' timing alone can't easily make 23%. Calendars that keep each storm's year and month, so that the seasons and the log's growth match, beat it 1.9% of the time, and shifts of the whole calendar 4.1%. Matching the month even helps the note, because most storms fall in spring and autumn, when the log has slightly fewer rows than in its July–August peak. The daily counts vary almost exactly as chance predicts, with no carry-over from one day to the next, so the tests that allow for more spread give the same answer (our analysis) [2].

Nor does the excess hang on a few storms. 79 of the 127 storms were followed by no logged explosion at all on days 1–3, and the 48 that were is about what chance gives, 43.8; the excess comes from storms followed by several. Dropping the five storms with the most rows leaves 1.04, but that removal lowers any ratio: on shifted calendars it leaves a median of 0.82, and only 4% of them keep 1.04 or more (post hoc). So the other 122 storms still run above chance. In the same way, leaving out any one of 27 storms takes the adjusted interval's lower end below 1, which says only that it is 1.01 to begin with (our analysis).

What moves the result is the note's own choices. Days 1–3 is the best of the 24 three-day windows in the chart above. Fixed in advance, as the note says it was [1 Methods, Chance], it is matched by 4.1% of shifted calendars; picked after looking, the best window of a shifted calendar would reach 1.23 half the time, and a third of the time among the 14 windows from the storm day on (post hoc). The note offers no dated record of the choice, so we rate "fixed in advance" Unverified. A window one day either side gives +10% or +14%. NOAA's G3 line is Kp 7 [5]; the note's 7− takes in a third of a Kp step more, and strict Kp 7 gives +18% (p = 0.13). The note's own dose table mixes the two conventions, with 5o and 6o for G1 and G2 but 7− and 8− for G3 and G4 [2]. Split at 2020/2021, both halves give about 1.2, and neither alone is clear of chance (our analysis).

The best-dated rows show an excess of the same size, with wider intervals. On the 677 rows with confirmed dates, days 1–3 give 32 explosions where 23.7 were expected, 1.35 (the note's permutation p 0.054, and 1.40, 0.97–2.01, adjusted), and without the work-related rows 1.43 (p 0.03) [2]. But their peak is day 2 (13 against 7.9), not day 3 (10 against 7.9), and in all rows the highest day of Pipeline A's profile is day −4, at 1.40, with day 5 at 0.48 and days 9 and 11 at 1.20 and 1.25, as the chart shows (our analysis) [2]. So the note's "day 3 is the loudest" [1 Claims 2] holds for Pipeline B but not for A. In Pipeline B, the confirmed rows give +8.8% on days 1–3 (p 0.24) [2].

The dating

The note says it analysed "UTC calendar days" [1 Methods, Inventory], and it lists re-dating to UTC among the tests that could prove it wrong [1 What would prove this wrong]. 87% of the rows give no time in UTC, so we can't see how each was dated, but the evidence we have says the note did what it says. The four sampled events whose local and UTC dates differ, three in the evening in the Americas and one early in the morning in Australia, were all on their UTC day, and of the seven rows we know of whose local and UTC dates differ, five carry the UTC date and two the local one (our analysis) [2]. At the worst, if every row without a time carried its local date, about 251 would fall on another UTC day; re-dating them at random by each country's offset from UTC gives a median of 1.20 (p = 0.07). Under shares that fit the evidence, a quarter of them or two in seven, the median stays at 1.23 (p = 0.04) (post hoc). The dating, then, is no reason to doubt the result.

OUR ANALYSIS

What the controls add up to. The excess is unlikely to be the storms' timing or the seasons, and it isn't carried by a handful of storms. But it passes them with one-sided p values of 0.006 to 0.047, and under choices as reasonable as the note's own, a window a day either side or NOAA's G3 line, it no longer stands clear of chance. That is the profile of a result worth testing again on new storms and an independent register, not of an established effect.

5 · PIPELINE B

The day-3 spike is the note's strongest number, at odds of about 1 in 24 to 1 in 43 among the tests it shows, and the two pipelines share their data

Pipeline B counts every storm day at Kp 6 or more, 341 of them. Its four lag tables show each day from a week before those days to two weeks after; at Kp 5 and 7 the note gives only days 1–3 and pooled windows [1 Methods, Storms][2].

Its day 3 is the note's strongest number: 73 explosions where 50.7 were expected, +44%, while days 1 and 2 give +2% and +7% [1 Headline values (B)]. Taken alone it is rare: 0.13% of shifted calendars match it (our analysis) [2]. But it is one of many numbers. The window the note set, days 1–3 pooled, gives +12% (p = 0.07), and the same tables show day 5 at −23%, day 7 at −22%, and days 9 and 11 at +18% and +19% [2].

The day-3 spike against the tests the note's tables show

FAMILY OF TESTS	TESTS	CHANCE OF A SPIKE AS RARE AS DAY 3'S
Days 1–3 at Kp 5, 6 and 7, all rows: the note's "1 in 70"	9	0.0087, about 1 in 115
Days −7 to +14 at Kp 6, all rows	22	0.018, about 1 in 57
Every all-rows test in the note's tables	**28**	**0.023, about 1 in 43**
The note's four lag tables, days −7 to +14 at Kp 6	**88**	**0.042, about 1 in 24**
Those, with days 1–3 at Kp 5 and 7 for the same four sets of rows	112	0.054, about 1 in 19

Our analysis, post hoc, over all 10,050 circular shifts of the storm calendar: the share in which the rarest test of the family is at least as rare as day 3 in all rows [2]. The four lag tables are all rows, those with no outside trigger named, those with one named, and those with no stated cause and no outside trigger, the note's 1,149. The rarest number in them, day 3 with no outside trigger named, is matched by 0.88% of shifted calendars across the 88 tests and 1.1% across the 112.

The note's "about 1 in 70" is its one-test 1 in 600 multiplied by nine, for three days and three thresholds [1 Claims 2]; allowing for those nine properly gives about 1 in 115, so for that family the note is fair. But its tables show 22 days at Kp 6, not three, and three more sets of rows. Allowing for the tests they show, a spike as rare as day 3's turns up about 1 in 43 times among all rows and about 1 in 24 anywhere in the four lag tables. The note's own rarest number, day 3 among rows with no outside trigger named, 69 against 44.9 (+54%), holds up better, at about 1 in 90 to 110 (post hoc).

The dip and the jump

The note adds that the three days before the storms ran 13% low, and that the jump from before to after, 1.29, happens in only about 2% of random calendars [1 Claims 2]. Both reproduce [2]. But a dip before storms has no proposed cause, and this one is within chance: 5.7% of shifted calendars dip at least as far, 11% counting both directions. Dividing by it doubles up the chance: days 1–3 alone give p = 0.07, the ratio p = 0.02. Had the days before been at their expected level, the jump would be 1.12 (our analysis).

Two definitions, one data set

Pipeline A's storms are Pipeline B's strongest days, so A sits inside B: all 173 of A's storm days are among B's 341, all 378 of A's window days among B's 809, and all 76 explosions behind A's headline among the 138 on B's days 1–3. Their day-by-day patterns correlate at 0.78 (our analysis) [2]. Their agreement on days 1–3 is built in: outside A's window days, B's days 1–3 give 62 explosions where 64.1 were expected. Their shared day-3 peak comes from different storms. About two thirds of the excess on B's day 3 comes from storm days that reach Kp 6 but not 7−, the moderate storms the note itself says leave days 1–3 "essentially flat" [1 Claims 4], and at Kp 7 Pipeline B has no day-3 peak at all, with days 1, 2 and 3 at 1.29, 1.24 and 1.13 [2] (post hoc). So the two agree where they share data and differ where they don't: not a second confirmation.

The control set

The note sets 125 explosions with a named outside cause, such as lightning, wind, vehicles or overload, against the rest [1 Methods, Cause sets][1 Claims 3]. On B's days 1–3 they give 9 where 13.3 were expected. That is "no bump", but 9 rows can't show one: the 95% interval runs from 0.31 to 1.29, which takes in the main result. And because the rest are 93% of all rows, their stronger figure, 1.18 against 1.12 in B, is nearly the same number (our analysis) [2].

6 · DOSE

Bigger storms give a bigger point estimate, inside intervals that all overlap

The note's fourth claim reads its tables as a dose response: minor and moderate storms leave days 1–3 about flat, strong and severe ones raise them by 23% [1 Claims 4].

Days 1–3 after storms, observed ÷ expected, by the storm threshold

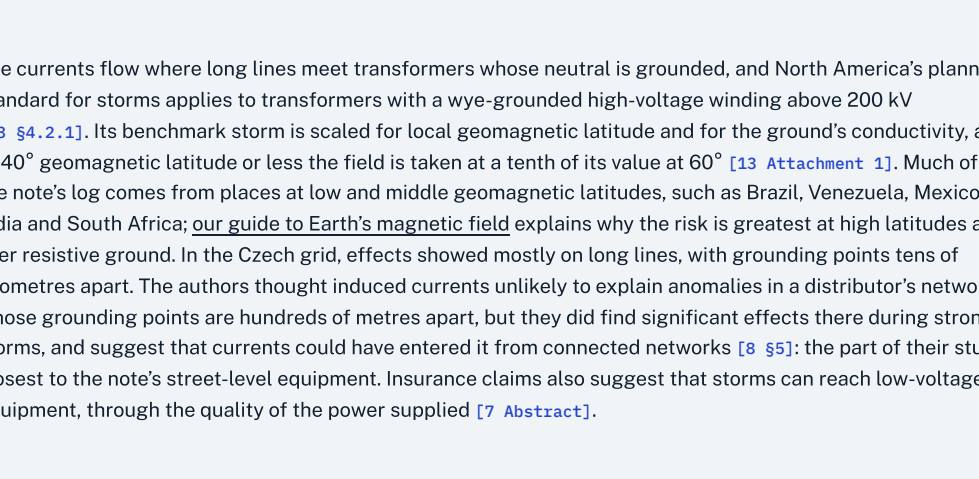

Our analysis, from the note's event table [2] and GFZ's Kp [4]: each row counts the runs of days at or above its threshold, timed from their peak day, as the note's own table does. The lines are exact Poisson 95% intervals with the expectation fixed. G-levels are NOAA's [5].

The step from about 1.06 to 1.23 is there, and it is the core of the note's case for Pipeline A. But every interval overlaps every other, the severe storms at Kp 8− (1.23) are no higher than the strong ones, and NOAA's strict G3 line, Kp 7, gives 1.18. Counted once each, by their peak alone, the G1 storms give 1.05 (0.91–1.21), the G2 storms 0.89 (0.66–1.19), the G3 storms 1.25 (0.92–1.67) and those of G4 and above 1.20 (0.80–1.73); in one model, the G4+ storms run at 0.96 times the G3 rate (0.60–1.53) (our analysis) [2]. Pipeline B's ladder, +5%, +12% and +18%, has p values of 0.13, 0.07 and 0.10 (our analysis). A dose response would strengthen the note's case; these numbers neither show one beyond chance nor rule one out.

7 · KNOWN PHYSICS

Storms do damage grid equipment, mostly high-voltage plant, and studies of grid and insurance records find effects, some of them delayed

Geomagnetic storms drive currents through long power lines and the grounded transformers at their ends. That is established; how often it shows up in failures of ordinary equipment is what statistical studies try to measure.

In March 1989 Hydro-Québec's grid collapsed in about 90 seconds, at 2:44 a.m. local time, and about six million people lost power for about nine hours [14][15]; in New Jersey the step-up transformer of the Salem No. 1 nuclear plant was found with "Charred winding series connections" [16 pp. 38–39]. In October 2003 the Halloween storms caused an outage of under an hour in Malmö, Sweden, for about 50,000 customers [17 p. 1782], and damage that led several large transformers in South Africa to fail in the following months [18][19]. At G2, NOAA's storm scale warns that "High-latitude power systems may experience voltage alarms, long-duration storms may cause transformer damage" [5]. Our guide to Earth's magnetic field and our page on the Sun's worst cases cover these storms.

The currents flow where long lines meet transformers whose neutral is grounded, and North America's planning standard for storms applies to transformers with a wye-grounded high-voltage winding above 200 kV [13 §4.2.1]. Its benchmark storm is scaled for local geomagnetic latitude and for the ground's conductivity, and at 40° geomagnetic latitude or less the field is taken at a tenth of its value at 60° [13 Attachment 1]. Much of the note's log comes from places at low and middle geomagnetic latitudes, such as Brazil, Venezuela, Mexico, India and South Africa; our guide to Earth's magnetic field explains why the risk is greatest at high latitudes and over resistive ground. In the Czech grid, effects showed mostly on long lines, with grounding points tens of kilometres apart. The authors thought induced currents unlikely to explain anomalies in a distributor's network whose grounding points are hundreds of metres apart, but they did find significant effects there during strong storms, and suggest that currents could have entered it from connected networks [8 §5]: the part of their study closest to the note's street-level equipment. Insurance claims also suggest that storms can reach low-voltage equipment, through the quality of the power supplied [7 Abstract].

What statistical studies of grid records found

STUDY	RECORDS	SIZE OF THE EFFECT	TIMING
Schrijver and Mitchell (2013)	1,216 US grid disturbances reported to the US Department of Energy, 1992–2010	Reports a day on the 2% most active days: 0.23, against 0.14 on random nearby days and 0.06 on nearby quiet days; about 4% of all disturbances attributable	The active days themselves [6 Abstract][6 §3]
Schrijver and colleagues (2014)	11,242 North American insurance claims for electrical and electronic equipment, 2000–2010	Claims rates: about 20% higher on the 5% most active days, 10% on the top third; about 4% of claims attributable	The active days; no clear difference between high and low geomagnetic latitudes [7 Abstract][7 §3]
Zois (2013)	Failures of Greece's 150 and 400 kV transformers, 1989–2010	19 failures within 3 days of 43 days with Ap of 100 or more; no expected count given	Within 3 days either side; yearly failures follow the solar cycle, some years behind it [12 Abstract]
Výbošťoková and Švanda (2019)	Czech maintenance logs, 12 years	More anomalies on high-voltage lines and substations in periods of tens of days around geomagnetic maxima	Hints of more just after the maxima than before [9 Abstract]
Švanda and colleagues (2020)	**The same Czech logs**	**Anomaly counts: about 5–10% more over the 5 days after a storm's onset, 5–20% after days with Dst below −50 nT**	**Power lines within a day; transformers 2–3 days later [8 Abstract][8 §6]**
Gil and colleagues (2021)	Polish transmission-line failures	More failures with intense storms	During and right after them [11 Abstract]
Výbošťoková and Švanda (2022)	Slovak transmission-grid logs	Failures follow temperature and rain most; with those removed, more around geomagnetic maxima	Periods of high activity, not single days [10 Abstract]

Read in the versions the source list names; abstracts are paraphrased. Ap is the daily planetary index of geomagnetic activity.

These studies give the note partial support. They measure different things, from disturbance reports and insurance claims to maintenance logs, and their effects range from 5–10% more anomalies over five days in the Czech logs to a same-day rise of about 70% over nearby days on the most active days in the US reports [6 §3][8 Abstract]. Some are delayed: in the Czech logs, transformer anomalies rose two to three days after a storm's onset [8 Abstract]. North American insurance claims show storms reaching low-voltage equipment too, though on the active days themselves [7 Abstract]. So the note's direction and size are plausible. Its own log just can't confirm them: it is a log of news reports, not a register kept for other reasons, and its excess is borderline (our analysis).

What the rows describe (exploratory)

Split by equipment, the point estimates run in the order the physics would predict: 1.67 for transmission and generating plant (13 explosions where 7.8 were expected), 1.49 for substations (21 against 14.1), 1.27 for transformers of unstated kind (21 against 16.6), 0.98 for distribution equipment (22 against 22.5), and none in the 35 other rows (0 against 1.5). But the classes don't differ beyond chance (p = 0.28) (exploratory; the test across classes is post hoc) [2]. Taken together, transmission plant and substations give 1.56, a combination we chose after seeing the split (post hoc). By region, the ratios are 2.37 in Africa (10 against 4.2), 1.55 in East and Southeast Asia (4 against 2.6), 1.30 in Latin America and the Caribbean (20 against 15.4), 1.15 in Europe (13 against 11.4), 1.13 in South Asia (4 against 3.5), 1.05 in the Middle East and Turkey (5 against 4.8), 1.04 in North America (20 against 19.3) and 0.82 in Oceania (1 against 1.2). These don't differ beyond chance either (p = 0.58) (exploratory). South Africa, which supplies most of Africa's rows, has a documented history of storm damage to its grid [18], but on these numbers that is a possibility at most.

OUR ANALYSIS

Through the catastrophism lens. Do storms quietly break the grid's equipment for days afterwards? In S0 News's log, which the note calls a floor [1 Methods, Inventory], the excess after strong storms comes to 14.5 extra explosions over 127 storms, about 0.1 a storm, and most storms were followed by none. That is a modest signal in a log built from news reports, not a grid failing for days on end. Utility records support a small and sometimes delayed effect on high-voltage transformers, one that the Czech authors judge "probably related to an exposure to GICs", geomagnetically induced currents [8 Abstract], and the rare extreme storm can do serious damage, as in 1989 and 2003 [16 pp. 38–39][18]. Neither the note nor this check finds evidence of more than that.

8 · LIMITS

What this check can and can't say

The verdicts rest on the note's own log and on tests whose rules we fixed before running them (our analysis).

One log, of unknown making

We can't say how or when the rows were found; the appended block and the "Catch-up" notes are the only traces [2]. A log kept for other reasons would be the real test.

No independent register

The US Department of Energy's OE-417 reports of grid disturbances are the obvious one, and the note names such a register as its own test [1 What would prove this wrong], but their host didn't answer us [20]. That test is still to be done.

A sample of rows

We read 31 of 35 sampled rows' sources; 4 refused our requests. A larger check could find more slips in the dates.

One record of storms

Every placebo reshuffles the same 28 years of storms and explosions; new storms are the stronger test.

Dating checked on a few rows

The sampled rows whose local and UTC dates differ were all on their UTC day, so the worst-case re-dating overstates the problem; but 87% of the rows give no time in UTC, and we could check only a few.

Splits on small numbers

The splits by equipment and region rest on 22 or fewer explosions per group, and are exploratory.

Authors' versions and abstracts

Three papers were read in full in their authors' versions on arXiv, [6], [7] and [8], two of them because the journal's site showed a bot check; four were read only as abstracts, [9], [10], [11] and [12]. The accounts of 1989 and 2003 are as checked for our earlier pages.

What would change these verdicts

01

An independent register, such as the OE-417 reports [20] or utility maintenance logs like the Czech ones [8 Abstract], showing more failures on days 1–3 after strong storms.

02

New strong storms after 3 November 2026, logged by a stated method that doesn't depend on storm dates, showing the same excess; the note sets this test itself [1 What would prove this wrong].

03

The rows re-dated to UTC from their sources' own times, with the excess still on days 1–3.

Our page on note 005 checks the four explosions it counts after this note's data were frozen.

9 · REPEAT OUR NUMBERS

The rules, the scripts and the derived numbers; the note's tables stay with S0 News

Every number of ours on this page comes from these files. The scripts fetch the note's tables and GFZ's Kp by URL under a plain client name [2][4]; we host only numbers we derived. No hosted file holds a name, a URL or any of S0 News's text, and dates appear only where the spot check concerned them.

TXT **RULES_reproduce.txt · RULES_inventory.txt · RULES_inventory_amendment1.txt · RULES_spotcheck.txt · RULES_controls.txt · RULES_controls_amendment1.txt · RULES_posthoc_top5.txt · RULES_posthoc_review.txt**

The rules of every test, each hashed before the computation it governs; the two post-hoc ones, on the five busiest storms and on the points of the fresh review, after the controls had run · 29 KB in all

TXT **analysis_log.md · HASHES.txt**

Our log, with the time and SHA-256 of every rules file, script, input and output as they came; and the checksums of every file here · 14 KB and 4 KB

PY **fetch_inputs.py**

Fetches the note's 18 Drive files, its text and GFZ's Kp file by URL under a plain client name, and compares each with the copy we used · 6 KB

PY **lib_ee.py**

Loaders, exact Poisson tests, and Poisson and negative-binomial models for designs of indicators, in the standard library · 10 KB

PY **01_reproduce.py · 02_inventory.py · 02b_dating_i6b.py**

The note's numbers and both storm lists, recomputed; the log by year, region, website, grade, equipment and dating; and the UTC day offsets implied by the rows that give a time · 13 KB, 15 KB and 5 KB

PY **03_spotcheck_draw.py · 03_spotcheck_fetch.sh**

Draws the 35 rows with fixed seeds; fetches their first sources under a plain client name, never retrying a refusal · 1 KB and 2 KB

PY **05_controls.py · 09_posthoc_top5.py · 13_posthoc_review.py**

The controls C1–C14 and the check of the "Catch-up" rows; the post-hoc check of the five busiest storms against shifted calendars; and the post-hoc checks after the fresh review: the families of tests the note's tables show, the dating under shares the evidence allows, the window picked after looking, and the tests across classes and regions · 33 KB, 3 KB and 6 KB

PY **08_chart_data.py · 10_charts.py**

The numbers behind the two charts, and the charts themselves, drawn into this page · 3 KB and 7 KB

JSON **reproduce.json · inventory.json · posthoc_top5.json · posthoc_review.json · chart_series.json**

Every number of the recomputation and of the inventory; the two sets of post-hoc checks; the charts' numbers · 21 KB, 5 KB, 1 KB, 2 KB and 5 KB

JSON **controls_C1.json · controls_C2.json · controls_C3.json · controls_C4.json · controls_C5.json · controls_C6.json · controls_C7.json · controls_C8.json · controls_C9.json · controls_C10.json · controls_C11.json · controls_C12.json · controls_C13.json · controls_C14.json · controls_L1L2.json**

The result of each control, named as in RULES_controls.txt · 14 KB in all

CSV **row_classes.csv · i6b.csv · spotcheck_findings.csv**

Our classes for every row by its sheet number, with none of S0 News's text, dates or grades; the day offsets implied by the rows that give a time; and the spot check's findings in our words, with a date only where the check concerned it · 66 KB, 27 KB and 6 KB

Run it

In an empty folder `<base>`, `python3 -I fetch_inputs.py <base>` fetches the note's 18 Drive files, its text and GFZ's Kp file, about 8 MB, and compares each with the copy we used. Then run `01_reproduce.py`, `02_inventory.py`, `02b_dating_i6b.py`, `03_spotcheck_draw.py`, `05_controls.py ALL`, `08_chart_data.py`, `09_posthoc_top5.py` and `13_posthoc_review.py`, in that order and each with `--base <base>`, to get the outputs listed here in `<base>/results/`; `bash 03_spotcheck_fetch.sh <base>` fetches the sampled rows' sources. Standard library only, with curl.

Related checks

PARTLY TRUE A "small" flare? The Disaster Cycle, Fact-Checked

Sources

Bracketed numbers refer to this list. For S0 News's research note [1], locators are its own section names: Title and Subtitle; Credit, the line under them; Claims 1 to 5; "What would prove this wrong"; Methods with the paragraph's opening word (Inventory, Cause sets, Storms, Expected counts, Chance); Headline values (A) and (B); and How to recreate. For papers, Abstract or § with a section number is the part we read, numbered as in the version the entry names; for NERC's standard [13], § and Attachment are its own; where only an abstract was read, the entry says so, and abstracts are paraphrased; p. and pp. are pages. Kp is GFZ's planetary index in thirds (7− is 6⅔); G1 to G5 are NOAA's storm levels; UTC is universal time. Links checked November 7, 2026.

1 S0 News (Space Weather News) (undated). Electrical explosions are more likely 1–3 days after a solar storm: +23% Explosions in the 3 Days Following Kp7 (G3) Solar Storms. "Space Weather News Research Note from internal GrokBot Team 003", a public Google Doc headed "SPACE WEATHER NEWS & RESEARCH PUBLICATIONS 2026". The note is undated; it says its inventory was frozen on November 3, 2026. Read in full, as text and HTML exports retrieved on November 7, 2026; SHA-256 of our text export: 9558eb05c827429249a3de95b0a8d8975e4be4eb313a452d063b878db9ddb394. Not peer reviewed. No licence is stated, so we quote and link it, redraw none of its figures and host nothing of it.

2 S0 News (Space Weather News) (undated). Derived data for research note 003. A public Google Drive folder of 18 files: 16 CSV tables, two of them identical copies of the event table EE_events.csv, and two identical copies of a MANIFEST.md, listed as modified on November 3 and 6, 2026. Downloaded on November 7, 2026, and again the same day with identical results; the tables agree with the manifest's row counts. No licence is stated: our scripts fetch the tables by URL, and we host only numbers we derived ourselves. No hosted file holds a name, a URL or any of the note's text, and dates appear only where our spot check concerned them.

3 SpaceWeatherNews (S0) (November 7, 2026). Solar Storm Electrical Explosions, Betelgeuse's Baby | S0 News Nov.7.2026. YouTube video, https://www.youtube.com/watch?v=WXp-m52Qftc, published at 12:50 UT. No transcript was available, so nothing from the episode is rated here; its title names the note's subject.

4 Matzka, J., Stolle, C., Yamazaki, Y., Bronkalla, O., and Morschhauser, A. (2021). The geomagnetic Kp index and derived indices of geomagnetic activity. *Space Weather* 19, e2020SW002641. Data: Matzka, J., Bronkalla, O., Tornow, K., Elger, K., and Stolle, C. (2021). Geomagnetic Kp index, V. 1.0, GFZ Data Services, CC BY 4.0: the Kp and ap file as of November 7, 2026 (SHA-256 20f204e2daef813016a3fae93a3d84363e1e66d4eb9f1a01ffb03755925e5068). Its own flag marks Kp as definitive to September 30, 2026 and as a nowcast from October 1. Daily maxima taken and storms counted by us.

5 NOAA Space Weather Prediction Center (undated). NOAA Space Weather Scales, read on November 7, 2026: geomagnetic storms from G1 (Minor), at Kp = 5, through G2 (Moderate) at Kp = 6, G3 (Strong) at Kp = 7 and G4 (Severe) at Kp = 8, including a 9−, to G5 (Extreme), at Kp = 9, with the effects on power systems expected at each level. US government work.

6 Schrijver, C. J., and Mitchell, S. D. (2013). Disturbances in the US electric grid associated with geomagnetic activity. *Journal of Space Weather and Space Climate* 3, A19. Read in full in the authors' version, arXiv:1304.5489v1, from its TeX source; the journal's site answered our request with a bot check.

7 Schrijver, C. J., Dobbins, R., Murtagh, W., and Petrinec, S. M. (2014). Assessing the impact of space weather on the electric power grid based on insurance claims for industrial electrical equipment. *Space Weather* 12(7), 487–498. CC BY-NC-ND 3.0, as deposited with Crossref. Read in full in the authors' version, arXiv:1406.7024v1, from its TeX source.

8 Švanda, M., Mourenas, D., Žertová, K., and Výbošťoková, T. (2020). Immediate and delayed responses of power lines and transformers in the Czech electric power grid to geomagnetic storms. *Journal of Space Weather and Space Climate* 10, 26. CC BY 4.0. Read in full in the authors' version, arXiv:2005.14448, from its TeX source; the journal's site answered our request with a bot check.

9 Výbošťoková, T., and Švanda, M. (2019). Statistical analysis of the correlation between anomalies in the Czech electric power grid and geomagnetic activity. *Space Weather* 17(8), 1208–1218. Abstract only, in the authors' version, arXiv:1907.01753; paraphrased.

10 Výbošťoková, T., and Švanda, M. (2022). Correlation of anomaly rates in the Slovak electric transmission grid with geomagnetic activity. *Advances in Space Research* 70(11), 3769–3780. Abstract only, in the authors' version, arXiv:2208.03813; paraphrased.

11 Gil, A., Berendt-Marchel, M., Modzelewska, R., Moskwa, S., Siluszyk, A., Siluszyk, M., Tomasik, L., Wawrzaszek, A., and Wawrzynczak, A. (2021). Evaluating the relationship between strong geomagnetic storms and electric grid failures in Poland using the geoelectric field as a GIC proxy. *Journal of Space Weather and Space Climate* 11, 30. CC BY 4.0. Abstract only, as deposited with Crossref; paraphrased. The journal's site answered our request with a bot check.

12 Zois, I. P. (2013). Solar activity and transformer failures in the Greek national electric grid. *Journal of Space Weather and Space Climate* 3, A32. Abstract only, as listed with the author's version at arXiv:1307.1149; paraphrased.

13 North American Electric Reliability Corporation (2020, February 6). TPL-007-4: Transmission System Planned Performance for Geomagnetic Disturbance Events. Reliability standard, adopted by NERC's Board of Trustees on that date and approved by FERC on March 19, 2020, as its version history says. Its applicability, section 4, and Attachment 1, on the benchmark storm and its scaling factors, read on November 7, 2026.

14 NOAA National Environmental Satellite, Data, and Information Service (2015, October 19). When solar storms attack: Space weather and our infrastructure. As cited on our guide to Earth's magnetic field.

15 Hydro-Québec. Understanding electricity: March 1989 blackout. As read for our page on the May 2024 storm, for the time of the collapse, 2:44 a.m. local time.

16 North American Electric Reliability Corporation (2012, February). 2012 Special Reliability Assessment Interim Report: Effects of Geomagnetic Disturbances on the Bulk Power System. As cited on our guide to Earth's magnetic field.

17 Wik, M., Pirjola, R., Lundstedt, H., Viljanen, A., Wintoft, P., and Pulkkinen, A. (2009). Space weather events in July 1982 and October 2003 and the effects of geomagnetically induced currents on Swedish technical systems. *Annales Geophysicae* 27(4), 1775–1787. CC BY 3.0. As read for our page on the May 2024 storm.

18 Gaunt, C. T., and Coetzee, G. (2007). Transformer failures in regions incorrectly considered to have low GIC-risk. *2007 IEEE Lausanne Power Tech*, 807–812. As cited on our guide to Earth's magnetic field.

19 Gaunt, C. T. (2014). Reducing uncertainty – responses for electricity utilities to severe solar storms. *Journal of Space Weather and Space Climate* 4, A01. As cited on our guide to Earth's magnetic field.

20 US Department of Energy. Electric Disturbance Events (OE-417) Annual Summaries. Not read: on November 7, 2026, the host's name did not resolve for our client, so this register is untested here.

METHOD

How this was checked

On November 7, 2026, we checked every claim of S0 News's research note 003, its title, subtitle and five numbered claims, and the statements in its methods that carry weight, against the note's own 18 derived files and against GFZ's Kp, from which we rebuilt both of its storm lists. We read the sources of 31 of 35 rows drawn from its log and checked the science against seven statistical studies; the accounts of the 1989 and 2003 storms are reused from our earlier pages, as checked there. Every other statement on this page was checked against its source, 151 references to 20 sources in all.

The verdicts and the items tagged "Our analysis" are our own assessment, built on the cited sources: our recomputation, the placebo calendars and the other controls, the spot check, the splits by equipment and region, and our reading of what the note does and doesn't test. The rules of every test were written down and hashed before it was run. Checks added after the results are marked post hoc: before the fresh review below, those of the five busiest storms against shifted calendars, of Pipeline B's day 3 by storm strength, and of two equipment classes combined; after it, the families of tests the note's tables show, the dating under shares the evidence allows, the odds of a window picked after looking, the tests across equipment classes and regions, and Pipeline B outside Pipeline A's window days. The splits by equipment and region themselves, fixed in advance as descriptive only, are marked exploratory.

Everything in the note reproduces: its row counts and flags, every day of both pipelines, its permutation and shift p values within chance, its adjusted rate ratio exactly, and its 127 storms and 341 storm days, rebuilt to the day from GFZ's file. Its negative-binomial check reproduces with a dispersion of 0.61, the moment estimate from the raw daily counts; estimated within the model, the dispersion is 0.08, and the answer is the simpler model's. A second fetch of all 20 inputs the same day gave identical files and identical results. Not verified: how the log's rows were found; four of the 35 sampled rows, whose sources refused our requests; an independent register, since the host of the DOE's OE-417 summaries didn't answer us; the journal versions of three papers read in their authors' versions; the full texts of four papers read only as abstracts; and whether the note's window was fixed in advance.

After the page was built, a fresh reviewer who hadn't seen the drafting recomputed everything in separate code: every headline number and control reproduced, and all 22 of our hosted outputs reproduced byte for byte from the hosted scripts. The changes that followed: the odds of the day-3 spike now come from the tests the note's tables actually show, about 1 in 24 to 1 in 43, not 1 in 16 from tests it never tabulated; one sampled row we had called mis-dated is on its UTC day, so all four date-crossing rows in the sample are, and the dating is no longer counted against the result; the note's dispersion is described as the moment estimate it is; the splits by equipment and region are reported in full and tested, and don't differ beyond chance; the best-dated rows, the peak day in Pipeline A and the odds of a window picked after looking are added; the studies' measures are stated; and about twenty smaller fixes.

This page was researched and written with Claude, an AI assistant made by Anthropic, and then re-checked against its sources in a separate verification pass.

The verdicts on this page

MISLEADING	Contains something true but leaves a false impression, usually by leaving out key context.
OVERSTATED	Has some basis but claims much more than the evidence supports.
UNCONFIRMED	Rests on a preliminary result that hasn't been confirmed.
TRUE	The evidence supports the claim.

All ten verdicts, and how we check

CORRECTIONS

Corrections and changes

If we get something wrong, we fix it and record the change below with the date. If a correction changes a verdict, the log says what the old verdict was.

Change log

Nov 7, 2026 Published on factchecked.science after a full check against the sources: S0 News's research note 003 and its 18 derived files, from which every number reproduces; GFZ's Kp, from which both storm lists rebuild exactly; a spot check of 35 rows against their sources; placebo calendars and other controls, each with its rules fixed in advance; and the published studies of storms and grid equipment. The rules, scripts and derived numbers are hosted; none of the note's figures is redrawn, and none of its tables is hosted. A fresh review before publication led us to base the odds of the day-3 spike on the tests the note's tables show, to correct the dating of one sampled row and drop the dating as a point against the result, to describe the note's dispersion as the moment estimate it is, and to report the exploratory splits in full.

Fact-Checked Science is independent. It isn't affiliated with the channels, presenters or authors whose claims it checks.

Last updated Nov 7, 2026