



JPEG-XS for Tactical ISR

TECHNICAL WHITEPAPER
MAY 2026

James Spriet
EW/ISR SME
Mission Systems Lead

Why the codec layer is now an architectural decision, and how DDV brings visually lossless, deterministic low-latency video to tactical ISR platforms today.

WHITEPAPER

Published by
Delta Digital Video, a Division of Delta Information Systems, Inc.
747 Dresher Road, Suite 125, Horsham, PA 19044

Audience. Program managers, system architects, prime integrators, and engineering leads for airborne, maritime, and ground ISR programs.

Featured products. Model 5480E UHD/HD/SD Video Encoder · Model VEC500 OEM Encoder Module · Model 9600 and 9610 Video Distribution Systems

Table of Contents

Executive summary	3
1. The bandwidth and latency problem in ISR.....	4
2. How JPEG-XS works, and why it matters.....	5
Intra-frame only	5
Wavelet transform with a fixed 5/3 filter	5
Slice hierarchy bounds latency in image lines.....	5
Stateless entropy coding.....	5
Latency in microseconds at the codec stage.....	5
3. Latency, decomposed	7
A. Codec-stage latency	7
B. Encoder appliance latency	7
C. Network and transport latency.....	7
D. Decoder, display, and exploitation latency	7
4. Where JPEG-XS already lives.....	9
5. How DDV implements JPEG-XS	9
5.1 Model 5480E UHD/HD/SD Video Encoder	10
5.2 Model VEC500 OEM Encoder Module.....	10
6. Mission applications by domain.....	11
6.1 Small UAS, gimbals, and seekers.....	11
6.2 Multi-INT mission recorders	11
6.3 Shipboard and ground station distribution	12
6.4 Cross-domain solutions and high-assurance review.....	12
6.5 Machine vision at the edge.....	12
6.6 Contested communications and DDIL	12
The honest framing	14
Three properties of JPEG-XS still apply on a tactical IP fabric under degradation	14
7. Timing and synchronization.....	15
8. Codec selection by mission thread.....	16
9. Procurement-ready requirement language	17
Verification.....	18
10. Lifecycle and adoption risk	18
11. Where JPEG-XS is not the right tool	18
12. How to integrate	19
Procurement notes	19
Sources	20
About the author	21
About Delta Digital Video.....	21

Executive summary

Modern ISR sensors produce more raw pixels per second than the network and the codec layer behind them were architected to carry. H.264 and H.265 solved the broadcast and surveillance problem of the last decade by trading latency, determinism, and forensic fidelity for compression ratio. That trade remains correct for analyst-fed PED, long-duration archive, and bandwidth-starved distribution. It is the wrong trade for the part of the ISR architecture where the consumer of video is a targeting solution, a machine vision model, a cooperative engagement loop, or a downstream re-exploitation chain.

JPEG-XS, formally ISO/IEC 21122, is the codec built for that second problem. At mezzanine compression ratios of roughly 6:1 to 10:1, it delivers visually lossless quality, with codec-stage latency measured in image lines rather than frames, intra-frame independence with no Group of Pictures, a deterministic, FPGA-native implementation, and direct support for raw Bayer transport. JPEG-XS is the compressed video essence defined by SMPTE ST 2110-22 and one of the principal mezzanine codecs used in professional broadcast IP video infrastructure. Defense IP video is being written on the same standards stack.

DDV ships JPEG-XS support in the Model 5480E commercial UHD encoder and the Model VEC500 OEM encoder module ahead of the broader military FMV encoder market. The same hardware family covers legacy MPEG-2 TS / STANAG distribution today and SMPTE ST 2110-22 tactical IP fabric tomorrow. The Model 9600 and 9610 distribution systems decode and re-encode across codecs at the operations center boundary. Program offices do not have to commit to a single codec generation. They can specify the codec by mission thread and field the same hardware against all of them.



WHAT THIS WHITEPAPER ARGUES

Codec choice is now an architectural ISR decision. It determines whether sensor-to-exploitation latency budgets close, whether forensic fidelity survives the recorder, whether cross-domain accreditation is tractable, and whether machine vision exploitation performs as designed. JPEG-XS is the best fit for threads where latency, determinism, forensic fidelity, and machine-vision quality dominate. H.264 and H.265 remain the best fit for the threads where compression ratio, link economics, and legacy interoperability dominate. The DDV 5480E and VEC500 carry all three codecs on the same hardware, so the codec selection becomes a configuration decision per mission thread rather than a procurement decision per platform.

1. The bandwidth and latency problem in ISR

Sensor capability has outpaced the codec architecture ISR has relied on for fifteen years. A representative sample of fielded and emerging sensor systems makes the gap visible.

Sensor / Platform	Native data character	Raw bandwidth class
L3Harris WESCAM MX-25	Up to six sensors on the MX-25, up to nine on the MX-25D, including HD EO, MWIR, SWIR, and laser designator. Multiple simultaneous HD streams.	Hundreds of Mbps to low Gbps per turret
Collins Aerospace MS-177 on RQ-4B Global Hawk Block 30	Long-range multispectral imaging sensor, SYERS-class. Initially demonstrated on MQ-9 Avenger in 2016; operational on Global Hawk Block 30 since 2021.	Hundreds of Mbps mode-dependent
Sierra Nevada Gorgon Stare Increment 2	WAMI pod pair on MQ-9. EO sensor derived from DARPA ARGUS-IS (368 five-megapixel imagers, 1.8 gigapixel composite at variable frame rate) plus separate IR pod.	Hundreds of Gbps raw focal plane throughput before compression and product generation. Operational downlink rates are configuration dependent and substantially lower.
RTX ASARS-2B SAR	Wide-area SAR, multiple modes, GMTI overlay.	Hundreds of Mbps mode-dependent
GA-ASI Lynx Multi-Mode Radar	SAR plus GMTI plus maritime modes on Group 4/5 UAS.	Hundreds of Mbps mode-dependent

Table 1. Representative sensor data rates that frame the codec problem.

The transport story is just as stiff. Standard Common Data Link return rates are selectable at 10.71, 137, or 274 Mbps in the Ku-band waveform specification. Tactical CDL extends range but tops out lower on most airframes. LEO SATCOM helps, but is still finite. Sensors produce in the gigabit range. The link carries in the high hundreds of megabits at best. The codec layer has historically absorbed that mismatch by compressing aggressively. H.264 and H.265 reach 100:1 or higher by exploiting temporal redundancy, reordering frames, and using stateful entropy coding. Those design choices cost latency, determinism, and forensic fidelity. For analyst-fed PED chains, the cost is often acceptable. For an engagement loop, a targeting solution, or a machine vision model, the cost shifts from an acceptable compression trade to an architectural constraint.

JPEG-XS absorbs less bandwidth pressure (6:1 to 10:1 rather than 100:1) in exchange for fixing every other property at the same time. It is the right answer for the part of the kill chain where latency, determinism, and forensic fidelity dominate.

2. How JPEG-XS works, and why it matters

JPEG-XS, formally ISO/IEC 21122, is a mathematically clean intra-frame wavelet codec. Every design decision exists to satisfy three constraints: deterministic ultra-low latency, FPGA implementability at low cost, and predictable rate control.

Intra-frame only

Each frame is encoded independently. No inter-frame prediction, no I/P/B Group of Pictures, no motion estimation, no temporal entropy state. A bit error or packet loss affects at most one slice of one frame. There is no GOP corruption, no decoder reset, no error propagation across frames.

Wavelet transform with a fixed 5/3 filter

The codec uses a discrete wavelet transform with the LeGall 5/3 lifting filter over small spatial tiles called precincts. Compared to the 8x8 DCT in JPEG and H.264, wavelets avoid blocking artifacts and preserve high-frequency content more uniformly across spatial bands. That matters for ISR because small objects, targets, and edges live in high-frequency content.

Slice hierarchy bounds latency in image lines

A frame is divided into precincts, then into slices. A slice is the smallest independently decodable unit. Profiles specify a maximum slice height, typically 8, 16, or 32 image lines. This is the codec layer latency budget, measured in image lines rather than frames.

Stateless entropy coding

JPEG-XS does not use the context-adaptive binary arithmetic coding (CABAC) that drives H.265 efficiency. It uses a minimalist multi-pass scheme over horizontal lines of quantized wavelet coefficients: a significance bit, a Golomb-coded bitplane count (optionally DPCM-predicted from the prior line), raw bitplane data, and a sign pass. No context model, no arithmetic state machine, no branch divergence per symbol. JPEG-XS accepts lower compression efficiency than CABAC-based codecs in exchange for an order-of-magnitude reduction in implementation cost and fully deterministic silicon. For FPGA implementation and cross-domain accreditation, this is decisive.

Latency in microseconds at the codec stage

Because each slice is encodable as soon as its last line arrives at the input, the codec latency is bounded by the slice height plus a small, fixed processing overhead. The ISO/IEC 21122 specification frames this as 1 to 32 image lines of added end-to-end latency over an uncompressed path. At 1080p60, where one image line is roughly 15.4 microseconds, that puts the codec stage between roughly 250 microseconds and 1 millisecond, depending on slice height. In quality-optimized long-GOP configurations common to bandwidth-constrained FMV distribution, H.265 encode-plus-decode latency is typically in the tens of milliseconds or higher. H.265 can be configured for lower latency using all-intra, low-delay, or short-GOP modes, but those configurations reduce compression efficiency and remain harder to bound deterministically than JPEG-XS. Note that an end-to-end appliance budget on a fielded encoder LRU adds packetization, protocol stack, and



network buffering on top of the codec stage. The specification of as low as twenty milliseconds reflects that total appliance path, not the codec alone.

Property	JPEG-XS	H.265 (HEVC)	H.264 (AVC)	JPEG 2000
Typical compression ratio	6:1 to 10:1	100:1 to 200:1	50:1 to 150:1	10:1 to 30:1 (lossy), 1:1 to 3:1 lossless
Codec-stage latency	Tens to hundreds of microseconds; bounded by slice height, lines, not frames	Tens of ms or higher in quality-optimized long-GOP FMV configurations; lower latency possible with compression-efficiency trade	Single-digit to tens of ms possible in low-latency modes; compression-efficiency trade applies	Tens of milliseconds are typical
Inter-frame state	None	I/P/B GOP, motion compensation, reference picture buffer	I/P/B GOP, motion compensation	None
Entropy coder	Bitplane/Golomb, deterministic	CABAC, context-adaptive	CABAC or CAVLC	Arithmetic (MQ coder)
Error containment	Slice-bounded (image lines)	GOP-bounded (frames); IDR resync required	GOP-bounded; IDR resync required	Tile-bounded
FPGA implementability	Low-cost FPGA implementation; can avoid external frame-buffer DRAM in appropriate profiles and implementations	ASIC preferred, external DRAM typical	Moderate FPGA cost, external DRAM typical	Moderate FPGA cost
Raw Bayer support	Yes, dedicated RAW Bayer profile group with Star-Tetrix transform	No native Bayer path	No native Bayer path	Limited
Mathematically lossless	ISO/IEC 21122-2 lossless profile available	No	No	Yes, established lossless mode
Standards anchor	ISO/IEC 21122; SMPTE ST 2110-22; GenICam (2024); AMWA NMOS BCP-006-01	ITU-T H.265 / ISO/IEC 23008-2; MPEG-2 TS / STANAG 4609 endpoints	ITU-T H.264 / ISO/IEC 14496-10; MPEG-2 TS / STANAG 4609 endpoints	ISO/IEC 15444; SMPTE ST 2042 (VC-3 lineage), broadcast contribution

Table 2. Codec comparison on the properties that matter for ISR.

3. Latency, decomposed

Codec latency does not exist in isolation. A sensor-to-exploitation budget has four distinct layers, and a codec change only addresses one of them. Conflating the layers is how vendor claims become indefensible under review.

A. Codec-stage latency

The time from a pixel arriving at the encoder input to the same pixel being available at the decoder output, exclusive of all other system effects. This is where JPEG-XS is materially different from H.264 and H.265. ISO/IEC 21122 frames codec-stage latency in image lines. At 1080p60, one image line is roughly 15.4 microseconds. A typical JPEG-XS profile adds roughly 8 to 32 lines, putting the codec stage in the hundred-microsecond to one-millisecond range. In quality-optimized long-GOP configurations common to bandwidth-constrained FMV distribution, H.265 encode/decode latency is typically in the tens of milliseconds or higher. H.265 can be configured for lower latency using all-intra, low-delay, or short-GOP modes, but those configurations reduce compression efficiency and remain harder to bound deterministically than JPEG-XS.

B. Encoder appliance latency

The time from the sensor electrical interface to the network egress at the encoder LRU. This includes ISP and color processing, encoder pipeline fill, packetization to MPEG-2 TS, RTP, SRT, RIST, or ST 2110-22, encryption, and PHY transmit. The DDV 5480E specification of as low as 20ms is an appliance-level figure, not a codec-stage figure. Once the codec is no longer the bottleneck, packetization, buffering, encryption, and the protocol stack become the dominant latency-control problems. That is a feature, not a defect, because those layers are addressable through implementation rather than codec choice.

C. Network and transport latency

Switching, routing, error correction, store-and-forward, link MTU, FEC overhead, queueing, and one-way propagation. On a managed 2110 fabric, this can be sub-millisecond. On a satellite link, physics dominates.

D. Decoder, display, and exploitation latency

Decode pipeline fill, frame buffer depth, display refresh phase, AI ingest queueing, operator console paint cycle. On commercial displays, this is often the largest single contributor and is bounded by the display's vertical refresh.



Layer	JPEG-XS contribution	H.265 contribution	Notes
Sensor readout	1 frame interval at worst, less with rolling shutter region access	Same	Codec-agnostic. Sensor and ISP problem.
ISP/image processing	Implementation-dependent. Often single-digit ms. The RAW Bayer path, using the JPEG-XS Bayer profile, can avoid full debayering before encoding.	Same. The Bayer path generally requires a full ISP before encoding.	Where the VEC500 MIPI CSI-2 plus Bayer-profile JPEG-XS path is differentiated.
Codec stage	~250us to ~1ms	Tens of milliseconds or higher in quality-optimized long-GOP FMV configurations; lower latency possible with all-intra, low-delay, or short-GOP modes at compression-efficiency cost.	Primary delta.
Packetization (RTP/TS/SRT/RIST /ST 2110-22)	Single-digit ms typical	Same	Implementation. JPEG-XS over ST 2110-22 has the cleanest path.
Network transport	Sub-millisecond on managed fabric; tens of ms on tactical CDL; hundreds of ms on SATCOM	Same	Codec-agnostic.
Decoder	~250us to ~1ms	Mirrors codec-stage behavior; lower latency possible with compression-efficiency trade.	Mirrors codec stage.
Display/exploitation/ AI ingest	8 to 16ms (one frame at 60 Hz) typical; AI ingest may be faster	Same	Codec-agnostic.

Table 3. Layered latency budget at 1080p60.

JPEG-XS reduces the codec stage and decoder from dominant latency contributors to bounded, sub-frame contributors. It does not remove the other latency layers. A program that wants a deterministic glass-to-glass or sensor-to-exploitation budget has to design the timing architecture, packetization, network, and display path with the same discipline applied to the codec. JPEG-XS makes that design tractable. It does not perform that design.

Engagement type	Total sensor-to-exploitation budget	Codec stage budget	Achievable with JPEG-XS	Achievable with H.265
Analyst-in-loop PED	Seconds, tolerant	Tens of ms acceptable	Yes	Yes
Sensor-to-shooter, fixed target	Several hundred ms	Single-digit ms preferred	Yes	Yes with low-latency config and compression cost
Moving target engagement	Roughly 200ms	Single-digit ms required	Yes	Possible but tight; harder to bound
Cooperative engagement with wingman / CCA	Tens to low hundreds of ms	Sub-millisecond preferred	Yes	Difficult to bound deterministically
Active protection, counter-UAS terminal	Tens of ms total	Sub-millisecond required	Yes	Not appropriate

Table 4. Engagement budgets by mission type.

The bottom two rows are where JPEG-XS is the better-aligned choice. The top three are where either codec can be made to work, and the decision is driven by compression ratio, link economics, and exploitation type.

4. Where JPEG-XS already lives

SMPTE ST 2110-22, ratified in 2019, defines JPEG-XS as the standard compressed video essence for IP video plant infrastructure. Major broadcasters have replaced SDI baseband infrastructure with 2110-based IP fabrics, with JPEG-XS keeping bandwidth inside rack-to-rack and inter-facility budgets. AMWA NMOS BCP-006-01 defines the discovery, registration, and connection management profile for NMOS deployments carrying JPEG-XS. Defense IP video is being written on the same standards stack. The transition runs alongside, not on top of, legacy STANAG 4609 distribution. STANAG 4609 endpoints are typically MPEG-2 Transport Stream with H.264 or H.265 essence and MISB KLV. SMPTE ST 2110-22 is the compressed IP essence path where JPEG-XS belongs. DDV products carry both transport stacks so program offices can field one hardware family across legacy and IP fabric mission threads. On the FPGA IP side, intoPIX TicoXS and Fraunhofer IIS reference implementations run on AMD Xilinx, Intel, and Lattice families. JPEG-XS was admitted to the GenICam machine vision standard in 2024, putting compressed essence directly at the imager interface. DDV shipped JPEG-XS in the Model 5480E and the Model VEC500 ahead of most other dedicated FMV encoder vendors.

5. How DDV implements JPEG-XS

DDV designs and manufactures video compression, recording, scan conversion, and distribution products for mission-critical applications across military air, land, and sea domains. The product line spans ultra-compact single-channel airborne encoders, OEM modules for camera and sensor integrators, multi-channel rugged encoders, and rack-

mount distribution systems. The JPEG-XS capability lives in the 5480E and the VEC500 today.




5.1 Model 5480E UHD/HD/SD Video Encoder

Compact, ruggedized network appliance for video-over-IP. A single device that encodes UHD, HD, or SD inputs in H.264, H.265, or JPEG-XS and streams over SMPTE ST 2110, SRT/RIST, RTSP/RTP/UDP, or MPEG-2 TS. One 5480E covers the bridge case most program offices face: legacy STANAG 4609 today, SMPTE 2110-22 tomorrow, same hardware, codec selection as a configuration choice.

Attribute	Specification
Video compression	H.264 (BP/MP/HP L5.1, 4:2:2 8-bit), H.265 (MP/M10P L5.1, 4:2:2 8/10-bit), JPEG-XS (4:2:2 8/10-bit)
Video input	SDI (SMPTE), HDMI, Composite NTSC/PAL
Resolutions/Frame rates	2160p, 1080p, 720p, 1080i, 480i, 576i / 25, 29.97, 30, 50, 59.94, 60 fps
Streaming data rate	64 Kbps to 12 Gbps
Streaming protocols	MPEG-2 TS, RTSP/RTP, SMPTE ST 2110, SRT/RIST
Encode latency	As low as 20 ms (JPEG-XS configurations)
Metadata	SDI VANC, RS-232, Ethernet. KLV or clear. MISB-compliant.
Encryption/Time sync	AES-128/AES-256 · GPS, IRIG-B (DC), PTP, NTP
Power/Environmental	+9 to +36 VDC, 15 W · Industrial -38 to +74 C, Storage -55 to +85 C
Size and weight	1.13" H x 5.50" W x 6.49" D, 1.1 lb
MTBF	Approximately 55,000 hours (airborne, inhabited cargo, +25 C)

Table 5. Model 5480E technical specifications.

Note: Encode latency is at the appliance level and is measured from sensor input to network egress, including codec, packetization, encryption, and the protocol stack. Codec-stage latency at 1080p60 is in the hundred-microsecond to one-millisecond range, depending on slice height.



5.2 Model VEC500 OEM Encoder Module

OEM module for camera and sensor manufacturers, gimbal vendors, and platform integrators who need to add tactical video streaming to their device rather than carry a separate encoder LRU. SDI or MIPI CSI-2 input directly from a sensor; H.264, H.265, or JPEG-XS encoding; GbE streaming with the same protocol stack as the 5480E. When the sensor, JPEG-XS RAW Bayer profile, and processing chain support raw Bayer transport, the VEC500 can avoid a full debayer-before-encode path, collapsing

several ISP stages into a single OEM board. Configurations that require stabilization, overlays, color correction, AI preprocessing, or visible-channel rendering for human operators may still require ISP functions in the chain.

Attribute	Specification
Video compression	H.264 (BP/MP/HP L5.1, 4:2:2 8-bit), H.265 (MP/M10P L5.1, 4:2:2 8/10-bit), JPEG-XS optional
Video input	SDI (SMPTE) and MIPI CSI-2
Resolutions/Frame rates	2160p, 1080p, 720p, 1080i, 480i, 576i / 25, 29.97, 30, 50, 59.94, 60 fps
Streaming data rate	64 Kbps to 1 Gbps
Streaming protocols	MPEG-2 TS, RTSP/RTP, SRT/RIST (SMPTE ST 2110 via DDV SDK configuration)
Metadata	SDI VANC, RS-232, Ethernet. KLV or clear. MISB-compliant.
Encryption/Time sync	AES-128/AES-256 · GPS, IRIG-B (DC), PTP, NTP
Power/Environmental	+9 to +36 VDC (12 VDC nominal), 10 W · Industrial -40 to +71 C
Size and weight	3.03" W x 2.36" D x 0.84" H, 3.6 oz (102 g)

Table 6. Model VEC500 technical specifications.

6. Mission applications by domain

6.1 Small UAS, gimbals, and seekers

ASIC encoders for 4K H.265 sit in the NVIDIA Jetson and Ambarella SoC class. Power budget runs five to fifteen watts, including external DDR. On a five- to twenty-pound aircraft, that is a meaningful share of the platform's power budget. JPEG-XS can be implemented in low-cost FPGA profiles with low fabric power and, in appropriate profiles and implementations, without external frame-buffer DRAM. That moves the codec into the gimbal or onto the imager carrier board rather than forcing a separate encoder LRU. This is the design point for which the VEC500 was built. MIPI CSI-2 direct from the sensor, plus the JPEG-XS RAW Bayer profile, can avoid a full debayer-before-encode path when the sensor, profile, and processing chain support it. In configurations that require image stabilization, overlays, color correction, AI preprocessing, or visible-channel rendering for human operators, the system may still require ISP functions in the chain. The architectural win is fewer mandatory ISP stages, not the universal elimination of the ISP.

6.2 Multi-INT mission recorders

Traditional tactical recorders ingest sensor data, apply aggressive lossy compression, and ship chipped distribution copies forward. That discard is forensic and irreversible. A JPEG-XS recording architecture ingests EO, IR, SAR product, and ELINT metadata at 6:1

to 10:1 visually lossless, frame-aligned via PTP under ST 2110-10, with KLV bound to frame. The result is forensically higher-fidelity than aggressively compressed inter-frame distribution copies and better suited to re-exploitation and ML training data generation. It is not mathematically lossless. Programs that require bit-exact reproduction for evidentiary or scientific purposes should specify either JPEG 2000 lossless or the JPEG-XS mathematically lossless profile defined in the next edition of ISO/IEC 21122-2. DDV encoders ship the metadata pipeline that this architecture requires.

6.3 Shipboard and ground station distribution

The Model 9600 and 9610 pair with the 5480E and VEC500 to form a complete encoder-to-distribution path for shipboard combat systems, ground stations, and tactical operations centers. The 9600 supports 15 channels of HD/SD in a 2U enclosure with IRIG-B time sync. The 9610 provides eight channels in 1U with optional integrated recording. Both decode and re-encode across codecs, which is the practical answer to coalition interoperability where some feeds arrive as H.264 over MPEG-2 TS, some as H.265, and some as JPEG-XS over ST 2110-22.

6.4 Cross-domain solutions and high-assurance review

Deterministic, intra-frame, slice-bounded decoding is structurally easier to analyze than long-GOP inter-frame pipelines. JPEG-XS has no inter-frame state, no arithmetic coder, and no context model. Each slice is a small, self-contained transform on a known pixel count. This makes JPEG-XS potentially simpler to validate through a data-diode chain and better suited to high-assurance review than stateful inter-frame codecs. Actual cross-domain accreditation depends on implementation hardening, parser robustness, protocol handling, metadata sanitization, test evidence, configuration control, and authority approval. The choice of codec removes one class of attack surface. It does not produce an accreditation.

6.5 Machine vision at the edge

H.265 quantization is matched to the human visual system. Modern computer vision models are not. Object detectors trained on uncompressed or lightly compressed imagery often degrade when fed aggressively compressed H.265 below a few Mbps at 1080p, with disproportionate degradation at small object scales. The wavelet quantization in JPEG-XS preserves high-frequency content more uniformly across spatial bands, more closely matching the data distribution most production CV models were trained on. JPEG-XS was adopted as a GenICam machine vision standard in 2024 for this reason. The honest qualifier is that JPEG-XS does not automatically improve every model. Any operational AI, ATR, or VMTI chain should validate model performance against the actual JPEG-XS profile, compression ratio, resolution, frame rate, sensor type, and mission scene. Models trained only on uncompressed or H.264 / H.265 imagery may require validation or fine-tuning when the compression regime changes.

6.6 Contested communications and DDIL

Bandwidth and link discipline matter more than codec marketing here. The right framing is bitrate math against link class, not codec tribalism.

Stream profile	Approximate bitrate	Notes
1080p60 4:2:2 10-bit uncompressed	~2.49 Gbps	Reference. Useful only on local SDI or 2110-20.
1080p60 4:2:2 10-bit JPEG-XS 6:1	~415 Mbps	Visually lossless quality.
1080p60 4:2:2 10-bit JPEG-XS 10:1	~250 Mbps	Visually lossless, typical operating point.
1080p60 4:2:2 10-bit JPEG-XS 15:1	~166 Mbps	Approaching the perceptual edge of visually lossless on demanding content.
4K60 4:2:2 10-bit uncompressed	~9.95 Gbps	Reference.
4K60 4:2:2 10-bit JPEG-XS 6:1	~1.66 Gbps	Requires 10 GbE or 2110 fabric.
4K60 4:2:2 10-bit JPEG-XS 10:1	~995 Mbps	Near 1 GbE line rate before transport overhead. 10 GbE preferred for operational margin, or reduce format/increase compression where acceptable.
4K60 4:2:2 10-bit JPEG-XS 15:1	~664 Mbps	
720p30 4:2:2 10-bit JPEG-XS 10:1	~55 Mbps	Practical for selected CDL-class operating points.
1080p30 H.264 STANAG 4609 ISR	~4 to 12 Mbps	Long-GOP, perceptual quality.
1080p30 H.265 ISR	~2 to 6 Mbps	Long-GOP.
480p30 H.264 narrowband	~0.5 to 2 Mbps	Tactical fallback.

Table 7. Bitrate reference by stream profile.

Bitrates are active-video approximations before transport overhead unless otherwise stated. Operational link budgets must include RTP/IP/Ethernet overhead, packet pacing, timing, ancillary data, encryption, FEC where used, and engineering margin.

Link class	Typical sustained rate	Practical codec fit
Managed 1/10 GbE IP fabric, SMPTE ST 2110	1 to 10+ Gbps	JPEG-XS at native resolution.
CDL Ku-band, full rate	137 or 274 Mbps return per JSC EMC waveform analysis	JPEG-XS may fit selected 1080p30 or 720p60 feeds depending on compression ratio, transport overhead, link margin, and mission tolerance. At 137 Mbps, profile selection and overhead control are critical. At 274 Mbps, selected HD JPEG-XS feeds are more practical. Full 4K is not supported at visually lossless ratios. H.264 / H.265 carries full 4K.
TCDL airborne return	Roughly 10.71 Mbps return is typical	H.264 or H.265. JPEG-XS only at substantially reduced resolution, frame rate, or region of interest.
Mil-Ka / Ku SATCOM	10s to low 100s of Mbps mission-dependent	JPEG-XS feasible for selected high-value low-latency feeds at reduced format. H.264 / H.265 for full sensor coverage.
Iridium Certus, L-band BLOS, jammed fallback	Sub-megabit to a few Mbps	H.264 or H.265 only. JPEG-XS does not fit at any operationally useful video format.

Table 8. Link class reference.

The honest framing

JPEG-XS is the right answer for managed high-bandwidth IP fabric, for selected high-value feeds on CDL-class links where latency matters more than full sensor coverage, and for any path where bounded codec latency and deterministic behavior are mission-critical. JPEG-XS is not the answer for narrowband links. The 5480E and VEC500 ship H.264 and H.265 because the DDIL fallback path is a real procurement requirement, and JPEG-XS at 6:1 to 10:1 does not fit on a 10 Mbps pipe at any tactically useful video format. A program-grade DDIL plan specifies the codec per link class, not per platform.

Three properties of JPEG-XS still apply on a tactical IP fabric under degradation

These properties apply where the link budget can already carry the selected JPEG-XS stream; they do not make JPEG-XS suitable for narrowband DDIL paths. Error resilience is slice-bounded rather than GOP-bounded, so a packet loss damages a few image lines, not a window of frames. Rate-adaptive intra-frame coding degrades visually rather than collapsing on a missing reference picture. Codec-stage latency stays bounded under buffer pressure because there is no GOP reset to wait for. These properties matter at the IP fabric layer, even when the codec is no longer the bottleneck.

7. Timing and synchronization

Deterministic video over IP requires a deliberate timing architecture, not just a codec. JPEG-XS reduces codec-stage uncertainty. The platform still has to control everything else.

■ Grandmaster

A program-grade ST 2110 deployment requires a PTP grandmaster, typically GPS-disciplined, with a documented holdover specification. Tactical platforms often run with a primary GPS-disciplined grandmaster and an IRIG-B (DC level shift) fallback for GPS-denied operation.

■ Boundary and transparent clocks

Switches in the fabric should be PTP-aware. Boundary clocks segment the timing domain. Transparent clocks correct for switch residence time. Mixing PTP-unaware switches into a 2110 fabric degrades timing accuracy in ways that are difficult to debug after the fact.

■ Holdover

Specify the required holdover accuracy under GPS denial in the requirements document. Class B and Class C holdover categories from ITU-T G.8273.2 can serve as starting points for tailoring requirements, but the final holdover requirement should be derived from the platform timing budget, mission duration under GPS denial, and KLV-frame alignment tolerance.

■ Frame-to-metadata alignment

KLV metadata bound to frame requires that the timing source feeding the encoder, the timing source feeding the metadata generator, and the timing source feeding the recorder are coherent. Drift between these sources produces a KLV-frame mismatch that survives the recording and surfaces in exploitation.

■ Degraded timing

Specify how the encoder, decoder, and recorder behave when the timing source is lost or degraded. A program-grade encoder should free-run with documented drift, not silently corrupt timestamps.

The DDV product line accepts PTP, IRIG-B, GPS, and NTP timing inputs. The timing architecture above is the integrator's design responsibility. JPEG-XS makes that design tractable. It does not perform the design.

8. Codec selection by mission thread

The following table makes codec choice a structured decision per mission thread rather than per platform.

Mission thread	Link condition	Consumer	Latency priority	Quality / forensic priority	Recommended codec	Recommended transport	DDV product fit
Analyst PED is over-constrained for the downlink	TCDL, narrowband SATCOM, Iridium-class	Human analyst	Low	Moderate	H.264 or H.265	MPEG-2 TS, SRT, RIST	5480E or VEC500 in H.264/H.265 mode
Pilot/operator visual feedback	Onboard wired or managed IP	Human pilot, operator	High	Moderate to high	JPEG-XS preferred	RTP, ST 2110-22	5480E or VEC500 in JPEG-XS mode
Gimbal control loop	Onboard high-bandwidth	Closed-loop controller, operator	Very high (sub-frame)	Moderate	JPEG-XS	RTP, ST 2110-22	VEC500 in JPEG-XS mode, MIPI CSI-2 direct
ATR/VMTI ingest	Managed IP fabric	Computer vision model	High	High (small-object fidelity)	JPEG-XS	RTP, ST 2110-22	5480E or VEC500 in JPEG-XS, validated per model
Multi-INT mission recorder	Onboard storage fabric	Post-mission exploitation, ML training	Moderate	Very high	JPEG-XS preferred; JPEG 2000 or JPEG-XS lossless when bit-exact required	ST 2110-22 ingest; native recording	5480E or VEC500 feeding mission recorder
Shipboard/ground station IP fabric	Managed 1/10 GbE	Multiple consumers	High	High	Mixed, codec per consumer	ST 2110-22 plus MPEG-2 TS legacy	9600/9610 decode and re-encode hub
Coalition STANAG endpoint	Coalition tactical IP	Partner force ground station	Moderate	Moderate	H.264 or H.265	MPEG-2 TS, MISB KLV	5480E or VEC500 in H.264/H.265 mode
Narrowband DDIL fallback	Iridium, L-band, jammed	Human analyst	Moderate	Moderate	H.264 or H.265	MPEG-2 TS, SRT, RIST	5480E or VEC500 in H.264/H.265 mode
Multi-INT re-exploitation archive	Storage	ML training, forensic	Low	Very high	JPEG-XS visually lossless or JPEG 2000 / JPEG-XS lossless for bit-exact	File or ST 2110-22 recording	5480E or VEC500 feeding archive

Table 9. Codec selection by mission thread.

9. Procurement-ready requirement language

Specify the codec in the requirements document, not as an aspiration. Specify it per mission thread. The following language is suggested as a starting point. Tailor to program, classification, and authority.

- Encoder shall support JPEG-XS per ISO/IEC 21122. Specify the profile group (Main, Light, Light-subline, High, RAW Bayer, or TDC), level, and sublevel appropriate to the mission thread's resolution and frame rate.
- Codec-stage latency shall be bounded by slice height. Encoder shall report the configured slice height and the resulting bounded codec-stage latency at acceptance.
- Appliance-level encode latency at the specified resolution and frame rate shall not exceed [program-specified figure] milliseconds at the network egress, measured per a documented acceptance procedure.
- Encoder shall support 4:2:2 at 8-bit and 10-bit. Where applicable, the encoder shall support the JPEG-XS RAW Bayer profile group with the documented sensor configuration.
- Where raw sensor transport is required, the encoder shall identify the JPEG-XS RAW Bayer profile group, supported Bayer pattern or sensor packing, bit depth, level, sublevel, and any required pre-processing or metadata needed to reconstruct the image chain.
- For ST 2110 deployments, requirements should address ST 2110-22 RTP encapsulation, ST 2110-10 timing, and applicable ST 2110-21 traffic-shaping behavior. Acceptance testing should verify the selected sender profile and traffic behavior on the target network.
- Encoder shall support H.264 and H.265 with MPEG-2 Transport Stream output for legacy STANAG 4609 endpoints, with full MISB KLV insertion and timestamp integrity.
- Encoder shall accept timing inputs from PTP IEEE 1588-2008 or IEEE 1588-2019, IRIG-B (DC level shift), GPS, and NTP. Holdover behavior under GPS denial shall be documented.
- Encoder shall support AES-128 and AES-256 encryption for the protected video and metadata stream, with the protected boundary, encryption mode, key-management interface, and performance impact documented.
- Encoder shall support SRT and RIST output for selected mission threads.
- Metadata binding shall preserve KLV-to-frame alignment under the specified timing source. Acceptance shall verify timestamp integrity end-to-end.
- Environmental and SWaP constraints shall be specified per the platform integration document. Encoder MTBF shall be documented per MIL-HDBK-217 or equivalent.

Verification

Every latency, bitrate, codec, and metadata claim above shall be testable at acceptance with documented procedure, measurement points, and pass/fail criteria. Vendor specifications that cannot be reproduced under a documented test should be treated as marketing, not as acceptance evidence.

10. Lifecycle and adoption risk

JPEG-XS adoption in defense is early but directionally aligned with the same IP video transition already visible in broadcast, machine vision, and managed video infrastructure. Programs adopting JPEG-XS should verify:

- Decoder availability across the program ecosystem, including coalition partners, exploitation tools, and CDS endpoints.
- Licensing model for the JPEG-XS patent pool and any FPGA IP cores (intoPIX TicoXS, Fraunhofer IIS, Lattice TICO-XS) used in implementation.
- Long-term IP-core support and roadmap from the encoder vendor.
- Interoperability with selected SMPTE ST 2110, AMWA NMOS, and BCP-006-01 tools.
- Fallback behavior to H.264 and H.265 endpoints, which the 5480E and VEC500 carry natively.

The fallback path to H.264 and H.265 is what makes the procurement decision tractable today. A platform that supports both stacks does not commit a program to a codec generation before the ecosystem requires the commitment.

11. Where JPEG-XS is not the right tool

JPEG-XS is not a universal solvent. The 5480E and VEC500 ship H.264 and H.265 alongside JPEG-XS because the right codec depends on the link, the consumer, and the storage budget.

■ Storage-bound archive

At 10:1, JPEG-XS saves an order of magnitude over uncompressed. H.265 at 200:1 saves two. For long-duration archival where perceptual drop is tolerable, H.265 wins on dollars per terabyte. Use H.265 mode on the 5480E.

■ Bandwidth-starved low-resolution links

For a 480p feed on a 1 Mbps link, JPEG-XS is too fat. H.264 or H.265 at 1 to 2 Mbps owns this slot. Use H.264 or H.265 mode on the 5480E or VEC500.

■ Mathematically lossless archive

JPEG 2000 lossless is still the right answer for bit-exact reproduction. The next edition of ISO/IEC 21122-2 adds a JPEG-XS mathematically lossless profile that closes this gap.

■ **STANAG 4609 endpoints that have not been updated**

If the ground station only speaks H.264, the 5480E and VEC500 ship H.264 with full MISB KLV support. JPEG-XS is the future path, not the only path.

12. How to integrate

Integration pattern	DDV product mix
Group 1-2 sUAS with a single EO/IR sensor	VEC500 OEM module inside the gimbal or sensor carrier board. MIPI CSI-2 direct from imager, JPEG-XS or H.264/H.265 selectable, 102 g, 10 W.
Group 3-5 UAS or manned ISR with existing CDL/TCDL	5480E standalone per channel, or VEC500 inside the sensor. Full SMPTE ST 2110, SRT/RIST, AES, KLV. Bridge from STANAG 4609 H.264 today to SMPTE 2110-22 JPEG-XS tomorrow.
Shipboard, ground station, FOB ops center	5480E and VEC500 as encoders, 9600 or 9610 as the distribution and re-encode hub. IRIG-B time sync, optional integrated recording on the 9610.
Multi-INT mission recorder for ML training data	JPEG-XS as recording essence, DDV encoders feeding over ST 2110-22. Frame-aligned via PTP, KLV bound to frame, 6:1 to 10:1, visually lossless.

Table 10. Integration patterns and DDV product mix.

Procurement notes

Specify the codec in the requirements document, not as an aspiration. Call out JPEG-XS profile, level, and sublevel where latency matters. Specify SMPTE ST 2110-22 as the IP transport for any new airborne or shipboard video subsystem, with H.264 or H.265 over MPEG-2 TS as the legacy path. Treat MISB KLV, AES-128/AES-256, and GPS/IRIG-B/PTP/NTP timing support as base requirements, not options. These capabilities should be verified against the selected DDV model and configuration during acceptance.

THE BOTTOM LINE

Codec choice is now an architectural decision. JPEG-XS is the best fit for ISR mission threads where latency, determinism, forensic fidelity, and machine vision performance dominate. H.264 and H.265 remain the best fit for threads dominated by compression ratio, link economics, and legacy interoperability. The DDV 5480E and VEC500 carry all three codecs on the same hardware. The 9600 and 9610 distribute and re-encode across them.

Sources

1. JPEG Committee, JPEG XS overview (ISO/IEC JTC1 SC29 WG1). <https://jpeg.org/jpegxs/>
2. Descampe et al., JPEG XS, a new standard for visually lossless low-latency lightweight image coding, Proceedings of the IEEE, vol. 109, no. 9, pp. 1559-1577, 2021. <https://doi.org/10.1109/JPROC.2021.3080916>
3. ISO/IEC 21122-1:2024, JPEG XS low-latency lightweight image coding system, Part 1: Core coding system. <https://www.iso.org/standard/85247.html>
4. Descampe et al., JPEG XS technical reference paper, Proceedings of the IEEE, vol. 109, no. 9, 2021. <https://doi.org/10.1109/JPROC.2021.3080916>
5. ISO/IEC 21122-2:2024, JPEG XS low-latency lightweight image coding system, Part 2: Profiles and buffer models. <https://www.iso.org/standard/85250.html>
6. NATO STANAG 4609 Ed 5, NATO Digital Motion Imagery Standard, Joint Capability Group on ISR, 2020. <https://nisp.nw3.dk/coverdoc/nato-stanag4609ed5.html>
7. Fraunhofer IIS, JPEG XS research and licensing. <https://www.iis.fraunhofer.de/en/ff/amm/content-production/jpegxs.html>
8. SMPTE, ST 2110 FAQ, Society of Motion Picture and Television Engineers. <https://www.smpte.org/smpte-st-2110-faq>
9. Descampe et al., FPGA, ASIC, CPU, and GPU implementation of JPEG XS, Proceedings of the IEEE, vol. 109, 2021. <https://doi.org/10.1109/JPROC.2021.3080916>
10. L3Harris, WESCAM MX-25D airborne targeting and designating. <https://www.l3harris.com/all-capabilities/wescam-mx-25d-airborne-targeting-and-designating>
11. GA-ASI, Predator C Avenger flies with MS-177 long-range sensor. <https://www.ga-asi.com/predator-c-avenger-flies-with-ms-177-long-range-sensor>
12. ASDNews, Collins Aerospace MS-177 operational on Global Hawk Block 30, 2021. <https://www.asdnews.com/news/defense/2021/09/21/collins-aerospaces-ms177-isr-sensor-now-supporting-global-hawk-operational-missions>
13. DARPA Director testimony, House Armed Services Subcommittee on Emerging Threats and Capabilities, 26 March 2015 (ARGUS-IS program reference). <https://www.darpa.mil/sites/default/files/attachment/2025-06/prabhakar-a20150326.pdf>
14. Headquarters United States Air Force, RPA Vector: Vision and Enabling Concepts 2013-2038, UNCLASSIFIED, February 2014. <https://www.af.mil/Portals/1/documents/news/USAFRPVectorVisionandEnablingConcepts2013-2038.pdf>
15. Forecast International, ASARS-2 archived report. https://www.forecastinternational.com/archive/disp_old_pdf.cfm?ARC_ID=103
16. GA-ASI, Lynx Multi-Mode Radar. <https://www.ga-asi.com/radars/lynx-multi-mode-radar>
17. Joint Spectrum Center, Common Data Link EMC analysis, JSC-04-044. <https://info.publicintelligence.net/JSC-04-044.pdf>

18. AMWA BCP-006-01, NMOS With JPEG XS interoperability specification. <https://specs.amwa.tv/bcp-006-01/>
19. JPEG XS Patent Pool, administered by Vectis IP on behalf of Fraunhofer IIS and intoPIX (co-licensors of ISO/IEC 21122 essential patents). <https://www.jpegxspool.com/home>
20. Lattice Semiconductor, TICO-XS JPEG XS encoder/decoder IP cores. <https://www.latticesemi.com/en/Products/DesignSoftwareAndIP/IntellectualProperty/IPCore/intoPIXCores/TICO-XS-JPEG-XS-Encoder-Decoder>
21. MISB ST 0601, UAS Datalink Local Set, NGA Motion Imagery Standards Board. <https://nsgreg.nga.mil/doc/view?i=5471>
22. OPSWAT, Data diodes in transfer cross-domain solutions. <https://www.opswat.com/blog/data-diodes-in-transfer-cds-securing-high-assurance-cross-domain-solutions>
23. BAE Systems, Data Diode Solution. <https://www.baesystems.com/en/product/data-diode-solution>
24. EMVA, GenICam Working Group updates including JPEG XS support in SFNC 2.8 and GenDC 1.2. <https://www.emva.org/news/genicam-working-group-updates/>
25. Modern War Institute at West Point, Reimagining Contested Communications. <https://mwi.westpoint.edu/reimagining-contested-communications/>

About the author

James Spriet is an EW/ISR subject-matter expert and mission systems lead with more than twenty years of experience across Electronic Warfare, ISR, Cyber Defense, and Space Control. Fifteen years in the United States Air Force across RC-135 Rivet Joint and Combat Sent, on numerous SAP special missions, C-130H and A-10C mission systems lead, and a final deployment to northern Afghanistan in 2021. Industry work also spans the DoD Satellite Control Network, the Shared Early Warning System, defensive and offensive space control, the DoD's first Mission Defense Team, and the first deployed MDT. Writes a column in the Journal of Electromagnetic Dominance, the official publication of the Association of Old Crows, which is the largest non-profit organization of EW/ISR experts in the world.

About Delta Digital Video

Delta Digital Video, a division of Delta Information Systems, Inc., designs and manufactures video compression, recording, scan conversion, and distribution products for mission-critical applications across military air, land, and sea domains, law enforcement, broadcast, intelligent transportation, and security markets. Headquartered at 747 Dresher Road, Suite 125, Horsham, PA 19044. Contact deltadigitalvideo.com.