Enacting Prototype Warfare

Despite some misconceptions, robotic and autonomous systems (RAS) technology is 'now technology'; it has immediate military utility, which will evolve over time. It offers two key strategic benefits for Land forces: first, it can help mitigate their combat mass (or effect) challenges and, secondly, it can provide a broad range of military capability which can transform how they 'fight' and 'operate' (below the threshold of conflict).

Notwithstanding the difficulties of tight acquisition and modernisation budgets, Western nations need a more comprehensive approach to adopting and exploiting RAS capabilities.

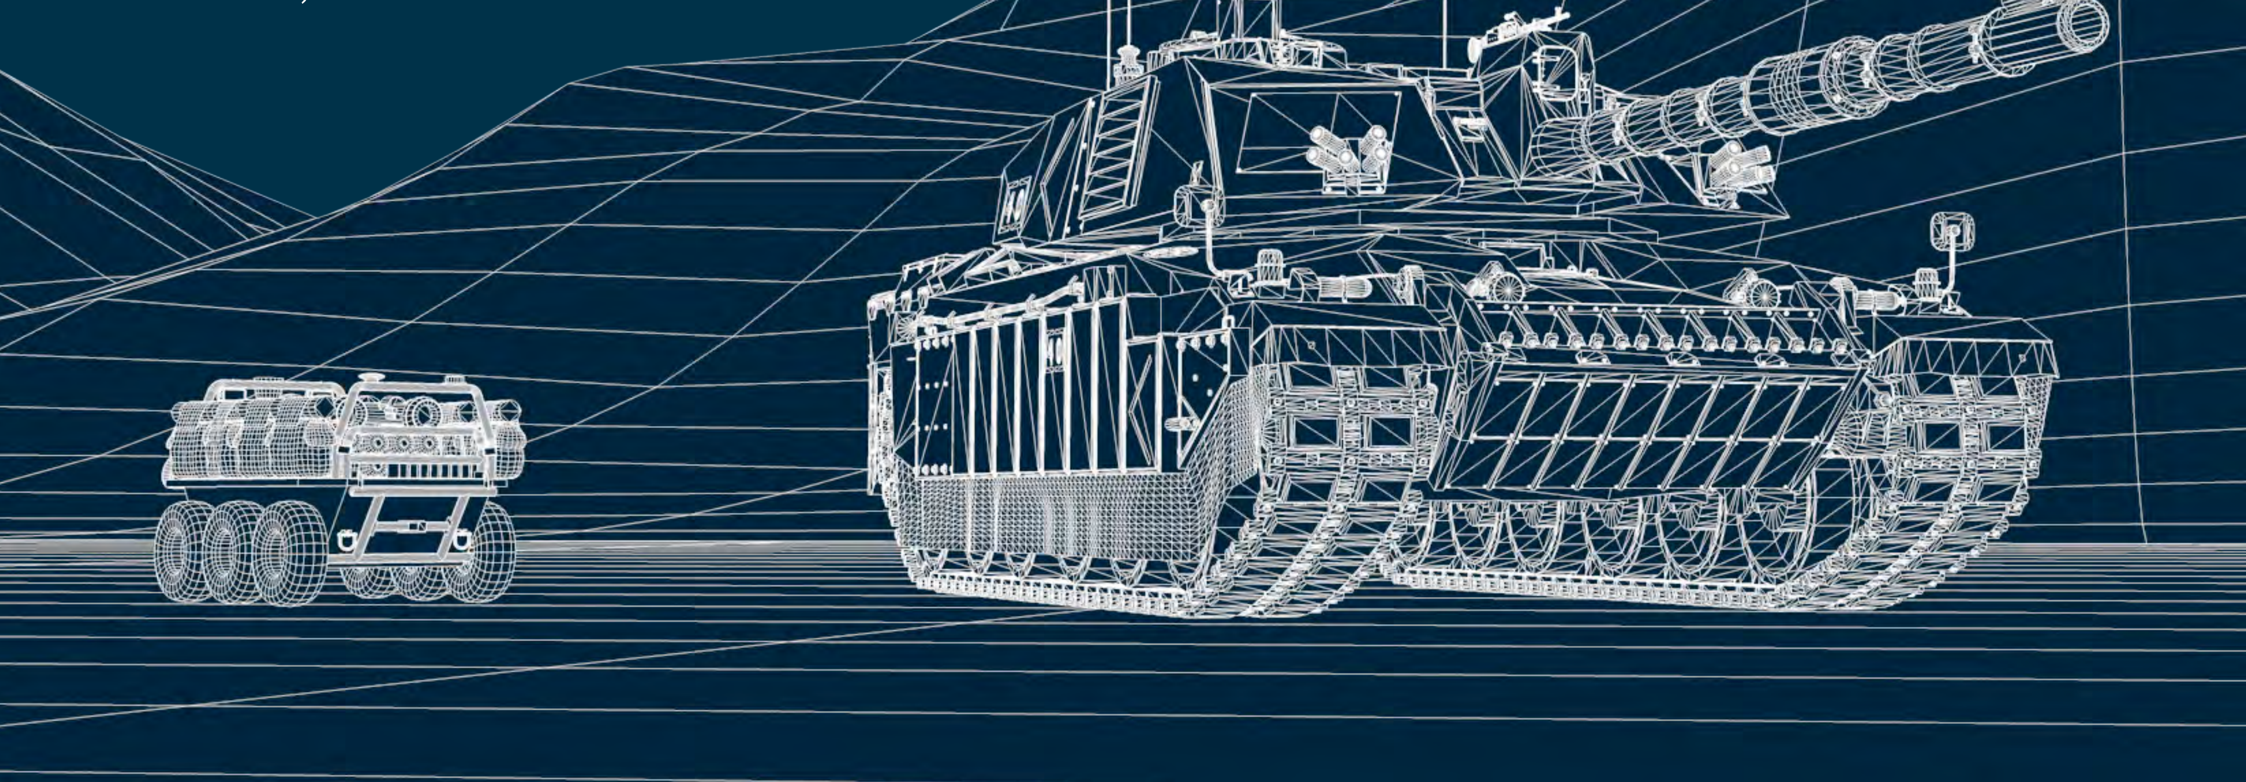

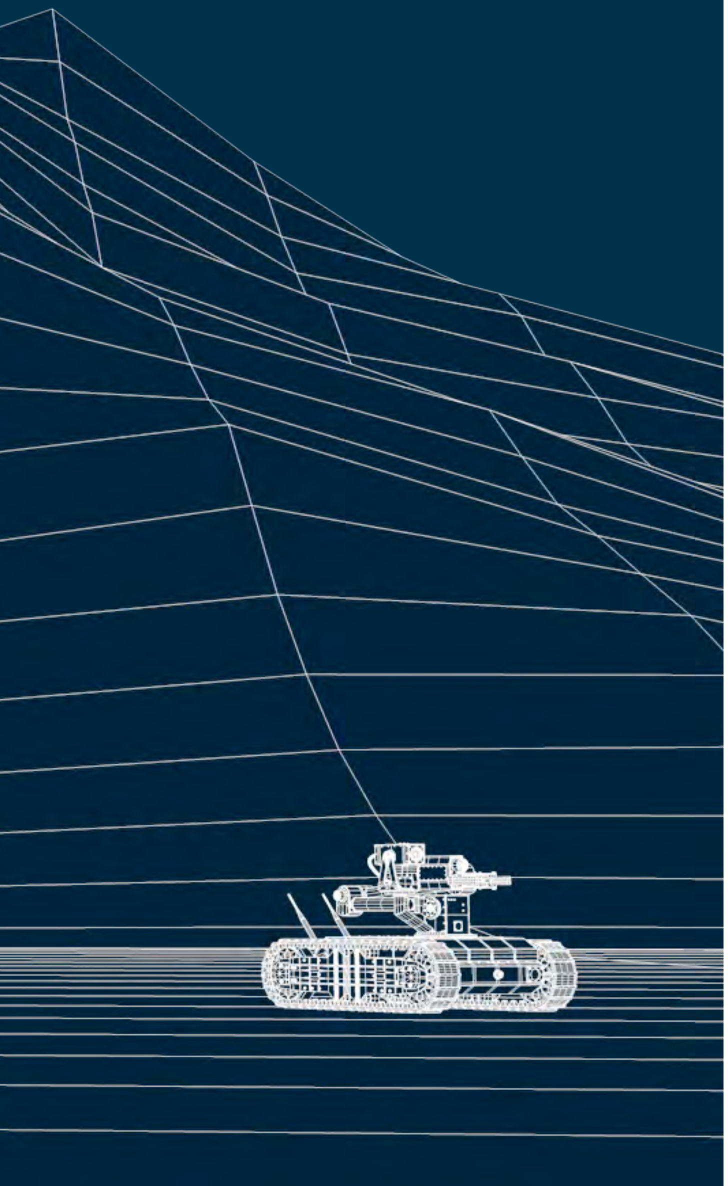

While important, we sense that much of Land forces' ongoing RAS experimentation remains peripheral activity to hedge some future as-yet undefined capability rather than early bold steps of a deliberate strategy to adopt more sophisticated manned/unmanned teaming. We observe that current development of RAS technology not only allows these steps to be taken now, but that Land forces need to start taking them now to address some key capability challenges. An understandable temptation to focus solely on the robotic platform as a measure of progress must be tempered with the reality that effective manned/unmanned teaming – even at basic levels – will only be enabled by addressing the non-equipment components of capability/lines of development early, particularly information.

For those who cry: "why now?" or advocate a fast follower approach, the RAS train has already left the station. Not only are many potential adversaries developing and fielding RAS capabilities – very publicly in many cases – the proliferation of civil sector RAS-related technology means that Land forces can harness viable military capability solutions now; there is also a unique opportunity to shape the way in which some of that technology is developed. The current opportunity (and risk) is neatly summed up by Australian General Angus Campbell who, when Chief of Army, observed that: "late adopters of technology usually struggle to catch up".

At the heart of any deliberate strategy to exploit the potential of RAS technology is the need for a vision to drive ambitious but realistic operational concepts. The US Army's RAS strategy, Australian Army's RAS strategy and a recently expressed British Army goal for a manned/unmanned teamed battle group in 2022 and brigade in 2024 are pivotal first steps which ought to drive a profoundly different view of experimentation, modernisation and acquisition. The operational concepts flowing from this need to describe how unmanned systems can augment manned capability and the manned/unmanned teaming complementarity which creates a combat system rather than a series of one-off robotic platforms or technologies. These concepts will need testing and adjusting through experimentation.

Since November 2019, QinetiQ has run a series of internal and external workshops with scientists, soldiers and other defence industry representatives to explore how RAS technology can improve and accelerate Land forces' manned/unmanned teamed capability to inform comprehensive Land RAS strategy. This work has led to this report which is designed to to stimulate early dialogue and debate of its key conclusions. At its heart is an exploration of the implications of RAS technologies for Land capability, and how to embrace them in a coherent way within a strategy, delivered through an experimentation and acquisition portfolio.

It demonstrates that RAS implementation is not simply a process of platform acquisition: it is a digital transformation which needs to be tackled as such. Without enabling information architecture and due consideration of other components of capability/lines of development, it will be near impossible to create coherent manned/unmanned teams to exploit the full potential of RAS for Land operations. A portfolio-based approach to spiral RAS acquisition, combined with concerted multi-year experimentation, would provide the necessary focus, coherence and integration to achieve the right capability outcomes in a flexible and agile way.

Strategic logic for the military use of RAS

The rationale for teaming humans and machines has been developing over many years. It has not been a linear path but fundamentally it is recognised that Land force generation in the digital era relies upon the complementarity of humans and machines, and the absorption of RAS technologies in a way that delivers a material difference to how Land forces protect, engage, contest and fight.

In the short term, it is crucial to set the technical foundations for RAS' long-term exploitation and establish the right pan-DLOD approach to enable challenges and opportunities to be identified and seized early as an enabler for longer term success.

In the medium to longer term, integrating RAS into any ground combat system is likely to reduce the cost and increase the effectiveness of maintaining (and when required, growing) combat mass to address inherently dynamic threats across the spectrum of conflict and confrontation.

At the heart of this are three key benefits:

Reducing risk

Effective adoption of RAS technologies offers the opportunity to reduce risks to life and military capability in a way that has not previously been possible in the Land domain. Stand-off capabilities have long been a feature of the Air and Maritime environments, offering a range of effects to influence policy outcomes with reduced levels of risk to own forces; this is less easy in a Land environment. Complex, congested, contested and constrained, the Land environment is rarely permissive. Consequently, whenever Land forces are committed to deployments, there are typically substantial risks of casualties or capability attrition.

Integrating robots within Land capability offers a way to degrade or avoid a threat before committing humans or expensive, finite manned platforms. RAS can replace humans in dull, dirty, dangerous and demanding tasks and allow them and other manned capabilities to be concentrated where they can achieve greatest impact, given their unique skills. This reduction in risk has been shown to have noticeable impact on operational performance. With a commander committing unmanned systems in place of manned systems, the option exists to commit mass and force upon an adversary with a vastly revised risk calculus.

Increasing tempo

RAS can expand the scale, direction and velocity of threats presented to an adversary, forcing multiple simultaneous dilemmas on them to the point of overwhelming their decision-making processes. Preventing the enemy from being able to observe the situation; orient themselves to their circumstances; decide what to do; and then act (often referred to as the 'OODA loop') is an effective way to neutralise them. RAS enables that by providing a way to very quickly increase the decision points an enemy needs to contemplate. Multiple RAS assets directing effects on an adversary present a far greater decision-making challenge than a single manned asset doing the same.

The more decision points forced upon an opponent, the quicker cognitive overmatch can be achieved. The value of RAS is to expand the capacity to present decision-making dilemmas to an enemy, without the need for the associated increase in manpower that traditional manned assets would demand. In this way, RAS drive an increase in the tempo of military operations and offer tactical opportunities less possible solely with manned capability.

Reducing the cognitive burden

Land forces are utilising technological advances in sensing to provide ever-greater volumes of information to warfighters. This has been employed to create an operational advantage; however, it also places complex demands on the information-processing capabilities of soldiers. Sensors are very good at gathering data and detecting anomalies but seldom confirm or deny anything themselves. Their output requires processing, exploitation and dissemination – frequently relying upon human cognition to turn the raw data into information or intelligence. That requires considerable judgement and requires soldiers to wade through multiple data points – often including false positives – in order to make decisions. It also means they have to spend time cross-correlating data to corroborate a situation. This is significant because these demands render soldiers vulnerable to cognitive overload, which in turn can result in errors. RAS can take on this responsibility – gathering and processing data either at the edge (within the equipment that has gathered the data itself) or via centralised data stores. These systems never get tired or cognitively overwhelmed, and they can present the human in the loop with analysed data, freeing them to make faster, better decisions. While these three benefits may seem to focus on Land forces' ability to 'fight' with greater success, RAS has very clear utility for other military activity: 'operate'. RAS technologies are becoming an essential part of how Land forces can evolve to meet the needs of more complex operational tasks and the needs of homeland security and resilience.

RAS also fit well with changing political and economic backdrops that expect Land forces to be able to boost combat mass rapidly while simultaneously working within considerable financial constraints. Recent circumstances further highlighted the importance to look innovatively at addressing the need for combat mass without unduly increasing expenditure. This remains essential if Land forces are to maintain their edge to 'fight' effectively – at scale. Implemented well, RAS technologies could be a fundamental part of this solution. By prioritising core information architectures, a scalable backbone can be created which can accommodate rapid expansion, as and when the operational environment requires greater combat mass.

The current paradigm of depending on an ever-decreasing number of capable but expensive multi-mission platforms which, if lost, would be impossible to replace quickly, is unsustainable.

By contrast, a system of low-cost and distributed autonomous systems can be acquired in large numbers to perform a number of the same tasks and be used flexibly to support a range of tactical functions; they can also be replaced at speed – all whilst working within a realistic capability budget.

Furthermore, augmentation of Land forces' most valuable asset – their people – to create human-machine teams enables the multiplication of combat mass without the need to alter significantly existing manning levels. In a time of fiscal and geopolitical uncertainty, the economic benefits of RAS technology are a vital part of the overall rationale for change.

Emerging technologies have the potential not just to change the character of war, but possibly even affect its immutable nature as a contest of wills. Creating and adopting a coherent RAS capability strategy will enable Land forces to adapt to exploit the opportunities inherent in this technology. It follows that an effective RAS strategy must consider the application of RAS technologies and their use and benefits across multiple Land tactical functions, from close combat to fires, manoeuvre, sustainment and intelligence. The remainder of this report explores the practical implications of developing and implementing a RAS strategy through an experimentation and acquisition portfolio to enact the tenets of Prototype Warfare.

-45.7836
-45.7836
AC [ ]
563457568674
563457568674
563457568674
01
02
RDE 056
RDE 056
RDE 056
70
52
84

Ethics and regulation

Rogue states and non-state actors are unconstrained by the Laws of Armed Conflict, enabling them to gain advantage using tactics that cannot be reciprocated by law-abiding nations. This asymmetry cannot and should not be addressed by lowering ethical standards to respond in kind, but can be mitigated to a large degree through smarter use of technology to create capability overmatch.

RAS offer unique opportunities to counter asymmetric threats – but as emerging technology and capability, ethical frameworks and regulations guiding its use are incomplete. It is therefore vital that these ethical guidelines and regulations evolve in tandem with the technology, to ensure RAS strategy remains compliant with the Laws of Armed Conflict.

The minimum standard is compliance with established international laws, such as the Geneva Conventions, Hague Conventions and International Humanitarian Law, which are non-negotiable. However, as legislation struggles to match the pace of technological progress, law alone may be insufficient as a guide to ethical conduct. RAS developers and users must exercise good judgement in developing operating concepts to avoid actions that could later be deemed unethical or illegal.

Organisations should take personal accountability for ethics, establishing 'red lines' and formalising and enforcing them using the following steps:

1. **Create a RAS-focused Ethics Committee** – comprising senior [executive committee-level] leadership, corporate responsibility professionals, legal advisors and other key stakeholders
2. **Agree an independent principles charter** – specifying what outcomes the organisation deems unacceptable for its technology to enable
3. **Through doctrine, tactics and procedures** – share the charter with all personnel and key stakeholders to direct ethical and legal use of RAS
4. **Monitor ethical use** – as RAS capability grows in use there are also opportunities for greater agility in the drafting and application of operational regulations, governed by such bodies as:

- Civil Aviation Authority
- Health and Safety Executive
- Maritime and Coastguard Agency
- Defence Safety Authority
- Military Aviation Authority
- Defence Maritime Regulator

New technologies have always pushed these regulatory boundaries and regulations have always adapted to accommodate them. However, this process must be accelerated through closer collaboration if RAS-related regulations are to keep pace with technological developments.


000067866756
FOLDER 02

Technology overview:

The ‘art of the possible’

Creating an effective RAS strategy requires an understanding of what technology is available today and what may be available in the future. Only by recognising the limits of existing technology and potential of emerging technology to underpin the adoption of autonomous platforms, can Land forces plot a realistic path to success.

This section explores the technology roadmap for RAS. This should inform a Land RAS strategy. First we establish a baseline from the civil environment, and then we look at the individual enabling technologies across three time horizons – 2025, 2030, and 2035.

Technology leadership vs. followership

The transfer of supremacy in the fields of technology and innovation from the public to private sector in the 21st century is well-documented. While the Cold War galvanised governments to compete for technological dominance in the 20th century, progress is now largely driven by privately owned commercial organisations. For the foreseeable future, the private sector will continue to lead development of RAS and the military will follow. While this may feel uncomfortable, it does not have to be a disadvantage. By closely watching – and helping to shape – developments in the private sector, defence can benefit from lessons being learned and adopt technologies as they reach sufficient maturity to suit their prototyping needs.

Furthermore, Western allies will need to track closely RAS developments of adversaries and be ready to respond to them with technologies that can be deployed quickly to neutralise the threat. They will also need to be ready to develop their own novel sovereign capabilities at the pace of relevance.

Commercial markets as a baseline

The scale of autonomy used by commercial vehicle manufacturers offers a useful baseline for how technology might develop for use in military scenarios. There are six levels of autonomy that are broadly agreed for civil vehicles. These run from zero – at which there is no autonomy at all, to level five – where a vehicle is fully autonomous and, as a user, there are no controls, the vehicle does everything you need it to do, and you have no option to interfere.

Most modern vehicles with elements of autonomy built in, such as the latest Teslas, are at level two. There are certain driver aids such as lane departure warnings, speed control, and automatic lights and wipers that help the driver. But the driver retains responsibility for the majority of the vehicle functions and is legally responsible.

At level three, the system now handles all safety management functions. The system is starting to take legal responsibility but is programmed to recognise its limits, and return responsibility to the driver when they are approaching. This is a potentially dangerous place to be because, as a user, the system may return control to you at a point of low situational awareness.

At level four you get more autonomous functionality, but it exists within a very clear set of boundaries – such as on a motorway, for example. So, this will probably be very relevant for logistics companies. Humans can drive through a complex urban environment and, when they reach the motorway, they can switch to greater autonomy because it is a relatively constrained environment. The autonomy capability would then drop to level three or perhaps even two upon exiting the motorway.

Whilst not a viable scale on which to base the military use of RAS, it does at least offer a broad starting point for understanding the levels of autonomy that should be considered, alongside the individual technologies that will enable them to be adopted.

Autonomy - the big picture

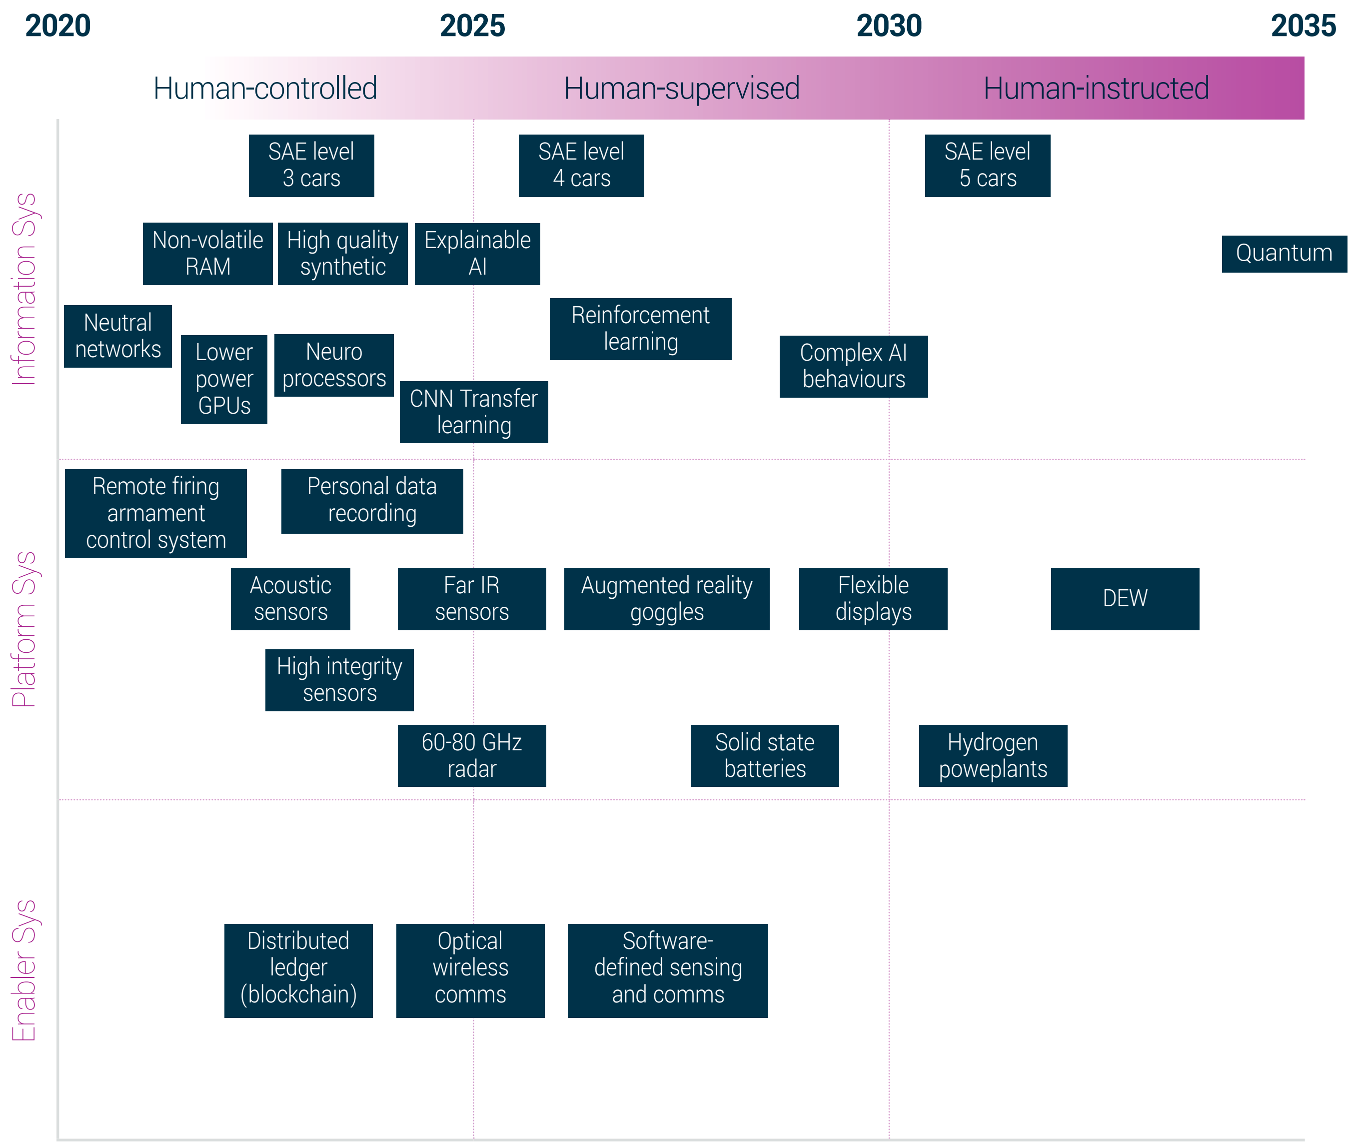

What is worth considering is the difference between expectation and reality. Five years ago, every single major automotive manufacturer was claiming level-five autonomy would be here in 2020. Over the last three years, all bar one has withdrawn that expectation from the market – the exception being Tesla.

The focus is now on level four autonomy, with major vehicle manufacturers now reluctant to even provide a date when they expect level-five autonomy to happen.

A lot of this is down to legal challenges and, in particular, the question of who is responsible if an accident occurs. The same challenge pervades the roadmap for military adoption of RAS technologies.

Three classes of autonomy

Across the three time periods outlined above there is an anticipated shift in the relationship between humans and machines within teaming structures. This shift will be fluid and will not necessarily be neatly compartmentalised into those periods. But it will broadly move from human-controlled, through human-supervised, and towards human instructed. There will be a blending throughout, with early systems having some supervised features but mostly controlled. Over time the proportion will change as supervised/instructed features take over.

Viable technologies by 2025

Between 2020 and 2025 we expect RAS technologies to develop but to largely remain human-controlled. Humans will be in the loop, controlling assets almost all of the time. The teleoperation of a ground vehicle or a Reaper drone are good examples of this in practice. Limited levels of autonomy will enable platforms to navigate towards specific waypoints within comparatively simple environments. This would allow humans to handover specific functions to machines, where carefully monitored autonomy is in control for a short timeframe.

There are a number of technologies already available to support the delivery of human-controlled RAS capability. These include neural networks; computer vision systems; low power graphics processing units; high quality synthetic environments; and non-volatile RAM. These are readily available within the computing industry today and would transform the way certain functions work in Land. But, because of the way militaries tend to procure new equipment, this type of technology is not yet in the hands of soldiers. One of the challenges in the 2025 epoch is therefore how to make the most of what is already in the commercial world.

Viable technologies by 2030

Beyond 2025 we expect RAS technologies to see some shifts into the domain of 'human-supervised'. In this scenario, a human will give an instruction, leave the system to progress the task, and check back in on the system frequently to ensure it is performing as expected, and that it is not encountering anything it cannot deal with.

The major transformation we expect to enable this shift is the availability of explainable artificial intelligence (AI). AI is already used in many different defence scenarios, including automatic target recognition systems, but it cannot be used in any safety functions. It is already possible to create a neural network that can navigate a vehicle, but there is no way to certificate it for use in defence environments because it is a 'black box' – there is no way for a user to see the way in which conclusions are reached and decisions are made. Until that is possible, an additional safety system will always be required and that limits any shift from 'human-controlled' to 'human-supervised'.

We can also expect to see the introduction of high-integrity sensors that can be relied upon to a far greater extent than current technology. So, not only will the AI software needed to drive RAS adoption be certified for safety functions, so will many of the sensors required to gather the data on which decisions are based. Combined – this allows the development of 'systems of systems' safety arguments that would otherwise be a hurdle to moving away from a human-controlled environment.

Viable technologies by 2035

By 2035, there will still be human-controlled and supervised technologies, but we expect to see more humans instructing RAS technologies, in the same way that they would instruct another person. They will be expected, based on their technology, training and competence, to deliver against a task requirement without human intervention, and come back to the operator if a problem is encountered.

One of the key technology changes to enable this will be around power sources to provide RAS with much longer operating times before the need to replace or recharge. A move to solid state batteries will provide up to 100 times the electrical power of existing power sources and hydrogen systems will also be an option for much longer availability.

These will be combined with a new range of sensors that are software driven – ensuring a single sensor can be used for multiple functions, depending on the immediate data need. This adds significant flexibility to a RAS asset and means each platform can use the power of information and software to reduce the amount of hardware on which it relies. It also means each asset can be used for a much wider range of functions, automatically switching from comms, to sensing, and analysis. This will vastly reduce the amount of human control and instruction required over each asset to deliver results.

Combined with the availability of explainable AI, we expect to see software, hardware and power sources coming together to form systems that can deliver true RAS capability, not simply RAS technology.

Use Case Scenarios

A good way to illustrate the roles RAS capability can play in today's operational environments is to look to the past. In this section we have created simplified maps of Operation Plunder – one of the most key operations in the closing stages of the Second World War in Europe. The operation was to seize and exploit bridgeheads across the River Rhine from which to advance quickly deep into Germany. A complex blend of 'deep', 'close' and 'rear' actions around a major obstacle crossing at scale, Op Plunder (and the associated air operation, codenamed Op Varsity) is an excellent vehicle to show the plethora of roles and effects RAS technology can have in a Land-based, multi-domain operation.

This map demonstrates a subset of Op Plunder – known as Op Torchlight. On that map we have highlighted where RAS technologies could have been deployed, and for what purpose, as an illustration of their potential to support as part of human-machine teams. Below we offer a more detailed explanation of how they would execute their functions in this scenario:

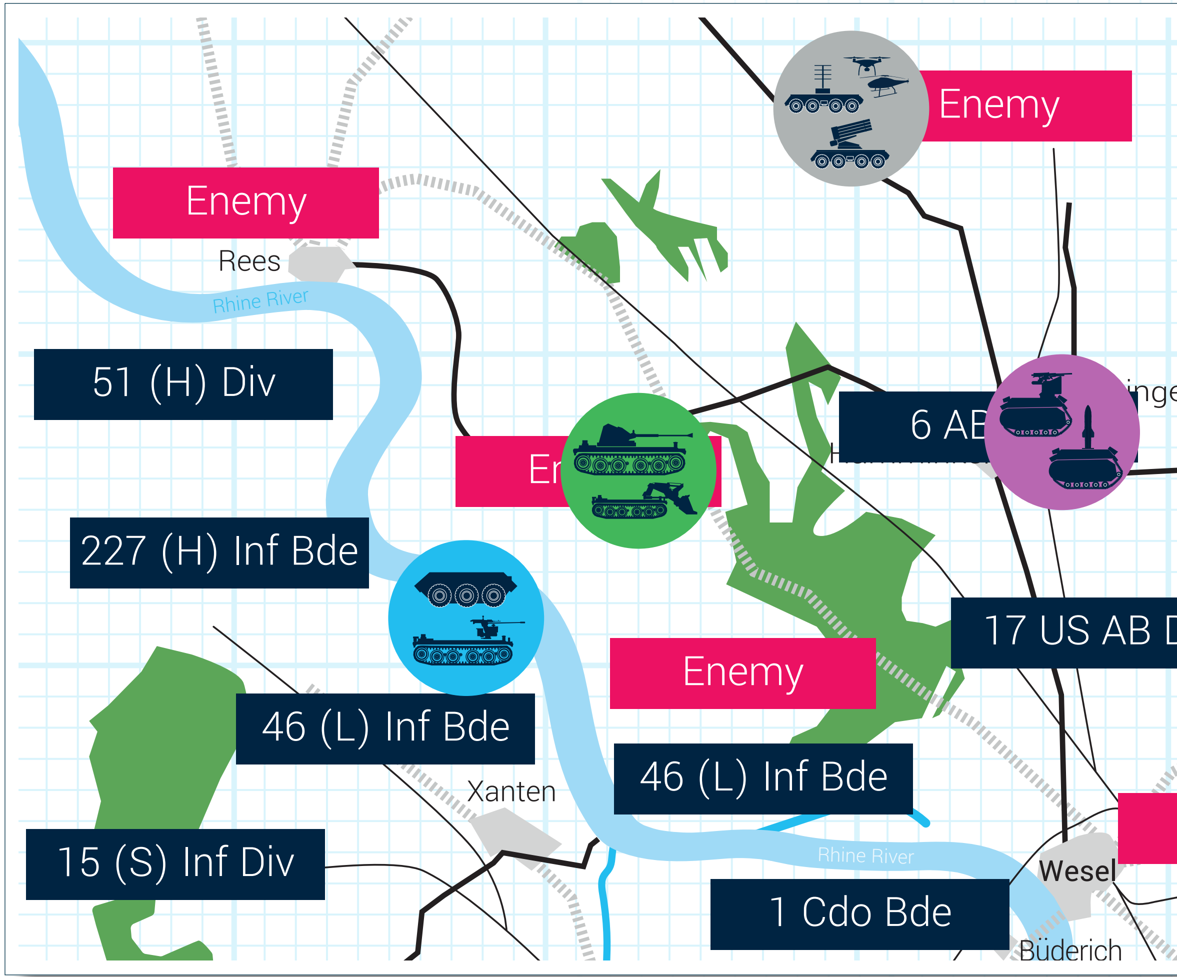

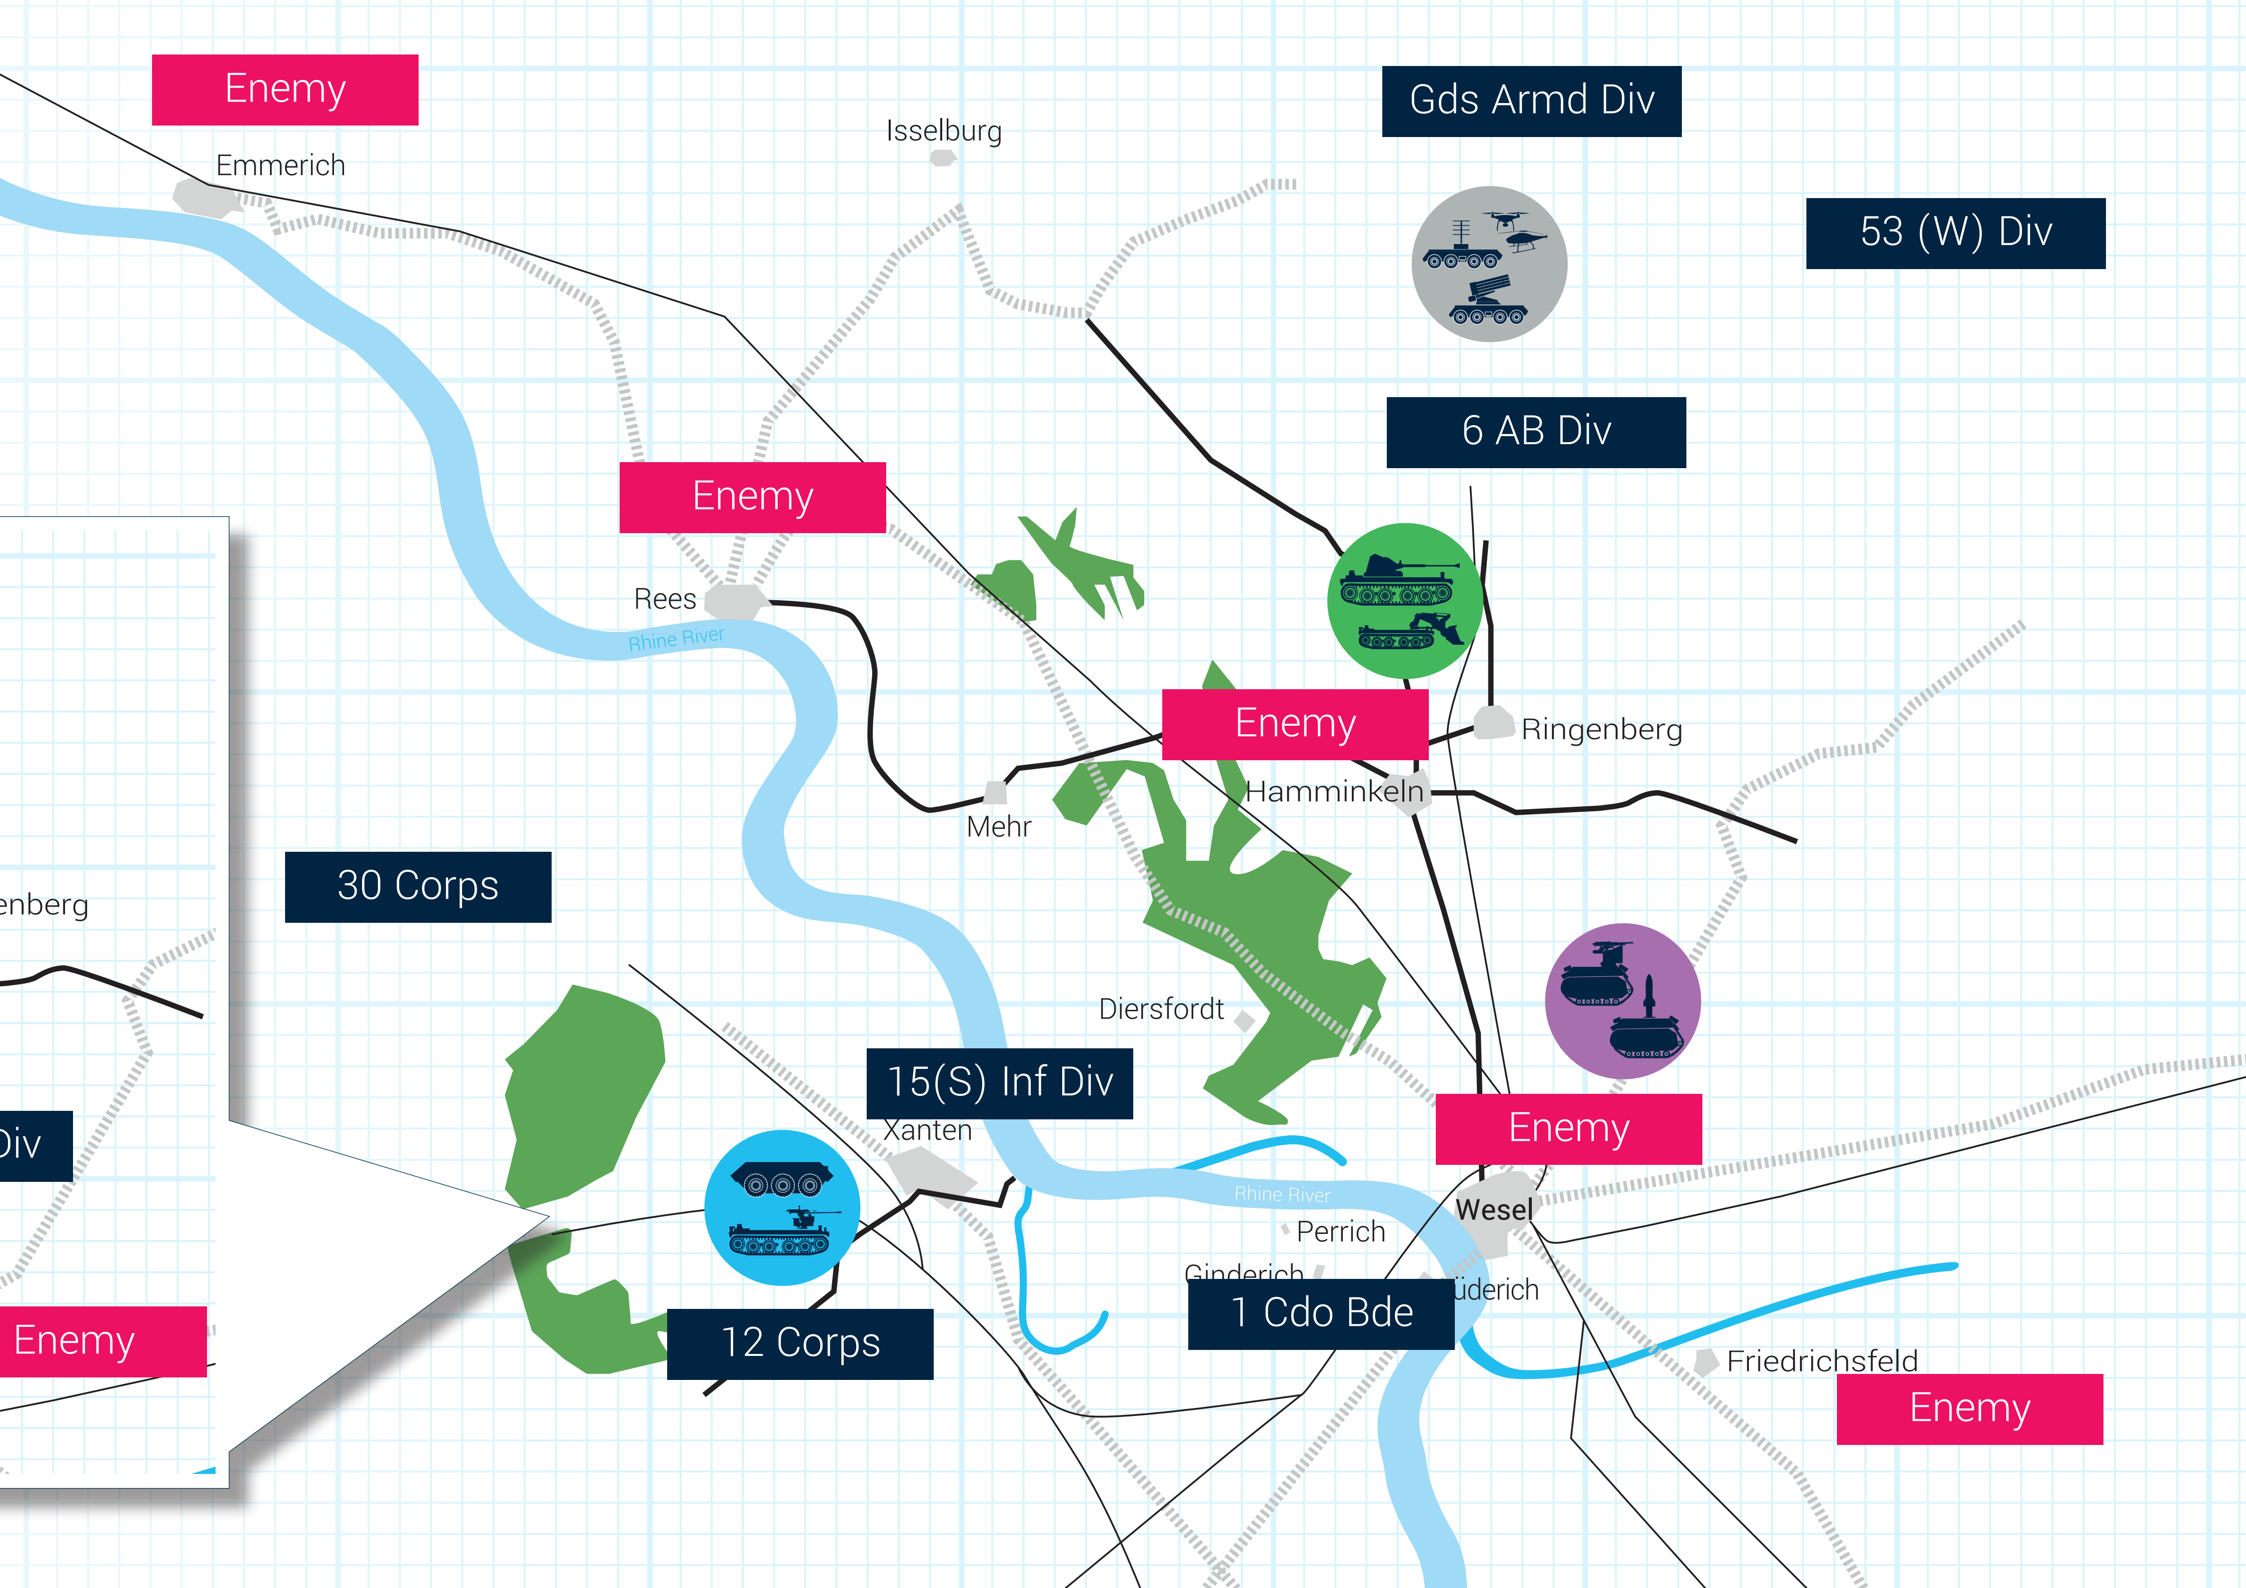
Enemy
Emmerich
Isselburg
Gds Armd Div
53 (W) Div
6 AB Div
Enemy
Rees
Rhine River
Enemy
Ringenberg
Hamminkeln
Mehr
30 Corps
Diersfordt
15(S) Inf Div
Xanten
Enemy
Rhine River
Wesel
Perrich
1 Cdo Bde
12 Corps
Friedrichsfeld
Enemy
Enemy

1. Close Operations – Fight

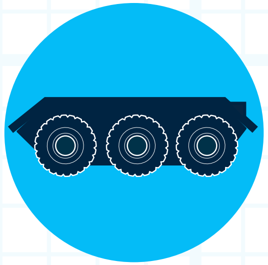

MEDEVAC UGV

A small sized wheeled UGV able to carry 1-2 casualties at a time conducts rapid evacuation. Once stabilised, casualties are packaged for transit and brought from point of injury to the nearest aid station. This minimises the physical burden on troops by reducing stretcher carries, and allows combat effectiveness to be maintained – keeping troops in the fight. Basic life support systems are integrated to the vehicle to sustain casualties in transit. The on-board autonomy is designed to select routes with as little undulation as possible in order to minimise the risk of further injury, whilst covering each journey with the appropriate speed.

Fire Support UGV

A light tracked UGV with high mobility and a payload of approx. 3 tonnes carries a variety of direct fire systems and integrated sensors. A suite of advanced information systems enable its performance. These include autonomous navigation that allows it to move in formation with manned vehicles and fuse situational awareness of opposing force locations, in order to take up suitable fire positions. Its sensors and effectors are collaboratively networked with others in the battlespace, meaning it can both acquire and engage targets as a member of a human-machine team.

2. Close Operations – Advance

Over-watch and Heavy Fire Support UGV

Teamed with other AFVS, this large tracked UGV (up to 20 tonnes) provides heavy direct fire support to its manned vehicle team-mates. Moving ahead of the manned vehicles, it takes up over-watch positions at vulnerable points, its primary weapon providing a means to engage enemy armour in a close fight. Its autonomy stack is designed to manoeuvre the vehicle aggressively when required – taking the fight to the enemy and accepting a higher risk profile in how it reacts to obstacles. Its collaborative targeting software means it can receive situational awareness from any other sensor in the battlespace and automatically suggest courses of action to engage threats to its human team-mate.

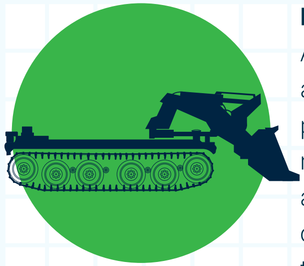

Mobility Support

A medium tracked UGV fitted with a combat engineering payload provides mobility support. As manned assets advance, they are likely to face counter-mobility obstacles. This system moves in to deal with those obstacles, but without exposing human operators to high threat areas. The on-board sensors and autonomy stack prioritise the gathering of local geospatial data, providing an engineering reconnaissance function as the platform advances. This data is shared with other platforms in the vicinity. The payload can complete a number of tasks automatically, but for highly specialised and dynamic tasks, it can revert to remote control - allowing the human operator to apply their specialist skills.

3. Close Operations – Light Forces

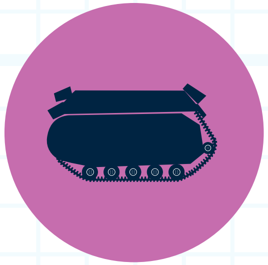

Resupply/Mule

A small tracked UGV (under 2 tonnes) minimises the load carriage burden on Light forces. The system can function in two modes – either following a patrol as a "mule" vehicle, or conducting resupply operations autonomously. The autonomy stack is calibrated to manoeuvre the vehicle in difficult and constrained environments. It requires minimal operator input and can function in "follow-me" mode for extended periods. It is networked to a logistic information system that means it knows where to go to collect supplies, and how to distribute them in priority order.

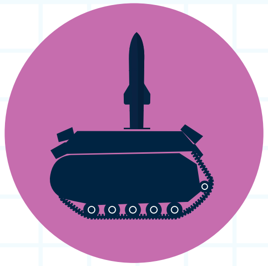

Fire Support/Anti-Armour

A small tracked UGV carries a range of fire support and anti-armour effects. This brings a range of capability to Light forces that would otherwise have to be carried by dismounted troops or on poorly protected vehicles. Its autonomy stack prioritises stealthy movement so that its signature does not alert enemy forces. This makes it suitable for close support of dismounted troops. Its targeting suite is linked to a Dismounted Situational Awareness system – meaning troops can rapidly pass it targets to engage and ensure de-confliction of fires when operating in close confines.

4. Deep Operations – Shaping

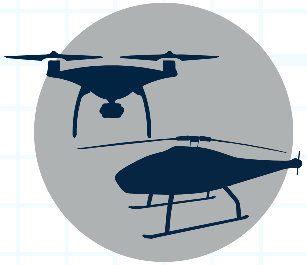

ISTAR

The FIND function in the Deep battlespace is conducted by a range of UAS platforms – both rotary and fixed wing. With varying range, payloads and signatures, a family of systems work collaboratively to perform multi-spectral ISTAR missions. Single human operators control multiple autonomously navigating platforms with information processing and fusion taking place at the edge – minimising the amount of data to be backhauled over the bearer network. Advanced analytics further minimise the operator burden, providing intelligence driven alerts that enable the cross cueing of both kinetic and non-kinetic effects.

Fires

A large wheeled UGV (payload of up to 8 tonnes) carries a GLMRS payload. Operating dispersed, it can deliver effects into the deep battlespace at range – providing a stand-off fires capability. Its autonomy stack is calibrated towards survivability and avoiding detection. Upon firing, it immediately performs a “shoot and scoot” manoeuvre so that it can rapidly vacate the launch site and avoid counter battery fire. It navigates autonomously to a resupply point for reloading, before dispersing once again. Targets are passed to the system via a network suite of ISTAR assets.

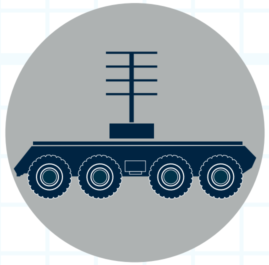

Deception

A wheeled UGV carries payloads that mimic the RF/electronic emissions of a larger formation. Further EW/EA modes are also available – functioning within a wider CEMA plan. The autonomy stack is configurable to replicate the movement of other categories of vehicle – e.g. to imitate a convoy of logistics trucks on a road move, or IFVs on the march.

Across the battlespace – information systems: the golden thread

All of the capabilities within the human-machine team function as part of an integrated and coherent information system. Every soldier, platform and sensor is both consumer and supplier of information. Wherever possible, processing, exploitation and dissemination of that information is done at the edge. Information is only backhauled via the network where relevant. But, what information is shared enables a host of information-driven activities – from the sharing of geospatial data to support autonomous vehicle navigation, to the integrated and collaborative engagement of the enemy via networked ISTAR. Information is prized with this force; it is the key driver of agile C2 and improved decision-making, creating better operational outcomes.

Autonomous Systems - Conceptual Architecture

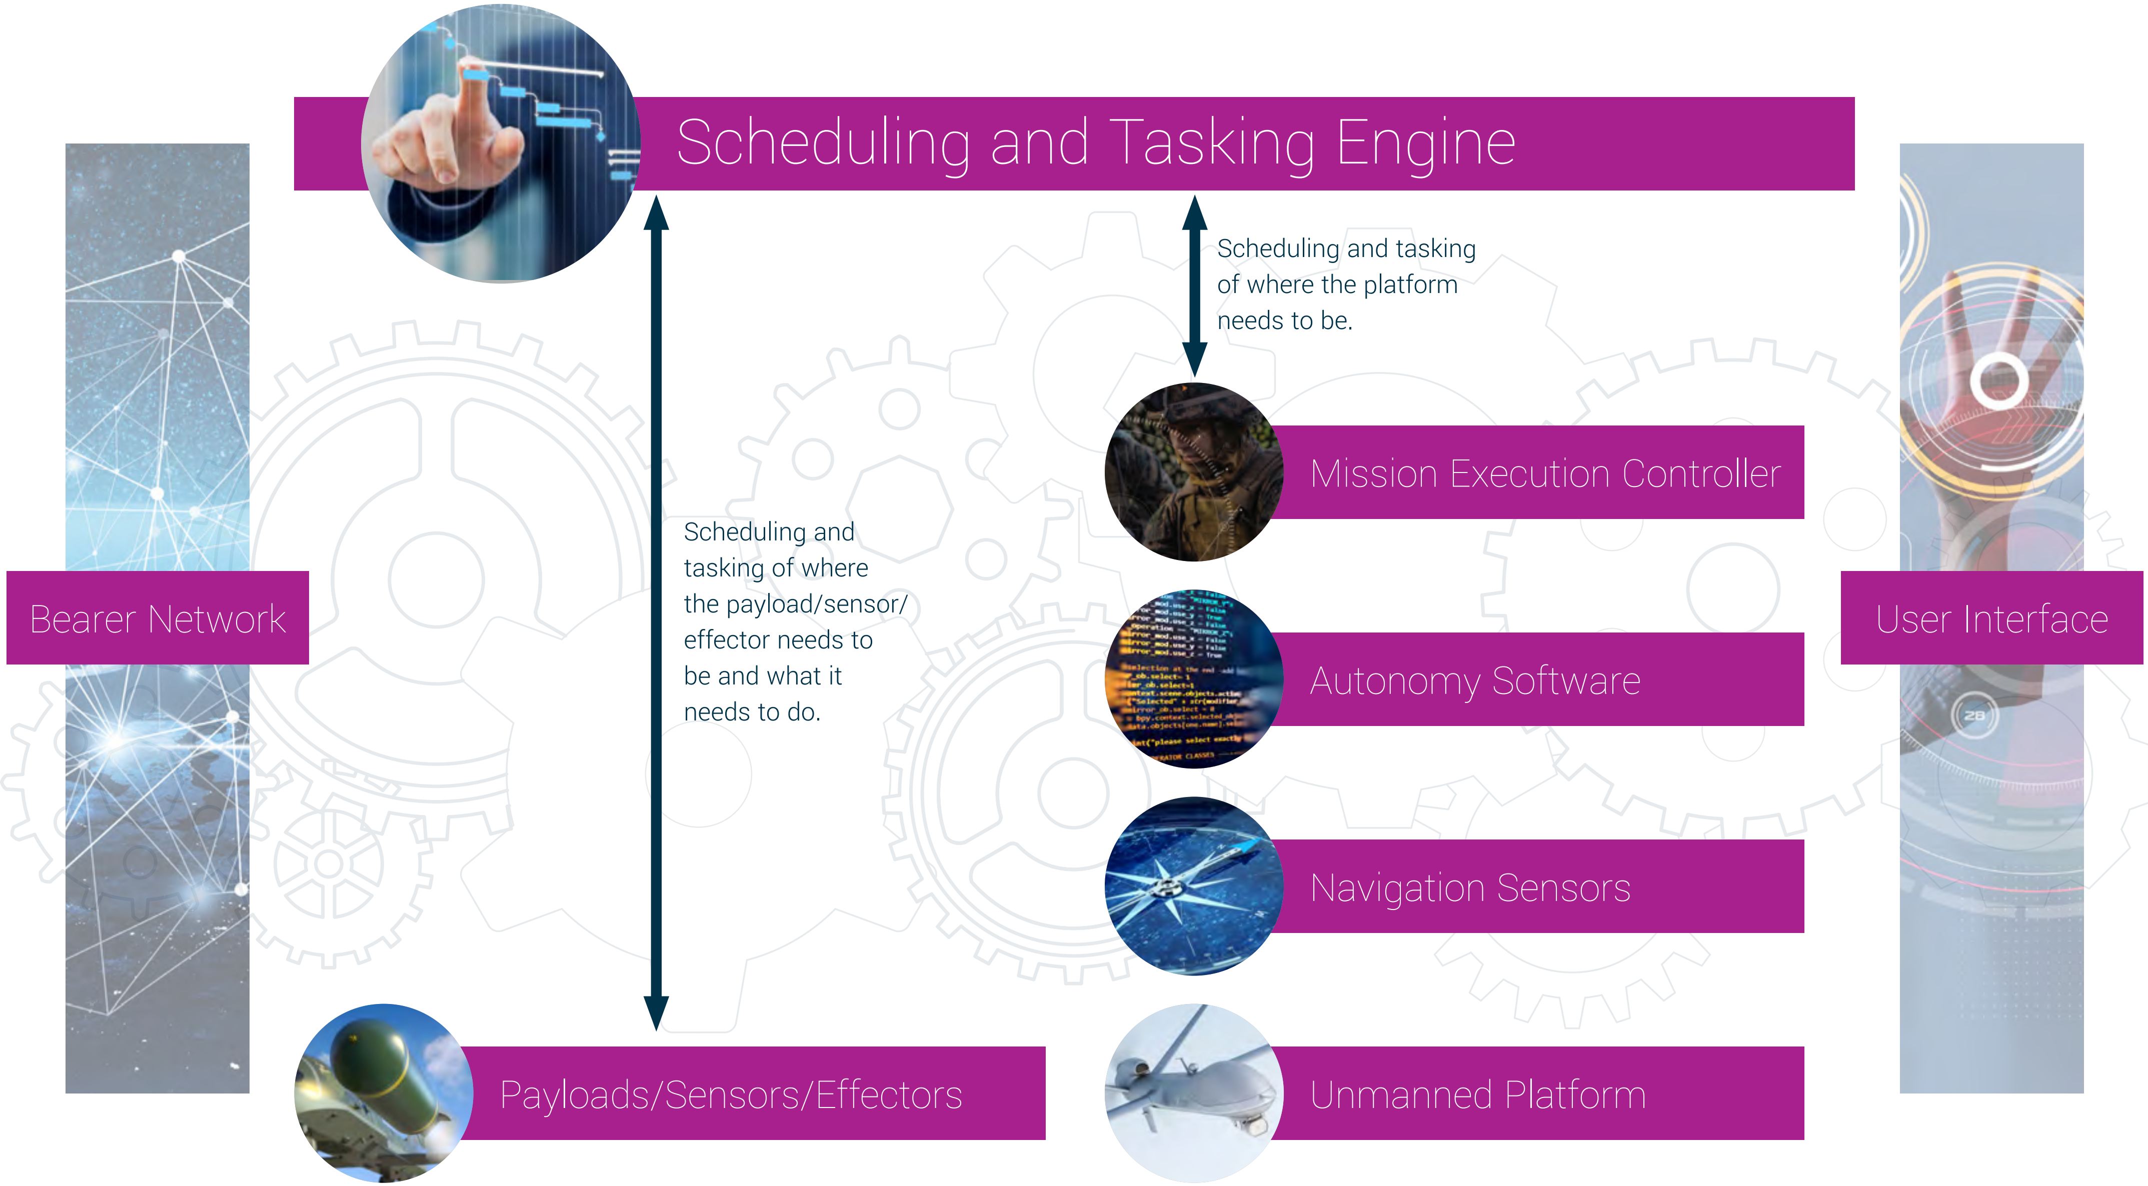

The implications of RAS

From physical infrastructure to human skills, the introduction of RAS will have several knock-on effects that must be taken into account when planning future Land strategy.

It’s 2030, and the new DCGS takes to the stage at the Land Capability Industry forum. Her career has been meteoric to date, rising through a succession of command and staff jobs, showing considerable acumen, judgement and leadership throughout. Well-versed in capability management and acquisition, she understands instinctively the challenges of creating a coherent fighting force from the Army’s tightly squeezed budget. She starts to outline a bold new agenda to transform the Army’s equipment and information systems.

Her PowerPoint slides show an ambitious programme of development and bold action that aims to accelerate the adoption of technologies which are now commonplace in civilian life. She pauses to reflect on the progress of recent years; her slideshow switches to images of recent acquisition programmes and evaluation procurements, displaying a range of UAVs and UGVs. These exciting technologies had promised so much at the outset, she reflected, but they have one major problem: despite all the money invested and good intentions of those involved in their development, requirements setting, acquisition and fielding – none of them work together. They don’t talk to each other; they are all on the same network but can’t share data or situational awareness.

DCGS turns to the audience and asks a simple question: “how did we allow this to happen and how do we stop it from happening this time?”

Implication:

Information – the heart of the RAS ecosystem

The movement of information and data between platforms and operators does not just enable RAS capability – it is RAS capability. The power of Land RAS comes from leveraging information shared between multiple systems. Robots with no infrastructure can be likened to mobile phone handsets with no network, no data architecture, and no operating system. Without connectivity and the ability to share and process data, they cannot operate. A secure, high-capacity, resilient network is therefore vital.

Communication between robotic systems will take place in a congested and contested digital space, competing for access to the electromagnetic spectrum. Infrastructure must be built with this in mind, and, once established, secured with a robust spectrum management strategy. This is covered in more detail in the 'security implications' section of this report.

Building the architecture to facilitate connectivity is the essential first step in deploying RAS in a land environment. But what is that architecture composed of?

The digital backbone

A robotic autonomous system is a collection of hardware and software elements that both consume and supply information.

At the front end of the system is the user interface – the door through which the human accesses the capability. Via this interface, the user tasks platforms and their subsystems (such as payloads, sensors and effectors), and receives information from those systems that informs subsequent tasking decisions.

At the far end are the multiple RAS carrying out their tasks. These systems consist of robotic platforms, navigation sensors, effectors, autonomy software, and the mission execution controller responsible for timing and sequencing. Every action, threat and attack is logged at this end, then processed and delivered back up the chain to the commander.

Linking the two ends is the scheduling and tasking engine, which is critical to the operation of the system as a whole. Without this vital intermediary, the operator would quickly suffer cognitive burnout trying to task and monitor multiple systems, while simultaneously interpreting and acting upon the incoming information. The role of the scheduling and tasking engine is therefore twofold. For the human, it serves to distribute the outbound commands to the systems,

and fuse the inbound data into a comprehendible and actionable tactical picture – prioritising that which is most important to minimise the cognitive load. For the machine, it translates the operator's instructions into packets of data that are relevant and understandable for the on-board software. This can include geospatial data for navigation and routing, situational awareness on the location of hostile forces or information of potential civilian population/traffic interactions to guide vehicle behaviours.

Incremental steps

Taken as a whole, the task of implementing brand new system of systems architecture can seem daunting – but it does not have to be completed in full at the outset. The initial system must contain all of the fundamental elements outlined above, but not to the highest possible degree. RAS architecture has the advantage of being eminently scalable, so it is entirely feasible to begin at a modest scale and gradually grow operations to become more complex over time. The first step may involve three small UGVs, with a longer-term aspiration to oversee a full-size tank with two UGV wingmen and a heavy UAV for aerial navigation and reconnaissance. The smaller-scale systems at the more accessible end of this spectrum are implementable today – we have already seen single-user, multi-platform control exercised in live demonstrations, with a lone operator commanding three to four light UxVs. This is enough to provide a solid foundation on which to build a more adaptable, diverse and powerful capability.

Taking the first steps

The information systems dialogue right now needs to focus on the data. Who owns it, how is it assured, what is the optimum format, and is it secure? These considerations should form part of the parallel conversations on the safety and interoperability challenges, outlined in the corresponding sections of this report.

From a practical standpoint, the most important activity is experimentation – but where information flows are prioritised rather than just focusing on physical platforms. The first step is to take existing data platforms and user interfaces proven in live demonstrations, plug in new platforms, develop new use cases and continually push the boundaries in pursuit of longer-term aspirations. This is already happening – but the technology cannot remain on the testing range forever. The next step is therefore to begin deploying the systems into active service, taking incremental steps forward in terms of complexity and risk. If the right information can be delivered to the right places, at the right times, in the right formats, the power of the capability will grow exponentially in response.

Implication:

The impact of RAS for the human

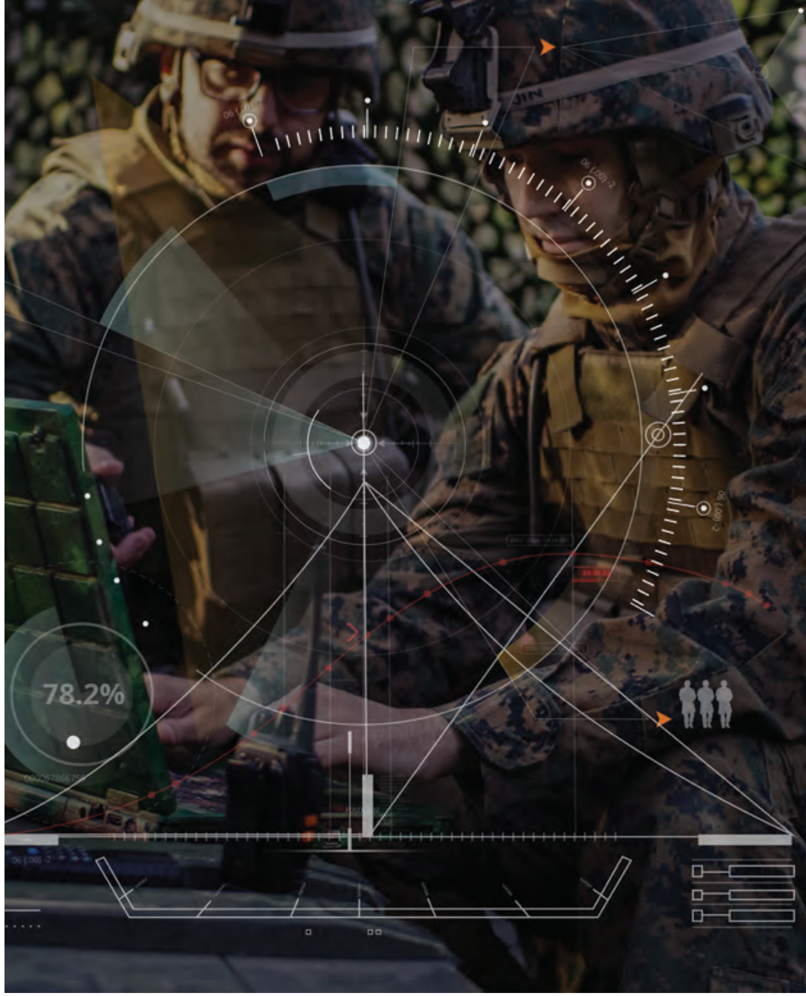

The human brain remains the most sophisticated computer on the battlefield. Any decision regarding the use of RAS should therefore focus on how the technology can benefit human soldiers in the most meaningful way. It is a mistake to begin forming concepts of operations based on what platforms and systems are capable of doing without first considering the human need, followed by an assessment of how RAS can help to fulfil it.

For example, there is a constant human need in warfare to reduce the information burden on the soldier. An excess of information can lead to cognitive fatigue, increasing the likelihood of fatal errors. Starting with this challenge, RAS use cases can be conceived to ease the cognitive burden. Similarly, there is a need to prevent human casualties, and so further use cases will focus on minimising soldiers' exposure to dangerous environments.

But perhaps the more urgent task is to determine the roles best suited to humans, and those best suited to robots. These will be divided down both practical and ethical lines, as there will be instances where the operational advantage of a particular course of action is outweighed by the moral case against it. The optimum outcome is for soldiers to be committed to roles that only soldiers can perform, while robots free up manpower by conducting the missions that do not require the human touch. But how are those duties assigned?

The role of the human in the RAS-equipped battlespace will be subject to significant change. The psychological effect of this should not be underestimated – a tank commander of twenty years may respond with a tribal mentality, seeing unmanned combat vehicles as a threat to the skills that have come to define them. The unconscious reaction may be one of self-preservation, causing them to fight back against the perceived threat. But resisting the change is not an option. Just as the tank superseded the cavalry during World War I, RAS will soon supersede the tank – and those who have not adapted will find themselves at a severe disadvantage.

While RAS capability will be designed around human requirements, in fulfilling those requirements it will displace humans and create new demands concerning training, skills, responsibilities and behaviours.

User responses

Understanding the behaviour of humans in their interactions with robots is critical in developing safe and effective capability. The need for trust in robotic team members is no less vital than for human colleagues, although the way in which this trust is earned is different. Trust between humans is forged on mutual understanding and a shared sense of duty, while a human's trust in RAS is based on an assurance that the reality of the robot's performance will match the expectation. Live experimentation provides the conditions in which to build that trust, but also helps to manage user expectations. This ensures users do not inadvertently trust platforms to perform tasks for which they are unsuitable.

There is also a psychological aspect to teaming humans with machines. Human teams are bonded through reliance, respect and camaraderie. Assuming these qualities do not extend to robotic teammates – what effect does that have on decision-making? Conversely, what if human soldiers do form emotional bonds with robots? In a 2015 study by researchers at Germany's Heinrich-Heine-Universität Düsseldorf, brain scans revealed that human participants felt empathy for robotic vacuum cleaners that were kicked or verbally abused. Could subconscious sentimentality cloud judgement on the battlefield?

The way in which information is presented to users is also important. As well as being optimised to minimise the cognitive burden, it must not become a distraction in its own right. A soldier constantly looking at a screen may lack sufficient awareness of physical threats developing around them, like a mobile phone user who walks into a lamppost. Data provision and interfaces must be developed so as not to demand too much of a user's attention.

Not only must humans understand and trust their robotic teammates, but the machines must also have some level of 'understanding' of human behaviour. For example, if during battle an UGV operating autonomously becomes stuck in a ditch, it may sound an alarm to alert the supervisor that human intervention is required to free it. However, freeing that specific UGV may not be the supervisor's priority in the wider context of the battle, and repeatedly sounding an alarm may unnecessarily add to the cognitive load and stress. Autonomous systems should therefore exercise 'judgement' about their interactions with humans to ensure they do not hinder operations. Where this is not possible, there should be an option for the operator to quickly overrule the alert.

Live and virtual experimentation, under the observation of human performance scientists, will amass evidence in relation to all of these factors.

Service and support

The introduction of RAS will place other demands upon human skills. As Land increases the volume of platforms in the battlespace, maintenance and repair requirements will heighten. The skills needed for mechanical maintenance already exist within the armed forces, but RAS brings a new demand for software developers, data scientists, and other expertise not currently prevalent in the military. The training and recruiting implications of this change are discussed below, and there are also questions about where service and support crews should be physically located, outlined in the 'Organisational Design' chapter of this report.

Training and recruiting

Training for operating RAS assets can be accelerated using virtual platforms operated in synthetic environments, reducing reliance on the availability of physical platforms and ranges. It will also equip operators with the skills to operate multiple platform types, as assigning specialist operators to specific platforms would place severe limits on overall operational flexibility, and the advantage that comes with it.

But that is not the limit of training requirements. Training will also be needed to support RAS technologies both in operation, and back at base. This skills gap will not be addressed overnight – it will take many years to foster new skills pipelines, build career structures, and train recruits to fill vacancies in the volumes required. In the near term, armies will rely on contractors, who will take up new positions within deployed formations to meet service and support requirements. Another possible avenue is to draw on the reserve forces – by both deploying existing reserves with the necessary skills, and by launching reservist recruitment campaigns similar to those previously targeted at cyber security specialists.

From augmentation to teaming

The transition of RAS platforms from tools to teammates will be incremental, evolving through three distinct horizons: human-controlled; human-supervised; and human-instructed. Human-controlled robots are already an established feature of Land combat, such as those used for ordnance disposal and scouting missions. Next, human-supervised robots will exercise degrees of autonomy, overseen by an operator who may intervene to perform certain actions. For example, a platform equipped with a weapon may navigate autonomously, but will require a human to take firing decisions. Weaponised platforms will not progress beyond this stage. Finally, human-instructed robots will be tasked by a human, carry out the mission autonomously, and return to base with no intervention. This will require the highest level of trust, acquired through rigorous experimentation and combat experience, and training informed by an understanding of human psychological responses to RAS.

Implication:

Moving from concept to capability

The extent to which RAS can be effective rests on exploring their incorporation into concepts and culture. This goes beyond finding ways to simply replicate existing concepts using robots instead of manned systems – an approach that has been dismissively labelled 'digitising the analogue' by senior military figures. Land concepts must undergo radical reform to provide the conditions under which RAS can offer genuine, transformative advantage.

Advantage will come from the ability of RAS to accelerate the 'OODA loop' of observe-orient-decide-act. Achieving this will depend on multiple factors, including:

- Most effective transmission, fusion and presentation of data
- Possession of the most capable platforms and subsystems
- Weapons and techniques to disrupt the enemy's OODA loop
- Security and resilience against OODA loop disruption by the enemy
- Most advantageous doctrine concerning RAS use

These raise technical, procedural and cultural questions that must be addressed to achieve capability overmatch against adversaries, and the resulting acceleration of the decision-making cycle. How are the platforms commanded and controlled? What tactical functions suit unmanned capability? And what roles demand manned capability?

Should it stay or should it go?

As RAS proliferate they will inevitably start to encroach on duties traditionally performed by legacy platforms, in which governments and industry have invested millions. At the point in the future where RAS can fulfil the roles of these legacy platforms with greater assurance and at lower cost, lower risk, and with less manpower, Land forces will be faced with a strategic decision – is some deletion of legacy capability necessary?

Forces must not fall victim to the 'sunk-cost fallacy' – believing that past investment justifies future expenditure. RAS should not be viewed as an additional expense on top of existing capability, but as a force multiplier and a means of achieving more with less. At the right time, decisions will therefore need to be made about which platforms forces can live without. This is a vital step and one that is traditionally uncomfortable for Land forces, where signature platforms and equipment can almost come to define the identity of an army, trade or regiment. However, such disruption should be positively embraced if RAS can be proven to provide superior capability.

RAS will also displace operational roles currently carried out by humans. The ratio of humans to RAS on any specific operation will depend on how many robots a single person can control, but at a force level the number of humans should remain broadly unchanged, with those displaced moved into other essential roles, that only humans can realistically perform. The potential of RAS can only be harnessed effectively by operating as part of a human-machine team, augmenting the fundamental human component of Land capability. This is discussed in more detail in the Human Implications chapter of this report.

From concept to capability

To make the progression from concept to capability, the following steps will be necessary:

1 **Articulate the concept:** To secure the required funding for RAS development, government defence departments and industry will need to make the operational and economic case and articulate the national security advantages to policy-makers. They will need to argue that RAS generates additional combat mass without increasing financial or human capital costs, and that it enables resource to be scaled up rapidly in times of conflict.

2 **R&D, trials and experimentation:** Articulating the advantages of RAS and making the economic case to treasuries will require supporting evidence, amassed through progressive experimentation. Modelling, simulation and virtual experimentation will help to determine which capabilities are genuine force multipliers – reducing risk; reducing the cognitive burden on human soldiers; and increasing operational tempo.

3 **Cross-DLOD consultation:** Requirements must be defined in consultation with experts across all defence lines of development (DLOD): concepts and doctrine; equipment; information; infrastructure; logistics; organisation; personnel; training; and interoperability.

4 **Drafting requirements documents:** R&D, trials and experimentation will identify use cases and shape concepts of operations. These will in turn inform requirements, enabling governments to make informed procurement decisions and build an acquisition portfolio. The inherently evolutionary nature of RAS technology and capability suits a more flexible approach to acquisition, whereby the capability is continuously evolved and updated.

RAS have clear utility for major conflict and for a wide range of sub-threshold activity, including forward presence, homeland security and assistance to civil authorities. But it is not enough simply to make these claims; the benefits and limitations of RAS for these circumstances must also be articulated clearly and supported by credible evidence gained through trials and experimentation.

Early consideration of cross-Service applications will inevitably economise on scale and effort, leading to greater value for money. RAS platforms can easily by adapted to support amphibious assault/landing force and airfield force protection roles and cross-Service utility and integration should be an early aspect of all experimentation and acquisition.

The long game

The expectation must be set that RAS will not enable greater efficiency from day one. In fact, implementation may consume more manpower in the early days, as humans, processes, culture and infrastructure adapt to accommodate the new way of operating. However, initial investments of time and effort are essential precursors to reap two significant rewards: first, the effectiveness of the fighting force will improve almost immediately; and secondly, efficiency will increase over time until true force multiplication (reduced risk; reduced cognitive burden; increased tempo) is achieved.

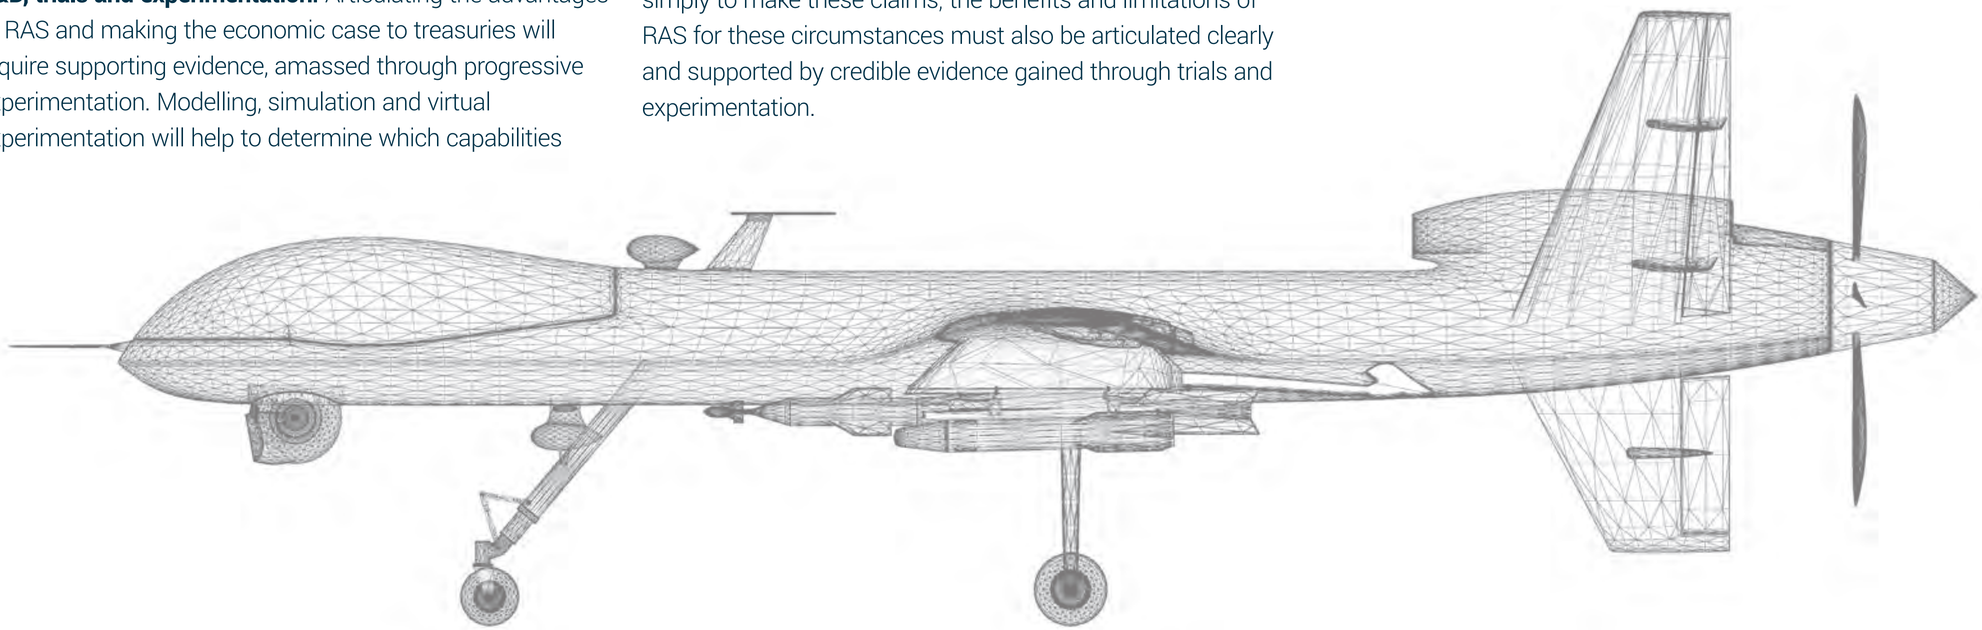

Implication:

Organisational Design

While much of the conversation around RAS is focused on technology, the way in which Land forces organise around that technology will be fundamental to success. Few other technological developments have required such seismic shifts in organisational design. One example is the progress from direct fire to indirect fire – the ability to strike from beyond line of sight – which saw the rapid expansion of artillery and created demand for new skills such as geometry and meteorology for targeting and ranging. Other examples include wireless communication, which enabled operations to be coordinated on an unprecedented geographic scale; and the adoption of flight into military endeavour, which ultimately led to the creation of air forces.

In each of these examples the change was not just a case of accommodating new technologies within existing operations, but radically transforming force structures (and even founding new ones) to place the technologies at the centre. Only by doing this could the capabilities be used to their full effect. The howitzer would have conferred no advantage if troops had continued fighting within line of sight.

The introduction of RAS will be similarly disruptive, but also equally game-changing if forces are optimally organised. To achieve this, militaries must consider factors such as the numbers, locations, functions and interactions of both human and robotic fighters.

Robotic wingmen: a thought experiment

Currently, a reconnaissance troop in a British battle group consists of eight manned vehicles, each containing a three-person crew. This configuration was conceived on the basis that it was needed to meet the battle group's reconnaissance requirement. Imagine now that each of those manned vehicles is augmented with two unmanned wingmen, which each increase the range of its coverage by 50 per cent. Assuming the original requirement was correctly defined and remains unchanged, the inclusion of robotic wingmen now creates surplus capability, raising a question as to whether either the manned vehicles and their human teams, or the unmanned vehicles, are best deployed elsewhere.

This illustrative example shows the type of decision that must be made when introducing RAS into Land operations. Does the aforementioned troop reduce its number of manned vehicles, removing human soldiers from harm's way? Or, does it retain the all vehicles augmented by the unmanned wingmen to bolster the existing capability?

Augmentation, not replacement

In the above scenario, reducing the number of manned vehicles from eight to four displaces twelve soldiers – but this should not mean they become redundant. With the reconnaissance requirement now met, the twelve can be redeployed into other roles that require uniquely human skills. This is critical, not just in mitigating the risk of future shortfalls in combat mass, but in compensating for the lack of mass that already exists today.

The ultimate aim of RAS is to generate advantage by increasing the combat mass and effect of the whole force, and this will not be achieved by simply replacing humans with machines to deliver the same outcomes. The opportunity in RAS is to deliver better outcomes using the same number of people. Relieving soldiers of dull, dirty, dangerous, demanding and difficult roles enables them to be deployed in a more flexible way and to greater effect. Losing them entirely creates distinct disadvantage at times when mass is needed for combat or humanitarian purposes. Similarly, there are elements of combat operations that will always require human decision-making. The aim of RAS is to give them extra capacity to enact those decisions.

Organising for RAS

In the early days of RAS deployment, autonomous systems are likely to be introduced as enhancements to the existing force structure. This may be at a platform level, using robotic wingmen; or at a subsystem level, with the automation of certain onboard functions in manned vehicles. In the earlier robotic wingmen scenario, the reconnaissance troop may remain at eight manned vehicles at the very start, augmented by two or three unmanned systems. This is a matter of building trust through experience. Live and virtual experimentation will develop new use cases, establish operating procedures, and ultimately equip users with the trust and knowledge of the unmanned technology to begin redeploying some of their manned assets elsewhere.

This section of the report has largely focused on the specific use case of unmanned wingmen within a reconnaissance troop, but similar decisions must be made at every level of an organisation.

On the front line, consideration must be given to factors such as the platforms' maintenance and sustainability requirements. For instance, a Warrior IFV is maintained by its crew – but who maintains the robots that are teamed with it? If all robots are required to return to a resupply base, where engineers and software developers are located, that base becomes an attractive target for enemy forces. Reorganisation must factor in the safety of maintenance crews. If they cannot be dispersed across the battlespace, their shared location must be adequately protected and fortified. Such decisions will go all the way to Division level, where the macro view will be critical in determining the geographical distribution of people and assets on the ground.

Finally, organisational design will be evolutionary and must remain continually under review, responding to factors such as technological advancements and changes in user trust.

Implication:

Interoperability and integration

Interoperability is the extent to which different platforms, people and systems are able to work together to fulfil a common mission. There are three recognised levels according to military strategists: deconflicted; compatible; and integrated. The UK Ministry of Defence (MOD) neatly defines these terms in a 2017 Joint Doctrine publication, in descending order of interoperability:

"Integrated means that forces are able to merge seamlessly and are interchangeable. Compatible means that forces can interact with each other in the same geographical battlespace in pursuit of a common goal. Deconflicted means that forces can co-exist but not interact with each other."

Joint Doctrine Publication 0-20, UK Land Power (UK Ministry of Defence, June 2017)

As interoperability of RAS technologies increases, so does their combined strength. If no consideration at all is given to interoperability, systems may clash with one another – not just negating advantage, but creating disadvantage. Deconflicting RAS is a neutral act. It prevents disadvantage without conferring a significant advantage at a system level. Interaction between compatible systems, however, does produce a positive effect. Unmanned platforms, instead of operating in isolation, can begin to share navigation data and contribute to a combined situational awareness picture. But the holy grail is full integration, which will allow multiple assets from multiple nations to work as one cohesive unit, multiplying combat force many times over. As adversaries' defence spending increases in comparison to that of NATO and allied countries, the ability to project greater combined force with less human resource will be vital. Allied countries already distribute conventional responsibilities at deconflicted and compatible levels. Some nations possess ground-based air defences; others have a nuclear deterrent; others have aircraft carriers – all of which can contribute to a fight against a common enemy.

RAS has a distinct advantage over conventional naval, aerial and land platforms, in that the rules have not yet been set. This provides an opportunity to meaningfully shape the capability and operating concepts of tomorrow – but to do so successfully will rely on understanding and overcoming the three major challenges of RAS interoperability:

Technical

The most obvious barrier to interoperability is technical incompatibility between systems. To function as a unit, each platform must at the very least interact with shared interfaces, such as command and control centres, but ideally with other platforms as well. They must use the compatible communications protocols, messaging formats, programming languages and software development standards. Other industries have achieved interoperability through the collaborative development of common standards and architectures, perhaps most notably telecommunications. Competing telecoms companies realised in the 1990s that web browsing using mobile handsets would depend on a common technical standard for accessing information over wireless networks. They formed a working group that led to the introduction of the Wireless Application Protocol (WAP) – and laid the foundation for the perpetually connected world we live in today. But while a similar requirement for RAS is obvious, it is not easy to implement. National priorities around machine learning and artificial intelligence will drive a diverse range of approaches to cyber security, information assurance, data-set management, safety accreditation and legal approval. All of these factors have implications for achieving technical interoperability and each factor will be further influenced further by the following challenges.

Doctrinal

If the technical challenge is about what RAS technology we want to use, the doctrinal challenge is about how we want to use it. How will capabilities and responsibilities be divided up between nations? Are all allies united around the same strategy and concepts of operations? How do nations want to use RAS within operations; are they prioritised around the same tactical functions? In what ways do they interact with manned assets? What are the appropriate human to machine ratios? What level of information sharing between platforms and systems is required to manoeuvre a RAS capability effectively? These decisions must be made early and cooperatively, and it is vital that they are subsequently stress-tested in collaborative exercises to expose unforeseen weaknesses. Finally, early collaborative experimentation is critical for fostering the right level of cooperation between nations' forces.

Cultural

The final challenge for interoperability is the operation of RAS technologies across different cultures. Human variables, such as belief systems, are harder to change than technologies and doctrines – and so the latter must be designed with the former in mind. This challenge manifests in the more metaphysical aspects of RAS, like how much decision-making autonomy is given to a machine learning-equipped platform, or the ethical basis for favouring one course of action over another. As an illustrative example from outside defence, driverless cars may be forced to prioritise whose lives to preserve in the event of an unavoidable collision. Country A's culture places greater importance on youth and potential; Country B's culture values age and experience. Should the car's decision on who to save be different depending on the country it is in at the time? Or should it maintain the standard of the country in which it was manufactured, regardless of its location at the time? These problems may sound abstract, but resolving them is fundamental to the realisation of RAS interoperability.

It is important to recognise that these three challenges are interlinked. Firstly, the cultural norms of a society influence concepts of operations. Secondly, RAS are programmed by humans, so the beliefs and cognitive biases of the programmers can become ingrained into the technology itself. For instance, some facial recognition tools developed in predominately white cultures fail to identify black features, having only been tested on white subjects. To avoid such oversights, technical, doctrinal and cultural challenges should be tackled collaboratively through user-centred development and joint experimentation.

Establishing interoperability

The widespread introduction of RAS technologies into the battlespace is only just beginning. Prohibitive technologies and practices are not yet entrenched, but there is a shrinking window in which to build in interoperability from the outset. We have identified three recommendations for getting it right first time:

1. **RAS capability must be international by design.** From the very beginning it must be conceived and built using the appropriate common standards and protocols, based on a mutual understanding of allies' doctrines and cultures.
2. **Nations must state their proposed levels of interoperability.** To what extent are they willing to share technology, tactics and intellectual property in order to benefit from the advantages of interoperability?
3. **Multinational experimentation is critical in proving concepts and troubleshooting operational challenges.** Collaborative training is also vital in exposing and bridging skills gaps, ensuring interoperability of human teams, not just technologies.

It is also important to note that as interoperability increases, security and assurance can be eroded if steps are not taken to mitigate the effects. We discuss this in more detail in the 'security implications' section of this report.

Interoperability is not just a key technical requirement, but a founding principle on which any RAS strategy should be based. There are varying degrees of interoperability, and numerous challenges in attaining it at the highest levels – but the greater the extent to which technologies, people, forces and countries can work together, the greater the effect.

Implication:

Platform Selection

Although building a RAS capability is chiefly an information architecture task, the effects of the capability are still ultimately delivered in a physical battlespace by physical platforms. The physicality of the Land domain is ingrained in the psyche of manufacturers, buyers and end users. At present, design and procurement decisions are made in response to the specific physical requirements of the current conflict – whether urban, rural, desert or jungle. When the nature of the battlespace changes, a new set of requirements is drafted and a new design and acquisition cycle begins.

There is a tendency to address each new use case with a brand new piece of physical technology – a platform or a sensor designed for a niche and specific task. This is unsustainable. Fielding unique robots for "edge cases" within reconnaissance, rescue, gap crossing, resupply, offensive strike, IED disposal and myriad other missions will multiply the cost and logistic burden without increasing the effectiveness of the capability. The aim is precisely the opposite – to multiply force without increasing complexity or cost.

Platforms must be designed for compatibility with a common communication and C2 architecture, not separate architectures developed in isolation to accommodate disparate platforms. The former encourages the integration of multiple platforms into a powerful unified capability, while the latter severs the connections on which RAS depends.

Shifting the mindset away from these established norms will be challenging, but it is crucial if RAS are to fulfil their economic and strategic promises. Platform selection as the primary acquisition decision at the outset cannot be done without diminishing the focus upon the other equally vital implications outlined in this report. But what is the alternative?

A new model

To use RAS in a truly transformational way we must think beyond specific use cases – even beyond individual missions and campaigns – and consider RAS platform selection at a whole-force level. This begins with using simulation and modelling to test potential use cases for likely conceivable scenarios (present and future) and mapping them against combinations of platform types and subsystems that could fulfil the demands. From this process, it will be possible to identify the smallest number of platforms capable of carrying out the greatest number of missions at the lowest cost. The selected platforms can then be adapted and plugged into the established architecture as needed.

Platform characteristics

Many RAS platforms will be small, inexpensive and expendable – manufactured and deployed in high numbers but at a lower cost than expensive manned platforms. But there will still be a need for more capable, sophisticated RAS – the cost of which will begin to encroach on budgets currently allocated to manned platforms. This will force difficult choices about the need to disinvest in legacy platforms, or else the cost saving benefits of a RAS strategy will not be realised – RAS cannot be funded in addition to the core equipment plan. All of these choices will be need to be informed by the simulation and modelling of scenarios mentioned above.

When reducing the number of available platform types it becomes more important that each one can adapt to fulfil a wider range of roles. The key criterion for platform selection is therefore adaptability. An unmanned fleet may comprise six or seven base ground platforms and range of air platforms, each built on modular architecture, so that subsystems can be introduced or swapped out as the mission dictates. A typical inventory of base platforms may consist of:

Tracked vehicles

Small Tracked Unmanned Ground Vehicle (UGV) (1-2 tonnes)

- Suitable for support to Light Forces with mobility to follow dismounted troops at speeds <24kph

Light Tracked UGV (3-5 tonnes)

- Carrying a range of payloads and effectors (e.g. up to 30mm) and capable of matching manned tactical vehicle mobility

Medium Tracked UGV (8-12 tonnes)

- Carrying heavier payloads, matching tactical vehicle mobility and with survivability against some ballistic threats

Heavy Tracked UGV (up to 20 tonnes)

- Payloads capacity to include direct fire weapons capable of engaging adversary armour. Survivability appropriate to its role

Wheeled vehicles

Light, Wheeled UGV (1-2 tonnes)

- Support to Light Forces with higher speed than tracked equivalents, <40kph

Large Wheeled UGV (4-8 tonnes)

- Effectively replacing existing logistics trucks with comparable mobility characteristics; flat load-bed that can be adapted to fulfil multiple roles – e.g. pallet-hook, ISTAR suites, indirect fire systems, drone-launch systems. Survivability limited to protection of critical sensors and sub-systems

Medium Wheeled UGV (4-6 tonnes)

- Mid-sized platform with adaptable load bed to take a variety of mission specific payloads. Suited for use in urban/semi environments where tracked platforms are not appropriate. Mobility sufficient to keep up with tactical vehicles and adaptable survivability dependent upon threat levels

Unmanned Aerial Vehicles

Nano-UAS

- Organic UAS capability for Light forces and other Close Combat personnel to provide tactical ISTAR

Class 1 Rotary Wing UAS

- Organic UAS capability for a range of combat, combat support, combat service support and CIS roles to carry specialist payloads aloft; where precision of control and flight characteristics are a priority

Class 1 Fixed Wing UAS

- Organic UAS capability for a range of combat, combat support, combat service support and CIS roles to carry specialist payloads aloft; where range and endurance are a priority

Medium Lift UAS

- Capable of lifting 25kg to 75kg with a variety of designs viable – including rotary wing, fixed wing and novel design (e.g. paramotor). Suitable for resupply and ISTAR tasks

Heavy Lift UAS

- Capable of lifting 75kg to >150kg with a variety of designs viable – including rotary wing, fixed wing and novel design (e.g. paramotor). Suitable for resupply, ISTAR and delivery of kinetic effects

Medium Altitude Long Endurance (MALE) Fixedwing UAS

- A well-understood class of platform already in service with most nations. Primarily used for ISTAR in the deep battlespace with sensing and kinetic effects integrated

In the case of ground vehicles, it is worth noting that many of the existing logistic support vehicles, tanks and armoured fighting vehicles currently performing these roles can be retrofitted with autonomous systems, meaning new fleets do not necessarily need to be designed, built and procured from the ground up. However, the full potential of RAS will be fulfilled in the longer term by platforms designed specifically for autonomy. Development of any platform (including those that fly) necessitates trade-offs between mobility, firepower and protection. In manned platforms, the balance tends to skew toward the protection of the human operator. With the human removed, new vehicles can be designed with mobility and firepower at the fore.

Advantages

Simplifying the portfolio of platforms reduces complexity in fleet composition and allows a greater focus on what enables those platforms to perform. Having a limited number of modular, multi-mission platforms supports the exploration of use cases through experimentation, innovation and acquisition. Instead of bringing multiple platforms to the testing range for a single mission scenario, the user can trial each one against a host of missions, increasing their availability and maximising the value of the time spent on the range.

The advantages of this approach go beyond those enjoyed by the end user – there is also a strong economic case for both manufacturer and buyer. The more niche a platform is, the fewer a manufacturer is able to sell. A highly adaptable base platform is manufactured and sold in higher numbers, allowing the business and its customers to benefit from the economy of both scale and scope. Achieving this economy is crucial to the underpinning logic of RAS; as soon as RAS become niche and exquisite, the investment case is much harder to justify.

What about manned platforms?

Manned platforms which operate within a future new RAS-enabled system of systems will inevitably require adaptation to allow them to operate effectively within it. Physical architecture of vehicles may need to change to be able to accommodate UxV operators and control stations. Information architectures are also likely to need to take into account RAS and any communication between human-crewed and robotic platforms, ensuring data can be transmitted, received and stored securely and safely.

Implication:

Security and assurance

Manned systems rely on their human operators to gather information, process it and act upon it. They have therefore evolved to safeguard their operators and provide backups (such as co-drivers in vehicles) that allow continued operation in the event of emergency deterioration of the security environment. As we increase the physical distance between machines and the human members of the team we must turn our attention to similarly protecting sensors, data links and other critical systems from attack, and backing them up in case of failure.

Cyber and Electromagnetic Activities (CEMA)

The most distinct threat for RAS is that posed by cyber and electromagnetic (EM) attacks. Autonomous navigation relies chiefly on sensors that harness different sections of the electromagnetic spectrum. A global navigation satellite system (GNSS) is vulnerable to signal interference, jamming and spoofing. Multi-constellation, multi-service receivers protect against these by identifying and tracking the strongest and safest available frequency from any satellite. Encrypted signals are used to deny unauthorised access and prevent spoofing.

In the visual part of the EM spectrum, cameras used for navigation or object characterisation are vulnerable to dazzling, which may be conducted using something as rudimentary as a laser pointer. Physical shielding can protect optical sensors from such attacks – but most importantly, no critical system should rely on any single sensor. An autonomous vehicle must adopt a multispectral approach, operating using multiple sensors on different parts of the electromagnetic spectrum. If GNSS access is denied, the vehicle can navigate using optical sensors. If its optical sensors are damaged, it can find its way using radar or LIDAR. The key is to build in multiple layers of redundancy, so that the failure of a sensor – or a combination of sensors – does not result in the failure of the whole system.

Sensors are not the only vulnerability. Data gathered by the sensors must be communicated wirelessly, which exposes another threat vector. Hackers may intercept or interfere with communications to gain access to the network, steal data, or give false commands. However, the risk is the same as for any networked system, and so the tools and techniques for mitigating it already exist. The platform must also have the ability to operate with degrees of degraded capability until a point where it is compelled to return to base, either autonomously or under the remote control of the operator.

Lastly, the collected and transmitted data must be processed, which may also be a target for attacks. If an attacker understands the machine learning model used to power a platform's visual classification system they can spoof it, causing it to misidentify objects. Algorithm and data security is therefore as vital as traditional cyber-security. Just as organisations employ penetration testers to stress-test their cyber defences, militaries must employ people to attack their AI systems and expose weaknesses before enemies can exploit them.

Physical threats

Cyber and electromagnetic activities are not the only security concerns when operating RAS. Vehicles and other platforms may be targeted with direct fire that destroys sensors or masts. Physical armour, the layers of redundancy provided by multiple sensors, and an understanding of the platform's ability to operate at reduced capability are all part of the mitigation suite.

Just as with a manned military vehicle, an unmanned vehicle may be subject to counter-mobility tactics, such as mines, ditches or roadblocks. Countering these requires a sophisticated sense-and-avoid capability, built on evidence acquired through testing about which obstacles the vehicle is able to tackle and which must be evaded.

Finally, there must be systems in place to guard against capture or physical interference. The first strategy for preventing capture is one of avoidance and stealth. A strong situational awareness picture assists the platform in avoiding enemy combatants, while silent manoeuvre and silent watch capabilities minimise the risk of discovery within close proximity of adversaries' locations.

In case capture cannot be avoided, the platform itself should store as little sensitive data as possible which the adversary could steal and exploit to their advantage.

Physical interference may include attaching improvised explosive devices or surveillance devices to the vehicle, which are then transported back into the unit's forward operating base. Onboard sensor systems can help to detect signs of physical interference, although procedures should also be adapted to factor in this threat, such as establishing checkpoints with scanners at the entrances to bases.

Securing RAS

RAS are not necessarily more or less vulnerable to attack than manned systems, but the threats are different and must therefore be countered using bespoke security strategies. While the vulnerabilities of human system operators are well understood, those of sensors and autonomous systems are comparatively new and constantly evolving. It is therefore vital to test and experiment with RAS to identify weaknesses, protect against attack, and build in multiple layers of redundancy.

Assuring RAS as a capability

It is vital to ensure RAS does not endanger users or civilians and their property; this also affects user trust in the technology. Any safety related incident will fuel suspicion of RAS technology, leading to its underutilisation. Fielding a safe capability will require building assurance through evidence which can be gathered via simulation and modelling, progressing to live trials, and finally a phased introduction into service. Earlier digital twinning and modelling can start to build safety case and regulatory body evidence early.

Regulatory implications are addressed in the opening sections of this report, but it is worth reiterating their importance, as getting regulations wrong can set progress back years. Regulations have always adapted to accommodate new technologies, but the modern pace of technological progress means regulators can struggle to keep up. This failure cannot be attributed solely to the regulatory bodies – developers must able to provide evidence of compliance with existing regulations, and engage fully and early with regulators to build a meaningful case for change.

Recommendations

Development and execution of a Land forces RAS strategy is the perfect opportunity to apply the philosophy of Prototype Warfare.

Prototype Warfare champions early, safe and well-thought-out experimentation, including on operations, as a way to drive pace and progress in capability development and adoption.

This approach will be key to the successful exploitation of RAS technology in future integrated military human-machine teams. Only by doing so will risks and opportunities be addressed effectively and the operational and cost effectiveness potential be realised.

This report has highlighted a number of areas that any military RAS strategy must consider. In summary, we offer four key recommendations for change, each with a practical first step:

1 Think holistically, based on operating concepts

Any Land forces RAS strategy needs to be comprehensive across all lines of development and components of capability. Envisioning how RAS can be incorporated into military use is essential and needs to be expressed in an operating concept hypothesis which forms the basis of experimentation and force development. Information architectures sit right at the heart of any effective RAS strategy but changes to organisational design and platform choices will impact how those systems are designed and deployed. A Land forces RAS strategy must avoid addressing issues in silos but plan to create a human-machine teamed system of systems from the outset.

2 Adopt a portfolio approach

Adoption of RAS capability and execution of a Land forces RAS strategy are significant undertakings and, over a 10 to 15-year time period, will represent a major shift from current capability. The early steps are vital and the complexities involved mean it cannot be delivered through a single programme. It will require investment and carefully planned, connected and integrated steps – some of which will be experimental – across a number of programmes concurrently, linked by a common and defined (albeit by necessity flexible) RAS outcome. That does not mean the speed of delivery should be slow; on the contrary, it needs to deliver useable Land forces capability at the pace of relevance. A 'portfolio' approach will ensure the required system of systems capability integration, accommodate cross-programme exploitation of learning from experimentation and can engender better system safety and value for money.

3 Seek to augment, not replace

RAS capability should not be seen simply as a means of replacing existing personnel, nor should it be a case of simply adding new robotics systems into existing 'human teams'. RAS capability should supplement humans in future more sophisticated operating concept-based human-machine teams. As the report outlines, some tasks fundamentally require humans and organisational design will need to adjust to better reflect the type of structures that best harness the potential RAS technologies can offer. The most difficult question is how to establish what those structures need to look like at the outset. Over time and with supporting evidence from experimentation and experience, judgments can be made on the respective roles of humans and machines.

4 Embrace experimentation

Successful adoption of RAS technology by Land forces will require a major element of experimentation to help determine optimum roles, tactics, structures and support needs across various tactical functions. As espoused in our earlier Prototype Warfare work, this includes a willingness to experiment on live operations. Experimentation must be carefully and effectively integrated, planned and conducted to ensure that useful conclusions can be drawn safety and securely.

To end, we have not sought to describe some distant hypothetical idea that can be ruminated upon in clubs and bars for a few years. RAS technology is 'now technology', in use by adversaries and with immediate military utility for Land forces; it will evolve over time. It is hard to envisage a situation in 10-15 years where RAS are not integral to Land forces capability, with their potential harnessed within sophisticated human-machine teams.

While the cross-lines of development/components of capability journey to get to that point may seem daunting, we advocate an experimentation and acquisition portfolio approach to ensure that investment of inevitably tight resources achieves maximum benefit and capability is fully integrated, while leaving the freedom to explore and adjust along the way. Only by adopting this approach – and in particular, recognising that RAS capability adoption is less about the platform, and more about the other lines of development/components of capability – do we believe that Land forces will be able to realise the full potential of RAS at a pace of relevance. The early steps are key and action is needed now.

Cody Technology Park
Ively Road, Farnborough
Hampshire, GU14 0LX
United Kingdom
+44 (0)1252 392000
prototypewarfare@QinetiQ.com
www.QinetiQ.com

QINETIQ/20/02976