

Summary of Questions I have Received About the EOL of the SBN

by

Stonie Cooper

I. Why not use Ku-band, as it is an existing commercial geosynchronous platform paradigm?

Ku-band *can* move bits fast, but it makes a **poor substitute for C-band NOAAPort** because the physics of Ku-band propagation, the regulatory environment, and the operational realities of meteorological distribution all work against it. The shortcomings fall into five hard categories.

1. Rain fade is severe and unavoidable

Ku-band suffers **orders-of-magnitude higher attenuation** during precipitation than C-band. A 2025 IEEE study shows Ku-band latency and signal degradation increase significantly under tropical rainfall, with link performance deteriorating as attenuation rises.

For meteorological data distribution, this is structurally self-defeating:

- The worse the weather, the worse the Ku-band link performs.
- The exact moments when WFOs most need deterministic, uninterrupted data are the moments Ku-band is least reliable.
- Heavy rain, wet snow, hail, and even dense cloud fields cause measurable attenuation.

C-band was chosen for NOAAPort precisely because it is **resilient to precipitation**.

2. Ku-band link availability is not high enough for COOP requirements

NOAAPort is part of NWS continuity-of-operations (COOP). Ku-band cannot meet the required availability because:

- Ku-band outages during moderate to heavy rain are common.
- Even short fades break multicast continuity and corrupt GRIB/BUFR streams.
- Ku-band antennas require tighter pointing tolerances, increasing vulnerability to wind and ice.

The OpenSat4Weather dataset (Ku-band link monitoring) shows frequent disturbances and anomalous patterns in raw Ku-band data, including precipitation-induced signal

drops and noise bursts. These are tolerable for opportunistic rainfall sensing, but **unacceptable for operational meteorological distribution.**

3. Ku-band transponders are not engineered for deterministic broadcast

NOAAPort relies on a **dedicated, deterministic, one-to-many broadcasts** with guaranteed bandwidth and priority. Ku-band capacity is typically:

- Shared with consumer broadband or TV services.
- Subject to dynamic bandwidth allocation.
- Not guaranteed for fixed-rate, latency-sensitive meteorological streams.

Meteorological data (especially GRIB2, BUFR, GOES-R, JPSS, and radar) requires **predictable throughput**. Ku-band operators do not provide the service-level guarantees NOAA requires.

4. Ku-band uplink/downlink power budgets are too tight for large antennas at WFOs

C-band NOAAPort dishes are ~3.8 m and operate with comfortable link margins.

Ku-band would require:

- Smaller dishes (due to higher frequency) but **much tighter pointing accuracy**.
- Higher susceptibility to ice loading and wind misalignment.
- Higher gain LNBs and more expensive RF chains.
- More frequent maintenance.

This increases operational cost and reduces reliability across 122+ WFOs, many in harsh climates.

5. Regulatory and spectrum-sharing constraints

Ku-band is crowded:

- Consumer broadband (VSAT)
- Direct-to-home TV
- Enterprise networks
- Government/military links

This creates:

- Higher interference risk

- More transponder contention
- Less control over long-term spectrum stability

C-band, by contrast, has historically been used for professional broadcast and is more stable for fixed-service meteorological distribution.

6. Ku-band is poorly suited for large-volume meteorological payloads

NOAAPort carries:

- High-resolution model output
- Radar mosaics
- GOES-R ABI products
- BUFR observational streams
- Emergency dissemination products

These are **large, continuous, and time-critical**. Ku-band's susceptibility to fade and jitter makes it unsuitable for:

- Continuous GRIB2 ingestion
- High-rate satellite imagery
- Radar super-resolution streams
- Time-sensitive warning products

Even small interruptions cause corrupted files, missing tiles, or delayed warnings.

7. Operational risk during severe weather

This is the fatal flaw.

When a hurricane, derecho, blizzard, or severe convective outbreak occurs:

- Ku-band attenuation spikes.
- Ku-band link availability drops.
- Ku-band latency increases.
- Ku-band packet loss rises.

Exactly when WFOs need the most reliable feed, Ku-band becomes least reliable.

C-band does not exhibit this failure mode.

Bottom line

Ku-band cannot meet the meteorological, operational, and COOP requirements of NOAAPort. Its rain fade characteristics, link instability, shared-spectrum environment, and lack of deterministic broadcast capacity make it fundamentally unsuitable as a C-band replacement.

II. What existing commercial geosynchronous platforms can transmit using a weather-resilient band?

There are *no* commercial GEO satellite platforms using a weather-resilient band comparable to C-band that NOAA could realistically adopt as a NOAAPort replacement. Everything commercially available today either uses Ku-band or Ka-band, both of which are *more* vulnerable to rain fade than C-band, not less.

In search results I saw the GeoXO program, confirming NOAA is planning **new weather-observation satellites**, **BUT NOT** new distribution satellites, and they do not introduce any new “weather-resilient” communications bands.

Below is the full breakdown.

1. Commercial GEO satellites today use Ku, Ka, or mixed Ku/Ka — none are weather-resilient

Every major commercial GEO operator (Intelsat, SES, Eutelsat, Viasat, Hughes, JSAT, Arabsat, Hispasat) uses:

- **Ku-band** (10–14 GHz) → moderate rain fade
- **Ka-band** (18–30 GHz) → *severe* rain fade
- **Occasional L-band feeder links** → but not for high-rate broadcast

None of these bands approach the precipitation resilience of **C-band (3.7–4.2 GHz)**.

There is **no commercial GEO service** offering wide-area broadcast in:

- S-band (2–4 GHz)
- L-band (1–2 GHz)
- UHF (300–900 MHz)

These bands *would* be weather-resilient, but they are not used for high-rate GEO broadcast because:

- They cannot support the bandwidth needed for meteorological data distribution.

- They are heavily allocated to aviation, mobile, military, and MSS services.
- GEO payloads in these bands are extremely rare and not commercially leased.

2. C-band itself is disappearing from commercial GEO

The 3.7–4.2 GHz downlink band historically used for NOAAPort is being:

- repurposed for 5G in the U.S.
- reduced in available transponder inventory
- deprioritized by satellite operators

Commercial operators are retiring or repacking C-band transponders rather than expanding them.

This is why NOAA's current NOAAPort lease is one of the last remaining large-scale C-band meteorological broadcasts.

3. Ka-band “high-throughput satellites” are the opposite of weather-resilient

Ka-band HTS platforms (Viasat-3, Inmarsat GX, SES-17, Eutelsat Konnect) offer huge bandwidth but:

- suffer **10–20×** the rain attenuation of C-band
- require adaptive coding/modulation
- cannot guarantee deterministic multicast during severe weather
- are designed for consumer broadband, not fixed broadcast

For meteorological distribution, Ka-band is *less* reliable than Ku-band.

4. L-band GEO exists (Inmarsat), but cannot carry NOAAPort-scale data

Inmarsat's L-band GEO constellation is extremely weather-resilient — but it offers:

- ~432 kbps per channel (classic)
- ~1–2 Mbps (BGAN)
- high cost per bit
- no wideband broadcast capability

NOAAPort requires **50–100+ Mbps sustained multicast**. L-band cannot do this.

5. S-band GEO is essentially nonexistent commercially

S-band is used for:

- NASA TDRSS
- some military relay systems
- a few experimental mobile-satellite systems

There is **no commercial GEO S-band broadcast** with the bandwidth needed for meteorological distribution.

6. GeoXO (NOAA's next-generation GEO satellites) do not solve distribution

The search results you saw describe **GeoXO**, NOAA's replacement for GOES. GeoXO provides *observations*, not *data distribution*. It does not introduce:

- new weather-resilient comms bands
- a replacement for NOAAPort
- a new GEO broadcast service

GeoXO is about sensors (imager + sounder), not communications.

7. Conclusion: No commercial GEO platform offers a weather-resilient band suitable for NOAAPort replacement

To replace C-band NOAAPort, NOAA would need:

- a GEO band with **low rain fade**
- **wide-area broadcast capability**
- **tens of Mbps sustained throughput**
- **long-term regulatory stability**

No commercial GEO band meets these criteria today.

The only bands that *could* (L-band, S-band) lack the bandwidth. The bands that have bandwidth (Ku, Ka) lack the weather resilience.

III. If you were to guess, Stonie, what would a future NOAA data distribution mechanism look like?

In a vacuum, conclusions can only be drawn from what is logical, and what is available.

1. High-level design

Core idea:

- Terrestrial IP (N-Wave + commercial carriers) becomes the primary distribution path.
- A GEO satellite feed (likely Ku/Ka) becomes a degraded-service fallback, not the main artery.
- Every WFO/RFC/center ingests data via multiple independent paths, with AWIPS and LDM aware of path health and priority.

Main components:

- National core: NOAA/NCEP data centers + N-Wave backbone.
- Regional hubs: Regional aggregation points with strong connectivity.
- Edge sites: WFOs, RFCs, national centers (SPC, WPC, NHC, etc.).
- Satellite uplink: One or more GEO transponders carrying a reduced, critical subset of products.
- Path-aware ingest layer: LDM/AWIPS logic that understands multiple feeds and failover.

2. Terrestrial primary path

2.1 Backbone: N-Wave + peering

- N-Wave provides the federal backbone: high-capacity, redundant, MPLS/segment-routed core.
- NOAA peers with:
 - Internet2 / ESnet for research partners.
 - Major carriers for last-mile diversity.
- Core services:
 - Multicast or replicated unicast for NOAAPort-like streams.
 - Message-oriented distribution (e.g., Kafka, AMQP, or LDM over TCP) for products.

2.2 Regional aggregation

- Regional hubs (e.g., each NWS region) host:
 - Regional caches of model data, satellite imagery, radar mosaics.
 - Local fan-out to WFOs via:

- Dedicated circuits (MPLS/VPN).
- Dual-homed commercial fiber.
- Fixed wireless or microwave in rural areas.

2.3 Edge connectivity

Each WFO/RFC gets at least two terrestrial paths:

- Primary: N-Wave or dedicated carrier circuit.
- Secondary: Different carrier, different physical route (fiber, microwave, or 5G fixed wireless).
- Optional third path: Business-grade Internet with encrypted tunnels (IPsec/SSL VPN).

Goal: >99.99% availability for core data, independent of any single carrier or route.

3. GEO satellite fallback

3.1 Role of GEO in the new architecture

Satellite becomes:

- Not the main firehose.
- A “lifeboat” channel (i.e. NWWS) carrying:
 - Critical warnings (TOR, SVR, FFW, etc.).
 - Key model summaries (e.g., select grids, not full fields).
 - Essential radar mosaics at reduced resolution.
 - Core satellite imagery (e.g., CONUS/Full-disk at lower cadence or resolution).

Think of it as: *“If everything terrestrial breaks, you still have enough to warn people.”*

3.2 Band and platform

Realistically:

- NWWS would have to switch from C-band to Ku or Ka GEO leased from a commercial operator.
- Uplink from:
 - A primary NOAA uplink site.

- A hot-standby uplink at a geographically separate facility.

Mitigation strategies for rain fade at uplink/downlink:

- Site diversity: Multiple uplink sites in different climates.
- Adaptive coding/modulation: To keep the link alive at lower throughput.
- Smaller, robust dishes at WFOs with:
 - Good radome design.
 - Automatic tracking/pointing.
 - Conservative link margins.

3.3 Content selection

Because GEO fallback has limited, less reliable bandwidth:

- Tier 1 (must-have):
 - Warnings and watches.
 - SPC meso discussions and outlooks.
 - NHC advisories.
- Tier 2 (high-value, reduced):
 - Coarse-resolution model fields (e.g., 0.5° global, key parameters).
 - National radar mosaics at reduced temporal/spatial resolution.
 - Key satellite imagery sectors.
- Tier 3 (optional, opportunistic):
 - Additional products when link margin is high.

4. Multipath ingest and AWIPS/LDM changes

4.1 Path abstraction

At each WFO/RFC, define logical feeds:

- Feed A: Terrestrial primary (N-Wave).
- Feed B: Terrestrial secondary (carrier 2 / VPN).
- Feed C: GEO satellite fallback.

LDM/AWIPS treat these as equivalent sources of the same product IDs, with:

- Priority order: $A > B > C$.
- Health monitoring: Latency, loss, completeness.

4.2 Product-level redundancy

For each product (e.g., GRIB2 file, radar tile, BUFR message):

- Ingest from whichever path delivers it first.
- Validate:
 - Size.
 - Checksum.
 - Sequence number.
- Discard duplicates from other paths.

This avoids double-processing and ensures first-arrival wins, regardless of path.

4.3 Degradation modes

When terrestrial paths degrade:

- Mode 0 (normal): Full product set via terrestrial; satellite mostly idle or used for monitoring.
- Mode 1 (partial terrestrial loss): Some products missing; satellite fills in Tier 1/Tier 2.
- Mode 2 (terrestrial outage): Satellite provides minimal but sufficient dataset; AWIPS switches to “degraded operations” profile:
 - Coarser grids.
 - Fewer imagery sectors.
 - Emphasis on warnings and situational awareness.

AWIPS UI could expose a “data health” indicator so forecasters know they’re in degraded mode.

5. Resilience, testing, and governance

5.1 COOP and drills

- Regular failover exercises:
 - Simulate terrestrial outage.

- Force AWIPS to rely on satellite + local caches.
- Measure:
 - Time to fail over.
 - Product completeness.
 - Forecaster usability.

5.2 Governance and SLAs

- Formal SLAs with:
 - N-Wave.
 - Commercial carriers.
 - Satellite operator(s).
- Clear RTO/RPO targets:
 - Recovery Time Objective for full data.
 - Recovery Point Objective for critical warnings (near-zero).

5.3 Security

- End-to-end encryption and authentication on terrestrial paths.
- Satellite feed:
 - Encrypted uplink.
 - Authenticated receivers.
- AWIPS/LDM:
 - Product signing/verification to prevent tampering.

6. Why this is realistic (and not sci-fi)

This architecture:

- I. Uses existing technologies:
 - N-Wave.
 - Commercial GEO.
 - VPNs, MPLS, segment routing.
 - LDM/AWIPS enhancements.

II. Respects physics:

- C-band is gone; Ku/Ka are imperfect, so they're fallback only.

III. Aligns with COOP reality:

- Multiple independent paths.
 - Satellite as last resort, not first line.
-

IV. Specific to the LDM, what is Stonie doing to prepare for a scenario like this?

I have already started to implement a new data insertion paradigm using a highly configurable file pick-up paradigm, capitalizing on the Linux kernel inotify subsystem. This would be a drop-in replacement or companion for the noaaportIngestor application already within the LDM suite of applications. This would allow for an easily configurable mechanism to insert data from multiple points and populate a future IDD essentially as crowd sourced.

The main obstacle now in that development is the productID assignment, along with sequence number assignment, in such a way as to maintain backwards compatibility for legacy databasing and display software.