01

DEVELOPMENT: 2003 - 2023

The McDonnell Douglas (now Boeing) F-15E Strike Eagle entered the twenty-first century as a Cold War, multi-mission workhorse pressed into action in a succession of wars it had not been optimised to fight. The Strike Eagle community meanwhile carried the weight of expectation from earlier successes in Iraq and the Balkans.

In those conflicts, the Strike Eagle had been proven lethal, adaptable, and indispensable. But it remained an aircraft built around old-world assumptions: that information would be scarce, that targets would be largely fixed, that little data would come from off-board sensors, and that the crew would be responsible for fusing together a picture of the battle space.

What followed over the next two decades was an uneven transformation, shaped by urgent operational needs and constrained by budget. Smart weapons, smart sensors, and even smarter crewmembers would come to define this period.

This chapter traces how the F-15E evolved from independent interdictor into the most networked, sensor-rich, and digitally mediated 4th Generation strike fighter in the US inventory. That's not hyperbole, either – it is a view widely supported by those from the F-14 Tomcat, F/A-18 Hornet and F-16 Fighting Falcon who had the good fortune to fly the F-15E during this period.

It is a story of incremental software builds, increasing computational power, and architectural expansion. It is also a story of stores that arrived before the jet was truly ready for them, and of stores the jet was ready for that were slow to arrive.

But as you read these pages, you will come to realise that above all else, it is the story of a community that was never allowed to rest and that learned to be "comfortable amidst complex chaos". You will also learn how the Air Force struggled to keep its "only long-range, deep interdiction fighter" relevant amidst a sea of competing priorities.

As you read these words, bear in mind, too, the aerospace industry context that the requirements planners, HQ Air Staff budgeters, and acquisition teams faced while this all took place.

By the early 2000s, the integration of the former McDonnell Douglas (McAir) all-star team into a hybridised Boeing team was complete. These McAir men and women had deftly brought the F-15 – the pinnacle of the 4th Generation fighter – into the inventories of both the USAF and international air forces. Now, the old guard found themselves intermingled with an unfamiliar and often combative peer and leadership team.

Since 1997, when Boeing had acquired McAir, the F-15's designers, engineers, and software developers had been slowly introduced to "the Boeing way". For many, it felt like indoctrination. Some were gradually replaced by Boeing draft picks. Others simply retired, satisfied with having witnessed both the genesis and

OPPOSITE

E227 (second from left), E231 (third) and E236 (right) on the ramp at RAF Lakenheath, England, in 2005. These were part of the final batch of ten F-15Es delivered to the US Air Force and were the most advanced Strike Eagles in the inventory at the time. (Author) the early transformation of the world's greatest air superiority fighter.

As if the loss of Eagle corporate knowledge were not enough of a blow, F-15E acquisition and in-service numbers generated an additional challenge. During this period, fewer than 230 Strike Eagles were in service with the United States Air Force (USAF), with as many as 48 of those earmarked for Replacement Training Units (RTUs).

With no prospect of adding more Strike Eagles to the ledger, the cost of any improvements or modifications to the aeroplane would be borne by a small fleet that was ageing and shrinking, to boot. This meant that Boeing's contributions would become inordinately expensive.

To further make matters worse, every dollar allocated to the F-15E would have to compete with similar development needs of the Lockheed Martin F-22 Raptor, and of the "new kid on the block" Lockheed Martin F-35 Joint Strike Fighter, the latter of which Lockheed, in a spectacular moment of hubris, had posited would end the need for the 4th Gen fighter altogether.

At the operator level, the scale of these costs became clear in 2003, when a Boeing technical representative, observing a Seymour Johnson crew debrief after flying a new software load, noted that Boeing would charge $20,000 for each line of code that needed to be changed.

But change was not optional: it was essential and it was all-consuming. Between 2003 and 2023, the F-15E did not merely *receive* upgrades. It *shape shifted* between wars because of them.

It developed to fight inside dense networks, to generate and share targeting data at speed, to survive in increasingly contested electromagnetic and surface-to-air threat environments, and to act as both shooter and sensor in campaigns that always demanded either persistence or precision, and often demanded both.

The Strike Eagle that emerged at the end of this period bore little resemblance – technically or tactically – to the jet that crossed the Saudi-Iraq border in March 2003.

The transformation was long, grinding, and driven by the relentless pressure to deliver over the course of twenty years at war.

This is how it happened.

2003–2007: Old Strike Eagle Meets New Wars

By the time the United States launched Operation Iraqi Freedom (OIF) in March 2003, the Strike Eagle was still, at heart, an analogue jet built for Cold War interdiction. It was fast and long-ranged, and it was lethal – but it soon found itself entering a conflict that increasingly demanded speed of information, dynamic targeting, and constant interaction with a broader network.

These demands were indubitably worsened by the overarching expectation that it would execute "precision strikes": expected by mission planners, expected by an American public growing weary of forever wars, and expected by a press corps eager to amplify the slightest hint of operational error.

Be Afraid of the Dark, The Strike Eagle Story, Part 1: 1978–2002 concluded with the final ten USAF Strike Eagles, ordered under a May 2001 contract worth $571 million, being prepared for delivery to the 48th Fighter Wing at RAF Lakenheath, England, in June 2002.

At that time, these airframes (E227 through E237, USAF serials 00-3000 through 00-3004, and 01-2000 through 01-2004) were the most advanced F-15Es in the fleet. They featured the new Advanced Display Core Processor (ADCP), replacing the twenty-year-old Multi-purpose Display Processor (MPDP); Fighter Data Link (FDL); the new Programmable Armament Control Set-45 (PACS-45); and an associated J-Series weapon data bus architecture that was compliant with MIL-STD-1760.

In the cockpit were active-matrix liquid-crystal, flat-panel colour displays; an Enhanced Engine Monitor Display; and all the functionality offered by the latest version of the Suite 4E+ operational flight program (OFP) software.

In the years since the F-15E had entered service, its avionics architecture had steadily moved away from the legacy H009 data bus (pronounced who-nine) – a proprietary McDonnell Douglas system inherited from the F-15A-D family – and consolidated instead around MIL-STD-1553B.

By the late 1990s, H009 no longer supported any operational mission function, its relevance fading as the Strike Eagle's weapons, sensors, and mission computers became entirely 1553-centric. The retirement of the AIM-7 Sparrow marked the disappearance of the final system that had meaningfully depended on H009.

The adoption of MIL-STD-1553B changed how avionics within the aircraft communicated, enabling a more scalable and upgradeable "central nervous system".

At the same time, a second standard – MIL-STD-1760 – defined the electrical and logical interface between the Strike Eagle and its externally carried stores: weapons, pods, and other payloads. 1760 used the 1553 bus as its underlying data path and became a prerequisite for the employment of modern smart weapons.

On these newest F-15Es, only two conformal fuel tank stations (LC/RC2) were wired for 1760 stores, supplementing the existing fuselage station 5 and wing stations 2 and 8, and therefore limiting full digital carriage to just five stations.

Notably, the structural and mechanical interfaces necessary to use wing stations 1 and 9 were in place, but these remained electrically and logically dormant in the USAF F-15E fleet, in contrast to their later activation in the F-15SA ordered by Saudi Arabia (see Chapter 5: International Operators and Chapter 6: Variant Briefings).

This constraint would ease only gradually as wiring modifications propagated across the fleet over the next two decades, but it is noteworthy that for only $50 million, the Air Force could have converted all fifteen stations across the entire inventory in one fell swoop. It was a decision, driven by budget, that would eventually prove a false economy.

Operational Flight Program (OFP) Suite 4E+ introduced the Strike Eagle's Link 16-based FDL to frontline squadrons and delivered early weapon management and software integration hooks for the Joint Direct Attack Munition (JDAM), Joint Stand-Off Weapon (JSOW), Wind Corrected Munitions Dispenser (WCMD), and Small Diameter Bomb (SDB) weapons that were to follow.

Thanks to a generic weapons page in the stores management system called Smart Unknown, the crew could also hand jam the GBU-38 500 lb JDAM into the Programmable Armament Control Set (PACS) as an interim measure until an upcoming OFP integrated the weapon natively.

Busses for Nerds

In modern fighters, the way systems talk to each other matters almost as much as what those systems actually do.

For the F-15E, two US military standards sat at the heart of that conversation: MIL-STD-1553 and MIL-STD-1760. They were often referenced interchangeably, but they existed for very different reasons and solved fundamentally different problems.

MIL-STD-1553 was the aircraft's central nervous system. Developed in the 1970s, it was created to standardise how avionics boxes communicated inside military platforms. Flight controls, navigation systems, sensors, and central computers all used 1553 to exchange commands and status information in a way that was predictable, reliable, and resilient to failure.

It was never designed for large data flows or high-speed sensor fusion. Instead, it existed to move commands, flags, and small packets of mission-critical information under the worst possible conditions – electromagnetic interference, vibration, temperature extremes, and combat damage.

It operated at a fixed data rate of 1 Mbps and was time-multiplexed and command-driven: a single bus controller dictated which system spoke and when. Nothing transmitted unless it was told to, which made the bus deterministic and immune to data collisions. Dual-redundant wiring added further protection, allowing the system to continue functioning even if one side failed. The result was a data bus that was extraordinarily robust.

In contrast, MIL-STD-1760 was not an internal avionics bus, but a weapons interface.

Introduced in the 1980s, 1760 was developed to standardise how an aircraft connected to the munitions attached to it. Where 1553 stayed inside the jet, 1760 carried data across the pylon and defined not only how data was exchanged, but also how power was delivered, how connectors were pinned, and how a weapon was physically and electrically integrated with the aircraft.

A single 1760 interface could provide electrical power, control signals, targeting data, weapon status, and built-in test feedback, all through a common, standardised connection.

And this is why 1760 was key to the modern Strike Eagle. A smart bomb or guided missile was not just a dumb store. It was a powered subsystem that had to be managed, monitored for health, configured for employment in real time, and sometimes updated right up until the moment of release.

1760 therefore carried both data and power, and supported higher data throughput, reflecting the demands of modern precision weapons.

Smart Unknown required the crew to enter the MIL-STD interface type, power interface and weight class, into the PACS for each store. It used existing JDAM logic and interface architectures, allowing rapid fielding of new weapons as long as the vendor designed the weapon communications logic to follow the JDAM ICD (interface control document).

It was generally used to get the aeroplane to recognise entirely new weapons that had not been fully integrated, but was also useful for new variants of existing weapons – the JDAM, JSOW and WCMD series in particular – that had also arrived between OFPs.

Even so, new weapons and deriviatives would still have to be properly and fully integrated into the aircraft's stores management system, a process that typically took years. That burden would not truly disappear until the arrival of the Universal Armament Interface (UAI – see below).

Smart Unknown had a major limitation: the crew would not receive a launch acceptable region (LAR: the area in which the weapon could be released and

Operational Flight Program: An Explainer

At the heart of the F-15E is the Operational Flight Program (OFP) software – the code that underpins every aircraft function, analogous to the operating system of a PC or Mac.

As modern combat aircraft became increasingly software-defined, the OFP provided the baseline code that enabled core functions while integrating an expanding array of sensors, weapons, and avionics. New software builds were developed alongside planned hardware upgrades, with the combined package released as a Suite. Software and hardware were inseparable, sharing the same development, test, and budget cycles, making each Suite a complex combination of engineering, testing, training, and funding that would repeatedly challenge the programme.

Between 2003 and 2023, F-15E OFP development continued steadily, introducing new capabilities while enabling the Strike Eagle's gradual digital transformation. The F-15 Systems Program Office (SPO) worked with Air Combat Command (ACC) to define requirements, fund software development and aircraft modifications, and oversee updates to the simulator fleet, ensuring crews could train on new capabilities before they reached operational squadrons.

During the same period, the Air Force significantly reorganised its test enterprise. Responsibility for developmental testing gradually shifted from Edwards AFB, California, to Eglin AFB, Florida, where the final Strike Eagle test aircraft, specialist equipment, and personnel eventually relocated. The 445th Flight Test Squadron handed responsibility to the 40th Flight Test Squadron, while Developmental and Operational Test activities became increasingly integrated through the combined efforts of the 40th and ACC's 85th Test and Evaluation Squadron. Operational Test and tactics development remained the responsibility of the 28th Test and Evaluation Squadron at Eglin and the 422d Test and Evaluation Squadron at Nellis AFB, Nevada.

Until Suite 9, the F-15C and F-15E each maintained separate OFP development paths. From Suite 9 onwards, both fleets shared a common software baseline, although subsequent revisions could still be type-specific, resulting in designations such as Suite 9E+ and Suite 9E+ Rev 1.

Separate OFP builds supported the Strike Eagle's nuclear mission, identified by the suffix "N" (for example, Suite 8.01N), while bespoke software variants were also developed for international Strike Eagle operators.

See Appendix A: OFP Suite Development and Chapter 6: Variant Briefings.

reach the target with the planned impact parameters) indication in the cockpit, so they would have to calculate the release point either on-the-go or during mission planning.

Sensor Fusion and Cueing: The Modern Strike Eagle's Roots

Suite 4E+ also continued the iterative but nascent implementation of both sensor fusion and of combinations of cueing weapons and sensors in the jet.

Sensor fusion was an overarching data sharing model that allowed sensors (including weapons that could act as sensors) to share their view of the battle space with the MPDP (often referred to simply as the central computer). The CC processed the data and presented it as additional symbology on the multi-purpose displays (MPDs).

It had been introduced to the Strike Eagle in Suite 3, but at that time data sharing was limited to the APG-70 radar and the AAQ-14 target pod. Even so, this early instantiation allowed the crew to look at the radar page and see half-intensity rectangles representing aircraft detected by the pod operating in air-to-air mode. The radar page also presented a dashed vertical line that showed the pod's line of sight, giving better situational awareness and allowing the crew to rapidly align a narrowed radar scan with the pod, if desired.

The ability to fuse data from both sensors on a single MPD page meant the crew could not only counter certain types of radar jamming, but also mitigate certain inherent limitations of the radar.

For example, if the radar was building a track file for a single contact, but the pod detected multiple aircraft in resolution cell formation – too close for the radar to make them out individually – then the CC would display a "multi-target" asterisk symbol on the radar page. As the name suggests, this indicated the presence of more than one contact at that location, and the symbol would replace the single full-intensity rectangle that denoted the track file being maintained by the radar.

It was sometimes not desirable to see this fused picture because it could be erroneous and misleading. To address this, Correlated mode (COR) acted as a filter that inhibited multi-target symbology, even when the CC considered it to be valid. COR was located at push button 1 (PB1) on the radar page and, when selected, was referred to as "boxed COR". In a later suite, COR would be renamed to the much more logical Declutter (DCL)

Boxed COR also removed low-fidelity track file symbology from the display, and this was most bothersome when the radar was in search mode. From Suite 2 onwards, the APG-70 created low fidelity

ABOVE

Chiefs and Lancer jets take part in an 'elephant walk' at Seymour Johnson AFB, North Carolina, in 2012. By that point in time, the Chiefs had been deployed to combat on and off since 1991. (USAF)

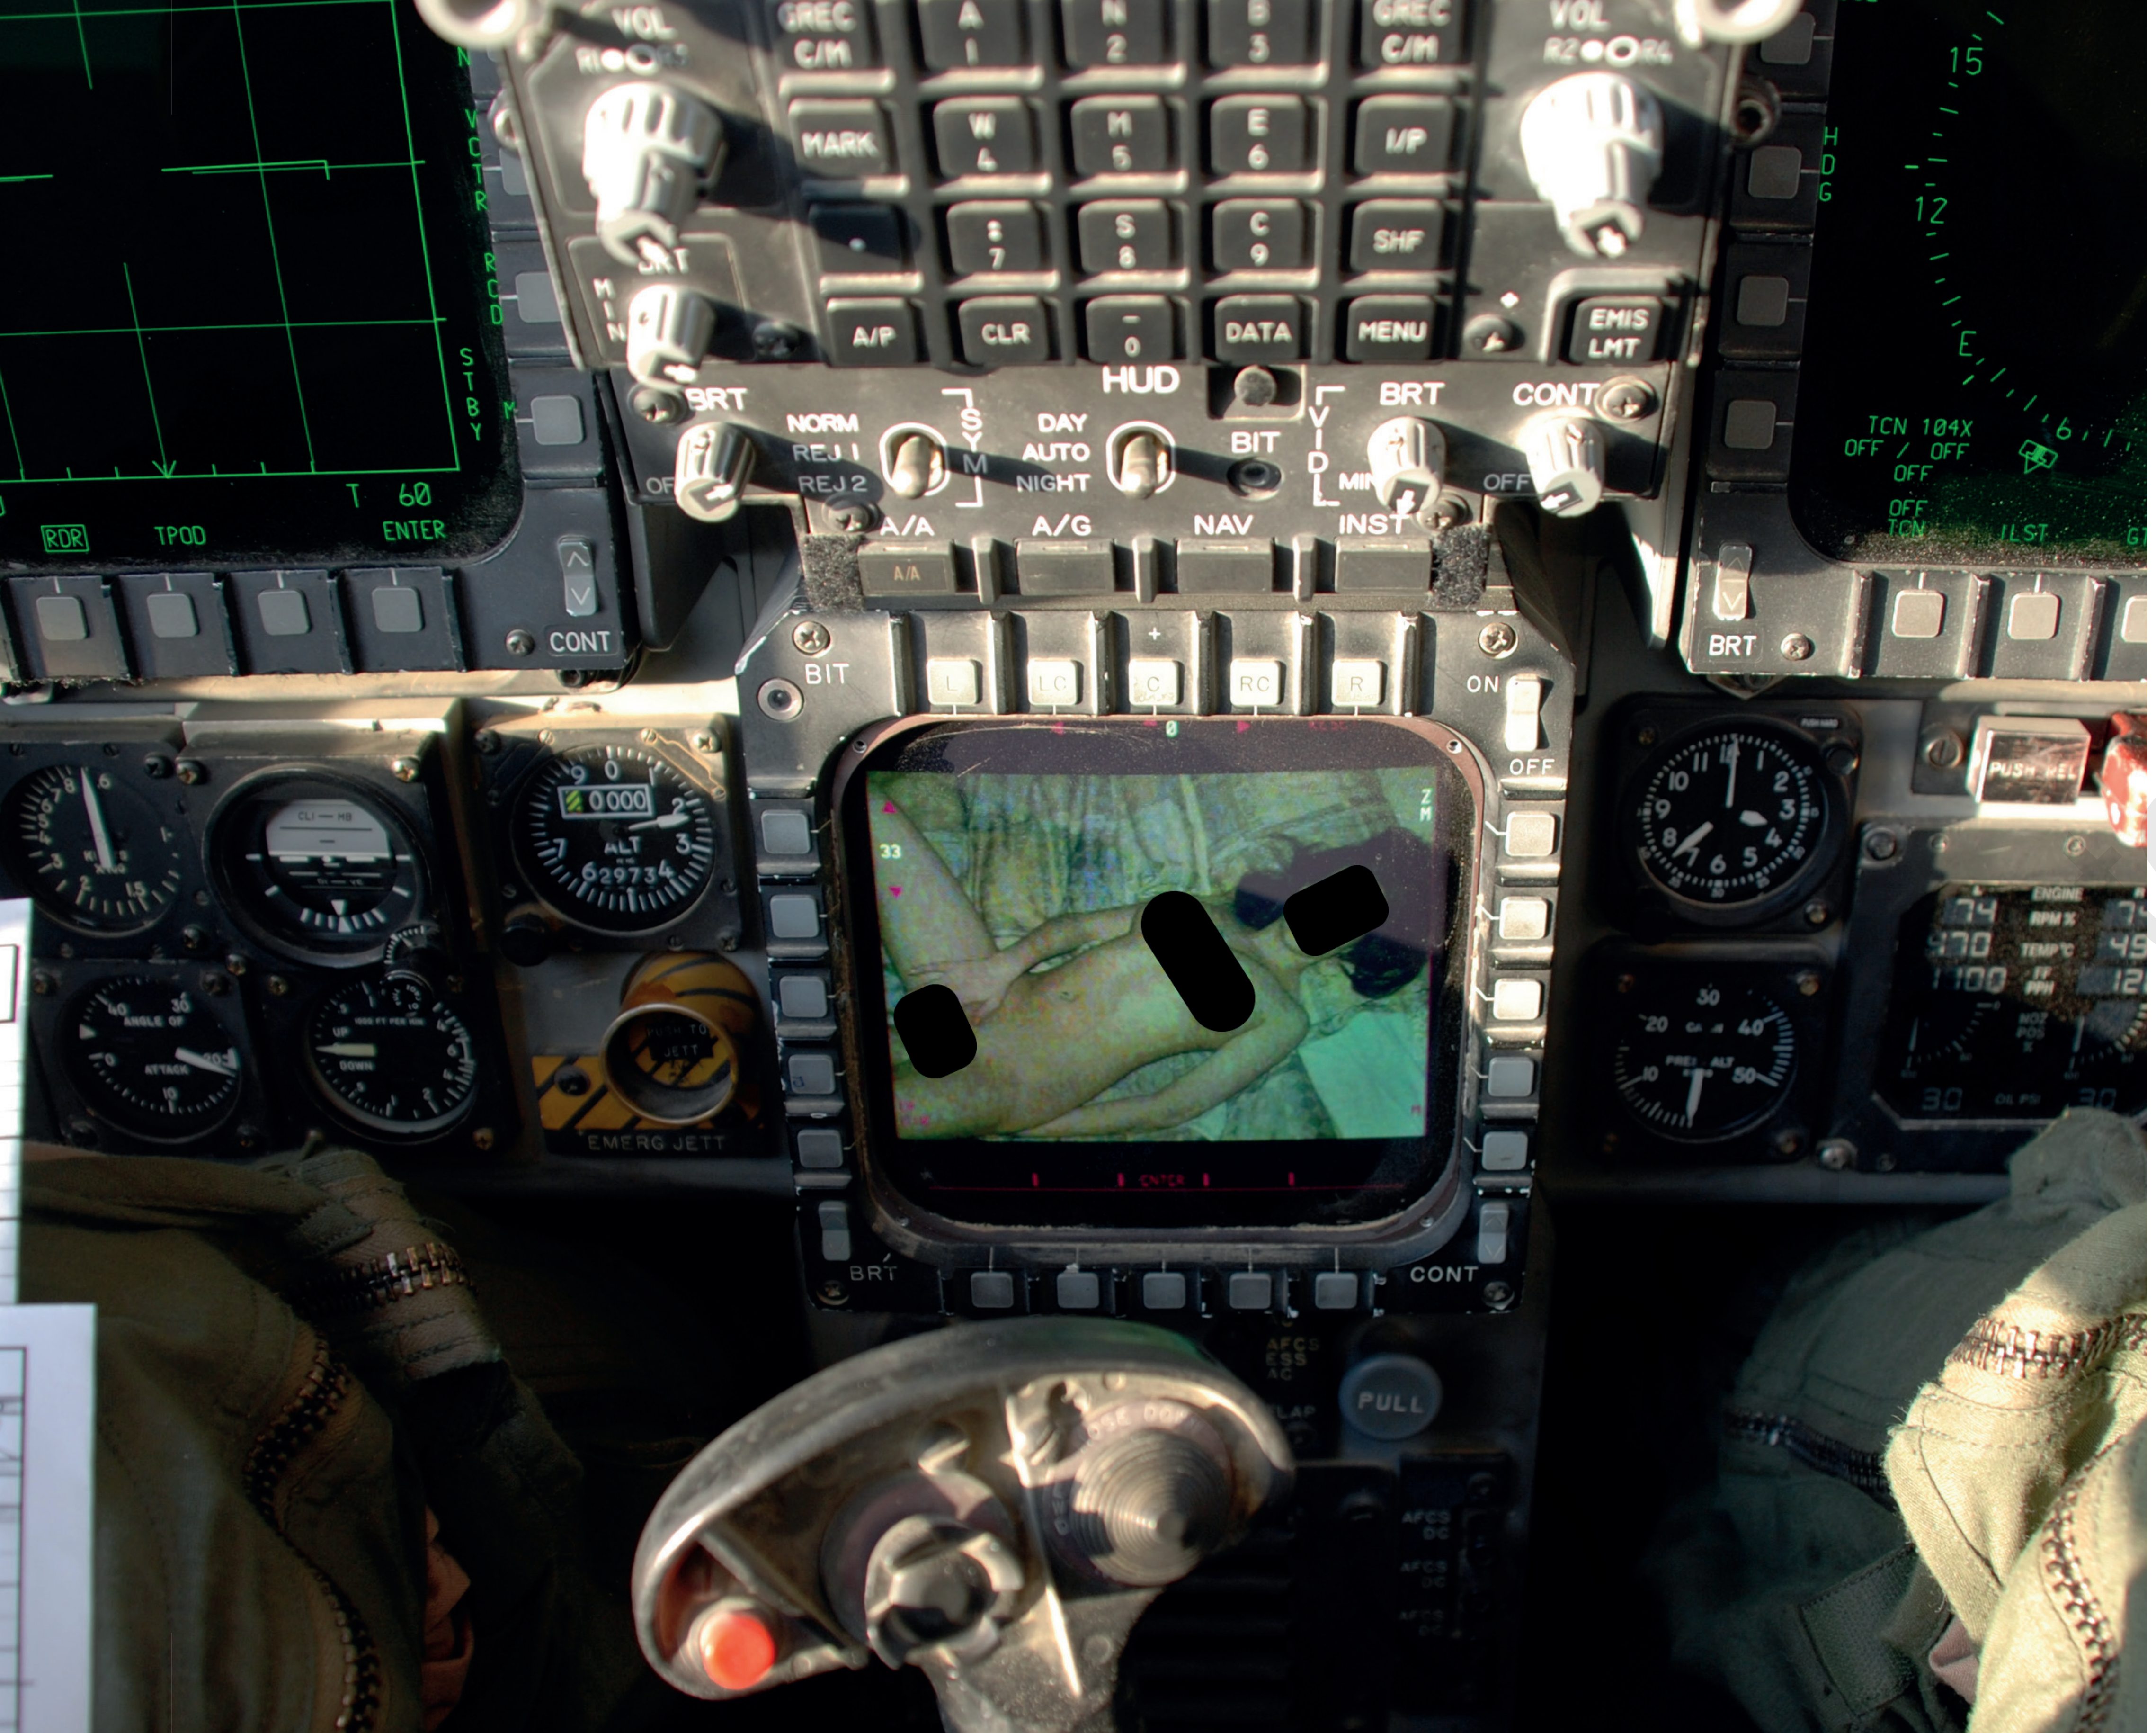

track files of targets in search mode – sometimes referred to as "background TWS" because the radar would spend some time processing the Doppler data from each return. On the radar page, the progression from low- to high-fidelity track file was characterised by a rectangular track file symbol that would mature after a few sweeps from an initial "filled in container" to one with a "short donkey dick", to one with a "long donkey dick", aligned to the heading of the hit. Such symbols could overwhelm the display, so boxed COR inhibited them and only established track files were displayed.

Suite 4E+ enhanced the data-fused landscape by adding FDL to the list of data sharing sensors. When this happened, COR really came into its own because when the crew were handed a target over data link, its position was often at odds with where ownship sensors assessed it to be. When this happened, erroneous multi-target symbols were generated.

Cross-cueing and cross-targeting were highly desirable by-products of this sensor fusion capability. While fused data created situational awareness, the ability to slave – or cue – one sensor to another (cross cue), or for any one of several sensors to cue a weapon (cross target), meant that it gave the crew much more than just enhanced situational awareness: they could now employ weapons directly off it.

Cueing methods had matured over time, driven in part by the development of the target pod's air-to-air functionality. In Suite 1, the crew had been able to slave the AAQ-14 to the radar line of sight and command an IR track to monitor the target, but the pod could not be used to generate a track file. By Suite 2, a quasi air-to-air mode had been introduced to the pod, which now used a pixel tracker, similar to Point Track in air-to-ground mode, that would generate azimuth-elevation rate data from the point tracking error-rate signals. With the radar in search, slaving it to the target pod would command a periodic short, low-power burst ranging waveform down the pod line of sight, thereby producing range and range-rate Doppler measurements that the radar could use to generate a track file. But the data this mode produced would become stale relatively quickly and paled in comparison to what was produced by a single target track lock of the kind necessary to employ an AIM-7 Sparrow missile in a guided mode (the Sparrow could also be employed in FLOOD mode, which used a 16°x40° radar beam and did not require a lock).

Suite 3 had taken advantage of continued progress in the development of IR detection algorithms and had finally introduced a truly robust air-to-air mode, but it had some important limitations. It was usually useful only in clear weather, and mainly looking up, because the target point tracker tended to "grab clouds" and other pixel-dense background objects. Azimuth rate data was also poor when the target was near head-on. Even so, it was a powerful utility: prior to this point, correlating the presence of multiple targets within a single radar hit had necessitated voice communication between formation members, further complicating an already jam-packed radio choreography. For example, an initial call identifying a relevant radar contact would start with "Hammer 1, group, bull 270 at 40". It could be followed by a wingman spotting a pair of aircraft in the targeting pod cued to that area, who then called "Hammer 2 has eyeball, two there."

Suite 4E+ now consolidated the F-15E's cross-cueing capabilities, most notably by allowing a slaved acquisition: the WSO could run the pod in the air-to-air mode, command an IR track on a target, cue the radar to it, then finally enter single target track (STT) with the radar so that an AIM-7 could be employed. This cross cueing technique also worked with the radar's dual target track (DTT), where the radar independently built the track file for one target and was cued by the pod for the other, but this mode was deemed unreliable.

Careful management of COR was required as the crew went through the cross cueing process. In the cueing transition from target pod to STT, boxed COR could prevent some misalignment between pod and radar. But if the target group was in or near the resolution cell, COR would filter out the multi-target cue, and that information could be useful if the formation was trying to "sort" the group (to divvy up who was going to target which bandit in the group). One technique for a side-by-side sort with COR unboxed, and with the multi-target cue therefore present, was for the lead crew to command STT just to the left of the asterisk denoting the radar group or hit (assuming lead was on the left side of the formation) and swing the APG-70's radar beam into the group from left to right, while the number two crew did the same with their radar, but from right-to-left. This

OPPOSITE

The digital map replaced the original film cartridge system and was introduced with Suite 4E+. The system could store reference photos for inflight access, and while these were typically of the target, it was a longstanding tradition for mission planners to sneakily insert the occasional nude photo into such storage systems. (Author's collection) offered no guarantee of a successful left-right sort, but it was a good start.

With the data fused and sensors cross cueing one another to the point that the crew could employ weapons, the final piece of the puzzle was cross-targeting. In practical terms in the pre-AIM-120 AMRAAM age, this meant slaving the radar dish to the pod and then using the pod to guide a Sparrow all the way to impact. Once the fully-active AIM-120 arrived, the possibilities expanded dramatically.

To cross-target an AIM-7, the crew would put the radar into a search mode, command a visual track of the target in the pod page, then cue the radar antenna to the pod line of sight by boxing SL (PB5) on the radar page. The radar would transmit a narrow beam of high PRF energy, illuminating the target and guiding the Sparrow for as long as the pod continued to track the target. The same technique could be accomplished against a jamming target. In both cases, range to the target was estimated by the central computer until it fell within the 13.2 nm reach of the AAQ-14's 100 millijoule tactical laser, at which time the true distance to target could be confirmed.

Early cueing methods were process intensive and required extensive manual input to execute. But in the years ahead, slaving would improve to the point where it could be configured to happen automatically. In its ultimate incarnation, just one or two HOTAS combinations was all that was required to employ an active radar-guided air-to-air missile, supported by the radar that was itself being cued by the RWR, which was passively sensing the emissions of an adversary air intercept radar. This level of autonomy was still many years off, though

ABOVE

The APG-70, Suite 4E+, in search mode. Note the COR and SL options at PB1 and PB5, respectively. (Author)

Automation in general became increasingly common over time. For example, the transfer of NCTR (non-cooperative target recognition) data from the radar to the AIM-120 AMRAAM had started off as a manual process, then became a semi-automatic process, and ended up as an option that was boxed before before the crew even taxied. The NCTR classification was used to determine when the AMRAAM should go active with its own radar (based on the estimated or known radar cross section of the target).

It doesn't take much imagination to see how cross cueing and cross targeting could be extremely useful for operating at a higher level of proficiency all round, and for operating passively with limited electronic emissions. These technical solutions also offered options against adversaries for whom electronic attack (EA) was a key part of tactical employment.

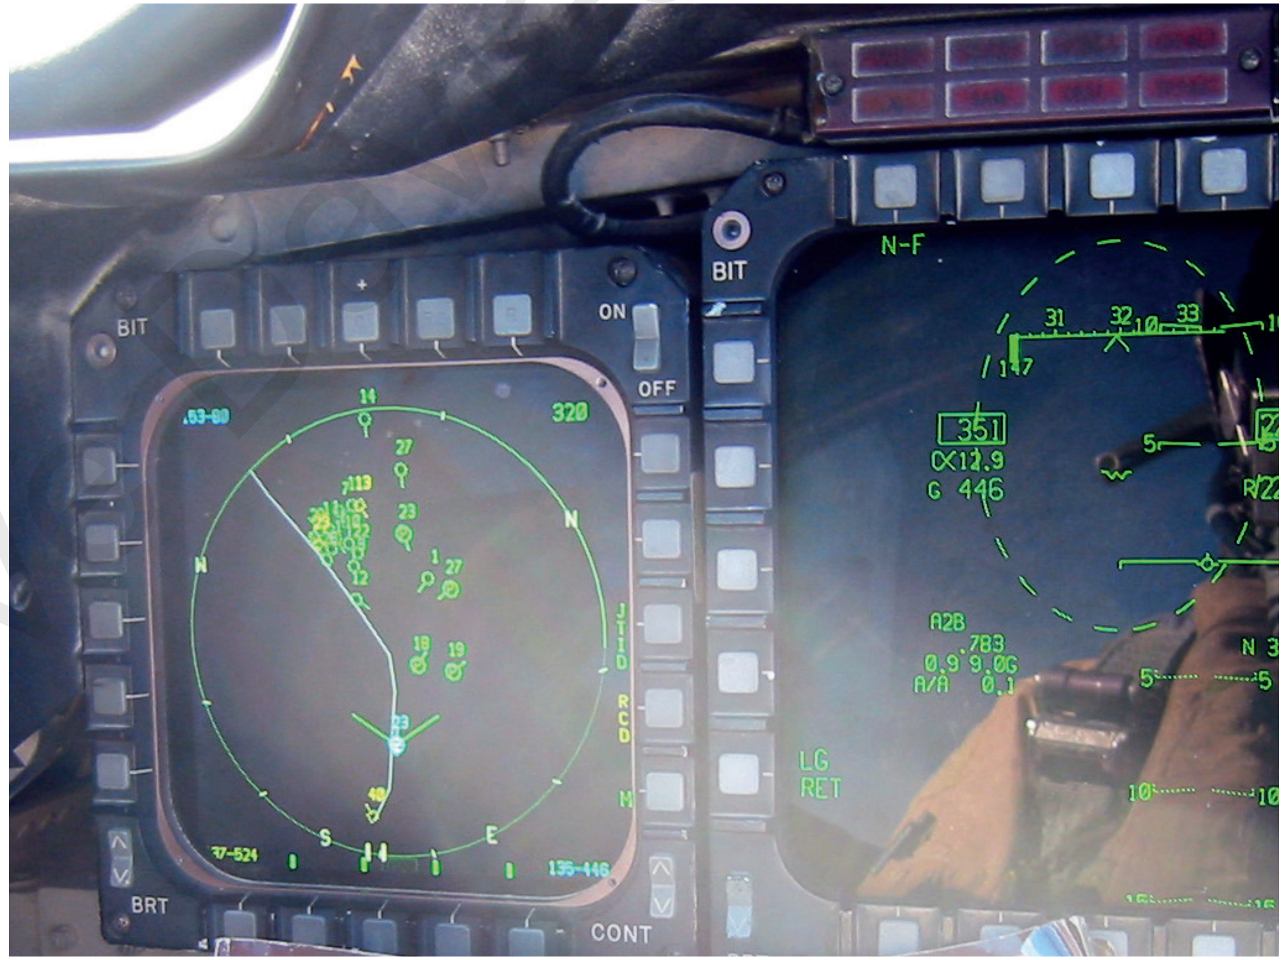

ABOVE

An early instantiation of the SIT display (left) set to 320 nm range scale, showing network contacts on the Link 16 data link marshalling in the distance. The HUD repeater is on the right. (Still taken from an unofficial end of deployment video)

But they were not yet mature enough to force doctrinal shifts in how the Strike Eagle was employed in the early 2000s (although FDL would change this, as we will see). There were several reasons for this.

First, at this stage, sensor fusion was still fragile. The CC could associate tracks and cross cue sensors, but without the processing margin and message discipline that would later be delivered by a new processor and data architecture (ADCP and MEDB, see below), these behaviours were workload-dependent and not always stable under heavy data link or sensor loads.

Second, cross cueing and cross targeting naturally lent themselves to passive engagement tactics and were favoured when the crew wanted to avoid STT, which would likely trigger the adversary radar warning receiver. But *at this time*, the United States still enjoyed a healthy shot advantage over Russian and Chinese air-to-air missiles: the AIM-120A/B could, Western intelligence agencies believed, outrange everything that threat nations had to offer (this too would change, as we will see). With this decisive advantage in place, and recognising that the Strike Eagle had a huge radar cross section, most tactics were built not around low probability of intercept (LPI) techniques, but around active radar use and intimidating the threat by *deliberately* lighting up their RWR. Using the APG-70's SNIFF mode (receive-only), its LPI mode (which scanned a very narrow azimuth and did so very slowly), or the Emissions Limit switch (which ceased all outbound transmissions), were seen as purely defensive options and were thus rarely if ever employed.

Thirdly, while it was known that adversary nations fielded passive, anti-radiation (ARM) versions of their missiles that could home in on the APG-70's emitted radiation, these were not considered an issue: the community was confident that they would be pitching out from their AMRAAM shots before an ARM became a real threat. And if it did become a concern, when their own missile timed out, the pilot went to the notch while the WSO boxed SNIFF mode to trash any ARM shots.

Finally, the very early passive capabilities of the pod-radar combination were just not mature enough. Burst ranging with the radar to create a track file from the pod might have offered an LPI way to engage an adversary, but it was also considered a low probability of kill (Pk) shot. And the tracking algorithm of the pod could sometimes grab other objects, as already described.