EARLY EVOLUTION AND OPERATIONAL REQUIREMENTS OF AEW AIRCRAFT

Original text by **Mike Hirst**, adapted by **Dimitris Dranidis**

Airborne early-warning aircraft became necessary almost as soon as radar was put into service. The reasons why are evident in simple physics. Radar sets would not, and still cannot, see round corners. Because radar waves travel in approximately straight lines, they are limited by the horizon and cannot see into airspace hidden by the curvature of the earth; therefore the lower an aircraft flies, the closer it can get to a radar site before it is detectable. Piston-engined heavy bombers of the type used in the early 1940s were visible almost an hour before they ever reached the outposts of a defender's territory. Even the high-speed, high-altitude bombers that came along later were visible as much as a half-hour before they reached the radars.

When aircraft take cover, by flying lower (but not NOE), nature plays into the hands of the radar user. Halving the altitude at which an aircraft penetrates only means a detection range reduction of barely 25 per cent. One message was clear, even before the final phases of the Second World War: to avoid radar detection, aircraft have to come in literally on the deck. A few figures illustrate this: A radar station located at sea level cannot detect an aircraft approaching at 500 ft altitude until it is within 28 nm, and if the target approaches at 200 ft it gets closer than 18 nm before its presence is evident.

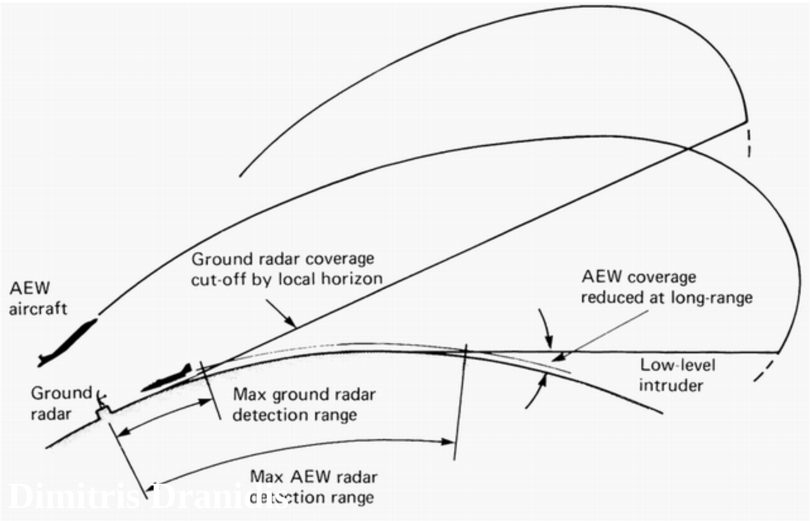

Operations in this band of altitude were already used across the English Channel, by both sides, after the clear advantages offered by radar had become evident around 1942. The targets more likely to face a low-level onslaught in those early days were coastal defences, including radar sites themselves. Over water, the attackers could fly flat out without fear of hitting obstructions, and the only way that defenders could claw back some detection range was by raising the radar above sea level. When a radar antenna is raised by just 50 ft, an extra 9 nm of detecting range is possible, and an aircraft approaching at 200 ft altitude will be detectable at up to 27nm range. This figure is a theoretical maximum: In general, radar screens will be cluttered, so targets will appear only intermittently – but on the other hand, factors such as surface ducting may actually increase the practical detection range beyond the theoretical numbers. With modern aircraft travelling 4 nm or more between successive radar scans *(assuming mechanically scanned radars; phased-array radars can re-scan an air sector almost instantaneously)* it is inevitable that aircraft will get much closer than the theoretical maximum detection range before they are evident to an operator.

Inland targets are more difficult propositions to analyse. First, attacking aircraft tend to fly slightly higher, between 200 and 500 ft typically, to clear terrain and other obstructions, and a wise crew will also make use of terrain masking by flying in the shadow regions on the side furthest from the radar.

The defender's best plan has always been to locate radars as close as practically possible to coastlines and borders, and at strategic points within one’s own borders. Even so, it has been painfully obvious, ever since plans to provide such methods of defence were first drawn up in the 1940s, that the number of radar stations required to protect an area of high-value targets means it is a costly business.

In the 1940s it was also time for sea commanders to question the vulnerability of their fleets at sea. If a radar set mounted atop a 50 ft mast could not see an attacking aircraft until it was within 20 nm, ship defence wasn't going to be an easy proposition. The world's leading navies began to question whether they could have a ship- and aircraft-detection system offering much better performance than the visual maritime-reconnaissance aircraft types then in use. Both navies and air forces were drawing the obvious conclusion that the best way of seeing low-level attackers was to get the radar sensors as high as possible. Airborne-early warning (AEW) was a buzzword even before 1945.

The logic adopted was concerned solely with the threat posed by attacking aircraft, and this has become the central issue in all AEW operational requirements. Radar has remained the primary sensor in all AEW aircraft for over 50 years now, and its main task has always been to detect aircraft at all altitudes, and at as long a range as possible[1]. Other capabilities have been incorporated as they have descended naturally from the primary requirements, an example being the ability to obtain highly accurate bearings on escort or stand-off jammers which are beyond the line-of-sight of the equipment they disrupt on the ground.

The detection of low-flying aircraft and surface objects is the primary role of all AEW development, but several generations of radar development have been necessary to reach the full specification.

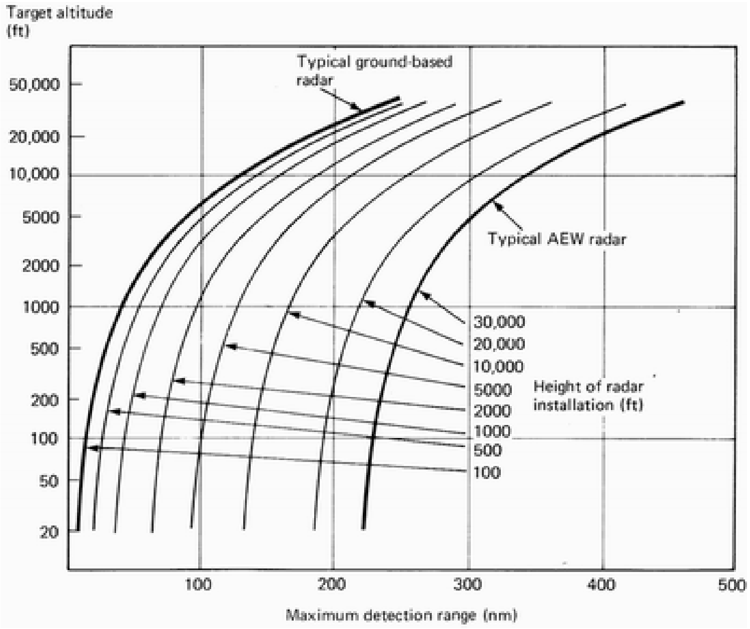

Detection range curves for various target altitudes depending on the height of the radar installation. Notice the vast superiority of AEW systems in low-flying target detection.

In the immediate post-war period, airborne radar sets were rudimentary. They usually operated at wavelengths between 3 to 10 cm, and such short wavelength radiation could be directed quite adequately using a parabolic dish, although high sidelobe losses were inevitable. At the time, only low pulse repetition frequency (PRF) sets were used, and when such radars were flown at high-altitude, around 18,000 ft, it soon became evident that if the radar beam was directed such that it did not reflect any energy off the ground, it was able to detect aircraft at and above the same altitude, and possibly slightly below. In this guise systems were used in the first night-fighters. It must be pointed out that earlier interceptor aircraft had to be vectored by ground-based radars, even if they had simple airborne radars, and they were nowhere near as successful as future generations of autonomous night-fighter types.

Pointing the radar down produced a particularly important effect. Reflections from the earth's surface caused a big target to appear at a range equivalent to the cruising height, and unfortunately, due to the relatively large amount of energy being lost in antenna sidelobes, this would sometimes happen even when the beam was directed immediately ahead.

When radiation is reflected from land, it is reflected back randomly by grass, hedges, trees, buildings, and so on. Reflections from flat surfaces, such as water, tend to be directed away from the radar, however. On a conventional radar scope (called a plan position indicator, or PPI) the result of scanning downwards at close range is a map of the earth's surface showing strong returns from the land, and weak returns from water. Even a simple airborne radar can pinpoint a coastline, river courses and lakes. During hostilities in the Second World War this attribute was used to good effect, providing navigation fixes from above cloud, or at night. As ancient considerations have located most cities and major industrial plants next to rivers, blind-attack target-fixing was possible too. Radars which can be used this way are categorized as 'ground-mapping' types, and it is a method of operation available on almost all airborne radars today, including the elementary weather-detection radars used in airliners and light aircraft.

---

[1] *The abilities to detect surface vessels and, more lately, land targets, while valid operational requirements in their own right, are less defensively oriented and more geared towards providing a complete tactical picture for both offensive and defensive purposes. For anti-air AEW systems, the most significant impetus was the need to counter the threat of sudden air strikes ie. a largely defensive stature.*

Unfortunately, simple radars cannot detect details in the sense that one might expect. Those unwholesome sidelobes, and effects similar to the diffraction scatter which can be experienced with light waves, make detailed resolution virtually impossible. Radar is rather like an eye watching a scene through frosted glass, in that it can detect large-scale features, observe movement, and gain an overall impression of the scene that it is surveying, but the actual details of the scene are near invisible. A consequence in the immediate post-war period was that overland AEW was out of the question; radars could produce a rough map, but objects such as low-flying aircraft could not be picked-out from the background. Low-flying aircraft detection was possible only over water, and only under favorable conditions.

Therefore, the enormous costs of constructing and maintaining vast chains of radar stations to defend land masses was accepted as inevitable. This produced several important defence systems. The NATO Air Defence Ground Environment (NADGE) is a still-operative chain of radar stations along NATO’s border with Warsaw Pact countries, and the American DEW-line stations (gradually replaced by the North Warning System from the mid-80s) were just one element of the immense North American Air Defence (NORAD) system set up by the US and Canada. The British-developed UKADGE (United Kingdom Air Defence Ground Environment) is a much more capable version of its original Chain Home system. Some of the capabilities & limitations of a ground-based air defence system will be analyzed in more detail a little later – meanwhile, attention should be devoted to overwater airborne-radar operations, where almost all early AEW operational experience was concentrated.

Over a flat area of sea an airborne-radar, with a fairly large antenna (an important precondition for mechanical-scan antennas to get long-range target-detection) has a reasonable chance of seeing ships, and so it was in the maritime-surveillance role that the first effective airborne-early warning (AEW) systems were introduced. Radar-equipped aircraft, flown from shore bases or off ships, and patrolling at high-altitude, promised to provide long-range detection of shipping, and so forewarn the captains of vessels operating nearby. Several aircraft types were pressed into service in the radar maritime-surveillance role in the late 1940s, and they were categorized as AEW types.

Coherent MTI

In the decade and a half after the Second World War AEW capability improvements were very limited, because radars’ poor performance when looking down at the earth. In the late-1950s some military commanders, realizing that low-level aircraft were virtually invisible even to AEW radars, regarded them as the most persistent threat, and convinced the industry to design aircraft specifically for such operations. It is intriguing to look back now and to realise that while the world's air forces still stuck to the development of long-range high-altitude strategic bombers, ignoring all commonsense about radar-detection capability and the missile-design improvements taking place, it was the naval arms that started the brunt of developmental work on high-speed, low-level attack bombers. In the early 1950s for instance, a project sponsored by the British Admiralty matured as the Blackburn Buccaneer and the US Navy wrote the specification for the Grumman Intruder. Both aircraft were designed to operate from aircraft carriers, and their specifications called for them to speed along, barely above wave crests, and attack targets at ranges of several hundreds of miles. As soon as these two significant bombers began to show their mettle, they reinforced the need for an AEW sensor capable of detecting low-flying aircraft at long ranges. A lot of interims have to be reported first, but it is surely significant that it was in the same two companies that developed those first-generation low-level strike bombers, (Blackburn at Brough in the UK, and Grumman at Bethpage in America) that the most enduring steps in AEW aircraft development were to take place in the ensuing decade.

AEW airframes were almost incidental. So long as aircraft were available which could carry the new types of radar-clutter rejection equipment becoming available in the 1950s, they were pressed into service and became the first truly effective AEW types. Exactly what these new clutter rejection systems, called moving-target indicators (MTI) involved, is an intriguing tale of engineering ingenuity.

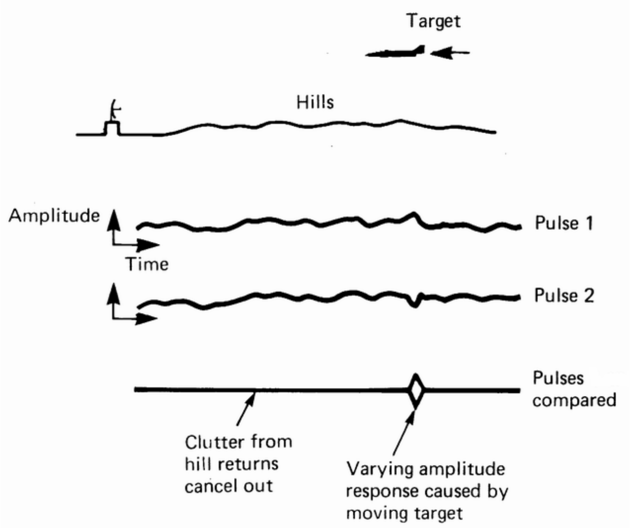

The operating principle of coherent MTI

Early enthusiasm about radar operations from ground sites had been overly optimistic, the enthusiasm masking several fundamental drawbacks. One was to do with the effects of terrain, which would poke into the base of radar beams and clutter radar screens with permanent echoes. To nullify its influence radars had been installed on mountain tops, or any other prominent and high location.

Solitary mountains however are rare. Geological origins more often than not dictate that, when mountains well up from the plains and oceans, they form formidable ranges and archipelogos. To the radar operators, those mountains made their job problematical. They would inevitably clutter the radar screens, so what was scenic to the eye was more than likely a nuisance to the radar.

When no high ground was available for a radar site, the result was disastrous, as antenna side-lobes produced reflections from the surrounding countryside out to 30 miles or beyond. To add a further injury, most early surveillance radars used 10 cm wavelength radiation, and these can detect weather almost as readily as aircraft. Consequently, on cloudy days, the radar screens could be totally cluttered by permanent echoes from mountain sides and reflections from rain. This made air-defence duties and air-traffic control, the prime roles of ground-based radar, nearly impossible.

Radar engineers reasoned that, if the information returning from two successive radar pulses included returns relating to approximately the same clutter region, it would be of virtually identical amplitude at any given range. Anything moving however would produce a varying amplitude return. If the amplitude pattern from two successive scans could be compared, or even better, if several successive scan amplitude patterns could be compared, 'clutter' returns would remain the same, but moving targets would show amplitude variations. The concept is simple enough, and this is indeed the basis of most moving-target indicator (MTI) systems.

At the time of MTI's development, methods of storing electrical signals, even for the 1/500th of a second or so between pulses, were woefully inadequate. It has to be kept in mind that in 1/500th of a second, radiation in space will travel out 300 km, and back again. In a wire it would cover a barely less significant distance, but even a resounding signal sent down a 600 km length of wire will emerge hopelessly weak, and the physical task of getting such a long wire into an electronics cabinet defies comprehension. An analogous but more elegant solution had to be found. It came in the form of a 'delay line', in which the output is virtually identical to the input, but delayed by the interval between successive radar pulses.

No energy is stored, but the equipment required to delay a signal by one PRF cycle is quite substantial, and a 'delay line' apparatus, as first produced in the early 1950s, was not the sort of thing to put in anything less than a furniture van. To all practical intents however, the method was satisfactory. It is widely used in ground-based radars today, and called coherent-MTI. One drawback of its use is that if any moving target happens to cover a distance equivalent to an exact number of wavelengths between successive radar illuminations, the amplitude of the returning signal will remain constant, so the MTI system does not recognize its movement, and it is cancelled. This happens at regularly-spaced speeds, a typical coherent-MTI radar being blind to aircraft travelling at 80, 160, 240, 320 kt, etc. This is called blind-velocity fade, and although an attacker might think he could bluff his way through by regulating his groundspeed to equal the known MTI blind-velocity, it is a virtually impossible task. An associated drawback of MTI is tangential-fade. In this case as a target moves abeam the radar, for a short period it does not move in the direction of expansion of the wave-front. Successive returns appear to have identical amplitude and are cancelled, causing the target to fade momentarily as it passes abeam the radar (this is an oft-used tactic in A2A combat and missile avoidance). Neither of these problems has been so severe as to prevent MTI being welcome wherever it has been installed, being as it is a tremendous improvement over what was hitherto available.

Non-coherent MTI

Coherent MTI cannot distinguish between slow-moving, large, targets - such as clouds - and an aircraft flying at high speed over or under the cloud. A more complex form of MTI was developed therefore: Several successive pulses are compared, and moving targets appear as 'noisy' portions. If the large moving object, which is causing this noisiness, has another object moving above or below it, the latter will cause a 'beat' frequency to be superimposed. This type of MTI is used in some AEW radars, where it is called airborne-MTI (AMTI).

Unfortunately, coherent-MTI is really too simple. A second form of MTI arrived, called non-coherent MTI. Instead of using amplitude variations exactly as received, non-coherent MTI compares amplitudes from successive sets of returning energy and determines 'beat' frequencies set up by fast and slow-moving targets, even if they are located at the same position relative to the radar. This will seem unnecessarily complex perhaps, but at least non-coherent MTI radar can identify blankets of moving weather, or areas of land, and aircraft above them. It cancels the weather returns, leaving on the screen targets that are moving relative to the earth and weather, i.e. aircraft.

Non-coherent MTI has to memorize where it can expect to find beat frequencies, and when these conditions are not present it uses coherent-MTI techniques to arbitrarily determine clutter. The requirement to memorise the conditions relevant to various areas of the radar coverage brought radar engineers into contact with data-storage techniques for the first time. Whereas digital computers with banks of data-storage are commonplace now, they were very rarely used for online applications in the early 1950s. This memory-only electronics application, carried out using analogue equipment only initially, was a significant step ahead for radar engineers. The AEW radar fraternity watched non-coherent MTI developments in ground-radar installations rather jealously.

When a radar is installed into an aircraft that is moving relative to the ground, it sees everything when it looks-down, the ground included, as a moving target. To add further to the AEW radar designer's dilemma, the world's oceans are continuously moving too, either driven by tides or pushed along by the wind. These effects cause waves to pile-up and reflect energy back when an airborne antenna looks along the direction of the wind or tides. The surface is textured, like land, and moving. Unless a non-coherent MTI radar is used there is little chance of detecting even a large ship, except at extremely long ranges when the vessel is virtually silhouetted, or at short ranges, dangerously close, where the sea texture can be differentiated from the ship.

Non-coherent MTI techniques therefore produced the first systems which got anywhere near meeting contemporary AEW requirements, and it is a great testimony to the skill of the engineers who made this important step forward, that the technical metamorphosis, from a furniture-van load of ground-based equipment to a carrier-borne aircraft load of black boxes, was made in the early 1960s. The system was called Airborne MTI, or just AMTI. It was developed first in the US, and the US Navy had a special aircraft designed to carry the system, the E-2 Hawkeye. Crude as the original version was by modern standards, it has been refined over the years and is now the most numerous AEW system in the world, with a significant export success to boot.

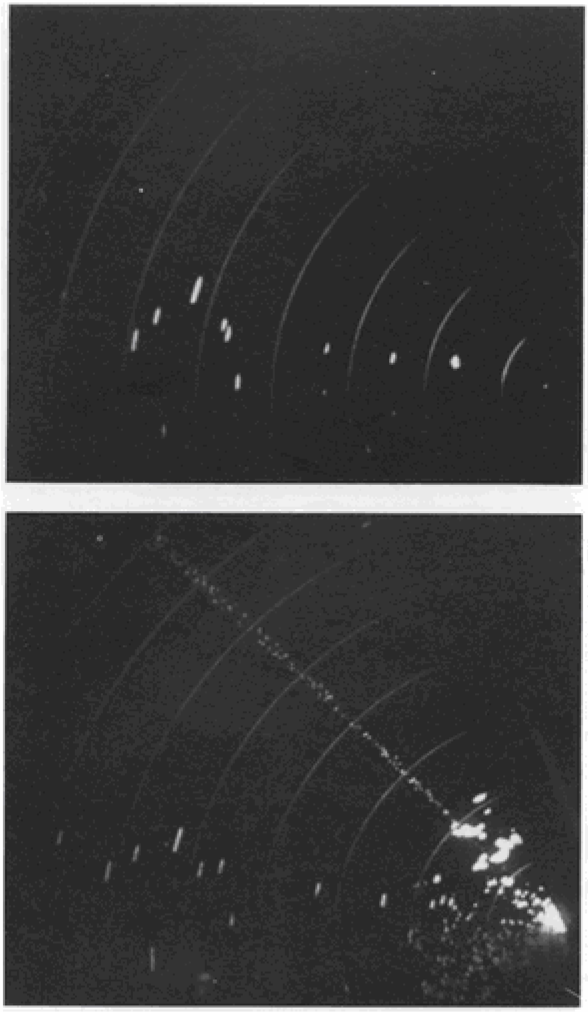

These photographs of a radar screen were taken a few seconds apart. The upper picture, taken when coherent-MTI was in use, shows aircraft clearly while the lower picture, taken when MTI was not selected, shows aircraft close to the station masked by clutter

Britain used earlier US-built long-range radars in its maritime AEW types, the Shackleton and Gannet, but later added an indigenously developed AMTI, so non-coherent MTI technology was the lynch-pin of early AEW capability. The US and Britain were alone worldwide, such was the tremendous technical challenge involved in producing AEW equipment to these standards.

There was still a hole in the best defences however. Even these sophisticated early AEW systems could not reliably detect low-flying aircraft, especially at long-range. By 1961 this was bad news. Aircraft such as the Buccaneer and A-6 Intruder were already close to production, and projects such as TSR-2 and F-111 were in prospect. It was obvious that Soviet designers must be tackling similar projects, and that Western defences had to find an antidote. The quest for a radar technique able to give good 'look-down' capability, over either water or land, was being reinforced daily.

Operational studies conducted in the early 1960s showed that an AEW aircraft was needed which would cruise at 20,000 ft or higher, so that it could see low-flying targets across a wide front at 160 nm range or more. If targets had not been detected around this range, and defensive activities got underway before the target had penetrated a further 60 nm, there was a chance that too many aircraft would get through to the defender's ground targets, or even to the AEW aircraft themselves. In the intervening decades these figures have scarcely been reduced, the advent of stand-off air-to-surface weapons having added at least a few more tens of miles to the advantage of the attacker.

To all the electronic wizards in the late 1960s, one trend was unmistakable, and it gave them reason to be hopeful of a new era in AEW radar development: Analogue electronics were going to give way to digital systems. The first integrated circuits, although almost trivial by today's standards, were about to enter production, and the promise of electronic reprogrammable logic in general-purpose computers was being viewed optimistically. At first, to engineers outside the inner clique of electronics experts, it seemed impossible to believe that the scale of what was being predicted would ever happen. Typical general-purpose digital computers in the early 60s had magnetic-core stores, which were relatively large and heavy, and the processing sections ambled through calculations at the rate of 10,000 or so additions or subtractions per second (0,01 MIPS). Even a computer with a 4096-word (4KB) memory still required a small office, a lot of air-conditioning, and suffered from mechanical and electronic headaches frequently enough to keep armies of technicians in business. They were also very expensive.

Within a little more than a decade the electronics engineers revealed the fruits of their optimism. They could provide as much as a million-words (1MB) of storage with a processor able to add or subtract in less than a microsecond (1 MIPS), and all this was in units barely bigger than a family suitcase. Equally as dramatic was a phenomenal improvement in failure rates, leading to a great fall in costs. The day of small general-purpose computers was suddenly a fact to live with, and with its arrival came the chance, if all other supporting requirements were met, to provide a radar with 'look-down' capability, in any weather, and over any terrain.

The pulse-Doppler radars, nowadays almost standard items for high-performance military aircraft, were what eventually emerged from this. Radars could now track low-flying aircraft, almost anywhere within their field-of-view, and irrespective of weather or terrain. Compared to what had been used before, these radars set remarkable standards.

Operational requirements and cost-effectiveness vs. ground-based systems

AEW aircraft are undeniably expensive to buy, maintain and operate. In leading aviation research establishments throughout the world, numbers have been put into trade-off studies which have attempted to show how ground-based radar networks and AEW aircraft compare in terms of cost-effectiveness. Only the major aviation industries have decided that AEW aircraft are an investment well worth the cost of development and operation. Subsequently, several other nations have decided that they cannot afford not to operate an AEW fleet, and have bought systems from abroad rather than face the cost of developing anything so complex.

Evaluating the relative costs of two aircraft detection systems which use different operating principles is an almost impossible task, but a simple example will illustrate the sort of analysis that is conducted.

A straight border, 500 nm long, is to be defended, and it is assumed that a radar located anywhere on the friendly side of the border can achieve the maximum possible detection range, irrespective of any terrain-imposed constraints. Each ground radar installation is assumed to be a site with the antenna raised 50ft above the ground, and enemy aircraft approaching at 200ft altitude are assumed to be detectable at 30nm from the site. These assumptions mean that radar stations must be provided at no less than 60nm intervals along the border, and consequently at least nine stations will be necessary to provide an unbroken line of radar coverage along 500nm.

It is clearly impractical to place these radars right at the border, where they would be easy targets. At 20nm from the border they would still be easy meat for tactical missiles launched by the enemy from his own side of the border, and at 30nm they will be just a little more safe. With the latter arrangement, geometric considerations suggest that there is a 50% probability of detecting attacking aircraft within 4 nm of the border, and of detecting all aircraft by 30nm. In actual fact these are optimistic assessments, as it has been considered that any target detected at the maximum range credited to each radar installation will be immediately recognisable.

Making a more realistic assessment will require too many subsidiary assumptions for inclusion here, and the clear consequence can only be that the number of radars installed must be increased. For the time being consider that nine installations close to the border will suffice.

It was quoted earlier that AEW assessments have suggested that attacking aircraft must be detected about 160nm before they reach their targets if all conventional defence tactics are to be employed, so we have to assume now that all likely targets within the territory being defended are at least 160nm from the border. To detect and track attacking aircraft over this distance requires three lines of radar installations, comprising 28 radar sites in total.

Those radars installed only 30nm from the border certainly need to be defended, probably by a SAM battery associated with each site. Indeed, it would not be unreasonable to assume that both forward lines of radar need to have defensive systems. It is not impertinent to note also that if all targets are at least 160nm from the border then the intervening area is likely to be water, or at best an inhospitable tract of desert, swamp or other form of no-man's-land. The question of how one locates no fewer than 19 radar sites and their defences in such places, and maintains them, deserves serious study.

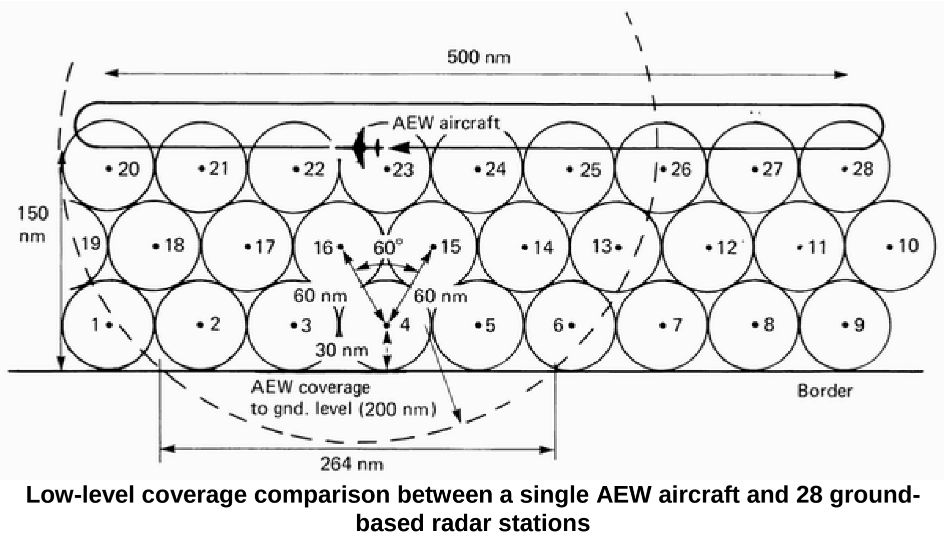

Low-level coverage comparison between a single AEW aircraft and 28 ground-based radar stations

If a nation has targets within 160nm of its borders (and this is far from a rare occurrence), it is clear that conventional ground-based radar systems will not be able to provide adequate warning of attacks for all precautions to be completed in adequate time. It would be impossible, for instance, to provide anything more than just a few minutes' warning to a populated area, whereas about 15 minutes is desirable to permit residents to take cover.

The only alternative available with current technology is an AEW aircraft patrol. If a patrol pattern is set-up along the full 500nm-length of the border, one aircraft will only be able to provide intermittent cover at any point, so the first priority is to discover how many aircraft are needed. First, assume an AEW patroller cruises around 30,000 ft altitude, and it can detect a target approaching at 200ft altitude some 200nm from its patrol line. If the AEW pattern is set up 150nm from the border, it will see approaching aircraft, at best, 50nm inside enemy territory, and at any time it will provide full coverage along a 264 nm stretch of the border. If several AEW aircraft are used, and they all cruise at

300kt in a clockwise pattern, it would be sensible to space them regularly so that each patroller's coverage sweeps behind the preceding aircraft. Ideally there would be sufficient aircraft for the coverage to remain uninterrupted along the border, in which case aircraft should fly at 264 nm intervals. This is regarded as unnecessary in real-life, as, by having reasonably small coverage gaps between aircraft, considerable savings in equipment are possible without a great increase in vulnerability. In the example of a 500nm pattern, if three aircraft are used, spaced regularly in the pattern, and cruising at 300kt, vulnerability gaps of around 15 minutes duration open between successive aircraft. Due to the racetrack type pattern aircraft in the rear of the pattern fill in between aircraft at the front of the pattern, and the diagram accompanying illustrates how gaps open only at the end of the pattern, and at two places along the border, each separated by 166 nm.

A bold enemy pilot might want to enter in one of these blind areas, during the 15 minutes that it existed, and if he entered as soon as the gap opened one might imagine that he could penetrate up to 150nm. In fact, he cannot, as at all times AEW radar coverage is complete within 40 nm of the border, and now he stands being detected by not just one, but by two AEW patrollers. His task would be extremely difficult, his timing must be absolutely accurate, and the penalty of getting detected is that with two witnesses to the incursion there can be little hope of continuing to a target which is still over 100nm distant.

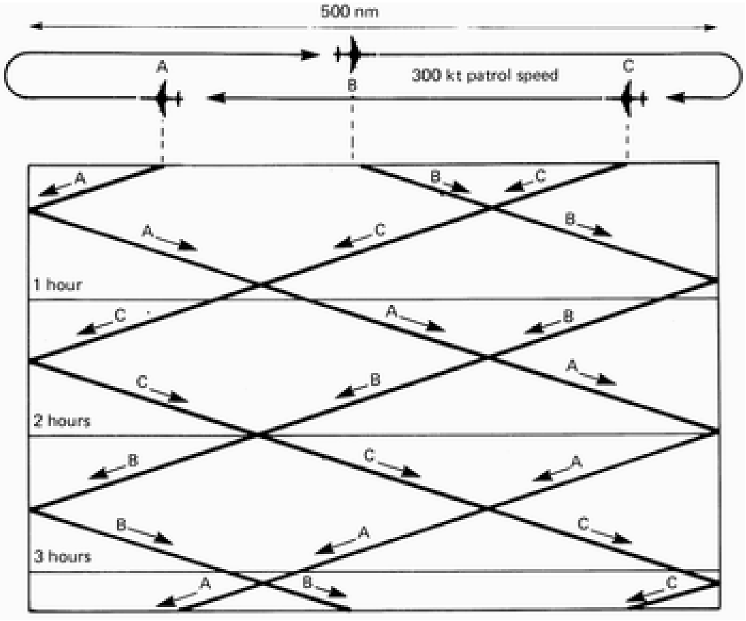

Using several AEW aircraft in a single patrol pattern

This analysis shows that a three aircraft patrol, which would probably require a fleet of between 8 to 10 aircraft to be operated continuously, is a viable alternative to a 28-site ground-based radar system. It is easier to protect than the ground-radar system, and the detection performance is predictable with far greater confidence. But this does not mean that the cost comparison will come up with an absolutely clear answer.

Compared to a ground-based radar system, which can use many identical installations (each of modest technical capability), an AEW radar system is very expensive to develop. On the other hand, with fewer installations, even though each of them is a veritable airborne command-centre, the AEW alternative can be less expensive to operate. Most analyses show that radar stations will usually be competitive in terms of costs, over about 15 years, if the forward sites are not located in difficult terrain. Bear in mind however, that many other related costs which must be taken into account are ignored in this evaluation. For instance, it would be essential to provide a comprehensive method of coordinating activities throughout the total defence system, and the command and control systems necessary for each alternative might bear considerably different magnitudes of cost. In a nutshell, cost comparisons are not easy.

Reverting to a more realistic viewpoint, it was clear in the example above that assumptions associated with the AEW system were considerably fewer than with the ground-based radar system. In practical terms, the conclusions from that experience are that if the border area is an ocean, possibly if it is inhospitable land terrain too, AEW has a great advantage, almost irrespective of the border length. This largely accounts for why the UK has always been a leading supporter of AEW technology. Alternatively, if a border much larger than 500nm has to be protected, the cost of a ground-based radar system can become extremely high, whereas a well-utilised AEW fleet will not cost proportionally more for each mile added to the patrol length. Again AEW can begin to show benefits here, and it was from conclusions in this direction that the US and the Soviet Union found sufficient evidence to warrant the development of AEW fleets.

Small island states can almost always make a strong case for AEW protection, and so not surprisingly Japan has become an early customer for a non-indigenous AEW system. NATO, with its long borders and multiple high-value targets massed relatively close to the line which divides East and West Europe, was obliged to invest early in an AEW fleet too, while a clear case of political necessity brought Israel into the ranks of AEW nations at a very early date. In general, as time passes, and as more long-range surface-to-surface and low-level strike penetration capability is brought to the hands of the leading armed forces of the world, the case for AEW protection gets increasingly stronger.

It might be erroneous to give the impression that military planners are moving towards putting all their eggs in one basket however. All countries which have the technical ability to develop an AEW system can also implement effective ground-based air defence systems, and a combination of AEW and ground-based radars makes a sensible compromise. Established radar border-guard systems such as NADGE, the NORAD DEW-line and UKADGE are all being strengthened at the same time as AEW operations are being introduced. Ground radars not only provide a fallback element if AEW coverage falters, but they are more amenable to having sections hived-off to take on local control and defence requirements. Nowadays, mobile radar sites are often used too. These can be relocated in a matter of hours, but to have this degree of flexibility does tend to increase costs, as more stations are required to ensure that full coverage is available while some units are on the move. A mobile radar system can recover from having a single element destroyed much more rapidly than a fixed-based system however.

The simple fact is that AEW capability does not come cheap. They are amongst the most expensive aircraft in the world to develop, manufacture and operate. When a nation sets its sights on developing an indigenous AEW capability it is embarking on a programme taking probably in excess of ten years from project definition to service entry, and in that period the programme will absorb the efforts of a combined team which will be in strong demand for uses elsewhere in securing the nation's defences. Yet the demand for these aircraft, even expensive as they are, can hardly be expected to diminish in the near future.
