



DARPA, America's Defense Research Lab

Kim Kyung-jin

AI Library · kimkj.com · 2026

Table of Contents

1. Preface

Part I Organization and Budget

2. Chapter 1 A Vessel for Mission: The Reorganization of MXO and IPTO

Part I. Organization and Budgets

3. Chapter 2. The Direction Budgets Reveal: From Basic Research to Advanced Technology

Part I. Organization and Budget

4. Chapter 3. The DARPA Way of Solving Problems: Clear Goals Over Good Science

Part II. AI and the Cyber Battlefield

5. Chapter 4. AlxCC and Automated Code Security

Part II AI and the Cyber Battlefield

6. Chapter 5 Trustworthy AI: AI Forge, SABER, and GARD

7. Chapter 6 Cyber-Physical Systems and Defense Cooperation

Part III Unmanned Systems and Hardware

8. Chapter 7 Intelligence Inside Hardware

Part III. Unmanned Systems and Hardware

9. Chapter 8. RACER and the Move Toward Combat-Ready Off-Road Autonomy

Part 3 Unmanned Systems and Hardware

10. Chapter 9 Aircraft Autonomy: Helicopter Transition, X-76, LongShot

Part III Unmanned Systems and Hardware

11. Chapter 10 Next-Generation Microelectronics and Photonics

12. Chapter 11 Space Robotic Servicing and Cybersecurity

Part 4 Quantum, Life, and Battlefield Medicine

13. Chapter 12 Quantum: Experiments That Separate Hype from Reality

Part IV Quantum, Life, and Battlefield Medicine

14. Chapter 13 Biological Computers and Optogenetics

15. Chapter 14 Battlefield Medical Autonomy: Triage, VITAL, RAPIID

16. Chapter 15 Strategic Materials and Bio-Based Manufacturing

Part V. Transition and Korea's Question

17. Chapter 16. Small Businesses and Startup Transition

Part V Transition and Korea's Question

18. Chapter 17 Academic Links and DARPAConnect

Part V. Transition and Korea's Question

19. Chapter 18. The Starting Point for a Korean DARPA

Preface

In a conference room near Washington, a one-page problem definition is placed on a desk. Some read it as a budget request. Some read it as a research announcement. Inside DARPA, it is read a little differently. It is an order to turn a capability that does not yet exist into an experiment that can be grasped within a few years.

This book follows the grammar of that order. A competition where AI fixes code. Unmanned vehicles crossing land without roads. Computing that keeps power grids and supply chains standing. Biological devices that try to calculate with cells. Plans to turn discarded biomass into strategic material. I did not collect only shining success stories. Plans are read as plans, results as results, and DARPA's own words are read together with the limits of their source.

What often goes missing when DARPA is discussed in Korea is not the organization's name, but the way it frames a problem. Before asking how much money was spent, there are other questions to ask. What failure did it accept? What goal did it define with almost unforgiving clarity? How did that goal change the behavior of researchers? If we hold on to these questions while reading, DARPA becomes not an institution in a distant country, but a mirror held up to our own research organizations.

Attorney Kim Kyung-jin

Chapter 1 A Vessel for Mission: The Reorganization of MXO and IPTO

What the renamed organizations actually changed

On May 20, 2026, a short post appeared on the DARPA website. Its title would read in Korean as something like "What's in a name?" The subject was an organizational reshuffling. Two of DARPA's six technical offices had changed their signs. The post was brief. Yet that brief post stands at the starting line for several programs this book will follow.

Changing a sign usually does not make big news. The world does not change because a company renames one department. Research institutions are a little different. At a place like DARPA, an office is not just a box on an organization chart. Several research programs sit inside it. Each program carries years of budget, a program manager, and outside research teams. Money flows along a direction, and the program manager stakes a name on that direction. An office is a vessel for mission. When the shape of the vessel changes, what it holds changes too. In such a place, a name comes close to a declaration of what the organization intends to do next. A promise about which problem it will solve sits inside the name. DARPA's decision to change two signs at once reads as a signal that it intends to move its center of gravity.

The first office to change was the Microsystems Technology Office, or MTO. For a long time, this office dealt with very small things. Semiconductor devices, tiny components, circuits placed on chips. Its stage was a world smaller than a fingernail, where the task was to draw out better performance. If equipment used on the battlefield is to become smarter, the chips inside it must become smarter first. If radar is to see farther and communications gear is to carry signals more densely, devices underneath must support that work. MTO handled that foundation. Its new name is MXO. Written out, it is the Multi-X Office.

The letter X in the middle is the point. DARPA explains this X as if it were a multiplication sign. It means multiplying scale, multiplying disciplinary boundaries, and multiplying effect in real operations. It may sound like wordplay. The direction inside it is fairly clear. If the old MTO focused on individual components one by one, MXO is turning its attention toward the integrated capability those components produce together. It has moved from the fight to squeeze another 1 percent of performance out of a component to the question of what several components can do when combined. Whitney Mason, the director who leads MXO, still called the office a place that deals with "small things that have huge impact." The office keeps its character of handling small things, but redraws how far those small things are meant to reach.

The starting point also moves farther down. From discovering new materials at the scale of atoms and molecules to turning those materials into capabilities that change future battlefields and the economy, the office is saying it will take in that entire long chain as its concern. It means placing the

very front end of the laboratory, where materials are found, and the far back end, where the results flow into weapons systems or industry, in one vessel. The calculation is that one office will take responsibility for the whole distance between discovery and application. From small components to large capability. That one line is the point of the first reorganization.

The second office to change was the Information Innovation Office, or I2O. This office had led research in artificial intelligence and information technology. Work under this office included having machines inspect code that no human could fully read, and building artificial intelligence that could explain why it made a given judgment. Its new name is IPTO, the Information Processing Technology Office. This name was not newly invented. It was revived. It is an old name used in the 1960s and 1970s. Early computing research that became the foundation of the modern digital world sprouted under that name. DARPA has taken that old sign back out and hung it again.

There is a reason attached to the return to the old name. Patrick Lincoln, who leads IPTO, explained that the renaming means looking again at the power and resilience of computing from first principles. To invent the next generation of technologies that will support national security, the office must not stop at touching up applications in front of it. It must go down to the foundations underneath. Improving an artificial intelligence service that appears on a screen so it feels a little smoother is one layer. Asking again about the principles of computation that support that service is another. The first widens a road already laid. The second opens a new road. The name IPTO aims at the second.

The target has widened as well. The office's work now includes foreseeing and reducing risk in large, complex, tangled systems such as supply chains and power grids, what DARPA calls megasystems. A power grid moves like one thing, with power plants, transmission lines, and substations tangled into thousands of paths. If one place collapses, it is hard to predict how far the shock will spread. IPTO's goal has become building the strength, or resilience, that allows such large systems to keep holding their mission even when attacked or broken. One more element is attached. Beyond placing security programs on top of finished systems, the office points to inherent cybersecurity, in which security is planted inside from the first stage of design. It is closer to the idea of building the walls so they cannot be breached, rather than adding locks after the house is finished.

Set the two reorganizations side by side, and on the surface their directions may seem opposite. One expands from small components toward large capability. The other digs down from visible artificial intelligence applications into the underlying principles of computing. One seems to widen upward, the other to dig downward. One thread runs through both. They look beyond individual pieces to the whole system those pieces form together, and they try to make that system hold even when shaken. From components to systems, from applications to foundations. The direction of the gaze is similar. When the two names are placed on top of each other, one can dimly see where DARPA now wants to put weight.

A pause is needed here. Everything described so far comes from DARPA's own explanation on its website. It is primary material showing how the organization understands its own reorganization. It is not an externally verified performance report. This distinction matters. Changing a name is easy. One press release can do it. The hard part comes next. Does the flow of money actually change? Is the list of research tasks rebuilt to match the new direction? Do people and organizations move that way? A new sign can go up while the work inside remains the same. Companies often change department names while the work stays unchanged for years.

So I will not count the renaming itself as an achievement. I am trying to keep, from the first scene, the promise not to draw DARPA only as a success myth. This reorganization is best read as a compass showing where DARPA wants to point. Whether the ship actually went where the compass pointed must be checked by following budget documents and individual programs one by one. That verification continues in the chapters that follow. The next chapter looks at the flow of money. After that, the programs are taken apart one by one.

For Korean readers, this scene is also a mirror. We too are used to making new organizations and giving them names. When we decide to change something, we first set up a committee, open a center, and hang a sign. Years later, we rarely look back calmly to ask whether the budget, people, and tasks behind that sign really moved in the new direction. The first day of a reorganization becomes an article. Whether the reorganization succeeded is quietly buried. Watching how DARPA explains its own reorganization makes us ask again how we should examine our own organizations. The first exercise in separating what is worth admiring from what should not be copied as is begins here.

What did the renamed organizations actually change? With the material now in hand, this is as far as one can go. DARPA has declared that it will put integrated capability ahead of small components, and the foundations of computing and the resilience of giant systems ahead of immediate applications. No one can yet guarantee what that declaration has led to inside laboratories. The answer is not in the name. It is in the money and tasks flowing beneath it.

The Research Method Indicated by MXO and IPTO

The previous section examined what the names declared. Now it is time to look at what shape the research is trying to take under that declaration. A name is a doorplate. The real content is what questions people ask inside the door.

DARPA is not one large laboratory. It is an organization made up of several technical offices with different missions. Each office has people called program managers. During terms of several years, they put their own ideas on the line and run programs. DARPA has long presented itself as a place that funds high-risk, high-return work if it can change the field, even when the chance of success is below half. A change in the office's name means that the question asked in that room, and the way it is asked, has been reset.

Look first at MXO. In the old MTO days, the question in this room was usually narrowed to one thing. How small can this chip be made? How much can the speed of this device be raised? The center of the stage was the work of squeezing hardware the size of a fingernail to its limits. In the new name MXO, the X in the middle widens that question. Director Mason describes this X as a multiplier, a value that multiplies. What it multiplies matters. The scale of research, the kinds of disciplines, and operational effects on the battlefield. It means making a larger result by multiplying those three together.

The core is contained in one sentence from Director Mason. Future advantage will not come only from advancing a single technology in isolation. However excellent a chip may be, it cannot exert power on the battlefield if it is not connected to other technologies. So MXO places multidisciplinary research at the front. It is abandoning the pattern in which materials science, physics, and engineering sit in separate rooms. The old method of pushing one component's performance to the limit and the new method of tying different fields together into one capability aim at different things. This is the point to watch. Such a direction changes the very reason a researcher receives funding. It is no longer enough to explain the work by producing the best component inside one's own field. The case for funding must also show how it interlocks with another discipline in the next room, and where the result will reach.

The range of research is also set wide. At one end is the search for new materials at the scale of atoms and molecules. It is a world invisible to the naked eye. At the other end are components, tools, and processes that can change the shape of future battlefields and economies. MXO says it will connect these two ends as one flow. DARPA calls this work of carrying foundational discovery at the atomic level to military advantage the making of architectures. It means not stopping at inventing components, but drawing the blueprint for how those components will connect and how they will be used. The move is from inventor to architect. From a policy viewpoint, this part is not ordinary. A country's ability to make excellent components and its ability to connect those components into a weapons system or an industry are different muscles. Strength in the first muscle does not automatically build the second. MXO's reorganization reads as a signal that weight will move from the first muscle toward the second.

Now turn to IPTO. History sits on this office's name. The new name IPTO revives the old name DARPA used in the 1960s and 1970s, when it was pushing computing research forward. The fact that a past name has been brought back is itself a message. It signals a return to the roots of information processing rather than to flashy applications.

Director Patrick Lincoln made the direction clear. The office will focus on the first principles of computing performance and resilience. The phrase first principles may sound difficult, but the meaning is not. It is the attitude of reexamining the principles at the bottom instead of placing one answer on top of someone else's answer. As artificial intelligence spreads quickly and competition around information becomes more refined, the judgment beneath this move is that research into

fundamentals, not technologies that polish the surface, will decide the next round.

The size of the research target has changed too. In the past, an individual artificial intelligence model or a single piece of software was the unit of research. IPTO raises that unit to national-scale supersystems. The examples DARPA directly cited in its official article are supply chains and power grids. The paths along which goods flow, the paths along which electricity flows. Foreseeing what happens when these vast networks are shaken by attack or accident, reducing that risk, and making sure that functions essential to the mission survive any blow. That is what resilience means. The direction goes beyond asking whether one program runs well. It asks whether the network supporting the whole country can avoid being cut.

Return to one concept passed over in the previous section. It is inherent cybersecurity. IPTO pushes aside the method of adding security on top of existing systems. Instead, it aims to plant security and privacy at the base of design so that the information processing structure itself repels attacks by nature or recovers on its own. It is research that goes beyond cleaning up after an accident and tries to remove from the start the gaps through which an accident can enter. This idea matches exactly with the IPTO name, which places performance and resilience side by side. The structure assigns one office both the work of making things fast and the work of making them hold.

Set the two offices side by side, and their directions overlap. MXO starts from small units of matter and moves upward. IPTO starts from the fundamental principles of information processing and broadens to national-scale systems. Their starting points are opposite, but both leave isolation behind and move toward connection and integration. On the hardware side, components are tied into architectures. On the software side, individual programs are expanded into national systems. DARPA's long-standing posture remains here: even if the chance of success is below half, it will place a bet if the work can change the field.

The question this turn poses to policy is fairly heavy. When a research organization moves its gaze from components to architectures, from individual software to national systems, the yardstick for evaluating that organization must change too. The performance of one component is easy to measure with numbers. How many nanometers it is, how many times faster it is, comes out at once. Whether an architecture tying several technologies together is well made, or whether the resilience of a giant system such as a power grid has actually increased, cannot be captured in numbers overnight. It is the kind of achievement whose presence or absence becomes visible only when crisis arrives. That makes research in this direction hard to boast about too early, and also hard to dismiss too quickly as unproductive. The evaluation timetable itself lengthens. How an organization endures the gap between short terms and long results is the real test of this kind of reorganization.

The line should be drawn clearly. The fact that the name changed does not itself mean achievement. The picture presented by MXO and IPTO is still mostly a plan and a direction. Future results will tell whether integrated architectures actually changed the battlefield, and whether the resilience of supersystems held in crisis. What can be read now is only a change in method. DARPA

has moved its gaze from components to architectures, from individual software to national-scale systems. Where that gaze will land and what it will make must be checked individually, starting in the next chapter.

How to Explain an Organizational Reorganization to Korean Readers

A policy official skims a morning briefing and stops at one line. "DARPA changed MTO to MXO." After a moment, the official turns the page. A changed office abbreviation does not seem worth noting. What weight could there be in three letters changing into three other letters? A common failure in explaining organizational reform to Korean readers begins in this scene. The fact of the name change is conveyed, but the reason for the change is left out. The reader gains one fact and loses the whole meaning.

So this section has a different texture from the two before it. Rather than adding new material, it looks from the position of the explainer at how the MXO and IPTO reorganizations should be unpacked for educated Korean readers and policy practitioners. It asks why this angle is needed. To Korean readers, a reorganization inside a U.S. defense organization can easily look like someone else's household affairs. It is the explainer's job to make those household affairs readable as our problem.

One point should be made in advance. I will not portray DARPA only as an object of admiration. No one yet knows whether the renamed organizations will actually produce results. A reorganization is a declaration of intent, not proof of outcome. The explainer's task has two sides. Explain faithfully the meaning contained in the reorganization, while not inflating that meaning as if it were already an accomplished result.

The first adjustment is to decode the alphabet. The X inside the name is the core. Director Mason described this X not as one component, but as a multiplier. It means a place that multiplies scale, academic fields, and operational effects. For Korean readers, it can be explained this way. If the old MTO asked how small and fast a fingernail-sized chip could be made, MXO asks how many times greater an effect that chip can produce on the battlefield when it meets other technologies. The level of the question has moved from component to system. It helps to read X like a multiplication sign. It signals that what will be multiplied by what remains open. When the question is about a component, the answer narrows to one path. How many nanometers can this chip be reduced to? When the question is about a system, the answer branches out. Where should this technology be placed, and with what other technology, to make a major difference on the battlefield? The name MXO writes those branches into the organization's name from the start.

One analogy may be used, but the boundary must be kept: an analogy helps understanding, but it is not fact itself. A good analogy brings readers to the door. Evidence must confirm what is actually inside the door. If the two are mixed, the explanation may feel successful while leaving no fact behind. There is one temptation the explainer must resist. It is to package the new name on the

screen as if it were already a major achievement. MXO is still a reorganized office. What that office has achieved is the next question. Explain the meaning of the name, but do not promise a result that has not arrived.

The second adjustment is to read the return of the old name. The change from I2O to IPTO is not the same as hanging a new sign. It is a case of taking out an old sign and putting it back up. IPTO was the name of an organization that actually existed inside DARPA in the 1960s and 1970s. Director Lincoln said he would return to the concerns of that period and focus again on the first principles of computing performance and resilience. When explaining this return to Korean readers, it is better to place the question first. Why go back to a name from half a century ago? Is it nostalgia for a time when things were going well? Read that way, it is only half right. It is closer to a company bringing back the mission from its founding days. It is a choice to return a scattered focus to its original place. That is why the phrase first principles carries weight. It means going down to the bottom questions of how fast computation is and how well a system can hold, instead of placing one more flashy application on top.

The word resilience also needs explanation. It may sound like an unfamiliar Sino-Korean term, but its meaning is not difficult. It points to the strength that keeps the whole from collapsing when one place fails, and that allows it to rise again after being hit. It is not visible in ordinary times. Its presence or absence appears only when an attack comes or one component breaks. IPTO's placement of performance and resilience side by side is a declaration that one office will take on both making systems fast and making them hold. The target has grown that much larger. Beyond one individual technology, the direction now points toward protecting supersystems tangled at the scale of a country, such as supply chains and power grids. If a power grid collapses, hospitals, communications, and water supply stop together. Making that chain endure has become the job of the new IPTO. The name has gone back to the past, but the size of the problem has grown. A Korean practitioner will naturally add a question of their own. How much collapse can our power grid and supply chain endure?

The third adjustment is the most delicate part of this section. It is to handle honestly the temptation of comparison. When explaining organizational reform, one quickly wants to draw a comparison table. The United States gathered scattered artificial intelligence projects into one place, another country split a large unit into several, and Korea expanded a small center into an institute. Such contrasts are powerful in a classroom. The story that each country reached a different answer because each had different circumstances falls neatly into place.

I will not draw that comparison table carelessly. There is a reason. The U.S. Department of Defense's case of integration, the division of units in China's People's Liberation Army, and the expanded reorganization of Korea's Agency for Defense Development have already been discussed in the author's earlier book. That book described those scenes on its own evidence. This manuscript begins with a different promise. New factual claims will be tied one by one to the official sources

secured for this work. Figures such as budget numbers, personnel size, and founding dates can be raised as central claims only when their sources can be checked again in the body of the text. The MXO and IPTO reorganization is grounded in a document DARPA released under its own name in May 2026. The detailed numbers for other countries, by contrast, sit outside the source ledger for this manuscript. The explainer should not hide this boundary from readers. A smooth three-part contrast of "the United States did this, China did that, Korea did something else" sounds good, but when the weight of evidence differs, it is hard to place them side by side in the same table. Numbers are weapons of persuasion. That is why only numbers that can be supported from behind should be brought out.

Putting away the comparison table does not make the Korean story disappear. It only moves its place. Instead of borrowing other countries' numbers to measure our condition, the explanation shifts toward drawing questions useful to us from the method of reorganization itself. The questions DARPA asked while changing its names can be moved into our hands as they are. What does the name of our organization ask? Does that name ask about components, or does it ask about systems? Are our supersystems built to endure even when one place collapses? These questions can be asked without the budget numbers of another country. And in fact, such questions last longer.

This honesty is useful to Korean practitioners. It is exactly the sense needed when judging organizational reform. Some reorganizations only change the doorplate. Others rewrite the mission itself. MXO's X and IPTO's return are closer to the latter. In changing the name, each office also changed what it asks. Of course, time must pass before we know whether those new questions lead to good answers. Giving that judgment in advance is not the explainer's role. When news arrives that some Korean organization has changed its name, the same order of questioning should be used. Has the problem it asks changed, or was only the sign replaced? And is it ready to answer the changed question? Redrawing the boxes on an organization chart and changing the questions that people inside those boxes ask every day do not carry the same weight. The first can be done in a day. The second takes years.

Explaining organizational reform consists of three adjustments. Decode the meaning of the name, read the return of the old name, and draw the boundary of comparison honestly. The last is the hardest of the three. It requires putting away a comparison table one wants to show off. Still, that place must be held if the question DARPA's reorganization poses to Korea is not to become blurred. Is this a name to admire, or a mission to interrogate? Holding that question in hand, the next chapter opens the budget documents. The check begins there, with whether money actually flowed in the direction indicated by the name.

Readers should look beyond names on an organization chart and follow the flow of questions. The problem held by the military is translated into the language of programs, and researchers break that sentence back down into laboratory tasks. The reorganization of MXO and IPTO was an event that reset that translation device. Such a sentence is not flashy. It quietly shows that when the way of war

changes, research organizations change first.

Chapter 2. The Direction Budgets Reveal: From Basic Research to Advanced Technology

Budget Flows from 2014 to 2023

Place two documents side by side on a desk. On the left is a thick budget volume issued by the U.S. Department of Defense in the summer of 2025. Its title is long: Fiscal Year 2026 Budget Estimates, DARPA Research and Development Master Justification Book. Tables and item numbers continue for hundreds of pages. The document on the right is thin. It is an article published in Science and Technology Review, a Chinese science and technology journal. It examines the budget of the U.S. Defense Advanced Research Projects Agency from 2014 to 2023. The two documents come from different countries and were written for different purposes. They ask the same question. Where did the money go?

One term needs to be unpacked first. DARPA's research and development budget is divided into several large categories according to the nature of the work. The U.S. Department of Defense calls these Budget Activities, or BA. The first stage is basic research. This is money spent on digging into the foundations of science with a view toward ten or twenty years later, even if it does not become a weapon right away. Physics and materials science in university laboratories belong in this category. The second stage is applied research. What has been dug out of the foundation is tested against a specific problem. Can this material make a lighter antenna? Can this algorithm hold up against real sensor signals? That is the stage. The third stage is advanced technology development. From this category on, tangible things are built. Prototypes are shaped, then run in the air and on the ground. Open the FY2026 justification book and DARPA's spending is separated into these categories. Even within the budget of the same agency, how much money goes into which category reveals what that agency now treats as urgent.

The first conclusion the Science and Technology Review article draws from ten years of numbers is this. The budget's center of gravity moved from basic research toward advanced technology development. Earlier DARPA budgets placed heavy weight on basic research. The article reads the past decade, and the later years in that decade in particular, as a period when basic research lost its first-place position in budget allocation. Money shifted toward the third stage, where existing ideas are pushed up into real systems. Put the other way around, DARPA's hand moved from seeds aimed at ten or twenty years later toward fruit that could be put onto the battlefield within a few years.

This point needs to be stated plainly. The picture of this ten-year trend relies in large part on the analysis in one Chinese article. DARPA's official budget documents show itemized figures for how much money went into each category every year. For those numbers to become a story about a shifting center of gravity, a layer of interpretation must be added. The shape of the curve changes depending on which items are treated as basic research and which items are grouped as advanced technology. When a program starts in applied research and moves into advanced technology

development, the accounting category changes, and it can look as if the center of gravity has moved. For that reason, this section speaks about direction, but it does not present precise annual amounts as settled fact. The direction is clear, and the decimal points require care. That is the posture this book tries to keep throughout.

Then where did the money point? The article identifies four technology areas where DARPA's budget was concentrated during this period, especially around 2022 and 2023: microelectronics, artificial intelligence, communications and networks, and biotechnology.

Microelectronics is the story of semiconductors, the brains inside every weapon. From smartphones to drones and radar, all of them contain a fingernail-sized chip. As the path of making that chip smaller reached physical limits, DARPA put money into preparing for what comes after smaller no longer works. Instead of competing to shrink transistors by a few nanometers, it turned to new ways to stack chips, remove heat, and join computation with memory. Artificial intelligence is the field where budgets grew quickly after deep learning broke through in the mid-2010s. The goal moved beyond machines that process only the data they are given, toward systems that can read context and make judgments even in battlefield situations they have never seen before. A battlefield is not a laboratory. Scenes never seen during training appear every moment. The goal became judgment that would not collapse in front of those unfamiliar scenes.

Communications and networks are the nerve fibers connecting the other three fields. In modern war, communication networks are the targets the enemy tries to hit first. Budgets went into creating connections that survive even where satellite communications or positioning signals are cut off. No matter how smart sensors and unmanned vehicles may be, if they cannot talk to one another, they are only scattered parts. Biotechnology's appearance on this list is itself a change. Biology, which until recently stood outside defense research and development, entered as one axis, and a biotechnology office under the Department of Defense led investments that overlap artificial intelligence and life science. The language of cells and genes entered budget tables that had been written in the language of missiles and radar.

A third feature ties these four branches together. The article reads the budget execution of this period as the prioritization of demonstration and verification. DARPA is known as an agency willing to bet boldly on things that could overturn the board, even when their odds of success are below fifty percent. That temperament was written directly into the budget. The agency guarded against research that stopped at papers or computer simulations, and it pushed money straight into demonstration stages where actual hardware was built and tested in the air and on the ground. It favored working prototypes over possibilities on paper. This posture has a cost. When steel is shaped and actually run, failure becomes visible. A design that always succeeded in simulation may stop in front of dust and dirt. The choice to put real objects first even while accepting that failure tells us the character of budget allocation in this period. Programs that appear in later chapters, from autonomous driving tests on roadless dirt tracks to demonstrations of autonomous air combat, grow

from these budgetary roots.

What endpoint did the direction of money point to? The FY2026 justification book gives a clue. One item that can be placed under advanced composite systems states its purpose on its own. The gist is this. Instead of depending on a small number of high-priced assets worth hundreds of billions of won each, DARPA seeks a new structure of weapons technology that gains lethality by combining overwhelming performance with overwhelming numbers. When one expensive decisive weapon collapses, the whole board shakes. That concern appears inside the sentences of the budget document. The picture is to scatter small, cheap, functionally spare sensors, communication nodes, unmanned vehicles, and strike systems, then bind them through an artificial intelligence communications network into one combat system. Why the four branches above moved together becomes clear in one line. If microelectronics mass-produces cheap brains, artificial intelligence commands them, communication networks connect them, and biotechnology is placed on the same board, then four different investments turn out to have been running along separate paths toward one picture of war.

This picture also offers a clue for Korean readers. Where the budget's center of gravity leans reveals what a country fears and what it treats as urgent. The fact that the United States pulled back slightly from the seeds of basic research and put more weight on real-world demonstrations is a signal that it judges the speed of pursuit to be that fast. A country planting seeds and a country picking fruit have different budget tables. When discussing a Korean-style research and development organization, this budget table first shows that the question of which category receives money is also the question of what the organization fears.

Before closing this section, one more line must be drawn. The direction described here stands on the narrow foundation of one official budget document and one external analytical article. The fact that both documents point to the same place gives weight to this direction. With only these two documents, however, it is impossible to assert the details of year-by-year increases and decreases, or the exact size of the amounts actually allocated to each program. The promise not to portray DARPA only as a success myth must be kept in the budget chapter. We can say where the money flowed. Whether all that money was worth it remains a separate question to ask, program by program, from the next section onward.

The Priorities Shown by the 2026 Budget Document

Open one budget book and you can see what an organization fears and what it is aiming at. A person's words can be dressed up, but the direction of money is harder to disguise. As you turn the pages of the U.S. Department of Defense's Fiscal Year 2026 DARPA budget justification documents, familiar items disappear one after another. This is a thick file in which the Department of Defense explains each program to Congress in order to receive funding. Boxes that once carried hundreds of millions of dollars now show zero dollars. At first glance it looks like a budget cut. Look more closely and the money has not disappeared. It has moved. Follow where it moved, and the organization's

future intent comes into view.

The first change that stands out is that the skeleton of the document itself has changed. Following recommendations from the PPBE Reform Commission, which reviewed the U.S. budgeting system, DARPA consolidated its program elements from sixteen to eleven. A program element is an accounting unit through which Congress allocates and tracks money. Reducing the number of these boxes means DARPA has set aside the old habit of dividing research by academic field. The budget document states the reason for this consolidation. The aim is to make procurement and execution simpler, make the budget more transparent to Congress, remove overlap, and bring in new technology more quickly. When the table of contents of a document changes, the organization's way of dividing up the world changes as well.

The result is zero dollars. The Defense Research Sciences item had 293.14 million dollars attached to it in 2025, but in 2026 it is empty at zero dollars. Names such as Electronic Technology, Network-Centric Warfare Technology, and Sensor Technology are the same. These names were long DARPA signboards. Electronics, communications, sensors. They overlap with the names of departments in university laboratories. Those boxes were emptied whole, and the money moved into new boxes named after missions. A budget document written in the language of academic disciplines was rewritten in the language of missions. The categories of money were redrawn around the problems to be solved, not the subjects to be studied.

Among the new boxes, one of the largest is Manufacturing, Sustainment, Supply Chain, and Logistics. It received 1.347 billion dollars. In Korean won, that is around 1.8 trillion. It is an unusual amount for a single item. The name sounds closer to a factory than a laboratory. Why did DARPA begin worrying about factories and warehouses? The judgment beneath it is that even if a fine weapon is created in the laboratory, it is useless on the battlefield if it cannot be produced cheaply and in large numbers inside the United States. An agency that once imagined future technology stepped into the problem of the hands and materials needed to actually make that technology.

Inside this box sits a next-generation microelectronics manufacturing program. The plan is to bring three-dimensional heterogeneous integration, which vertically stacks and joins multiple semiconductor layers, into production lines inside the United States. The stated goal is to build an open prototyping line used jointly by academia, government, and companies, shortening the cycle for making new chips. It means narrowing the distance between drawing a good design and turning that design into an actual chip. The concern that relying on others for semiconductors would tie America's hands on the day war breaks out has settled into the middle of a research agency's budget document. The same box includes a program called Burn n' Go. It is a project to produce solid rocket motors faster and more reliably. The diagnosis is that missile shortages come not from design failure but from production bottlenecks. The program says it will create a standard motor design usable across many weapons, reducing the time suppliers need to change production lines.

The Quantum Benchmarking Initiative, which deals with quantum computers, also entered this supply chain box. In 2025, Congress added 135 million dollars to it separately. The purpose is to measure coldly how fast commercial quantum computers are scaling, so that a strategic surprise does not erupt one day from another country. What stands out is that the name focuses on verification, not boasting. Measuring whether technology made by others is real was placed in a major budget category, alongside the work of making new technology first.

The second current is a changed way of looking at weapons. The newly created DARPA Advanced Technology Development item received 1.64346 billion dollars, and the Advanced Composite Systems item received 350.69 million dollars. In the mission descriptions for these two boxes, there is a sentence worth turning over for a long time. Instead of relying on one expensive high-value asset, DARPA seeks to open a new structure that gains lethality by mixing overwhelming performance with overwhelming numbers.

Put plainly, this means DARPA will not stake victory on one weapon worth hundreds of billions of won. It means connecting many cheap weapons like a net. Once one expensive weapon is lost, that is the end of it. If several cheap weapons are lost, the rest continue the mission. This arithmetic shifted the weight of the budget. Box Kick, included here, is a project to shoot down slow incoming cruise missiles or unmanned aircraft by cheap means. The calculation is to preserve expensive friendly weapons for more dangerous targets. Archer starts from the concern that friendly forces now rely too heavily on high-value large assets. It aims to create substitute strike systems so those large assets can be released for other missions. From protecting one expensive thing to spreading many cheap things. The place where the center of gravity of weapons philosophy has moved is printed in numbers in the budget document.

The third current is protecting AI and software from the inside. The Advanced Implementation Technology item received 335.64 million dollars. Constellation is a program that creates technology and prototypes so the military can conduct defense and counterattack together in cyberspace. By tying development, security, and operations into one method, it aims to deliver new cyber technologies quickly into actual operational settings.

SABER begins from a modest premise. It starts by acknowledging that AI placed on the battlefield is vulnerable to the enemy's physical and digital deception. This program seeks to create technologies that find and reduce those weaknesses, and to establish AI red teams inside the Department of Defense. It means the friendly side will strike its own AI first. SciFy is a program that uses computers to test whether scientific claims are truly possible. It is a device meant to prevent being fooled by inflated technical achievements. Place the names of these three programs side by side, and you can read an organization more worried about the weaknesses in its own technology than about the success of that technology.

The final box comes down to the human body. The Warfighter Performance item received 278.12 million dollars, and BEST inside it is a project to create sensors attached to wounds. It aims at a

closed-loop smart bandage that continuously watches pathogens and immune responses at the injury site, combining diagnosis and treatment. The picture imagines not a laboratory, but a forward emergency scene where a soldier has fallen. The fact that budgets for making weapons and budgets for saving people sit side by side in the same document shows another face of this file.

There is one more sentence that catches the eye alongside the direction of money. The Embedded Entrepreneurship Initiative inside the Advanced Technology Development item is a project that moves DARPA technology into the market, and the 2026 plan clearly includes an item on blocking problematic foreign investment in strategic technology. The document does not name which country's capital it has in mind. What is clear, though, is that protecting technology has now moved beyond patents and security into the matter of investment screening.

Put the five boxes together and the picture connects. Make cheaply and in large quantities inside the United States, bind cheap weapons together through numbers, protect the AI and cyber systems that move those weapons from the inside, revive wounded soldiers at the front, and prevent the technology from leaking out through another country's capital. Make, connect, protect, save, and police. A budget document once divided by academic field starts to read like a single war story.

Set this budget document beside the trend from 2014 to 2023, and the direction becomes clearer. A Chinese science and technology outlet analyzed ten years of DARPA budgets, and the center of gravity in that period differs in tone from the center of gravity in 2026. The face of this budget document is that the names of categories that once held basic research have disappeared, while manufacturing, supply chains, and volume have moved to the front. It is a signal that an agency once drawing the distant future has turned its eyes toward the problems of making and protecting now.

There is a point where we must step back. The fact that a large number appears in a budget document does not mean the program has succeeded. The supply chain project assigned 1.3 billion dollars and the weapons projects that speak of overwhelming volume are, for now, plans and statements of intent. Allocated money is not always spent, and spent money does not always bear fruit. Many DARPA programs are folded along the way, and the records of folded programs quietly disappear from the next edition of the budget document. Where the money flowed can be confirmed, but what that money created cannot be known until later budget documents and test results appear. This section can speak only to direction. The report card is not yet in our hands.

The Relationship Between the Direction of Money and Program Goals

Open a U.S. defense budget document. The cover says it is the Fiscal Year 2026 DARPA Research, Development, Test, and Evaluation budget justification document. Across hundreds of pages, program names and numbers line up, and beside some items there is only the lonely number zero. These are places that had received hundreds of millions of dollars until the previous year. Stare at those zeros long enough and something becomes visible. They are not traces of a modest budget cut. They are traces of a full turn in the direction the money had been facing.

Where an organization puts its money is an honest indicator of what it fears and what it wants to achieve. Words can be decorated as much as one likes. The sentences in a press release become smoother the more they are polished. The numbers in a budget table are harder to decorate. Accounting rules force the document to reveal how much was put into which box. That is why, when we set side by side where DARPA has moved money over the past decade and how it rebuilt its budget structure in 2026, the war this agency is now preparing for comes into fairly sharp view.

Start with the long flow. An article published in *Science and Technology Review*, a Chinese science and technology journal, examined ten years of DARPA budgets from 2014 to 2023. Why would Chinese scholars track another country's defense agency budget for ten years? That itself is a scene. It means a budget book can read to researchers in a rival country like the blueprint for future weapons. The shift in the center of gravity found by the article is this. The budget's axis moved clearly from basic research toward advanced technology development. Basic research is money that waters seeds whose final form is not yet known. Advanced technology development is money that chooses usable seeds, transplants them into the field, and makes them bear fruit. The character of the money changed from money that makes papers written to money that makes things built.

The same article points out that in the final two years, the budget was gathering into four branches: microelectronics, artificial intelligence, communications and networks, and biotechnology. Semiconductors, algorithms, networks, and the body. The fact that money concentrates in these four places draws a map on its own. The article summarizes the final goal of these four fields this way: accelerate the demonstration and verification of new technology, and assemble mosaic warfare as an actual method of combat.

The term mosaic warfare may sound difficult. Let us loosen it for a moment. It means abandoning the approach of putting every capability into one expensive fighter aircraft, and instead connecting a large number of cheap unmanned systems through a communications network so that each is weak alone but together they cover the enemy. One shard of glass breaks when stepped on. Join hundreds of shards together and they become a mural. Even if the enemy breaks a few pieces, the picture does not collapse. What must be held onto is the use of money. The money is not being spent to refine the theory of this approach. It is being spent to show that approach running before one's eyes. Not one article, but one working demonstration. The character of the budget divides at this fork.

This direction appears in the 2026 budget document in a much rougher and bolder form. DARPA merged long-separated budget items on a large scale. A budget item is a box that holds money. In the past, the boxes were divided closely by academic field. Electronics in this box, sensors in that box, networks in the box beside it. When divided this way, the accounting is clean. It is easy to see at a glance how much went into each discipline. But one problem appears. Walls arise between boxes. Electronics money and network money have trouble crossing over. In 2026, those walls fall.

Look at the numbers. The Defense Research Sciences item contained 293.14 million dollars through 2025. In 2026, it is zero dollars. The Network-Centric Warfare Technology item had 866.53 million dollars, and it also becomes zero dollars. The Sensor Technology item's 267.96 million dollars is likewise emptied to zero. Boxes led by technology names such as electronics, sensors, and networks all closed their doors at once. These three boxes alone total more than 1.4 billion dollars. That much money lost its name tag overnight.

The money did not disappear. It flowed into large newly created boxes. One that stands out is the item named Manufacturing, Sustainment, Supply Chain, and Logistics. This one box drew 1.34704 billion dollars, roughly 1.8 trillion won. The item called DARPA Advanced Technology Development was also newly created and received 1.64346 billion dollars. The ten-year flow seen earlier, the axis that moved from basic to advanced, has hardened into the names of budget boxes.

Let us unpack what this reorganization says. The old boxes were divided by type of technology. Money was divided according to what was being studied. The new boxes are different. Making, fixing, and carrying. The very names manufacturing and supply chain reveal the goal. This is a declaration that DARPA will not stop at inventing good parts. It has placed at the front of the defense problem the question of how quickly, steadily, and in what volume those parts can be produced inside the United States and sent to the front. Behind these numbers lies the judgment that in war, what truly trips you up is not a brilliant idea but quantity and delivery time, whether the item is a semiconductor or a rocket motor. Even an outstanding weapon is useless in total war if only ten can be made in a month. It is a painful calculation.

Move down to specific programs and this judgment becomes tangible. Look at Burn n' Go in the budget document. This program is tasked with designing standard solid rocket motors that can be inserted across many weapons. The goal is blunt. It seeks to reduce the time required to change production lines for different weapons, lower processing costs, and break through bottlenecks that keep blocking weapon production. This is not a story about discovering a new principle of propulsion. The program exists to answer how existing technology can be supplied steadily, cheaply, and quickly. Even one rocket motor carries the word supply, not invention.

In the same budget document, one can also see programs with different kinds of goals. The Quantum Benchmarking Initiative received 358.11 million dollars in related 2026 funding. The goal of this money is not to make a weapon. It is to measure rigorously how fast commercial quantum computers are developing, separate hype from reality, and prevent an unexpected technological surprise in advance. This is not money for creating something to brag about, but money for filtering out empty boasting. NASCENT received 17.44 million dollars in 2026. This program joins computer science with experimental economics to design economic incentives and foresee unexpected consequences from U.S. economic measures. This money targets advantage on the stage of economics, not guns or missiles. Even within the same budget document, some boxes point toward volume, some toward verification, and some toward economic competition. The fact that the money

points in several directions shows that this agency does not imagine a single battlefield.

The goal is also embedded in the way money is given. DARPA often uses contracts that fix success criteria in advance and release money whenever those criteria are reached. The next tranche is paid only after a defined step is visibly confirmed and passed. If performance is not observed, the money stops too. A device that pulls researchers toward the goal is built into the contract ahead of time. As much as the size of the budget, the conditions under which money is released determine the direction of the program. When money is given changes a researcher's day as much as how much is given.

There is something that should be stated honestly for Korean readers. Looking only at the numbers in the budget document, DARPA appears to be an organization with clear direction and quick decisions. That impression is only half right. The official materials secured for this book show where the money went, but they do not show whether that money actually produced results or how much was wasted. The budget table alone cannot yet judge what results the programs moved out of zero-dollar boxes produced after several years, or whether the 1.3 billion-dollar supply chain bet was right. A budget document is a document of plans and intent, not a report card. This is why the book promised earlier not to read DARPA only as a success myth. Boldness of direction and success or failure of results are different questions. The budget document testifies to the first, but it cannot testify to the second. The moment we blur this distinction, we mistake someone else's plan for someone else's achievement.

Still, one line remains to be held onto in this section. The numbers in a budget document are closer to a doctrine of war than to a ledger. Read which boxes were emptied to zero and which boxes were newly filled, and you can see what the organization intends to fight with for the next ten years. DARPA wrote that answer not on the side of invention, but on the side of demonstration and supply. The first question Korea should ask when discussing defense research budgets also divides at this fork. It is not how much more to spend, but into which box the money should go and what should be made to run before our eyes. Increasing the total budget can be done with one decision, but rebuilding the boxes of money means rewriting the organization's fears and ambitions. DARPA's 2026 budget document remains as a trace of that rewriting.

Chapter 3. The DARPA Way of Solving Problems: Clear Goals Over Good Science

A Story That Begins With One Arrow

In early 2026, one slide appeared in the presentation materials for a DARPA Proposers Day. Its title was "How DARPA Thinks About Your Project." On the left side of the screen, there was a single point labeled "Today." From it, several arrows spread outward like a fan. Interesting questions, academically valuable side paths, discoveries that could become papers. The presentation gave those fan ribs a name: "Good science."

Most researchers love that fan. It is what they learned in graduate school and what conferences reward. When you dig into one question, a more interesting question often beckons from the side, and if you follow it, a new branch opens. Knowledge grows that way. The point of the slide was to erase that fan. On the far right of the screen sat a target labeled "Program Goal." DARPA asked researchers for exactly one thing: start from today and go to that target in one thick, straight arrow. Do not drift sideways.

The person who prepared this presentation was program manager Michael Feasel, and the program he leads is called Protean. Across the materials, one sentence is fixed in capital letters. All quantitative metrics and deliverables "must be useful for measuring progress toward the overall program goal." Experiments built only to fill the research period, observations that are interesting but unrelated to the goal, do not count as progress in DARPA's arithmetic. The very fact that the sentence is written in capitals reveals a tension. It means DARPA has long seen researchers slide back toward the fan unless the agency stresses the point that hard.

This chapter is about that arrow. It looks at how DARPA frames a problem and how it turns the hands and feet of researchers toward that problem. Good science and clear goals often collide. In that collision, DARPA chooses the goal without hesitation. We will trace why, and whether that method is always right, through two programs and several documents. One program is a chemistry and biology story about protecting proteins inside the body. The other is an economics story about recalculating the world economy as a weapon. The subjects are far apart, but the way the two programs set goals is strikingly similar.

Protean: How Do You Stop Tens of Millions of Poisons?

Let us begin with the problem Protean is trying to solve. As artificial intelligence entered drug design, it did not make only good drugs easier to build. It made bad substances easier as well. The presentation shows the shadow of this dual use in numbers. AI-based molecular design can produce tens of millions of potential chemical and nerve agents. That number deserves a pause. In the past, defense often matched one antidote to one threat material and tried to delay the time to death. If there was sarin gas, prepare sarin treatment. If there was another nerve agent, prepare the

treatment for that agent. When the other side has ten or twenty types, that arithmetic can still hold. The moment the other side has tens of millions, the method collapses. You cannot make that many antidotes. No matter how diligent the defender is, there is no way to keep pace with the algorithm that designs new toxins.

Protean turns the problem around. Instead of blocking each poison that comes after us, it shifts toward protecting the targets inside the body that those poisons attack. That is where the change in thinking lies. Nerve agents, opioids, and ion-channel toxins ultimately have to bind somewhere in the body before they can harm a person. Those binding sites are receptors, enzymes, and ion channels. Then instead of stopping every poison one by one, make the protein itself, the place where the poison binds, strong enough to keep functioning no matter what poison arrives. Rather than guarding each door, build a house whose doors do not break. The program's target number makes a person stop. Build preventive and therapeutic molecules that allow proteins to preserve their function even after exposure to 10,000x LD50. In other words, even if the body is hit with ten thousand times the amount that would usually kill a person, the target inside the body must endure.

At first glance, this goal looks fanciful. So DARPA presents a hypothesis with it. Nature has already done this. Mongooses can resist cobra venom by 200 times, the Arg181Cys mutation shows 30-fold resistance to opioids, and cattle ticks can withstand insecticides. They share one trait. They gained resistance to a new toxin without losing the life-sustaining function they already had. The mongoose resists cobra venom and still uses its nerves normally. Nature knew how to place resistance in a protein without ruining the protein's original job. Feasel's hypothesis is that the secret lies in protein structure, protein motion, and the binding pocket. Instead of competing head-on at the active site where the toxin binds, act on another site in the protein, preserve normal function, and block only the toxin's binding. The presentation calls this non-orthosteric protection. In plain terms, it does not wrestle at the front gate. It changes the structure of a side door so that only intruders are filtered out.

That is the direction. What separates DARPA from other research funders comes next: the way it breaks deadlines and numbers into pieces. Many organizations announce an impressive direction and ask about results three years later. DARPA cuts those three years into smaller pieces.

The Stairway to Ten Thousandfold

Protean divides the full 33 months into two phases. Phase 1 lasts 18 months, Phase 2 lasts 15 months. Inside each phase, it cuts the points researchers must reach by month. This calendar is the program's real face.

Look at Phase 1. By month 6, researchers must identify at least one new distal regulatory site in the target protein. In other words, find at least one side door, not the front gate. By month 12, they must reduce the binding affinity of the chosen threat mimic by more than tenfold while keeping the effect on normal endogenous function below fivefold. That means leaving the body's original function

almost untouched while cutting the toxin's binding force to one-tenth. By month 16, they must either restore or prevent damaged protein function in vitro, and they must raise the effective dose by more than tenfold against at least three threat mimics. Then, at month 18, the end of Phase 1, an evaluation team from the U.S. government, not the researchers, comes in directly. It verifies in vitro efficacy against at least three real chemical threats designated by DARPA. Only teams that pass receive Phase 2 funding. A researcher cannot end the matter by reporting, "It worked." The grader comes from outside.

Phase 2 raises the intensity another level. By month 24, the effective dose must rise one thousandfold in vitro, there must be no acute toxicity, and the government evaluation team must confirm in rodents that LD50 has risen by at least one thousandfold. By month 30, the optimized lead compound must reach ten thousandfold in vitro. At month 33, the end of the program, success means rodents survive exposure to real threats with survivability improved by at least ten thousandfold over the baseline. From tenfold to one thousandfold, then to ten thousandfold. The target is laid out like stairs, and the government's hand enters at each step. When you see these stairs, you understand why DARPA erased the fan from that slide. A person who must hit a specific number every six months has no room to wander. The calendar pins the researcher's eyes to the target.

We will return shortly to what this tight structure does to researchers. Before that, we need to note how clearly DARPA writes down what not to do when it sets a goal. Half of goal setting rests there.

What Will Not Be Done

The Protean proposal has a separate section for exclusion criteria. It is a list of what this research must not touch. Strategies that only pursue inhibitors or activators competing head-on at the active site are excluded. Sponge-like materials that absorb or break down the threat substance itself are excluded. Research that directly edits human genes as an intervention is excluded. Incremental approaches that slightly improve existing medical countermeasures are excluded. Black-box AI molecular design that produces only results without understanding the principle is excluded.

Why does this list matter? If the goal is only stated broadly, researchers drift toward the easy path. One path tweaks existing technology by 10 to 20 percent and reports better performance. Another pours data into fashionable deep learning and presents only accuracy numbers. These paths carry less risk of failure, turn easily into papers, and help win the next grant. The exclusion criteria close those comfortable detours in advance. Protean wants a new principle that protects the target like a mongoose, not an improved version of an existing antidote. So it locks the escape route toward incremental improvement. [Inference: This reads as a signal that DARPA will not spend defense innovation money on safe and obvious research. The presentation itself does not state this interpretation in black and white, so this much is an inference drawn from the document's structure.]

One item in the exclusion criteria deserves attention: black-box AI. Protean pushes away AI design that offers performance without knowing the principle. This attitude appears again when we look at other programs later. DARPA does not dislike artificial intelligence. It only sees an inability to explain why AI produced a given answer as unacceptable in defense. A judgment that may decide human lives on the battlefield cannot be handed wholesale to a machine whose reasons are unknown.

To be honest, Protean is a plan, not an achievement. The sources secured for this manuscript do not include validation data showing that the program achieved ten thousandfold defense after 33 months. What we are reading now is the grammar of goal setting, not a victory report. Whether the ten thousandfold number will actually be reached, or whether the program will quietly fold after 33 months, cannot be known now. The promise not to read DARPA only as a success myth applies here as well. Skill in setting a goal and the result of achieving it are different matters.

NASCENT: Recalculating the Economy as a Weapon

The second case has a completely different texture. It is not chemistry or biology. It is money and economics. The program is called NASCENT, short for National Security Economic Theory. It is handled by a DARPA office connected to information, and it is led by program manager David Dewhurst.

Dewhurst begins by placing two scenes side by side: 1956 and today. In 1956, when Egypt nationalized the Suez Canal, Britain and France moved troops. The United States did not fire a single shot. Instead, it pressed on the fact that Britain's gold and dollar reserves were rapidly falling and threatened a devaluation of the pound. That one threat made Britain and France end their military intervention. The United States shook exchange rates with words and stopped a war. It was the moment when exchange rates, not artillery, delivered the decisive blow.

The scene seventy years later is humbler. In the Russia-Ukraine war, the United States imposed strong oil and financial sanctions. The timeline in the presentation, with reference points in February 2022, November 2022, and March 2026, shows that Russia maintained or even expanded territorial control, while Russia's oil and gas revenue since early 2022 held up as if mocking the sanctions. The sanctions failed to change the direction of the war. The same United States, holding the same economic weapon, produced such different results over seventy years.

NASCENT's diagnosis is sharp. The world economy of 1956 and today's world economy are too different to compare. Yet when the United States uses the economy as a weapon, it still relies on old tools. To borrow the wording of the U.S. Department of Defense Fiscal Year 2026 budget document (PDF-001), today's geoeconomic approach leans heavily on qualitative ideas from international relations and law. Such theories are often hard to falsify and narrow in scope. Scholars disagree, and there is no standard that can tell in advance in which situations a theory will be right or wrong. As a result, there is no way to know in advance when sanctions will work, when they will spin uselessly, or what unintended side effects may jump out. The budget document also treats this as a

problem of cost inefficiency. Because there is no institution where geoeconomic theories can be discussed neutrally and refined through repetition, building tools for economic warfare has itself become expensive and slow.

So NASCENT declares its aim loudly. It will pull economic sanctions down from the realm of intuition and experience and rebuild them as testable science.

The Unfamiliar Goal Called Quantitative Geoeconomics

The method is to join three disciplines: economic mechanism design, formal methods from computer science, and experimental economics. NASCENT tests the technical hypothesis that these three can be tied together to build a scientific skeleton for geoeconomics. Mechanism design is the branch of economics that designs rules so that desired outcomes emerge when people pursue their own interests. Formal methods are ways for computers to verify systems mathematically. Experimental economics observes economic behavior by gathering real people. NASCENT brings three tools with different characters into one place to redesign the old craft of economic sanctions. The program's picture of ten years later is a new standard, in government and the private sector, of economic statecraft that is theoretically rigorous and evidence-based. The program has already given this new discipline a name: Quantitative Geoeconomics. It is bold enough to name a field that does not yet exist, then fill that name with content.

The order of work is also reversed. Usually, policy is executed and the results are observed. Impose sanctions, keep them if they work, withdraw them if they do not. NASCENT goes the other way. It first defines the security outcome the United States wants with clarity, then works backward to calculate what incentives and mechanisms are needed to reach that outcome. In the wording of the presentation, it is to "work backwards" from the goal. Fix the destination and draw the path in reverse. The grammar is the same as Protean drawing the stairs backward from the destination of ten thousandfold.

To work backward, NASCENT must split a complex world into layers. The presentation divides the world into three layers. At the top are national goals: security, growth, social stability, and political stability. Below that, national security goals split inward and outward. Inward, it views domestic institutions through a security lens: research and development funding, public-private cooperation, subsidies, education and standards, certification, taxation, labor markets, housing policy, financial services, and election security. Outward, it uses alliances, trade agreements, commercial hubs and shipping-route security, strategic reserves, export and import restrictions, strategic acquisitions, sanctions and tariffs, and capital controls. At the bottom are the goals of firms in the market: market share, profit, and risk mitigation. [Inference: This breakdown appears to assume that the state's strategic goals and firms' profit motives are bound to collide. The state wants to stop capital from flowing to hostile countries, while firms want to go where money can be made. What NASCENT seems to seek through modeling and experiments is a mechanism that lets firms act naturally in pursuit of profit while U.S. security goals are also achieved. This interpretation is drawn from the

structure of the document.]

The line NASCENT draws around itself is also clear. It does not try to model the entire world perfectly, and it does not directly make policy. It goes only as far as building theories and tools that work backward from a target state and establish trust between action and result. Drawing the program's own limits in advance is also part of goal setting. If it claims it will do everything, nothing can be verified. The work is divided into a 4-month Phase 1 and a 24-month Phase 2. In Phase 1, the program builds an ontology that decomposes geoeconomics. It defines the world's economic and security elements one by one and maps their relationships. In Phase 2, research sprints and evaluation events repeat. The theories developed are tested through experimental wargames, executive tabletop exercises, and application to real operations. The point is not to leave theory on paper but to gather people, run simulated situations, and try to break it. The Fiscal Year 2026 plan says government and industry will hold a symposium to define security goals and economic incentives precisely, then build mechanisms that produce both security outcomes and economic benefits, and finally verify their empirical validity. It also includes a plan to build a long-running public-private feedback forum. That means leaving a standing channel through which policy and market signals can keep passing back and forth.

If Protean is a problem of protecting proteins inside the body by ten thousandfold, NASCENT is a problem of redrawing the world economic network with mathematics. The objects are opposites, but the grammar is the same. Define the threat clearly, work backward from the result, and end with verifiable metrics. DARPA does not ask for a good paper. It asks for a point that must be reached. This repeated grammar is how DARPA thinks as an organization.

How Goals Hold a Researcher's Hands

Now we arrive at the real question of this chapter. Do goals designed this way actually change researchers inside the laboratory? When documents from several programs are laid on top of one another, goals are not wishes. They work like devices that steer researcher behavior. The examples below come from programs other than Protean and NASCENT, but they are placed here to see whether the same handles appear again and again. These programs, too, are mostly represented by planning documents and FAQ-stage documents. This is a reading of goal documents, not a completed scorecard.

One common device is banning the black box. The VITAL program, which predicts the physiological responses of battlefield casualties, states in its FAQ that black-box machine learning relying only on statistical relationships between inputs and outputs, without validated biological and physical mechanisms, is out of scope. Even if accuracy is high, the program will not accept it if physiology cannot explain why the prediction emerged. Another item in the same document blocks a trap that researchers easily fall into. Building a perfect personalized model is not itself the goal. Personalization is only a means to meet required performance, such as 90 percent agreement with expert judgment or predictive lead time. The work of making a personalized model more precise is

academically attractive. The document holds researchers back in advance so that they do not lose the objective in that attraction.

There is also a way of demanding mathematics. MATHBAC, which studies AI's capacity for scientific discovery, says in its FAQ that better performance is welcome, but the central goal is to establish the fundamental mathematical theory that will push AI science forward. It warns that proposals claiming better performance by applying existing AI tools to a specific field do not fit the goal. The pressure is to move beyond coding and performance tuning and toward proving why something works. It is a way of pushing computer scientists to work with mathematicians.

Sometimes explainability is forced through a number. CLARA, which combines automated reasoning with machine learning, states in its FAQ that explanations must be hierarchical and granular natural deductions, and that the expansion between layers must be 10 or less. It turns the vague word explainability into the number 10. [Inference: This appears to be a device for preventing logical developments so complex that a human cannot follow them, forcing the system to be designed from the start in a structure that a commander can read.]

Another method is to redirect work through physical limits. O-Circuit, which attempts ultra-low-power computation with biological cells, sets a daily energy consumption target of 10 milliwatt-hours. That includes both standby and computation time. Purely silicon-based solutions are out of scope. In front of this extreme power limit, researchers have no choice but to put aside familiar silicon approaches. The evaluation method is also unusual. The biological computing device they build must learn behavior in a Pac-Man game environment and, when carried by a drone, must follow the smell of three specified chemicals to find a target. A condition is added that memory must remain even after 24 hours of rest without training. The demand is to design not cells that produce a one-time reaction, but a structure that remembers for a long time and changes. O-Circuit, too, is now represented by planning and FAQ-stage documents, not the scorecard of a completed biological computer. The principle of reading a plan as a plan holds here as well.

There is also pressure to invent the process itself. Fleetwood, which aims to turn discarded biomass into strategic materials, pushes common pyrolysis and gasification methods outside the scope. By blocking the familiar path, it tries to make researchers invent new chemical processes such as cell-free biocatalysis. Another program sets a quantitative target of raising reaction speed to one million turnovers per gram per second. That pressure forces researchers to move beyond the laboratory beaker and model, in mathematics, the amplification mechanism of mass production. This manuscript reads Fleetwood only within the range of planning documents. It is not a result showing that waste has actually become strategic material. It is a way to see what the goal forbids researchers from doing.

When these cases are overlaid, shared handles appear. Exclusion criteria block easy detours. Quantitative numbers disallow academic compromise. Demands for proof of mechanism block blind data learning. When goals are designed this way, researchers set aside the ivory tower's good

science for a while and turn their hands and feet toward the mission demanded by the state. The moment when the fan in the slide disappears and only one arrow remains is actually taking place throughout the documents.

A Mirror for Korean Readers

Now we add Korea carefully. One earlier book, *AI Defense Revolution*, discussed the culture of Korea's defense research field. It argued that because failures in budget execution are punished harshly, researchers tend to crowd into safe topics with low failure risk rather than source technologies that could change the board. They reach for tasks that cannot fail, such as improving existing radar performance by about 10 percent. This is not a newly verified fact in this manuscript. It is used only as Korean-language context offered by an earlier work. It will not be counted as independent evidence.

When the two cases above are placed on that background, the parts worth admiring and the parts that should not be copied separate. What is worth admiring is the grammar of goal setting. Define the threat at the scale of tens of millions, define success with the number ten thousandfold, and have the government intervene every six months to verify. Use exclusion criteria to state what will not be done and lock in advance the doors through which researchers might escape to easy paths. This grammar can be transplanted beyond defense into almost any research planning. Clarify the destination, draw the stairs backward from it, and place the grader outside. Korea can learn at least that much.

Yet some things become dangerous if copied as they are. Such tight milestones and conditional funding work only when a culture that accepts failure as learning and an institution that does not punish researchers when they collapse midway come with them. The structure in *Protean*, where the government evaluation team verifies results at month 18 and only passing teams receive the next funding, works in a healthy way only when teams that do not pass can step away quietly. If elimination leads straight to blame and disadvantage, researchers will not take risks toward the ten thousandfold goal. They will propose only low targets that are sure to pass. Without that foundation, the number ten thousandfold becomes not encouragement but a threat. A device meant to call forth challenge turns into a device that kills challenge.

One more point must be made. *Protean* and *NASCENT*, the two programs used in this chapter, are directions at this point, not results. The sources secured so far cannot say whether the two programs truly reached their goals. The documents we read are Proposers Day materials, program descriptions, and budget documents. They clearly reveal how the goals were set, but other materials are needed to confirm whether those goals were achieved. DARPA's strength lies in how it sets goals, not in always achieving the goals it sets. There are many DARPA programs in the world that put up dazzling goals and then quietly folded. The moment we blur this distinction, we are writing a success myth again. This manuscript stands against doing that.

Not good science, but a clear goal. The way that goal holds researchers' hands. And what must come with that method if it is to work in Korea. The hand that erases the fan from the slide can exist in the United States or in Korea. Whether an institution that can endure failure stands where that hand points determines whether the arrow becomes a weapon or an instrument of harm. From the next chapter, we follow how this grammar collides with the real battlefields of AI and cyber operations.

Chapter 4. AlxCC and Automated Code Security

The Design of the AI Cyber Challenge

It was the summer of 2025, in a competition hall in Las Vegas. On the stage of DEF CON 33, the festival of hackers, the final outcome was decided. The main actors on that stage were not human. Automated systems were competing against one another, reading software on their own with almost no human help, finding hidden weaknesses, and fixing them on the spot. This was the final scene of the AI Cyber Challenge, a competition DARPA had carried for nearly two years. People shorten its name to AlxCC.

To understand how this scene came to be, we first need to look at the problem the competition was built to face.

The lights we turn on in the morning, the traffic signals on the way to work, the water from the tap, and the communications in our hands now run almost entirely on software. Inside that software sits more code than a person could read in a lifetime. The problem is that code is not perfect. Human-written code always leaves gaps. Attackers push into those gaps. Defenders must find and close them first, but at the speed of humans checking line by line, they cannot keep up with code that grows every day. Code increases. Security staff fall short. As that gap widens, social infrastructure becomes more exposed. Can the work of finding and blocking weaknesses itself be entrusted to machines? That was the question AlxCC posed.

DARPA is not an agency that asks researchers for fine papers. It poses a problem first, then gives money and a stage to those who actually solve it. So it did not hand out this question as a paper topic. It turned it into an open competition with prize money. Nor did DARPA, a defense research agency, design this competition alone. ARPA-H, the Advanced Research Projects Agency for Health, joined in because it was also worried about code in hospitals and medical facilities. The judgment that protecting code is both a military problem and a problem of civilian life placed the two agencies on the same stage.

The method of competition is the skeleton of this event. Each participating team brings an automated cyber reasoning system to the contest. It is not a format in which humans give instructions in real time from the side. The system must analyze code on its own, find vulnerable points, create repair code, and attach it. The human role ends with designing that system well in advance. What actually competes on the stage is the automated defense program built by the participants. DARPA did not create a contest between human experts. It created a contest between human-made defensive AI systems.

That difference is larger than it may seem. In an ordinary security contest, skilled hackers stay up all night pulling code apart and solving problems. Human talent and concentration become the score. But the core of the problem we saw earlier was that human hands cannot keep up with the volume of

code. A contest that merely puts faster humans to work cannot answer that problem. To answer it, there must be machines that scan code without rest while humans sleep. The design of AlxCC wrote that demand directly into the rules of the competition. Humans step away from the stage, and the automated systems they built wrestle with code on it. The nature of the problem set the rules of the competition. The rules did not come first with the problem forced into them.

Why did it have to take the form of a prize competition? That question contains DARPA's own calculation. The agency could have written and released a defensive program itself. It chose not to. Instead of one arm of government creating and distributing a single answer, it opened the door to hackers and AI researchers around the world and set a goal: win, and receive a large prize. The plan was to let each of them build the best automated defense technology in their own way, then pull that private-sector skill toward automated national cybersecurity. Rather than naming the answer and issuing an order, the government laid out a field where people with good answers would gather on their own. If different teams are allowed to solve the problem by different paths, solutions the government could not have imagined alone appear on the stage. Competition was the device that brought those varied attempts into one place.

The office in charge of this competition is the Information Processing Techniques Office, IPTO. It is the same name that appeared in the previous chapter when we discussed organizational reform. The fact that an office with a long history in computing and information processing took on cyber defense also means AlxCC was not just a security contest, but a program placed within the larger current of AI and computing research. The competition is also classified in the cyber field. Reading and fixing code touches the work of AI handling language, that is, programming language. The fact that the office long responsible for that research ran the contest tells us something about its design. This was less a side experiment in cybersecurity than one branch of a larger question: how far can AI handle human language and machine language?

The scale of the money also reveals the design. Public budget materials show that the competition received \$34 million in fiscal year 2024 and \$40 million in fiscal year 2025. In the budget documents, it is grouped under information technology items. The total prize pool across the competition was \$29.5 million, including \$7 million in small business innovation support. A large share of that money flowed into prizes and operating the competition. It was not money spent for the government to buy one defensive program. It was money spent to run an entire field where several teams could attempt the problem in different ways. By size alone, it is hard to call this one of DARPA's very largest programs. But if we look at where the money went, the aim becomes clear. DARPA was trying to draw private talent with prize money and open a path where machines could take over code defense that human hands cannot reach. The direction of money is an honest indicator of what an organization sees as urgent.

DARPA itself called this competition a turning point in cyber defense. It is a large phrase. Such words need to be read carefully. An agency calling its own program a turning point is one thing.

Whether the program actually became one is another. Still, the phrase contains the design intent. It means DARPA wanted to create the point where the speed of humans and the speed of machines change places in the work of protecting code, the moment when defenders first receive real help from automation. The winners were announced on the DEF CON 33 stage in 2025. That stage was the final, and nearly two years of preparation and preliminary rounds stood before it.

If we follow the design, the question this competition poses to national defense cyber policy comes naturally into view. Will a country build its cyber defense only inside government organizations, or will it call in outside talent and build it together? The design of AlxCC chose the latter. It also chose to put defense capability inside automated machines instead of filling it with headcount. Neither choice has yet been proven as the answer. But DARPA showed through the form of a competition that the field on which a problem is placed already determines half of the answer. The ability to design the field becomes as much a national competitive strength as the ability to write good code.

What we have seen in this section is not the result but the blueprint. How well each automated system repaired code on the stage, and how we should read those scores, belongs to the next section. For now, three points should be held in place. First, DARPA framed the problem as a contest not between people but between defensive AI systems made by people. Second, instead of having the government create the answer, it chose to call in private answers through prize money. Third, IPTO ran that field as a cyber program, with \$34 million and \$40 million spent across two fiscal years.

Let us draw the line in advance. A well-designed competition does not itself mean success. It is too early to say that national defense cyber operations have already been automated because of AlxCC. From the materials available, the structure of the competition, its budget, and DARPA's stated goal can be clearly confirmed. What remains is whether this design actually changed code defense on the battlefield. That can be judged only by looking at the stage results and what followed. Reading the design accurately as design is the first step toward that judgment.

The Code Security Competition Revealed on the DEF CON Stage

Las Vegas Convention Center, August 2025. At the thirty-third DEF CON, one of the world's leading hacker festivals, a somewhat unfamiliar stage appeared. There were no hackers typing at keyboards on that stage. Code and scores moved across a large screen in real time. Machines, not people, were finding gaps in software and fixing the gaps they found on their own. Security experts in the audience watched the judgments of AI compete on the screen, not the work of their own hands. It was a familiar contest hall turned inside out.

In the previous section, we read the blueprint of this competition. Now it is time to see what numbers that blueprint produced on the actual stage.

The name of this stage was the AI Cyber Challenge, or AlxCC. Its target was somewhat unexpected. It was not the weapons system of an enemy state, but the open-source software we use

every day. Hospital medical records, power plant control devices, and water and sewage management systems mostly run on public code that anyone can inspect. In that code, holes no one has found may be hidden. The volume is too large for human eyes to cover. DARPA's question was this: Can AI read and fix first the code that humans cannot fully read?

There was calculation behind choosing open source as the target. Because not only Defense Department weapons but the whole of civilian infrastructure sits on the same public code, technology that protects that code becomes useful at once for both the military and civilians. The idea was to protect both worlds through one competition. The format may look like spectacle from the outside, but inside it is a method of research management. It turns vague hopes into numbers and claims into verification. The judging criteria are fixed in advance, scores are calculated in real time, and the results are made public.

After the semifinals at DEF CON in 2024, the finals were held from August 7 to 10, 2025. Seven teams reached the final. The systems they each brought are called cyber reasoning systems, or CRS. They are automated defense systems that analyze code, find vulnerabilities, and create repairs on their own without human instructions from the side. The key word is autonomy. Finding a vulnerability and fixing it used to belong to different people. A security researcher finds a hole and reports it. A developer reads that report and fixes it. The condition of this competition was to connect those two stages inside one system without passing through human hands. The final prize pool was \$8.5 million, and the grand prize for first place alone was \$4 million. By scale, it far exceeds the prize tables of ordinary hacking contests.

Let us walk through the numbers from the final stage. The systems of the seven teams read more than 54 million lines of code in total. The organizers had deliberately created vulnerabilities inside 63 tasks. The systems found 54 of them. That is 86 percent. They did not stop at finding them. For 43, they produced repair code that actually worked. That figure is 68 percent.

How fast this leap was becomes clear when we compare it with the year before. In the 2024 semifinals, the vulnerability discovery rate stayed at 37 percent, and the patch generation rate at 25 percent. In twelve months, discovery rose by more than double, and repair by nearly two and a half times. The range of languages also widened. In the semifinals, the systems mainly succeeded only on C code, but in the finals they showed similar success rates in both C and Java. That means the skill was not useful in only one language. The fact that the field moved this far in such a short time shows that it is still early and that much larger shifts may still come.

Another point drew attention. The systems did not catch only the fake vulnerabilities planted by the organizers. They also found real vulnerabilities that no one had known about. The finals uncovered 18 real vulnerabilities in all: 6 in C code and 12 in Java code. The participating teams did not secretly exploit those holes. They responsibly reported them to the maintainers of the relevant open-source projects. Some delivered repair code with the reports. An exercise held in a competition hall made real-world software a little safer. It was a moment when an experiment did not end as an experiment

but left a mark outside the room.

The speed was striking too. The average time for teams to submit repair code was around 45 minutes. At some moments, three teams scored on different tasks within one minute. The repair styles varied. One team solved a problem with a precise fix that changed only one line. Another blocked the issue by writing more than 300 new lines of code. Even before the same problem, each system showed a different character. It is a kind of variety similar to seating several human programmers in one room. The average cost to solve one task was about \$152. In Korean currency, a little over 200,000 won. This cost figure is an estimate from the organizers, so it should be read as approximate. But compared with the cost of having people find and fix one security hole, it sits in a different order of magnitude. It was a number that unsettled the arithmetic of protecting code.

Not all of this ran on the strength of a few teams alone. Frontier AI companies such as Anthropic, Google, and OpenAI gave each team \$50,000 in credits for using their large language models, for a total of \$350,000. Microsoft supplied the cloud resources needed for computation, and an open-source security organization under the Linux Foundation contributed expertise. It was a picture in which the latest private-sector technology and public concern met on one stage. Without such a coalition, the participating teams would have had to spend their energy securing expensive computing resources before even thinking about defensive algorithms. The fact that companies, not the government, lowered the entry barrier also shows the character of this competition. A national problem was not carried by the nation alone. It was placed on private companies' infrastructure and solved together.

The final winner was Team Atlanta. It was a multinational coalition of researchers from Georgia Tech, Samsung Research, KAIST, and POSTECH. The names of Korean universities and a Korean company appeared at the very top of the winners list in a U.S. defense agency competition. Trail of Bits, a small security company from New York, came second and received \$3 million. Theori, with Korean and American personnel, came third and received \$1.5 million. Even a quick look at the winners list shows that this competition was not an internal American affair, but a contest among talent crossing national borders.

DARPA did not intend to let this competition end as a one-day spectacle. Some of the finalist teams' cyber reasoning systems and related tools were released, and a plan was attached to move the rest through public channels. The choice was to avoid locking the winner's methods away with the prize and instead leave them open so the whole field of defense technology could rise together. Attack tools are often more useful when hidden, but defensive tools make the world safer as they spread. For a competition that began by trying to protect open-source code, thinking through the release of its results was a fitting close. Program manager Andrew Carney said that security experts who had been skeptical of AI before the competition had now become advocates and users. It was a place where technology changed people's minds.

We should step back and ask how to read this scorecard. An 86 percent discovery rate looks impressive, but flipped around, it means one in fourteen was missed. A 68 percent repair rate also means one out of the remaining three was not fixed. In a real defense system, leaving one out of three unfixed can make that one place the door through which the whole system is breached. Most of these figures came from materials DARPA released about its own competition. Since the organizer reported the competition's achievements, it is more honest to treat them as records until independent outside verification confirms the same results. Whether systems that worked on stage will show the same performance before old, tangled code in the real defense field remains an open question. Competition code is a clean sample arranged for judging. Military field code is full of traces patched and extended over decades. Victory on a screen is not the same as safety on a battlefield. In the next section, we will continue with the tasks this automated defense has left for national defense cyber operations.

The Tasks Automated Defense Leaves for National Defense Cyber

The stage in the previous section ended brilliantly. What remained after the stage was not only applause. Questions remained with it. This section takes those questions out one by one.

First, let us recall the scene. August 2025, Las Vegas. Seven team stations were set up on the DEF CON 33 stage. People sat at each station, but the hands that actually read and fixed the code were not human. The program each team brought, the cyber reasoning system, scanned software on its own. It found weak spots and created repair code to fill them. Humans were there to watch. The name of the competition was the AI Cyber Challenge, or AIxCC. Within DARPA, it was a cyber program handled by the office responsible for information processing technology.

We have already seen why a defense research agency held such a competition. Looking through massive code line by line with human eyes to find weaknesses is close to impossible from the start. Skilled security experts are few, and code grows every day. If that gap cannot be closed by hiring more people, the remaining path is to hand the reading to machines. AIxCC was the place where that possibility was tested on stage.

Let us unpack the process by which a machine repairs code. A cyber reasoning system links three tasks. First, it reads a program and points to suspicious places. Then it tests on its own whether those places can really be broken by an attack. Last, it writes and attaches repair code to close the hole. For a human, these three stages might take days. The system handles them in sequence without human instructions. That autonomy is why participants on the stage could take their hands off the keyboard.

The performance figures are as we saw in the previous section. The final systems read more than 54 million lines of code without passing through human hands, found 54 of the 63 weaknesses the organizers had secretly planted, and fixed 43. They also found 18 real weaknesses no one had known about, and for 11 of them they created repair code that could actually be used. The average

cost to solve one problem was around \$152, and the average time to repair was 45 minutes. The winner was Team Atlanta.

If we stop here, the story is clean. It is a picture of machines guarding code through the night in place of humans. The problem comes next. What this competition left for national defense cyber was less a finished shield than a bundle of tasks to solve. If we turn the numbers over and read them again, those tasks appear.

The first task is the missed places. The systems failed to find 9 out of 63 altogether. Among the 54 they did find, only 43 were actually fixed, so more than ten were found but not repaired. That was in a test hall where the answers had been set in advance. A real defense network is different from a test hall. No one has marked where the weaknesses are, and the way code is tangled together is far messier. Finding 86 percent on a controlled stage does not mean the same performance will appear in a real battlefield network. To bring automated defense into national defense, humans must keep watching where and how a missed piece might be breached.

The second task is trust. Suppose a machine creates repair code in 45 minutes. Can that code be attached immediately to a weapons control network or key infrastructure without human review? It is hard to see inside an answer produced by AI and know why it fixed the code that way. A repair that looks fine on the surface may create a new hole elsewhere. The logic that because it was fixed automatically it should be applied automatically is hard to accept where human lives and national infrastructure are at stake. The task remains to decide how far to keep the principle that machines may repair, but humans hold the final judgment. The faster defense becomes, the more important the rule becomes for when to stop that fast defense and bring humans in.

The third task is who verifies. The numbers listed so far, 54 million lines, \$152, and 45 minutes, all come from DARPA's own announcements about its own competition. That means the organizer of the contest counted the contest's achievements. Many other programs discussed in this book face similar circumstances. There is still not enough material from independent outside institutions repeating the work under the same conditions and confirming it. This is not a call to doubt the performance. It means national defense cyber still lacks the scale that separates performance from self-praise: a third-party verification system. The more defense is entrusted to machines, the heavier the question becomes: who checks whether the machine really did the work, and how?

The fourth task sits on the other side of the scale. Saying that one problem can be solved for about 200,000 won and 45 minutes means defense has become cheaper and faster. The same tools can be held by attackers too. Technology that finds weaknesses in code tells defenders which holes to close in advance, but it also tells attackers which holes to breach in advance. As the price falls, both sides will use it more often and across wider ground. As automated defense moves ahead, automated attack follows. A country that builds a good shield does not become safe at once. It enters a race where spear and shield both speed up. If one country leads in automated defense, its adversary will hurry into automated attack.

Fifth, there is the distance between stage and real operation. The DEF CON final was a place to solve set problems within a set time. A real defense network runs not for a day but 365 days a year. Attacks arrive without warning, and the code of yesterday is not the code of today. Whether a system that performed well in one competition can be attached to a real network that runs without rest and trusted for a long time is another question. A competition shows possibility. Moving that possibility into daily operation begins anew after the competition ends.

DARPA itself did not call this result the endpoint of defense. In official materials summarizing the competition, it described the moment as a turning point in cyber defense. That means less a declaration that everything has been achieved than a point where direction changes. The scene of machines fixing code on the stage was vivid, but what it left was closer to questions than applause. Running defense automatically also means deciding who watches the automatic device and who stops it. The work of making technology and the work of making rules to govern that technology must always move together.

For Korean readers, this scene can be read in two layers. Team Atlanta, the winner, included researchers from KAIST and POSTECH. The name of Samsung Research appeared with them. Korean hands reached the front row of this stage. But the framework that designed the stage, set the problems, counted the performance, and moved the results into national defense policy belonged to the United States. Doing well in a contest of technology and setting that technology as a national defense task to solve on one's own are different problems. A country that trains good players is not necessarily the country that designs a good arena. The list of tasks left by automated defense belongs not only to the United States. It has been handed to us as well, because our names appeared on that stage. More than who wins the next competition, the longer-lasting contest will be who first writes out, in their own language, the five questions that remain after the match is over.

Chapter 5 Trustworthy AI: AI Forge, SABER, and GARD

Opening: Smart AI and trustworthy AI are different problems

Let me open with two rooms placed side by side. The first is a conference room at a frontier AI company. A performance graph for a new model is on the screen. Accuracy has risen since the last quarter, users have grown, and the next investment round is close. People are discussing where to sell the model. Advertising, customer support, code writing, document summaries. The conversation moves toward money. It is a familiar scene. Most of the AI advances that have stunned the world over the past few years came from rooms like this.

At the same time, there is a completely different room. AI chiefs from the Department of Defense and intelligence agencies have gathered. The questions they ask have a different texture from those in the conference room. Will this model make the same judgment on a battlefield where the power is out and communications are blocked? If an adversary has secretly poisoned the training data, can it notice? At the moment when a soldier's life depends on it, can a human immediately understand why it made that decision? Commercial models were not built to answer these questions. They had no reason to be.

The three programs in this chapter were all born from the distance between those two rooms. AI Forge tries to narrow that entire distance into a national research arena. SABER presses into the problem of actually testing AI that will go to the battlefield. GARD tries to explain, in theory, the foundations of why AI is deceived. They aim at different places, but the question beneath them is one. Making smart AI and confirming whether that AI can be trusted are separate tasks. The U.S. Defense Advanced Research Projects Agency, DARPA, has begun to take up this second problem directly.

AI Forge and the National Security Problem

DARPA's AI Forge program, introduced in 2026, began by giving a name to the distance between the two rooms just described. DARPA called this distance the 'deployable AI gap.' Commercial AI has advanced at dazzling speed, but that progress does not transfer directly into national security settings. AI Forge was not created by DARPA alone. It is a research and development ecosystem joined by the U.S. National Science Foundation and the Center for AI Standards and Innovation under the National Institute of Standards and Technology. The word 'Forge' in the name shows the character of the plan. This is not a place that buys finished goods. It is a place where metal is heated and hammered, where needed tools are forged directly.

Why did this gap appear? DARPA's explanation is plain. AI innovation today is led by a small number of frontier AI companies. These companies focus on commercialization and returns on investment. That is not wrong. It is natural for companies to pursue profit. The problem is that the qualities national security demands usually do not make money. Assurance, predictability, reliability,

resilience, transparency. These five qualities do not need to be perfect in consumer apps. If a recommendation misses once, or a summary is slightly wrong, no one dies. Military operations are different. They unfold in high-risk environments where life and death are at stake, move at speeds no human can check one by one, and take place where adversaries constantly try to bring systems down. The Department of Defense's fiscal year 2026 budget documents contain the same diagnosis. While a small number of frontier AI companies have driven innovation in pursuit of profitability, the qualities required by national security have been pushed aside. AI Forge begins from the state's decision to take on these hard problems that the private sector has left aside.

What does it intend to do? DARPA brought together AI leaders from more than fifteen defense and intelligence agencies, frontier AI companies, and representatives from academia to identify AI tasks for national security. It released the results as an official report and connected them to the White House's 'America's AI Action Plan.' That means the program was placed inside a larger national picture. The selected tasks were grouped into three branches. Let us take them one by one.

The first branch is interpretability. It means humans must be able to know why AI produced a given answer. DARPA is not aiming at an explanation that researchers can slowly take apart in a laboratory. It demands a level at which an operator can immediately understand the system's behavior in the field, at the very moment a judgment must be made. DARPA calls this 'operational interpretability.' One step deeper, the term 'causal interpretation' appears. A correlation such as 'this feature affected the judgment' is not enough. The goal is to verify which mechanism inside the AI actually caused the decision and how that mechanism will move under pressure. This branch also includes the work of determining whether scientific outputs produced by AI are correct and of building evidence packages that can be traced and audited.

The second branch is control. As AI's military capabilities rise, devices that keep it from leaving human intent must be built into the system. DARPA wants to embed control mechanisms inside the model rather than attach them outside it. The design asks the system to present evidence for its own uncertainty and to know how to pause judgment when goals or situations change. When AI creates other code or data on its own, or modifies other models, provenance records must trace the origin all the way back so errors do not spread through the software supply chain. Containment structures that can safely stop or roll back AI when it leaves its expected path are part of the same design. The demand to abandon the weak claim that 'it works well in the lab' and build evaluation methods that estimate performance before fielding also belongs to this branch.

The third branch is adversarial resilience. Among the three branches, this one is closest to the battlefield. The opponent is not natural noise but a thinking adversary. Large models trained on internet-scale data are vulnerable to data contamination attacks, in which an adversary secretly plants malicious information in the training data. The traces do not show up easily in statistics. They hide out of sight, then pull the model in the wrong direction at a decisive moment. AI Forge seeks to build the scientific foundation for systems designed from the start to withstand such attacks. Its tasks

include protecting AI that learns and evolves in real time during a mission so it does not inherit new vulnerabilities, and creating standards that measure resilience with enough precision to be used for national security certification rather than in fragments. This third branch is where GARD and SABER, discussed below, directly meet the program.

The operating model is also worth attention. AI Forge is not just a window that distributes research funds. It is a platform that connects many actors. After DARPA issued a request for information, roughly 130 top U.S. universities responded. The number shows how thirsty academia was for this problem. The program is scheduled for official launch in the summer of 2026 in a form managed by a nonprofit organization, and it seeks to connect the huge models and computing resources of frontier AI companies with the basic research talent of universities. It aims at both scale and depth, which neither side can easily reach alone. Companies have models and computing power but lack enough personnel to carry out deep mission-oriented research. Universities have minds but not the resources to run giant models. The design of this arena is to mesh those two shortages. It is also notable that the fifteen tasks are not fixed in place but are set to be revised every six months. In a field where technology can change in half a year, the task list will try to keep pace.

That is what the official documents say. Let me draw a line here. AI Forge is a program just beginning in 2026. What we have in hand now is a plan, a report, and the responses from universities, not verified results. The fact that 130 universities raised their hands shows academic interest, but no one yet knows whether that interest has produced AI that will not fail in actual defense settings. Earlier, I promised not to read planning documents as if they were achievements. AI Forge belongs under that rule as well. What technology this arena, opening in the summer, will place in the field several years from now must be judged then, not now.

For Korean readers, the question this story poses is clear. The United States is refusing the path of simply buying finished commercial AI. It is trying to build its own defense-only AI foundation by tying together national capital, university talent, and Big Tech infrastructure. Beneath that judgment lies a sense of crisis: even a well-made AI becomes a hole in security if it collapses under enemy attack. Looking again at the three branches of tasks, this is not a contest over whose AI is smarter. It is a question of who first obtains AI that can endure under human control amid battlefield extremes and enemy hacking. It is easy to envy this structure. There is something to ask before envy. Have we ever redefined the vulnerabilities of commercial AI in the language of national security? If we postpone that answer and only import the names of other people's programs, we will be left with an arena that resembles the shell and nothing more. When the third of AI Forge's three questions, adversarial resilience, is brought down into the actual battlefield, the next program appears.

SABER and the Verification Problem for Battlefield AI

An AI mounted on a combat vehicle sees an object on the road. The screen says 'civilian vehicle.' Is that judgment right? If someone has attached a barely visible pattern to the surface of the vehicle, or mixed a tiny amount of noise into the signal entering the sensor, the same object may be read in a

completely different way. These tricks have been demonstrated many times in laboratories. The problem is that most of those demonstrations took place under conditions in which researchers could freely inspect the data and the inside of the model. Deceiving a system while holding everything in your hand is different from an ignorant adversary deceiving it on a battlefield filled with dust.

Let us look more closely at this difference. In the laboratory, a researcher who deceives AI usually holds three things: what data it was trained on, what structure the model has, and which point will shake its judgment if touched. An attack that succeeds under these conditions shows that 'this weakness exists in theory.' The enemy on the battlefield usually does not have that information. What the enemy has is the outward appearance of the friendly system and the will to shake it. Department of Defense budget documents also point out that most AI attack models studied so far are difficult to realize in the physical world or rely on unrealistic assumptions that the enemy already has some degree of access to friendly systems. So success in the lab does not immediately mean a battlefield threat, and safety in the lab does not guarantee battlefield safety. Verification becomes difficult in this gap.

This is the gap in which DARPA introduced the program called SABER. Its full name is Securing Artificial Intelligence for Battlefield Effective Resilience. The name is built from English initials, and it is run by DARPA's Information Processing Techniques Office. Its subjects span artificial intelligence, autonomy, and cyber. The very fact that these three branches are bound together inside one program tells us the character of the problem. AI's weaknesses do not exist only in code. They become truly dangerous when physical manipulation and cyber intrusion overlap. A pattern attached to a surface is physical manipulation, noise mixed into a signal is cyber intrusion, and misreading it is AI judgment. If the three are checked separately, each may look fine, yet they may collapse before a combined attack. The Department of Defense expects adversary states to develop exactly these combined AI, cyber, and physical attacks, refine them into clear tactics, techniques, and procedures, and use them to target U.S. military AI.

DARPA frames the problem this way. It is already known that adversaries have many ways to drive friendly AI into degraded performance. It is also known that AI contains weaknesses adversaries can target. What has not yet been firmly grasped is how adversaries will combine such attacks to create actual harm on the battlefield. U.S. Army Lieutenant Colonel Nathaniel Bastian, who leads SABER, says the program must stand at the leading edge of those combinations. He adds that warfighters have the right to know whether the AI they use is safe and can withstand enemy threats. If the person holding a weapon does not know whether that weapon may turn back on him at some moment, the weapon is not an asset but a burden. Knowing individual weaknesses and knowing the order in which those weaknesses will erupt in actual combat are different kinds of knowledge. The first is a list. The second is a scenario.

So SABER's approach is not to build one more defensive wall, but to create the attacker inside the friendly side. The idea is to establish a standing 'AI red team' inside the Department of Defense.

Borrowing the definition from the National Institute of Standards and Technology, a red team is a group authorized and organized to imitate the attack and exploitation capabilities of a potential adversary. If defenders are called the blue team, the red team deliberately attacks friendly AI and shows which defenses actually work. It resembles a soccer team playing hard against itself in training so it will not make mistakes in a real match. The difference is that this red-versus-blue exercise does not end after one round.

That is where SABER parts ways with existing checks. Existing AI verification has been closer to a method of checking once before deployment and stamping approval. The logic was that the weapon passed the test before being sent out, so it was fine. SABER says it will not remain at that one-time check. It will keep bringing in newly appearing counter-AI tools and hacking techniques, repeatedly attempt attacks, and in that process measure the scale of the risks an actual battlefield will face. It treats verification not as an event but as a state. An AI judged safe once is not safe forever. It must stay on the test stand while enemy methods change. Just as installing antivirus software once is not the end and it must keep being updated, AI safety is being redefined as something that must also be updated.

The targets placed on the test stand are clear. They are not imagined technologies from the distant future, but systems likely to be deployed on actual battlefields within the next one to three years. These include AI-driven autonomous ground and aerial unmanned systems, in plain terms, robotic vehicles and drones. Fixing the time frame this way has a meaning. Rather than worrying in advance about technologies decades away, the program will test weapons that will soon stand beside soldiers. The operating method is also public. It is a single-phase program running for 24 months, structured with four phases and experiments. Already built AI weapon systems are thrown into battlefield-like environments, attacked repeatedly, and observed to see whether they endure. Abstract submissions and proposals to gather performers each had deadlines. That means the documents state when the program opens and what it will handle.

We must step back. SABER is a plan, not a result. Program documents and announcements tell us what it intends to do, not what it has already achieved. We do not now have external materials showing how many vulnerabilities the red team found, or how much safer which weapon system became as a result. That is why SABER must be handled carefully. If we assume that a solid design will produce solid results, we fall into the contradiction of trusting a program about verification without verification. As is often the case with DARPA programs, time lies between the goals written in documents and the results in hand. SABER's question is clear, but we must wait for that time to pass before its answer appears.

Even so, what this program leaves for Korean readers is clear. When we talk about defense AI, our eyes usually go toward making it faster and smarter. More accurate target identification, faster judgment, more autonomy. SABER turns in the opposite direction and first asks whether the AI already made can avoid being played by the enemy. It shifts the weight from a race in development

speed to survival and verification. There is a place where the two collide. The force that wants to send weapons out quickly and the force that wants to verify their safety to the end pull against each other. If verification runs as a standing process, deployment slows. If deployment is rushed, verification thins. SABER does not remove that tension. It tries to absorb it into an institution. The standing red team is the form that fixes that tension inside the organization.

Any country that seeks AI weapons will stand before the same questions. Have we really tested the AI we plan to send to the battlefield? Will that test end once before deployment, or will it continue throughout the period in which the enemy changes methods? And who will conduct the test? Will the people who built the defense check it themselves, or will a separate dedicated organization be created to deliberately attack friendly systems? SABER has not yet proved its answer. But it has clearly moved the location of the question. Before the problem of making good AI, it has placed the problem of confirming whether that AI can be trusted. To do this work called verification, one more foundation is needed. It is a theory of why AI is deceived and by what principle. The third program presses into that place.

GARD's Account of Adversarial AI Defense

A photograph is on the screen. Nothing looks wrong to the human eye. A few tiny points of noise, too subtle for people to notice, are placed on top of it. Then the answer from the AI classifier flips completely. A machine that had just answered that it was one object now says it is a completely different object after the noise is added. Humans feel no difference, but the machine sees another world. DARPA gave this phenomenon a name: adversarial deception. GARD is a program aimed directly at this deception. If SABER asked how this deception should be verified on the battlefield, GARD tries to identify the nature of the deception itself one level below.

The name itself contains the purpose. GARD stands for Guaranteeing AI Robustness against Deception. In Korean terms, it means guaranteeing AI's resilience against deception. The program is run by DARPA's Information Processing Techniques Office, and its subjects are classified under algorithms and cyber. Both the object of defense and the goal are inside the name. What is to be stopped? Deception. What is to be protected? AI's resilience.

Why is such a program needed? Because as machine learning rapidly enters weapons and surveillance equipment, weaknesses that did not exist before have entered with it. The better an autonomous system is, the more it entrusts judgment to the machine's eyes. If those eyes are deceived, every judgment attached behind them goes wrong. If AI that recognizes targets is deceived, the result does not end as a wrong answer to one picture. On the battlefield, life-and-death judgments go wrong with it. Cases in which a barely visible sticker placed on a stop sign can cause AI in autonomous vehicles or target recognition systems to read it as a different signal show this danger well. The same principle can produce situations in which a certain pattern painted on a friendly tank causes enemy AI to ignore it as background, or a pattern attached to a civilian vehicle causes it to be mistaken for a target. The GARD program description notes that, despite the short

history of adversarial attack research, dozens of attack and defense methods have already appeared. The problem is that the balance tilts to one side. Attackers need to know less and less about the target system. Yet attacks keep growing stronger and easily pass defenses already built. The defender is always one step late. This asymmetry is where GARD begins.

Why have defenses so far fallen behind? GARD's description points to two reasons.

One is that defenses are too narrow. Today's defense techniques are tuned to one specific attack. If the attacker changes the form even slightly, the shield is pierced at once. The real problem is that there is no theory that gets to the bottom of why machine learning is deceived. A method that blocks holes one by one whenever they appear cannot keep pace with the speed at which new holes open. The defender can deal only with attacks that have already appeared, while the attacker can create attacks that have not yet appeared. From the start, this fight is unfavorable to the defender.

The other is that evaluation methods are detached from reality. The ways of measuring how sturdy AI is remain thin indicators unrelated to actual security situations. A model that looked sturdy in the testing room can collapse weakly before a real threat. We have mistaken models with high scores but poor field endurance for sturdy ones. If you hold the wrong ruler, you pass the wrong object.

GARD's answer was to change direction. Stop patching one hole at a time whenever it is pierced, and first build the theoretical foundation for why machine learning is deceived. On this foundation, GARD aims at three things: finding system weaknesses before they are attacked, identifying the properties that make models sturdy, and creating defenses that block a broad class of attacks rather than a single attack. The third is the core. The goal is not to defeat one attack, but to find a property that renders a whole family of attacks ineffective. The evaluation method was also rebuilt. It does not end with one benchmark score. It chooses a method that checks whether defenses truly endure inside scenarios modeled on real situations. The attitude is to measure field survival, not good laboratory grades.

Theory alone changes nothing. To turn this idea into actual tools, GARD formed a research coalition. Researchers from Two Six Technologies, IBM, MITRE, the University of Chicago, and Google Research joined together. And they did not lock the results in a Department of Defense vault. They placed them on GitHub, a public repository used widely around the world. Instead of hiding defense technology, they opened it so researchers worldwide could use and improve it together. In an arena where attack techniques spread faster than defense techniques, one can read the judgment that many eyes are more reliable than a small number of secrets.

Two results stand out. One is a virtual test range called Armory. It is a space where newly made defense algorithms can be tested again and again inside harsh virtual scenarios to see whether they are truly sturdy. It is a testing ground for researchers who want evaluations that can be repeated, scaled, and pushed hard. The other is the Adversarial Robustness Toolbox, ART for short. It is a software package that helps developers and researchers protect their models from various adversarial threats and measure how strong those defenses are. A test range and a toolbox were

opened together.

The concern raised by GARD does not end with this program alone. Other AI security programs DARPA later introduced inherit the same concern. In AI Forge, one of the three core research areas is adversarial resilience. AI Forge materials warn that deception has become more subtle. These include data poisoning and backdoor attacks that hide tiny manipulations inside the training data itself, and interactive AI attacks that move, search, and repeat without rest. GARD's concern with theoretical foundations and scenario-based evaluation continues in AI Forge as the challenge of building metrics and standards suited to national security environments. SABER goes one step further. Beyond laboratory conditions, it organizes an official AI red team to test defenses against ground and aerial unmanned systems likely to be deployed within the next few years. The idea is to deliberately deceive friendly AI weapons and find their gaps in advance.

Although I have connected the three programs, there is a point that requires caution. It is hard to conclude that AI Forge and SABER directly inherited GARD's tools as they are. Public documents show only that the three programs share the same concern. There is still not enough externally verifiable material on how much GARD's Armory and ART are actually used in SABER's red team, or what their results have been. I will not inflate that gap into an achievement. I will leave it as a gap. It is clear that the three programs share one philosophy, but evidence that this philosophy extended into tool succession is a separate matter.

For Korean readers, GARD's question is fairly sharp. Korea is also placing AI in weapons and surveillance equipment. How far has the effort come to accept that this AI can be deceived, to explain in theory why it is deceived, and to verify it together through open tools? Boasting of well-made AI and testing for ourselves how that AI is deceived are different things. GARD's real lesson is not in the boast that it built sturdy AI. It lies in the attitude of redesigning defense from the premise that AI is deceived, and in the attitude of opening that defense rather than hiding it so it can be verified together. The remaining question is how much of that attitude we have.

Putting the three programs together: trust is not made, it is verified

We have passed through AI Forge, SABER, and GARD one by one. Now let us place the three in one room. AI Forge tries to narrow the distance between commercial AI and national security AI through a national arena. SABER tries to keep testing AI headed for the battlefield through a standing red team. GARD explains in theory why AI is deceived and released tools for that work. Their layers differ, but the three programs point to the same place. High-performing AI and trustworthy AI are not the same object, and trust does not arise at the moment of creation. It is maintained only while verification continues.

In handling these three programs, we stopped three times at the same place. AI Forge is a plan awaiting launch. SABER is a design in progress. The links among the three programs still have thin external verification. The promise not to portray DARPA only as a success story carries weight in this chapter. Nothing is more inconsistent than trusting a program about verification without verification.

This chapter has repeatedly marked the time that lies between goals written in documents and results held in hand.

What this chapter leaves for Korean policy practitioners is not a list of things to envy. It is closer to three questions. Have we ever rewritten the vulnerabilities of commercial AI in the language of national security? Do we have an institution that holds AI we plan to send to the battlefield under standing verification rather than a single test? And are we prepared to open, rather than hide, the reasons that AI is deceived so they can be examined together? The three U.S. programs have not yet provided finished answers to these questions. But they have placed the questions in institutional form. Instead of copying someone else's answers, we can rebuild the questions others have raised in our own language. At that point, Korea's defense AI discussion can begin again.

Chapter 6 Cyber-Physical Systems and Defense Cooperation

Place three documents side by side. One is a program announcement that says it will overturn the way fighter aircraft are taken into the sky for testing. One is a cooperation announcement that says safety will be proven from the moment unmanned aircraft code is written. One is an argument that protecting satellites is, in the end, protecting soldiers on the ground. The stages of the three documents are far apart. A hangar, a ground control station, orbit. There are three stages, but one frame. How should we build and protect systems in which physically moving equipment and the code that moves it are bound into one body, cyber-physical systems? We will follow that question across three stages. We will not celebrate results in advance. All three sources are still plans, cooperation efforts, and claims standing somewhere near the beginning.

CyPhER Forge and the Digital Test Environment

Suppose a new fighter aircraft is being built. Once the drawings are fixed and the prototype rolls out of the hangar, the real work begins. A pilot takes that aircraft into the sky. The pilot changes altitude little by little, raises speed one step at a time, and reads each instrument to see where the airframe begins to shake. Each point recorded this way is called a test point. To verify the performance of just one aircraft, thousands of such points must be recorded. A person holds the stick and flies through them one by one. This is a scene that has repeated for decades at Edwards Air Force Base in California.

This method is slow and expensive. Open the CyPhER Forge program document that DARPA released in March 2026, and a frustrating number appears on the first page. Major U.S. defense programs take an average of twelve years to field a new capability. The place where that long time and enormous cost show up most clearly is flight testing. This testing consumes one quarter of the total cost of developing an aircraft and more than half of the development schedule. The final step of taking the aircraft into the sky to confirm it is the heaviest burden in the long journey called development. It is a structure in which even after holding a good design in hand, years are spent proving that the design is right.

There is another frustrating point. Computer performance and machine learning have become a different world over the past twenty years. The way complex cyber-physical systems are tested has barely moved. The situation is the same whether the system is an aircraft, a ground vehicle, or a ship. The term cyber-physical system may sound difficult, but it refers to equipment in which a physically moving machine and the software that controls it are joined as one body. A fighter aircraft is such a system. Of all the data poured out when this kind of equipment is tested, less than 5 percent is actually used. The remaining 95 percent sleeps inside storage devices. Engineers are tied to the labor of matching heavy, precise physical models against thousands of manually flown conditions one by one. CyPhER Forge aims at this old way of working. Its goal can be summed up in

one sentence: reduce the number of test points needed for systems under development by more than a factor of ten.

How does it intend to do that? The program places two axes side by side. One is the digital twin, and the other is the AI test agent. The program calls the automated test system that combines them CyPhER. The name is unfamiliar, but the idea is unexpectedly clear. Finish by calculation what can be known in advance by calculation, and confirm in the sky only what must be confirmed there. The judgment is that the way to save costly real flight is to do much more calculation in advance.

Start with the digital twin. Think of it as a computational twin built inside a computer by copying the real aircraft. This twin must have three abilities. First, it turns a heavy, slow original physical model into a lighter surrogate model that runs close to real time. Next, it calculates in numbers how trustworthy or untrustworthy its own predictions are. If it lacks confidence, it has the ability to admit that itself. Last, it immediately absorbs data pouring in during flight and keeps correcting itself. The structure keeps three actions running without pause: prediction, measurement of the uncertainty in that prediction, and relearning from new data. Only when these three actions mesh well can things that do not need to be confirmed in the sky be filtered out on the ground in advance.

The AI test agent is the side that receives the twin's predictions and makes decisions. It finds which maneuver the test aircraft should perform next to gain the most information while staying as safe as possible. It judges safety by statistics, not by human rules of thumb. It calculates a path that extracts as much data as possible from a single flight. It then executes the selected maneuver with machine precision. There are three axes: safety assurance grounded in statistics, maneuver design that raises information yield, and machine-level autonomous execution. Combined with the three pillars of the digital twin, six technical pillars support CyPhER. One side predicts the world, and the other decides movement based on that prediction.

What deserves attention in this picture is what DARPA builds on behalf of participating researchers. So researchers can focus on the algorithm itself, it provides the entire integration, validation, and verification environment. Over the full program period, high-performance computing comes with billions of CPU hours and a set amount of reserved GPU capacity. A controlled unclassified information data repository containing past flight test records also comes with it. Five high-fidelity models that can imitate an experimental aircraft are prepared as well. The five models cover airflow around the airframe, elasticity and inertia of structures, engine thermodynamics, flight control laws, and the movement of control-surface servos. On top of this, the researchers' algorithms train and are evaluated endlessly inside a virtual environment. The method lowers the barrier to research participation. Small teams that would have had to give up their ideas for lack of costly test infrastructure can now borrow data and computing resources from a shared ground and compete. The way the field is opened is itself a policy. Infrastructure decides who can compete.

The schedule is thirty months, arranged in three phases. The first six months are groundwork. The test environment, interfaces, and data structures are built. The next twelve months are the period for

creating prototype CyPhER systems. Challenge cycles begin at this point. Researchers compete on model accuracy, learning speed, and uncertainty calculation ability using simulation data and real flight data that has been hidden in advance. The use of hidden data matters. It means the program will evaluate not a student who has already seen the answer, but a student facing a problem for the first time. Only the top teams that pass objective metrics move on. At the end of the final eighteen months, the surviving algorithms are loaded onto real experimental aircraft at Edwards Air Force Base. A human remains in the loop as supervisor, but the maneuver commands are issued by CyPhER during real flight. Then the traditional test plan and the test plan led by CyPhER are compared side by side to see how many test points were saved.

DARPA's success criteria are harsh. Looking at the final-stage targets, the surrogate model's error must be below 0.25 on a normalized metric. A single calculation must not take more than 0.1 seconds. Updating the model with new data cannot take more than one minute. Under the old method, the same work would have taken months. The safety side is even tighter. The false alarm rate for risk must be below 0.1, and the rate of missing real danger must be below 0.005. Information search efficiency must exceed a factor of ten. Only if all these numbers line up does the final promise of reducing test points to one tenth hold. If even one slips, the banner of a tenfold reduction wavers. When metrics are nailed down this tightly, researchers spend their time not on the beauty of a paper but on crossing those thresholds. This is a scene in which goals change human behavior.

We need to step back. What has been described so far is not an achievement but a plan. The CyPhER Forge document is a revised version of a request for proposals issued in March 2026. The error of 0.25 and the safety miss rate of 0.005 above are not records that have been achieved yet. They are targets that teams must compete to meet. Real flight over Edwards led by CyPhER sits at the very end of the thirty-month schedule. Whether those numbers will hold in the real sky, and how many teams will pass the challenge cycles, has no answer yet. Reading this program as if it were already a victory misreads the source. What the document promised and what the document proved are different. What we hold in our hands is the former.

Seen through a policy lens, the question raised by CyPhER Forge comes before the technology. The United States is laying a new field where the contest is not the performance of weapons but the speed of verifying weapons. No matter how good a fighter design is, if confirming it in the sky takes twelve years, those twelve years become a strategic loss. DARPA is trying to move that bottleneck out of human hands and into calculation and autonomous judgment. This approach also carries weight on the other side. Although a human remains as supervisor, much of the judgment that selects and issues maneuvers moves to machines. The controlled unclassified information data repository and multi-physics high-fidelity models accumulated through testing may become a common ground through which U.S. defense companies and academia develop cyber-physical systems in the future. This is the point Korean policy practitioners should watch. It is not the ability to build one type of weapon quickly. It is a design that tries to make the method of testing weapons itself into a shared base. What we can confirm is how far the United States has placed something on

the test stand. Whether that test stand will actually hold remains for the sky over Edwards to answer thirty months later.

The problem CyPhER Forge targets is not confined to the test stage. If long confirmation in the sky is expensive, there is also a path to reduce problems from the stage where code is written in the first place. Next to the story of moving testing earlier sits the story of making development safe from the start. The stage moves from the hangar to the ground control station.

The Air Force and DARPA's Cyber Resilience Cooperation

Somewhere in the sky, an MQ-9 Reaper is flying. Inside the ground control station, a pilot sits before a screen, changes direction, turns sensors, and confirms a target. Those hand movements do not touch the aircraft directly. Code stands between them. The force that pulls the stick and the angle toward which the camera turns move only after software receives and interprets the command. A quiet question arises. If someone secretly touched that code, what could the pilot see to recognize it?

In June 2025, DARPA announced that it would work with the U.S. Air Force to address this question directly. The starting point was the concern that if the software moving weapons and support systems collapses, overall combat power shakes with it. DARPA's diagnosis was cold. The U.S. Department of Defense has run systems on nearly thirty years of accumulated old security policies and aging information technology foundations. Attackers push through those gaps. They steal sensitive military code and reverse engineer important equipment to find weaknesses. DARPA's picture was that older structures are more widely exposed to this kind of intrusion. If the flight testing in the previous section was about belatedly confirming the performance of a completed aircraft, this story shifts our gaze to the moment the code that moves that aircraft is created.

DARPA found the root of the problem in one practice. Steven Kuhn, the program manager leading this program, called it 'patch-and-pray.' Put plainly, it means this: finish the software, release it, plug holes as problems arise, and hope the next accident does not happen. For a smartphone app, this may not cause serious trouble. If an error appears, users can just download one more update. Kuhn stated firmly that this approach, which depends on luck, is unacceptable in weapons systems where human lives are at stake. The moment the code of a fighter aircraft or an unmanned aircraft goes wrong can lead directly to the loss of lives and the failure of a mission.

So the alternative the two organizations put forward is 'formal methods.' The name is unfamiliar, but the idea is unexpectedly clear. Instead of testing for vulnerabilities after software is finished, mathematical proof is used during the making process itself to examine the behavior of the code in advance. Before the code is sent into an actual battlefield, proof confirms that it works only as intended and does nothing else. Rather than building a house and then tapping the walls while hoping it will not collapse, the load-bearing capacity of the columns is proven by calculation at the blueprint stage. The work moves from plugging holes after the fact to fixing the design stage so holes

do not appear in the first place. If CyPhER Forge in the previous section brought confirmation in the sky forward through calculation, formal methods bring the correctness of code forward to before deployment. Both are movements in the same direction: pulling verification from the back to the front.

The substance of the cooperation is a program called Resilient Software Systems Capstone. Its targets are not example code inside a laboratory, but operational platforms actually used for missions. Each project is jointly funded by the Air Force and DARPA, with the goal of producing results in roughly 24 months. Two years is a short time for weapons-system development. Within that span, the Capstone aims at four things: making code safer from the beginning, advancing the time needed for the authorization to operate process, removing waste from development and testing, and creating a best-practices guide that other organizations can follow. The last goal stands out. It means the effort will not stop at fixing one weapons system, but will leave behind documents so the method can be moved elsewhere. The authorization to operate process is the stamp that says new software may be placed on an actual mission network. The time it takes to receive that stamp is the speed of fielding. The calculation underneath is that if safety is assured in advance by proof, the stamp can be received faster.

The first test target is the MQ-9 Reaper. The Reaper is a medium-altitude, long-endurance unmanned aircraft that moves between reconnaissance and strike missions deep inside enemy territory. What draws attention is that this choice was not an order handed down from above. The Air Force itself identified the pilot weapons system and first reached out to DARPA. The aircraft was built by General Atomics. Behind that decision was the judgment that for an unmanned aircraft that circles for a long time over enemy territory, the safety of the code that moves it is the success or failure of the mission. A single line of communication or a single command gone wrong can turn a reconnaissance asset into a threat asset in an instant. The fact that the military field recognized the problem first and brought in a research agency shows that DARPA is not an organization that only pushes technology downward from above.

Kuhn summarized the method of this cooperation with the phrase 'correct by construction.' It means not finishing something and then hoping it passes inspection, but making the construction itself guarantee correctness. He said he hoped this effort would become an opportunity to deliver resilient software tools to both the military services and industry partners. He added that he expected it to become a starting point and an example for other defense contractors or military agencies to bring DARPA's tools into their own platforms and development processes. The idea does not stop at making one tool. It tries to design the channels through which that tool will spread.

This reads like a neat success story. Inside DARPA's own explanation, there is an honest blank space. Some tools made with formal methods have already moved to several military services and are awaiting further development and actual deployment, it says. That is good news. In the same sentence, DARPA states firmly that much broader and faster adoption across the Department of

Defense remains an unfinished task. This cooperation is 'an attempt that began with one weapons system.' It is not a declaration that 'the whole defense system has already become safe.' The MQ-9 Reaper is the first pilot, and how far its results will spread to other aircraft and systems remains a question to watch. The 24-month deadline is not a promise of success but closer to a schedule for an experiment. Having a good tool and having that tool take root across the military are different stories. DARPA did not hide that distance.

For Korean readers, the question this case leaves behind is not the size of the achievement. What we should watch is the method. Will we stand in the place of plugging holes after problems erupt, or will we demand proof at the moment of making? Formal methods are not cheap. Building proof takes time and specialized people, and it is hard to apply them immediately to all code. That is why even DARPA began with one unmanned aircraft. If our defense software procurement still depends on inspection after completion, this case becomes one coordinate for changing direction. The first question is not what to introduce, but where to begin demanding proof. The fact that this experiment has not yet reached its ending is the reason we should read the approach before the results.

The two stages so far have been stories between ground and sky. The union of code and equipment reaches even higher. A satellite in orbit is also, in the end, a cyber-physical system moved by code. If that satellite wavers, the unit on the ground loses its eyes and ears. The third document lifts our gaze to space and then pulls it back down to the ground.

What Space Cyber Debates Ask of Ground Defense

Imagine a unit that has never launched a single object into space. Even such a unit loses both eyes and ears when satellites stop. Reconnaissance satellites capture enemy movement, and communications satellites connect platoon to platoon. When equipment in orbit wavers, the soldier lying on dirt pays the price first. The phrase 'space cyber' sounds like a conversation exchanged among scientists beyond the atmosphere. What the phrase actually targets is the problem of the person on the ground holding a rifle. Let us follow this fact through three official documents. One is a short argument, one is an announcement for an ongoing program, and one is an announcement that cooperation is beginning. When three sources of different character point in the same direction, we can read that direction carefully.

A short document released by DARPA in April 2026 addresses this point directly. Its title is 'Winning the New Space Race Starts with Cyber.' The document first summarizes the situation this way. The cost of going to space has fallen beyond comparison with the past, and as a result, more countries and companies have entered space. It states firmly that this is not science fiction but something happening now. As the threshold for launching rockets has fallen, orbit has begun to grow crowded. That crowding itself is a signal of new competition. The document calls this competition the 'new space race' and offers one judgment. The side that holds lasting advantage will not be the country that launches more rockets, but the country that first learns how to build and operate systems safely in this unfamiliar environment.

The author of the document is a DARPA 'service chief fellow.' It is a position held by someone dispatched from the military to work inside the research agency. This person says his mission focuses on a piece that people often miss: the place where space and cyber overlap, space cyber. If space is seen only as a problem of rockets and orbits, this piece does not come into view. Attention is easily drawn only to numbers such as launch success rates, satellite weight, and orbital altitude. The moment one remembers that each satellite is, in the end, a computer moved by code, the story changes. Anything that runs on code always has openings to probe. Who sees those openings first is the fellow's concern. It is the same kind of question as the unmanned aircraft code and fighter test software seen in the previous two sections.

Here comes the part that should pain ground defense authorities. We cannot carry the information security methods used on the ground into space as they are. On the ground, when a problem occurs, security software can be newly installed and patches can be downloaded often. A person in charge sits before a control screen, watches for abnormal signs, and turns equipment off and on if needed. Equipment that has entered orbit is hard to treat that way. No person can go fix it by hand, and there is little room to load heavy programs. Once a satellite has been sent up, it must operate in that state for years. So the document warns against treating security as something added later. In English, it says security must be 'baked in,' not 'bolted on.' This means it should be mixed into the dough from the first stage of design, not tightened onto the finished system later with screws. Sprinkling salt after bread has been baked is different from adding salt while kneading the dough. After a satellite has been baked, there is no way to add salt. When Kuhn's 'correct by construction' from the previous section rises into space, it becomes this 'baked in.' Different documents are expressing the same belief in different words.

What scene unfolds if this is missed? The document uses the expression 'messy patchwork.' It means hardware and software with different standards have been joined in a rush. If pieces are stitched together like a patchwork quilt, the system may appear to run on the surface, but the defending side struggles even to know where the holes are. Each joint where one piece meets another creates a new opening. The attacker only needs to find one joint. The defender must protect every joint. This asymmetry makes defense exhausting. That is why the document offers a 'government reference framework' as its prescription. It proposes putting orbiting vehicles, ground terminals, and control stations into one bundle and creating a common set of standards that lets them withstand and survive even as threats change. Instead of building to different standards, it proposes sharing the same underlying drawing from the beginning. There is another point to watch. The document proposes taking what has been learned from cyber tools once made for offense and returning it to space defense. It is a story of refashioning the shield with experience gained from sharpening the spear. The person who knows where the opponent's hand enters locks the door best.

Up to this point, the issue may look like a space story. Reduced to one line, the document's prescription is this: plant security in the design in advance. Do not make a patchwork quilt. Protect sky and ground as one bundle. These three sentences do not apply only to satellites. They apply to

every cyber-physical system on the ground. If physical equipment and software move together, weapons systems, control stations, and communications networks are no exception. Control devices that open and close power plants, combat systems on ships, and ground stations that launch drones all fit inside these sentences. The space cyber document raises the real question for ground defense. Do your ground systems keep these three sentences? Have we neglected the joints of ground systems precisely because they are not somewhere as unreachable as satellites?

DARPA's ground-side programs moving at the same time reveal the overlap. In 2026, DARPA issued a revised announcement for a cyber-physical program called CyPhER Forge. Its weight lies in striking systems where physical equipment and software are entangled inside a digital environment in advance. The approach is to find weaknesses first on a virtual test stand before sending real objects into the battlefield. Instead of learning by breaking actual equipment, it experiences attacks in advance against a model built the same way inside a computer. This shares the same root as the space document's baked in idea, placing security in the design stage. Instead of counting holes after the patchwork quilt has already become thick, it moves toward closing holes while drawing the blueprint. The earlier testing is moved, the lower the cost of repair becomes. The cost of fixing a weakness exposed on the battlefield and the cost of fixing it on a drawing are in different orders of magnitude.

There is one more case. In June 2025, an announcement said the Air Force and DARPA had joined hands to raise cyber resilience. Resilience is the key word. It is not a promise to block every attack perfectly, but a goal to withstand impact and stand up again even after being hit. The idea begins by admitting that the shield can be pierced, then leaves a path to continue the mission even after the breach. It is the same language as the 'resilient, survivable' demanded by the space document. Three different documents did not happen to choose the same word by chance. Whether in orbit or on a runway, it is closer to a signal that the grammar of defense is converging into one. The center of gravity is moving from an era that dreamed of perfect walls to an era that builds bodies able to stand again after collapse. That is why the three sections of this chapter, though they deal with different stages, gather into one story: pull verification forward, plant safety in design, and make systems withstand damage even after being hit.

We need to stay cold-eyed. What we hold in our hands now consists of three things. One is an argument written by one fellow, one is an announcement for an ongoing program, and one is an announcement that cooperation is beginning. They are not performance reports. We do not yet have a record that this framework blocked an actual attack on a satellite, or numbers showing that CyPhER Forge filtered out defects in some piece of equipment. Nor can this announcement alone tell us what results came from the cooperation with the Air Force. This is where the promise not to draw DARPA only as a success myth must be kept. A claim must be read as a claim, and a plan as a plan, so readers are not deceived. A good framing of a problem is not the same as a good result. There is a long distance between posing a problem clearly and actually solving it. These DARPA materials still stand near the beginning of that distance.

Then what question remains for Korean readers? Was the security of our military's satellite communications, ground control stations, and weapons-system software mixed into the dough from the beginning, or tightened on with screws after everything was built? How thick has the patchwork already become? We must ask what joints now connect equipment brought in at different times under different standards, and whether we have a single drawing that lets us see those joints at a glance. DARPA's documents do not answer this question. Instead, they change the direction of the question. Protecting space and protecting the ground are not separate tasks, but the front and back of the same design drawing.

We have walked through this chapter across three stages, but the lesson left by the three stages gathers in one place. What the United States is now contesting is not stronger weapons. It is the method for verifying weapons faster, designing them more safely, and making them hold even after being hit. What we need is not envy but a mirror. The very fact that these programs still remain at the stage of plans, cooperation, and arguments also means we have time to learn the approach before waiting for the results. When we unfold that design drawing will decide whether we can protect our eyes and ears in the next fight.

Chapter 7 Intelligence Inside Hardware

The Shift Called Physical Intelligence

A reconnaissance robot climbs a dusty hill. Its front wheel catches on a stone. A camera mounted on top of the robot captures the scene. The image passes at once into the computing unit inside the body. The unit examines the image. It decides that there is an obstacle. It sends a command to the motor that turns the wheels. Only then does the robot change direction. Written out in words, the sequence feels long, but in reality it happens in the blink of an eye. DARPA believes that even that blink is too long on the future battlefield. On a peaceful test range, that delay would not matter. Where bullets are flying, the situation changes.

On April 29, 2026, DARPA published a piece on its official news page calling for a new way to think about robotics. The title included an unfamiliar phrase: “physical intelligence.” At first glance, it can read like another fashionable term. But the question behind the piece is fairly basic. Where should intelligence reside? Is it really best to pack intelligence into a single computer chip in the robot’s head? DARPA puts a question mark there. In the agency’s own terms, meeting the demands of future dynamic environments requires changing the very “location where intelligence resides.” This is not a proposal to make a familiar answer a little better. It is a proposal to question the vessel that holds the answer.

To understand why this question appeared, we first have to look at the method that has governed robots and weapon systems until now. That method has been divided into three stages. First, sensors detect the outside world. Next, they send that information to the brain in the middle, the computing unit, for processing. Finally, the processed result is sent back down to arms and legs, wheels and wings, and turned into action. Sense, process, act. These three stages form a loop that keeps turning. Anyone who has studied computer architecture knows the picture. There is input, calculation happens in the middle, and output comes out. For the past half century, machine intelligence has mostly grown inside this frame. The frame has been used for so long that people forget it was a choice at all.

The weakness of this method lies in the fact that information has to keep moving from place to place. Program manager Julian McMorro points out that this traditional loop, which requires data to be moved continuously, holds back the speed and performance of robots. From sensor to brain, then from brain back to limbs, the short time spent sending signals back and forth adds up and slows the response. Under ordinary conditions, it is not a problem. But in the face of a threat flying in at hundreds of kilometers per hour, that slowness becomes the difference that decides survival. A delay of a few milliseconds can determine whether a weapon misses or hits its target, whether a robot dodges or breaks apart. The road that carries the data is the bottleneck. It is not a bottleneck that can be solved just by widening the road. It can be solved only by removing the road itself.

Physical intelligence is an attempt to solve this loop in another way. It does not divide sensing, computing, and movement among separate parts placed apart from one another. It plants those functions directly into the materials, components, and structures themselves. Instead of carrying information to a central processor, the physical design of the object itself reacts immediately. McMorrow describes this as “collapsing the loop.” The goal is to put intelligence directly into hardware and build systems that respond in real time without relying on the nonstop movement of data.

If the words feel difficult, the human body offers a useful comparison. When our hand touches something hot, we do not send the signal all the way to the brain, think about it, and then pull the hand away. The body reacts first at the level of the spinal cord. The brain recognizes “hot” only after the hand has already moved. The picture drawn by physical intelligence is similar. The idea is to make the material that forms the robot’s surface or frame respond reflexively the moment it senses danger. It removes the round trip in which a judgment is sent up to the center and then received back down again. Think not of a robot that thinks, but of a robot that reflects.

There is another point to watch in this shift. It is where DARPA has moved its attention. The focus drops from the finished “machine” to the “material” that makes up that machine. DARPA’s request for information in this area, the official document asking researchers for ideas, states clearly that the focus is on materials, not machines. It does not ask how to build one finished robot well. It asks what basic advances are possible at the level of materials, parts, and the kernel layer beneath them. It shakes the old assumption that computation must happen only on a silicon chip. When the problem is placed at such a low layer, the people searching for answers also change. It is not enough to call in robot designers. Materials scientists, chemists, and physicists have to sit at the same table. That is why this article carries three categories side by side: materials science, robotics, and sensors. The question is not set inside one field. It is placed where three fields overlap.

We need to pause and state this honestly. What I have introduced here is a concept contained in the news item and request for information that DARPA released in the spring of 2026. It is a document meant to gather ideas, not a result already proven on the battlefield. There is still not enough public material to confirm which research efforts under the name of physical intelligence have succeeded and which have hit a wall. So this section does not describe physical intelligence as a completed weapon. It speaks only about what DARPA is now asking and where that question points. The same is true of biological computing programs such as O-Circuit, which we will examine later. The idea of removing the silicon chip and computing with living cells is bold, but it too should be read as a planning document, not as a result in hand. We will leave the detailed story for Chapter 13. For now, it is enough to remember that even such extreme experiments sit on the same current: pushing intelligence into hardware.

Why, then, did DARPA choose this difficult path? The clue lies in battlefield conditions. Think of places where electricity is scarce, remote areas where refueling is hard, and missions that must

enter the middle of an enemy air defense network. In such settings, carrying a heavy, hot computing unit is itself a burden. A large computer consumes a great deal of electricity and gives off heat. Heat becomes an infrared signal and reveals position. If intelligence can be dissolved into material itself, machines can be built that judge with their bodies without carrying a separate large computer. Responses become faster, power consumption falls, and the room for enemy hacking or electronic interference also changes. Of course, all of this remains a goal, not proof. The first rule for reading this chapter is not to read goals ahead of time as achievements.

For Korean readers, the question raised by this shift goes beyond a show of advanced technology. DARPA is not telling anyone to buy a finished weapon. It first asks where intelligence should reside, then gathers materials scientists, roboticists, and biologists around that question. Perhaps what we should envy is not the advanced robot itself, but the ability to lower the layer of inquiry and frame the question again. The way an institution asks a problem is itself its strength. The next section follows what changes when this idea comes down into actual design, and what traces reducing central computation leaves on the body of a robot.

Design That Reduces Central Computation

If the previous section asked, “Where should intelligence reside?”, this section is where that question comes down onto the drawing board. However elegant the idea may be, someone eventually has to draw circuits and place devices. Imagine a small reconnaissance drone hovering over a dusty ridge. The video captured by its camera travels by satellite to a rear data center. Servers there identify the target, then send the command back to the drone. In normal times, it works well. The problem begins the moment the enemy cuts off communications. With one act of radio jamming, the drone becomes a lump of metal with its eyes open and no ability to judge. Meanwhile, its battery keeps draining just to send video to the rear. A body whose eyes are fine but whose head is far away. That is the robot we have had until now.

This scene is the starting point of the problem DARPA has seized. Until now, most machine intelligence has been gathered in one place. Sensors collect information, send it to a central computer, and the computer sends the calculated result back down to components. Sense, process, act. Robotics has long moved on circuits that keep these three stages apart. It is a convenient structure. The problem is the round trip that data has to make to and from the center. It takes time. It consumes electricity. If the path is cut, the whole system stops. When McMorow in the previous section spoke of “collapsing the loop,” this round trip was the loop he meant.

In the same April 29, 2026 article, McMorow states the core point in one sentence. The location where intelligence resides must change. Until now, intelligence has been inside the robot’s “brain,” the central processor. The direction McMorow proposes is to move that intelligence into the hardware itself, into the machine’s structure, materials, and parts. Instead of carrying information to the center, the material itself senses changes in the environment and reacts physically at the same time. McMorow described this as collapsing the loop that had run from sensing through processing

to action. The “design that reduces central computation” covered in this section is one set of faces of that collapse. It is not the name of a single invention, but a direction shared by several programs.

The first point lies inside the semiconductor. Today’s computers separate the device that calculates from the device that stores data. To calculate, the machine has to call numbers stored in memory into the processor, then put them back into memory after the calculation is done. It repeats this round trip billions of times. This structure is called the von Neumann architecture after the person who designed it long ago, and the point where data traffic creates a bottleneck is called the von Neumann bottleneck. DARPA’s fiscal year 2026 budget justification document lists one program aimed at this bottleneck. Its name is OPTIMA. The name carries the idea of optimal processing inside memory arrays. The idea is this. Do not call data into the processor. Finish multiplication and addition directly inside the memory device where the data already sits. If matrix multiplication, the backbone of artificial intelligence computation, is handled inside memory, the power and time lost in moving data can be saved. For a person, it is like reading the answer directly from the bookshelf in one’s hands instead of walking to the study and pulling out a book. If the “reflex” discussed in the previous section is pulled down to the semiconductor layer, it looks like this. Calculation and memory become one body.

A similar idea is moving on the sensor side. The same budget document contains a program called ScAN. Hyperspectral imagery used when an unmanned aircraft scans the ground produces a huge volume of analog signals. That is because it divides and captures tiny differences in the color of light across hundreds of bands. Until now, computers could handle these signals only after converting them into digital values of zero and one. ScAN seeks to skip that digital conversion step itself. It aims to infer directly on the analog signals emitted by the sensor. If the long passage of converting, sending to the center, and computing can be cut away, front-line equipment can make judgments on the spot. DARPA wrote in the document that it aims to greatly reduce the power needed for inference through this method. If OPTIMA reduces the round trip inside memory, ScAN reduces the conversion round trip at the sensor threshold. The targets differ, but the direction is the same. Finish it there before sending it to the center.

We need to draw a clear line. OPTIMA and ScAN are still designs and targets inside budget justification documents. They are not completed weapons, nor are they battlefield-proven achievements. The power savings numbers written in the documents are points the programs seek to reach, not results already in hand. Given the nature of a budget document, that is not strange. Such a document tries to persuade people why tax money should be spent on this research, so it tends to paint a plausible future in bright colors. Readers have to keep the two apart. This is also where the promise of this book, not to read DARPA only as a success story, must be kept. The moment a program’s target and whether that target has actually been achieved are lumped into one sentence, a plan disguises itself as an accomplishment. That is why, even when copying one number, we first have to ask whether it is a promise or a result.

There is also an extreme plan that shows how far the idea of moving intelligence into hardware can go. One example is O-Circuit, a concept that would culture living neurons and use them for computation. The idea is to remove silicon semiconductors and have a cluster of cells smell and choose its own direction. The program document states that purely silicon-based solutions are outside its scope, and instead sets a strange target: a biological processing unit. It even lists a small number, 10 milliwatt-hours, as the daily energy target. It is interesting, but at this point it is not a record of performance. It is an idea written in a planning document. This manuscript will read a plan only as a plan. There is still no public basis for calling it an achievement. The line from OPTIMA through ScAN to O-Circuit can be seen as a set of steps pushing the same direction with different degrees of boldness. Inside the semiconductor, at the sensor threshold, and beyond silicon. The layers differ, but the question is the same.

Why, then, try to scatter computation despite all this inconvenience? The answer lies in battlefield conditions. On the future battlefield, neither stable power nor an unbroken communication network can be guaranteed. If one central data center is attacked or communications are paralyzed, every machine dependent on that center stops with it. One vital point becomes the vital point of the whole. If computation is scattered to the edges, the situation changes. Whether drone or sensor, each continues to make its own share of judgments even where communications have been cut. A design that saves power increases the time a machine can last on a single battery. It is not a design for boasting about performance, but for surviving under the worst conditions. It wears the face of a flashy new technology, but at its root is an old military concern: how to build a body that does not die when cut off.

For Korean readers, this is not someone else's story. Our defense industry and research institutions are also thinking hard about placing artificial intelligence on weapon systems. But those discussions usually lean toward larger servers, faster communications, and more data. DARPA's question points in the opposite direction. Not how much more to add to the center, but what to remove from the center. The moment we bring this question into our own setting, semiconductor design, sensor design, and power budgets that have been treated separately begin to bind into one problem. The question of what a drone over a ridge can do on its own after communications are cut eventually returns to how that drone's body was designed from the start. It means this is not a problem to solve later by adding software, but one that has to be solved from the moment the metal is cut. The next section examines why this design principle appears so sharply in battlefield robots, and what separates them from robots that move boxes in a warehouse.

Why Battlefield Robots Must Compute with Their Bodies

Imagine a robot that carries goods in a warehouse. It stands beside people, moves boxes, stops when someone blocks its path, and displays a smiling face on its screen. Such machines put a familiar appearance and easy operation at the front. They are designed to work alongside people and help them. If they make a mistake, nothing catastrophic happens. If the robot stops for a

moment, a person can walk over and turn it back on. If communications are cut, a manager can call in by radio. If the battery runs low, the robot can be sent back to its charger. The need for the design principles discussed in the previous two sections becomes immediately clear if we place this warehouse robot unchanged on a battlefield.

A robot entering a battlefield faces different demands from the beginning. In late April 2026, DARPA stated in an article on rethinking robots through physical intelligence that commercial robots and national security robots need different abilities. Battlefield robots must operate where communications are easily cut, where humans have almost no room to intervene, and where the environment is hostile and unpredictable. This is a different kind of world from the one a warehouse robot faces. The concern of this section is contained in that one sentence. Where a robot works changes what counts as intelligence for that robot.

The measure of performance changes first. For a warehouse robot, handling large amounts of data quickly can be a point of pride. A battlefield robot is different. DARPA wrote here that performance is determined not by “how much data it handles,” but by “how quickly and reliably it responds.” If the response is even half a beat late, the robot becomes a target. Throughput may be something to boast about, but it does not protect survival. If this contrast feels unfamiliar, it is because we have been unconsciously applying the warehouse yardstick when talking about robots. A smarter brain, larger computation, more data. The battlefield turns that yardstick upside down.

To understand this difference, we have to recall once more the structure seen in the previous section. Until now, the intelligence of robots has generally been gathered in one place. Sensors perceive the outside world, send the signal to a central computer, and the computer judges before giving orders to the limbs. Sensing, processing, and action are separated, and data moves constantly between them. This structure has its own virtues. Gathering computation in one place makes systems easier to build, easier to repair, and easier to improve in performance. In a quiet laboratory, this structure causes no trouble. The path along which data travels is short and safe.

The problem is the battlefield. The short time spent sending a signal far away and waiting for an answer makes the robot sluggish. In places where radio waves do not reach well, such as between city buildings, underground, or in valleys, that round trip itself becomes fragile. If the enemy disturbs the radio signal or cuts communications, the whole structure stops at once. Intelligence gathered in one place becomes a single vital point. Strike that one place and the whole robot is paralyzed. Expensive equipment collapses in an instant into a lump of metal on the road. What was a virtue in the laboratory becomes a weakness on the battlefield. The same design receives the opposite evaluation because of one change in environment. The earlier scene of the drone on the ridge, stopped with its eyes open under radio jamming, is the physical form of this vulnerability.

DARPA’s renewed question begins here. Where should intelligence reside? The proposal is to stop relying on a method that keeps carrying signals around, and to make the physical design itself react immediately. Intelligence is planted directly into hardware, breaking down the boundaries

among the three stages of sensing, processing, and acting. The robot's body becomes the computing device. This is what DARPA calls physical intelligence. The words are difficult, but the target is clear. Remove the distance between the place that judges and the place that moves.

In terms of the human body, it looks like this. When a hand touches a hot pot, we pull the hand away before the brain judges that it is hot. The spinal cord handles it as a reflex, and the signal does not make a full trip to the brain at the center and back. Physical intelligence is close to an attempt to give robots this reflex. Instead of concentrating judgment in the center, it tries to place sensation and response together from the level of the materials and parts that make up the body. The picture is to reduce the time from sensing to action almost to zero. A robot whose fingertips decide for themselves instead of a brain deciding everything. Roughly put, that is the image.

That is why DARPA describes this direction as focused on materials, not machines. It means the work is not about assembling one finished robot better. It is a plan to rebuild intelligence from the level of materials, parts, and the foundation beneath them. Materials that possess their own sense, structures that respond without commands. These are the bottom-layer advances DARPA is targeting. That is also why the article is classified under materials science, robotics, and sensors. The finished product is not the protagonist. The material is. This is not a story about assembling a robot, but about changing the materials that make up the robot itself. If OPTIMA and ScAN from the previous section were answers at the component level of semiconductors and sensors, the material-centered design discussed in this section is the higher-level concern that threads those answers together.

It is also worth noting how this method fits the way the organization works. DARPA is less a laboratory than an institution that poses problems. It does not ask scientists to write good papers. It first makes public questions whose answers may not exist, then demands that researchers bring back the materials and structures suited to those questions. This time, the question placed there is: "Why must intelligence be only in the head?" This question does not order a particular weapon. It tells people to change the frame through which they look at robots. It asks them to set aside the familiar picture of a central brain and imagine a robot with intelligence scattered throughout its body. The point made earlier in this manuscript, that DARPA puts clear goals ahead of good science, repeats itself in this robotics program.

The reason for computing with the body can be gathered into one statement. It is to remove communication delays, remove the center that becomes a vital point, and allow robots to judge by their own power even in environments flooded with radio interference. If intelligence is scattered throughout the body, striking one place does not kill the whole. If the heavy, hot burden called the central computer is removed, the time a robot can last on a front line where resupply is difficult may also change. These are problems that the yardstick of the warehouse robot could not see. Communications, vital points, delay, and power. These four form the wall that separates battlefield robots from commercial robots.

I will avoid exaggeration. What DARPA has presented is not a completed weapon, but a direction for research. There is no evidence anywhere in this article that robots computing with their bodies are already roaming the battlefield. Physical intelligence now stands at the entrance to a concept that would redesign materials and components. Whether this direction will lead to real equipment, and how many years it would take if it does, cannot be known from this document alone. DARPA's preferred method has always been like this. It poses a clear question first and checks later whether answers emerge. It places the problem before the achievement. That attitude is the key to reading the agency. This is also why the previous section warned that the bright numbers in a budget document must be read separately as targets and achievements.

For Korean readers, this scene can be read in two ways. One concerns the direction of technology. When we discuss robots, we often talk about smarter brains and larger computation. DARPA is knocking on the opposite door. It is scattering the brain throughout the body. When we design unmanned systems, this means we must first ask what a robot can do alone at the moment communications are cut. The other concerns the way an organization works. Long before achievements appear, while the issue is still at the material stage, the institution publicly asks the question under its own name. It means opening the field early enough for failure to be allowed. When we think of the way Korean research and development is often chased by performance indicators and cannot allow such early failure, the weight of this attitude feels different. Perhaps what we should envy is not a finished robot, but the posture of putting forward an unanswered question first and enduring it. At this point, where the old assumption about where the robot's brain should be placed is being shaken, Korea has to ask itself whether it will remain a country that buys technology or become a country that frames questions.

At the end of intelligence inside hardware, what we should look at again is not the novelty of a technology name. DARPA's programs turn military questions into tasks that researchers can grasp, then attach schedules and methods of verification to those tasks. That is why the story of one chapter does not stop with an introduction to equipment or models. What remains is who defined the problem, what risk was accepted first, and what was learned when failure occurred. Following this current shows why DARPA is closer to an institution that selects battlefield questions than to a laboratory. When readers think about Korea's defense research, the same question follows. Before increasing the budget, what is needed is deciding which problem to hold to the end.

That choice does not end as a sentence in a conference room. A program manager gives researchers wide freedom while narrowing the goal so that the work does not drift into papers detached from the battlefield. This tension is the uncomfortable strength of DARPA-style research. If freedom goes too far, only papers useless on the battlefield pile up. If the goal goes too far, the work stays with predictable improvements. The hand that holds the space between them is the program manager. A chapter like this is not born if everyone only tries to avoid failure. If failure is not recorded, even the direction of the next experiment grows unclear.

Chapter 8. RACER and the Move Toward Combat-Ready Off-Road Autonomy

It starts with a scene in the dust.

In October 2025, a large unmanned vehicle rolled across the dry ground of Fort Hood, Texas. The driver's seat was empty. The steering wheel a person would have held and the pedals a person would have pressed were all in place, but no one sat in front of them. The vehicle was not following a paved road. It chose its own way across land with no road and turned as needed. Its wheels pressed through soft soil, cut through brush, and climbed over low rises. Then it stopped in front of a mined area. It laid out a long line of explosive charge, fired by rocket, across the ground and detonated it. Where the blast passed, a passage opened for people to walk through.

This had once been a soldier's job. Soldiers had to crawl low into mined ground, checking one spot after another with their hands and tools to make a path. Among combat engineer missions, it belongs to the category where the risk of death is high. This time, a machine went in instead of a person. If it failed, the loss would be expensive equipment, not a human life.

Behind this scene stands one DARPA program. Its name is RACER. Written out in full, it is Robotic Autonomy in Complex Environments with Resiliency. DARPA's Tactical Technology Office, TTO, runs it, and the official program description sums up the work in three words: algorithms, autonomy, adaptability. The words sound stiff, but the problem is concrete. Build a machine that can drive across roadless ground, as fast as a human driver, without a human.

I will draw the line clearly in advance. Most of the evidence used in this chapter comes from DARPA's own official program page and press releases. Press releases, by nature, put achievements first. Failed experiments, budget leakage, and demonstrations that fell short rarely appear in them. So this chapter will introduce what RACER achieved while keeping close at hand the fact that the account is DARPA's own voice. The sources used here do not contain enough independent third-party material verifying the same numbers. Later, when I discuss figures such as 'retraining in one day,' I will also say whose mouth that number came from. This book's rule is not to turn success stories into applause alone.

Autonomy on Roads, Autonomy Without Roads

First, we need to see why this is difficult. Civilian autonomous vehicles have grown remarkably over the past decade. Taxis now drive themselves through cities, and passenger cars that let drivers take their hands off the wheel on highways are on sale. That progress happened under specific conditions. The RACER program description states those conditions plainly. The commercial autonomous driving industry has developed algorithms in road environments where obstacles are limited, the structure is highly organized, and outcomes are relatively easy to predict. Lanes are marked, traffic lights exist, and everyone is presumed to follow the same traffic rules. Companies

improved performance through a fast cycle of driving vehicles, gathering data, fixing systems, and driving again. That method was possible because the road was already an ordered stage.

The battlefield destroys that premise from the start. The tactical technology section in the U.S. Department of Defense fiscal year 2026 budget documents describes in detail how harsh military maneuver environments are. There are hundreds of kinds of obstacles. Things to climb over and things to avoid are mixed together. Soil is soft enough for wheels to sink and spin. Grass and brush tangle in three dimensions and block the view. Above all, there is no fixed road network. There are no driving rules to obey. The promise that you must not leave your lane and the promise that everyone will obey signals do not exist there.

That is not the end of it. On the battlefield, terrain becomes a tool for hiding the body. You survive by moving behind hills and vegetation to avoid the enemy's eyes. The budget documents make this point and state that terrain features must be used for concealment and cover. This is not merely a space without rules. It is a space where it can be advantageous to break rules on purpose. If road autonomy stands on the belief that everyone follows rules, battlefield autonomy starts on ground where that belief does not work.

Today's operational concepts add one more layer. Under the Multi-Domain Operations, MDO, concept mentioned in the RACER description, ground combat platforms do not move as one mass. They must move widely dispersed to gain tactical advantage and improve survivability. If they are clustered in one place, one strike can bring them down together. If they are dispersed, that becomes harder. For dispersed unmanned vehicles to break through rough terrain on their own, central human control of every vehicle is not workable. One operator cannot drive ten vehicles at the same time. So RACER's question is blunt. On land with no roads, no rules, and no preplanned map, can a robot drive itself at a speed comparable to a human driver?

For Korean readers, the weight of this question can be brought home by thinking of the Korean Peninsula's terrain. There are many mountains, deep valleys, and a mixture of fields and low hills. The moment a vehicle leaves a well-paved road, it enters rough ground. If unmanned vehicles are used in an emergency, the ground they meet will not be smooth asphalt. It will be terrain like this. The problem RACER tried to solve is not only a story about another country's desert.

Build the Stack, Not the Vehicle

Once the problem is framed this way, the solution also takes a different shape. DARPA did not choose to build one special unmanned vehicle. In the June 2026 press release announcing RACER's completion, program manager Stuart Young explained that the goal was an Autonomy Stack. A stack is not a specific chassis. It is a bundle of software that brings together algorithms, datasets, and neural network models. We can also say this means building the vehicle's brain, not its body.

The strength of this approach lies in its ability to move from one platform to another. If the right sensors are installed, the stack can be placed on any ground vehicle and turn it into an autonomous driving machine. Without relying on GPS signals and even without any preplanned route, it can push across off-road terrain at mission-relevant speed. The program began in 2021, and the press release says the technology grew quickly from there to this point. Borrowing Stuart Young's phrasing, RACER did not merely duplicate existing military capability. It came from reimagining how the mission itself is carried out.

Why does it matter to build a stack instead of one vehicle? Let us unpack that a little more. When a weapon is completed, it ages together with that weapon. If the brain is fixed into the chassis, the brain must be discarded when the vehicle becomes obsolete. A stack is separated from the chassis, so when a new vehicle appears, it can be mounted again. In practice, RACER placed the same brain on both heavy tracked vehicles and light buggies. In defense procurement, platforms change often, while software lives longer. If a combat vehicle lasts for decades, the software riding on top of it may be replaced every few years. RACER's early focus on portability looks like a design that understood this biology.

For people used to the way weapons are bought and sold, the idea may feel unfamiliar. We often understand defense in terms such as 'a certain vehicle was acquired' or 'a certain robot was deployed.' That is object-centered thinking. RACER asks from the opposite side. What happens if what we buy is not an object but a capability, and if that capability can be carried across several objects? The answer is flexibility. Even if a new threat demands a new chassis, the capability itself can be moved over, so there is no reason to start again from the beginning.

A Machine That Knows Country Roads

Among RACER's technologies, the part Stuart Young explains with care is the Perception Architecture. This is the structure that determines how the machine reads the world and makes judgments. He compares it to driving on a country road. Since he shows a difficult concept first through a scene, we can follow it as it stands.

Imagine driving down a paved country road and cresting a gentle hill. The road beyond the hill is not yet visible. Even so, a person does not slam on the brakes. The body carries the belief that the road will not suddenly end behind the hill. Young calls this a priori insight. It is a sense born of experience, one that does not panic over what it cannot yet see. This is where the difference between a novice driver and an experienced driver appears. An experienced person keeps the flow even through a section that cannot be seen.

There is also the opposite case. Imagine that a car is awkwardly parked on the shoulder of that country road, or that orange cones are lined up. A warning sounds in the mind. There may be road workers ahead, or there may have been an accident, so the foot comes off the accelerator. Even without knowing exactly what is wrong, the mere feeling of uncertainty changes behavior in a more

cautious direction. The two senses point in opposite directions. One lets you move forward even when you cannot see. The other makes you slow down at an unusual signal.

RACER's Perception Architecture imitates these two human senses together. It does not merely react to obstacles right in front of the vehicle. It combines the evidence it has seen so far with prior information and anticipates what may happen next. Even when information is missing or blurry, it forms a prediction about the world and adjusts movement accordingly. This is a different level from earlier autonomous systems that only reacted to a rock one meter ahead. The difference appears as speed. A machine that sees only what is right in front of it keeps hesitating. A machine that looks ahead keeps the flow. The goal of speed comparable to a human driver comes from this predictive ability.

There is also the matter of retraining speed. According to the press release, early autonomous systems took weeks to retrain models so they could adapt to unfamiliar environments. RACER is said to have finished that process in one day. This number needs attention. 'One day' is a value reported by the program side. The sources used here do not include material showing that an independent third-party test organization reproduced the same figure. So it is more accurate to treat 'one day' as a result DARPA reported from its own experiments.

If that figure is true, it is an advantage that is hard to price for units that must quickly send robots into ground they have never crossed before. Operations are a fight against time. If it takes weeks to learn the nature of new terrain, that robot is useless in an urgent situation. Preparation will not finish before the situation changes. One day is different. If a robot can be sent today into terrain that received a deployment order yesterday, the equipment keeps pace with the operation. If this agility is real, it is not one performance metric. It is a condition that changes the entire way the system is used.

Back to the Minefield at Fort Hood

Technology often looks fine only inside the laboratory. Defense history is full of equipment whose graphs were beautiful in papers but failed in the field. What sets RACER apart from other autonomy projects is the fact that it went outside twice with real units in the fall of 2025. It was not a controlled test-course exercise. It was placed in the hands of living units.

The first was the scene that opened this chapter, breaching a minefield at Fort Hood, Texas. DARPA worked with the 36th Engineer Brigade under the U.S. Army III Corps and conducted the demonstration as part of a program called Machine Assisted Rugged Soldier. The equipment used was a large unmanned vehicle called the RACER Heavy Platform. Carnegie Robotics LLC built it on Textron's M-5 chassis. As the word 'Heavy' in the name suggests, it was not a light reconnaissance vehicle but a large, heavy machine.

The Army mounted an M58 MICLIC on this autonomous vehicle. It is a long line-shaped explosive charge pushed out by a rocket, used to open a lane through a minefield at once. The principle is this.

If a long line containing explosives is stretched across a minefield and detonated, the shock of the explosion detonates or neutralizes nearby mines. The narrow, long strip passed by the blast becomes a passage for people and vehicles. The problem is that this equipment must be brought close to the minefield. If people do that work, people must approach the front of the minefield.

In the demonstration, the autonomous vehicle drove itself off-road into the middle of the mined area, detonated the charge, and opened a passage for combat forces to follow. Stuart Young assessed that the demonstration did more than show the technology's promise. It proved how force protection could change in the future. Force protection means protecting troops from threats. It is hard to find a more direct form of protection than placing a machine where a person would otherwise stand in danger.

What deserves attention in this scene is not speed or firepower, but who is in the dangerous place. In minefield breaching, the person who dies is the person who steps on a mine. Putting a machine in that position means that, if it fails, the loss is limited to equipment cost. Unmanned vehicles are expensive too. A large autonomous vehicle such as the RACER Heavy Platform will not be cheap. Even so, it cannot be compared with a human life. The first value of unmanned technology on the battlefield is not firepower. It is this exchange of places.

Fort Irwin's 15 Kilometers

The second demonstration took place a month later at the National Training Center, NTC, at Fort Irwin, California. This is a major U.S. Army site for opposition-force training close to real combat. Its character differed from the first demonstration. It was not a one-time mission to detonate explosives, but a force-on-force exercise in which live units faced each other. That means it was not a scripted display. The opponent was moving.

Soldiers from a reconnaissance battalion under the U.S. Army's 11th Armored Cavalry Regiment used RACER technology on the opposing force side. This vehicle was not heavy. It was a RACER Fleet Vehicle, RFV, built on a Polaris RAZR buggy, and it carried integrated intelligence, surveillance, and reconnaissance, ISR, equipment. It is worth recalling again that the same stack sat on two different bodies: the heavy vehicle in the first demonstration and the light buggy in this one. This is a visible confirmation of the design choice to build the brain, not the car.

The soldiers assigned missions to the vehicles and sent them toward enemy territory. Without a person attached, the RFVs penetrated 12 to 15 kilometers deep into enemy territory and carried out long-range reconnaissance. The press release states that this was a form of autonomous reconnaissance the regiment had not had before. Staff Sergeant Gavin Ros, who watched it in the field, said it was an excellent system that fit well with the purpose they needed, and that he looked forward to seeing how the software would grow. Because this short quotation contains the assessment of a soldier who actually handled the equipment, not just promotional language, it carries some weight. It is the word of a user, not the boast of a maker.

Stuart Young summarized the meaning of this demonstration this way. By reducing dependence on GPS and preplanned routes, warfighters could send autonomous assets into off the grid environments where communications and power networks were cut. Instead of sending human scouts into enemy territory 12 to 15 kilometers away at great risk, people remain in a safer place while robots take on that danger.

Let us pause over the phrase 'off the grid.' Modern units depend deeply on GPS and wireless communications. They get position through GPS and exchange orders by radio. The stronger the enemy, the sooner it targets both. It jams signals, distorts location information, and cuts communications. Such environments are often called communications-denied and position-denied environments. A robot tied to preplanned routes and GPS becomes powerless at that moment. A map is useless if the machine does not know where it is, and a machine that moves only after receiving orders stops when the orders stop. RACER aims at the ability to read terrain and move forward even then. A machine that judges on its own after signals disappear. That is what off-grid autonomy means.

Uran-9 as a Mirror

When explaining to Korean readers why RACER moved in this direction, it helps to place one contrasting failure beside it. This is the story of Russia's Uran-9 unmanned ground combat vehicle being sent to Syria, which I discussed in my earlier book *AI Defense Revolution*. I do not use this case as a new factual claim in this book. I borrow it only as background. Verifying the detailed status of Russian equipment is not this chapter's task, so I will use it only as a mirror that shows the direction.

Uran-9 was an armed robot that looked solid from the outside, equipped with a 30 mm autocannon and even anti-tank missiles. In photographs, it looks threatening. In Syria's complex urban terrain, its remote-control communications repeatedly cut out within a few hundred meters. Buildings blocked the radio waves. When communications were lost, the robot stopped in place or slipped out of control. Its weapons were loud, but because it lacked the intelligence to find its own way, one loss of communications turned it into expensive scrap. No matter how many gun barrels it has, they are useless if the judgment to aim them is not inside the robot.

Put Uran-9's weakness beside RACER's direction, and the picture becomes clear. A robot that depends only on remote control is useful only while communications remain alive. The moment communications are cut, that robot becomes a weak object on the battlefield. The enemy aims at that communications line. RACER's effort to reduce dependence on GPS and prior routes and to put off-grid autonomy first can be read as a design aimed straight at this weakness. If you build a machine that keeps working after the communications line is cut, there is less to lose when the enemy cuts that line.

We also need to take one step back. Uran-9 is Russian operational equipment, while RACER is an American experimental program. One has already operated on a battlefield, and the other has completed demonstrations on training ranges. It is too much to place them on the same scale and judge which is superior. RACER has not driven 15 kilometers in a real war while facing active enemy interference. I will leave Uran-9 only as a mirror that shows 'what problem needed to be solved.' Anyone who has seen the fate of a robot tied to communications can guess why RACER's direction, breaking free from communications, appeared.

Drawing the Boundary of Achievement

To assess RACER's accomplishments, we need to separate two things: what has been confirmed and what has not. If that distinction is blurred, an experimental program reads like a finished weapon, and readers get the wrong picture.

What has been confirmed is this. DARPA conducted two demonstrations with real units in the fall of 2025. One was minefield breaching, and the other was autonomous reconnaissance at a depth of 12 to 15 kilometers. The stack was mounted on both a heavy tracked vehicle and a light buggy, and a positive assessment from a soldier in the field remains on record. These facts are recorded in DARPA's official press release with unit names, equipment names, and personal names attached. The virtue of this press release is that it is not vague self-praise. It is an account tied to verifiable proper nouns.

What has not been confirmed is also clear. Phrases such as 'retraining in one day' and 'speed comparable to a human driver' are explanations from the program side. The sources used here do not state clearly what terrain and conditions produced those values, or how many failed runs occurred. The demonstrations took place on training ranges. Fort Hood and Fort Irwin are training sites close to real combat, but they are still more controlled than an actual battlefield. These materials alone cannot tell us whether the same performance would appear in combat when an enemy actively jams radio signals and spoofs location.

That is why I think it is more honest to call RACER not a 'finished weapon' but an 'experiment that crossed the finish line.' The title of the press release itself used 'finish line,' and that supports this balance. A finish line is an arrival point, but it is also the starting point of the next race. It means one race has ended, not that every race has ended.

DARPA itself left clues about the program's limiting conditions. The press release stated that RACER's driving speed is limited only by three things: sensor performance, mechanical limits, and safety. At first glance, that sounds like a boast. Read in reverse, it also means these three remain the real walls. If sensors cannot see, the machine cannot drive. If the chassis cannot endure, it cannot gain speed. If safety judgment is not settled, it must stop. Without better sensors, a tougher chassis, and more reliable safety judgment, speed stops there. However smart the software is, it is useless if the body cannot keep up.

A Stack Beyond the Battlefield

The story does not end inside the military. The off-road problem RACER addressed is moving directly into the civilian economy. The press release says the program developed dual-use technology. That means it serves military and commercial needs at the same time. It also says private equity money has begun to attach itself to this flow. When private investment gathers around technology created by defense research, it is a signal that some people judge the technology can make money in the market.

Specific company names also appear. They are Field AI, spun out of DARPA-funded research at NASA's Jet Propulsion Laboratory, JPL, and Overland AI, spun out of the University of Washington Robot Learning Laboratory. These companies are moving the off-road autonomy stack demonstrated by RACER into fields such as agriculture, construction, mining, and transportation, opening new markets. There are more industries that must deal with roadless ground than one might think. Tractors in fields, dump trucks in mines, and heavy equipment on construction sites all move through spaces without lanes. There are many places outside the battlefield where technology that drives itself across non-road terrain is needed.

A story from 20 years ago appears again. The press release says DARPA's Grand Challenge opened the era of on-road autonomous driving 20 years ago. A race of unmanned vehicles across the desert became the seed of today's urban autonomous driving industry. RACER positions itself as the off-road version of that path. If autonomy on paved roads created one generation of industry, the narrative goes, autonomy on land without roads could create the next generation.

This is an attention-grabbing point, but it also needs to be read from one step back. The fact that spin-off companies exist and the question of whether those companies will actually build large markets are different issues. No one can guarantee what will happen a few years later to a newly spun-out company. The claim that the Grand Challenge 20 years ago produced today's achievements is also DARPA's own genealogy of the field. Where to place the roots of the autonomous driving industry can differ by viewpoint. The sources used here do not provide enough material either to refute that genealogy or to fully support it. I will present it as DARPA's self-narrative and leave verification for later, when more materials gather.

The Question Korea Should Take From This Chapter

Finally, let us set out what this experiment asks of Korean policy practitioners. The point is not to make a list of things to envy. This chapter does its job only if it makes clearer what should not be copied wholesale.

First comes how the problem is framed. RACER was not 'let us do good autonomous driving research.' It narrowed the problem to 'let us build a stack that can drive as fast as a person on land with no roads and no rules.' Because the goal is clear, the boundary between success and failure is also clear. Did it breach the minefield? Did it travel 15 kilometers? The answer comes as yes or no.

The point where Korean defense research often stumbles is that goals are set broadly, in terms such as 'securing capability,' making success and failure hard to distinguish. Broad goals are also good at hiding failure. It becomes hard to ask what was not achieved. RACER shows the opposite direction. A narrow and clear goal gives neither success nor failure a place to hide.

Next is the sense of investing in software, not the chassis. The decision to build a stack that can move across several vehicles instead of one showcase unmanned vehicle reads the fact that platforms change often and software lives longer. Korea's task is how to plant this sense on top of the inertia of buying weapon systems around hardware. A project to acquire one excellent vehicle is visible and easy to present as an achievement. A project to build a brain that will ride on many vehicles is less visible. The latter is what lasts.

The last point is the attitude toward verification. Throughout this chapter, I used DARPA's press releases as evidence while repeatedly attaching the fact that they are DARPA's own voice. The number 'one day' came with the condition that it was a reported number. The desert genealogy came with the condition that it was a genealogy. This attitude is often what goes missing when people talk about a Korean-style DARPA. It is easy to make a book of success stories. It is hard to record failure and waste along with them. Boasts remain in reports. Mistakes go into drawers.

RACER is striking not because of the demonstration itself, but because the demonstration was done outside with real units. The experiment was taken to the field, placed in soldiers' hands, and the soldiers' assessment was left in the record. Proper nouns and conditions were written together so that achievements and limits could be separated. The real question that the story of roadless battlefield autonomy gives Korea is not when we will build an impressive unmanned vehicle. It is how narrowly we can define the problem, whether we can separate capability from the body, and whether we can be honest enough to add conditions to our own achievements. The question rests on those three things.

Chapter 9 Aircraft Autonomy: Helicopter Transition, X-76, LongShot

This chapter asks how far people can step back from the sky. One story is about taking hands off the cockpit controls. Another is about being freed from the ground called a runway. A third is about separating the pilot from the launch point. Three scenes press the same question from different places. They are DARPA's transfer of autonomous helicopter technology to the Army in 2026, the X-76 experimental aircraft, and the LongShot unmanned aircraft. One point should be clear from the start. Of the three, only the autonomous helicopter technology has actually passed into the hands of an operational force. The other two remain on test ranges and in plans. I want to read them without blurring that line. The moment we write a plan as an achievement ahead of time, we buy a victory that has not yet arrived.

The Transfer of Autonomous Helicopter Technology to the Army

The cockpit is empty. Beyond the glass, instrument lights blink, but no one is sitting in front of them. The rotor begins to turn, and the aircraft slowly rises. From preflight checks to route flight and landing, no human hand touches it. In 2022, a Sikorsky Black Hawk helicopter performed that whole sequence by itself, without a pilot. Reports of helicopters flying on their own had appeared before. This flight had a different character. Even when a system failure was deliberately inserted during flight, the aircraft assessed the situation, responded on its own, and landed safely. A person did not intervene to save it. The machine crossed the crisis by itself.

The program that created this scene is called ALIAS. In English, ALIAS means a system that automates the work of the crew inside the cockpit. The U.S. Defense Advanced Research Projects Agency, DARPA, pushed this project for years. DARPA's aim was not to produce a new type of helicopter. It wanted to add automation to aircraft already used by units, so they could carry out missions more safely and flexibly even on battlefields where enemy threats are dense and conditions are complex. Instead of building an aircraft that did not exist, it inserted intelligence into aircraft that already existed. This difference later explains why the Army welcomed the technology. If a technology rests on top of maintenance systems and pilot training that units already use, the burden of learning new equipment from the beginning falls that much. Buying a new weapon and making an existing weapon smarter receive very different treatment in budget review rooms. The latter usually faces a lower threshold.

Why this approach matters becomes clear when we look at helicopter piloting itself. A helicopter pilot avoids terrain at low altitude while watching weather and enemy air defense threats at the same time. Hands and feet hold the aircraft, eyes move between instruments and the outside view, and the mind keeps making mission judgments. Human concentration has limits. When many tasks rush in at once, something will always be missed somewhere. If an automated system shares the heavy burden of holding the aircraft, the person can spend more attention on judgment. It releases the

human being, step by step, from the physical labor of piloting and lets the human focus on the mental work of mission command. That is the point ALIAS aimed at. The idea is not to erase the pilot. It is closer to redrawing the pilot's role.

In March 2026, DARPA officially announced that this technology had transferred to the U.S. Army. The item transferred was an experimental H-60Mx Black Hawk. It is a fly-by-wire aircraft, which sends and receives flight control signals through wires, and it integrates the Sikorsky MATRIX autonomy package completed with DARPA funding. The delivery took place under a technology transition agreement between DARPA and the Army's utility helicopter program manager. After the stage of proving the concept inside the lab, a real object that an actual force could hold and test crossed over to the Army. The weight of the news lies in the fact that this object is not a paper design or a simulator, but a real Black Hawk that flies in the sky.

At the time of the transfer, Program Manager Stuart Young, who led ALIAS in DARPA's Tactical Technology Office, explained it this way. ALIAS had developed and demonstrated a resilient and flexible automation architecture, and it was now ready to give the Army significant operational advantages. He compressed DARPA's mission into two parts: developing breakthrough technology and reducing the risks that come when the military adopts that technology. The second point deserves attention. DARPA is not a place that builds technology and then walks away. It sees part of its job as lowering the risk until the technology settles into real units. Offering a good idea and making that idea survive in the field are different tasks. DARPA takes the distance between them as its own responsibility. We will hear the same idea again with X-76 and LongShot. Reduce the risk first, then hand it over. That sentence runs through the whole chapter.

The receiver of the technology is the U.S. Army Combat Capabilities Development Command, known as DEVCOM. DEVCOM uses this experimental autonomous helicopter as a flying laboratory. The next stage begins here. One task is integrating advanced mission-specific sensors into the aircraft. The eyes a helicopter needs differ depending on whether the mission is reconnaissance or transport. Another task is examining what operational flexibility comes from reduced crew flight or flight with no pilot at all. One point must be marked. This DEVCOM work is not a finished achievement. It is a test plan to be carried out. It is true that the autonomy technology completed development and demonstration and moved to the Army. That does not mean the technology has become standard equipment in operational units. If we blur this point, we mistake a plan for an achievement. We have to read a planning document as a planning document if we do not want disappointment later. This is also where this book keeps its promise not to gather only success stories and applaud them.

The word transition may sound flat. In the world of defense research, it is heavy. Many research projects leave striking demonstrations behind, yet stop before reaching units. Between success in a lab and adoption in the field lies a wide valley. More technologies disappear because funding ends, they do not fit military procurement procedures, or the burdens of maintenance and training cannot

be carried. They leave one impressive video and are forgotten. That is why the fact that autonomy technology made by DARPA crossed into the Army's hands on a real Black Hawk is an achievement in itself. It means a bridge across the valley has been laid. A bridge does not mean the destination has been reached. The remaining question is how much cargo will actually cross that bridge from now on.

The question this case poses to Korean readers is clear. The organization that makes a technology and the organization that receives it and learns it in the field are different, yet a transition agreement and division of roles between them are written into an institution. DARPA reduces risk and hands it over. DEVCOM receives it as a flying laboratory. Who makes it and who grows it are set down in documents. This overlaps with a familiar frustration in Korean defense research and development: a demonstration succeeds, but units do not accept it; the development agency and the operating unit treat it as someone else's work. Before envying the performance of one autonomous helicopter, we have to ask first. Do we have a bridge that lets technology cross from the lab to the unit, and do we have an agreement that manages that bridge? The transition of the autonomous Black Hawk shows, in one scene, what that bridge looks like. The thing to envy may be not the aircraft, but the bridge.

The Flight Freedom X-76 Aims For

If the autonomous helicopter in the previous section was a story about lightening the burden inside the cockpit, X-76 now asks again about the ground on which an aircraft stands. The center of the problem moves away from human hands and toward the place where aircraft take off and land.

For a fighter pilot, a runway is ground with two faces. A heavy jet needs that ground to rise into the sky. At the same time, that ground is fixed on a map with exact coordinates. Friendly forces are not the only ones who know those coordinates. Enemy reconnaissance satellites and enemy missiles look at the same coordinates. An aircraft needs one long, flat strip of concrete to take off, but the moment that strip is broken, even a state-of-the-art aircraft sits down where it is. Aircraft live only if the base lives. This old pairing has troubled pilots and commanders alike.

DARPA does not skirt this problem. It confronts it directly. DARPA's official 2026 explanation includes this idea: runways have long been helpers that give aircraft speed, but at the same time they have been ropes tying aircraft to one place. DARPA says that rope has created a fatal weakness. The bargain is this: aircraft gain speed at the cost of losing freedom of place. Ending that bargain is the starting point of the X-76 story.

Let me briefly set the background for why this bargain has become a problem now. In the past, runways sat in safe rear areas. The enemy's eyes could not reach that far. Now it is different. Satellites look down from the sky, and precision-guided missiles strike those coordinates from far away. The wider and better paved a runway is, the easier a target it becomes to erase. This is an age in which the ground the aircraft stands on collapses before the costly aircraft does. In this game,

whether an aircraft can take off from anywhere becomes as important a question as how fast it can fly.

X-76 did not appear out of nowhere one day. Behind it is a program called SPRINT. The name stands for Speed and Runway Independent Technologies. The name itself contains the aim: to capture technology that can fly fast while not being bound to runways. DARPA is not carrying this program alone. It is building it with U.S. Special Operations Command. That is worth remembering. The fact that special operations forces are a joint sponsor tells us about half of where this aircraft is meant to be used. Places without runways, deep inside enemy territory, places where forces must enter and leave quietly. Infiltration and extraction, putting troops in and pulling them out secretly, are close to the uses SPRINT aimed at from the beginning.

The experimental aircraft born from this program is X-76. Such aircraft are called X-planes in English. They are challenges marked with the letter X, built as real aircraft to test ways of flying that have not existed before. Bell Textron is now building this aircraft. It has passed the paper design stage. It has cleared the critical design review, the gate that reviews the core design. It received the judgment that the design is mature enough to start cutting hardware.

The name has its own story. The 76 in X-76 points to 1776, the year the United States declared independence. The year 2026 marked 250 years since then, and DARPA chose the number to honor the spirit of that revolution. It engraved the founding year on a new aircraft. The name carries confidence that it will overturn old rules that have governed flight.

What freedom, then, is X-76 trying to gain? DARPA's goals are clear in three ways. It must take off and land on rough ground that has not been prepared in advance. Operations must be able to begin not from a smooth asphalt air base, but from unprepared terrain such as dirt or gravel. It must hover like a helicopter in austere environments without infrastructure. That is the ability to stop and hold in the air. And its cruise speed must exceed 400 knots. That means an aircraft that rises vertically like a helicopter, then flies like a jet once it is in the sky.

Here is why these three goals matter only when they are tied together. Helicopters already have the skill of taking off vertically. But helicopters are slow. The structure of a rotorcraft itself creates a speed limit. Being slow means staying longer inside enemy air defenses, and it also means arriving late. Fast speed, by contrast, belongs to jets. But jets need long runways to take off. The freedom to rise vertically and the speed to fly fast have rarely lived together in one aircraft. Gain one and you usually give up the other. X-76 tries to put both into one body. It aims to hold freedom of place and freedom of time at once.

Let us pause and say this honestly. What I have described so far is not an achievement, but a goal. X-76 had not flown as of this reference date. In May 2025, Bell was finally selected as the contractor, SPRINT moved into its second phase, and the program's effort is now focused not on flight tests but on the stages before them: making parts, bringing them together, assembling them, and testing them on the ground. The number 400 knots and takeoff and landing without runways are,

for now, target values written in design documents and test plans. A specification on paper and an ability proven in reality are different things. X-76 tests the promise made at the beginning: read planning documents as plans. The moment we call it a game changer before the experimental aircraft even flies, we buy the success story in advance.

How DARPA plans to handle this aircraft also shows its character. DARPA is not trying to hold X-76 to the end and make a completed operational aircraft. The demonstrator's role is to check whether key technologies work in representative environments and reduce, in advance, the risks that later operational systems will face. Will the technology work? Will the schedule hold? How much will it cost? This aircraft is the runner that carries that uncertainty first. When that work is done, the aircraft is expected to move to another military organization for the next stage. It is the same path as the autonomous helicopter that moved to the Army in the previous section. The research organization germinates the seed, and the services take over operational use. The difference is that the helicopter has already crossed the bridge, while X-76 is standing before it.

The weight this question has for Korean readers continues from the previous section. We tend to envy a new weapon. What is harder to copy is not the aircraft, but the structure of handoff. One organization stands in front and carries risk, then an operational organization takes over behind it. Even if an experiment fails, the failure is not an organizational shame but an asset for the next design. That division of labor is harder to copy than the number called a 400-knot vertical takeoff and landing aircraft. Korean research and development does not endure failure well. If audits and reprimands come first, no one will attach a budget to an experimental aircraft that has not flown. The real question X-76 raises is not whether we can reach 400 knots. It is whether we have an institution that can set 400 knots as a goal and accept the risk of failure.

If we reduce the flight freedom X-76 seeks to one line, it is this: an attempt to break free both from the rope called the runway and from the speed limit called the helicopter. Yet that freedom still sits on a ground test range. If LongShot, which we meet in the next section, aims at another rope called distance, X-76 aims at two old ropes together: place and speed. We will know whether the ropes have truly come loose only on the day this aircraft first rises vertically from a dirt field.

LongShot and the Distance Problem in Air Combat

If the helicopter lessened the work of the pilot's hands and X-76 redrew the ground under the aircraft's feet, a third rope remains. It is the distance between the person and the target. LongShot aims at that distance.

A fighter flies toward enemy territory. The pilot has not yet seen the enemy aircraft with the naked eye. Only a small mark floats far away on the screen. Air combat today begins this way. It is not old-style close air combat, where pilots see each other and chase tails. The side that shoots first with radar and long-range missiles from beyond visible distance wins. Soldiers call this kind of fight beyond-visual-range combat, or BVR.

The problem appears at a certain point. No matter how far a missile can fly, it only goes as far as its rocket fuel allows. After it flies for a long time following launch, the missile loses energy. By the time it barely reaches the enemy aircraft, it may not have enough energy left to follow if the other aircraft suddenly turns away and runs. To hit with confidence, the missile must be fired from a shorter distance. Getting closer means a crewed fighter goes inside enemy air defenses and risks the pilot's life. To extend range, one must push a human being into danger. This is the distance problem air combat has carried for a long time.

This knot cannot be untied by technology alone. A costly fifth-generation fighter flies with the training and life of one pilot inside it. For a commander, deciding to push such an aircraft within range of enemy missiles is always a weighing act. How firmly we can hold the target and how many of our pilots we are prepared to lose sit on the same scale. Range and safety have long pushed against each other. Gain one and you had to give up the other.

DARPA's LongShot program rebuilds that scale from another direction. The idea is this: let an unmanned aircraft fly the dangerous final segment instead of a person. LongShot's official goal is to create an unmanned, air-launched aircraft that can carry current air-to-air weapons as they are, greatly increasing engagement range and mission effectiveness.

The sequence is easy to understand if we picture it. First, this unmanned aircraft is launched from a fighter or bomber. After launch, it does not rush to accelerate. It saves fuel and moves slowly, entering toward enemy territory ahead of friendly crewed aircraft. Then, when the final engagement moment comes and the distance to the enemy aircraft has narrowed enough, the unmanned aircraft fires the high-energy air-to-air missiles it carries from that location on its own. From the missile's point of view, the starting point has moved much farther forward. Instead of spending its strength flying a long distance, it is launched right in front of the target and uses the rocket's power to chase the enemy aircraft. The unmanned aircraft, in effect, carries the missile to the enemy's chin.

In this structure, the position of the crewed fighter changes. The pilot can remain far behind the forward edge. The unmanned aircraft takes on the dangerous entry, and the person commands the overall picture from a safe distance. DARPA explains that this approach greatly improves pilot safety while expanding the force's overall reach and strike effect. The balance point of the scale shifts merely because one aircraft with no person aboard enters the space between the person in the cockpit and the target. If X-76 in the previous section pushed the runway backward, LongShot pushes the person backward. The idea of placing machines forward in the space people have vacated repeats throughout this chapter.

LongShot has another design principle worth watching. It is not tailored to one specific fighter. The developers conceived this unmanned aircraft so it can depart from many kinds of host aircraft. It can be hung outside a fighter, carried inside a bomber's weapons bay, or dropped from a transport aircraft as palletized munitions. They even leave open the idea of a cargo aircraft pushing out multiple unmanned aircraft from the rear. The developers call this quality host-platform independent.

It means that even without enough advanced fighters, existing transports and bombers can become launch platforms for long-range air combat. The image of throwing many shots forward from cheaper launch platforms rests on a different calculation from betting everything on a small number of elite aircraft.

The concept has moved beyond paper and into a real aircraft. This air-launched unmanned aircraft has received the official experimental aircraft designation X-68A. General Atomics Aeronautical Systems, or GA-ASI, was selected as the industry lead for development, production, and demonstration. When a small unmanned aircraft fires a large missile, recoil can easily shake the aircraft or make it lose attitude. This problem of stability and control has been regarded as one of the program's major technical hurdles. The reaction at the moment one missile is pushed away can flip a light aircraft as a whole.

It is also striking how many hands gather to solve one problem. DARPA has tied several services and agencies into this program. Air Force flight test units and laboratories, Navy aviation-related centers, NASA, and Army test ranges all participate. Among them, the Navy and Air Force have been designated as transition partners who will actually receive and use the technology if it is completed. The goal of handing it to the military rather than ending with research appears directly in the list of participating organizations. It also means this is not a single laboratory's achievement, but work that can move only with agreement across several services. This is the third handoff structure we meet in this chapter. The helicopter to the Army, X-76 to another military organization, LongShot to the Navy and Air Force. The habit of pairing the maker and the receiver in advance is the backbone of the DARPA method.

The progress so far is managed by a team led by Colonel John Casey. The team said it had passed full-scale wind tunnel testing, testing of the parachute system that recovers the aircraft, and testing of the weapon release system that had been considered difficult. The next goal is a flight test in the real sky. DARPA has presented a flight test plan aimed as early as late 2026. The sequence is to build and verify test aircraft, integrate software, receive airworthiness certification, and then carry out captive-carry tests while flying together with tactical aircraft such as the F-15. At the end, the program will confirm that the unmanned aircraft can safely release and fire its own missile during flight.

One more step back is needed for a cool reading. As of this reference date, X-68A had never flown. It had passed wind tunnel and ground tests and set a flight test schedule. It had not shot down an enemy aircraft in combat. Colonel Casey's statement is also about direction, not completion. He said LongShot reduces technical risk and shows a viable path for the military to move toward unmanned air launch. A viable path. That means a passage, not an arrival. Plans must be read as plans, and tests as tests. Across three sections, we have repeated the same caution three times. The helicopter has transferred but has not become a standard. X-76 has passed design review but has not flown. LongShot has passed ground tests but has not risen into the sky.

Seen through policy, the question deepens by one layer. What LongShot targets is not the performance of one weapon, but the calculation that makes up air power itself. Until now, a strong air force has been measured by how many costly crewed fighters it has. If a layer of unmanned launchers is placed in front of them, the standard of calculation moves from the number of pilots to the number of launch bodies that can be thrown forward. The side that can put many threats forward with less money may gain an advantage. This trend is not only an American story. For a country like Korea, which does not have an abundant number of elite aircraft, the idea of increasing range while exposing people to less danger does not sound remote. This is not a call to rush into changing our force structure around a concept that has not yet been proved. Yet the question of how far we can move the person in the cockpit backward will soon appear in budget rooms as well. To answer it, we first need an institution that can fund a concept aircraft that has not flown and endure failure. The handoff structure and the culture that endures failure, repeated three times in this chapter, are lessons that will last longer than the three weapon names.

LongShot stands in line with the autonomous helicopter moving to the Army and X-76 asking again about freedom of flight. The helicopter lightens the pilot's hands. X-76 redraws the limits of flight itself. LongShot separates the person from the launch point. The three programs push the same question from different places. How far can people step back, and how far can machines move forward? This chapter's question about aircraft autonomy appears sharply at this point, on the distance problem in air combat. The answer is inside the test range. What we need to hold on to now is the honesty not to write that answer as success in advance, and the question of whether our own institutions can endure such experiments.

At the end of Aircraft Autonomy: Helicopter Transition, X-76, LongShot, the point to look at again is not the novelty of the technology names. DARPA's programs turn questions held by the military into tasks researchers can grasp, then attach timelines and verification methods to those tasks. That is why the story of one chapter does not stop at introducing equipment or models. What remains is who defined the problem, which risks were accepted first, and what was learned when something failed. Follow this flow, and DARPA looks less like a research institute than an organization that chooses the battlefield's questions. When readers think about Korean defense research, the same question follows. Before increasing budgets, we have to decide which problems we will hold to the end.

Chapter 10 Next-Generation Microelectronics and Photonics

JUMP 2.0 and the University Research Ecosystem

Imagine a semiconductor laboratory at an American university. Late at night, several graduate students place a thin wafer under a microscope and peer into it. Their adviser is revising slides for a conference presentation next month. There are no photographs of weapons on the wall, and no one in uniform walks down the hallway. From the outside, this room seems to have nothing to do with the battlefield. Yet if we trace, one by one, the sources of the research money flowing into such a laboratory, a familiar name appears. The Defense Advanced Research Projects Agency of the United States, DARPA. There is a point where basic semiconductor research and the future battlefield connect in a single line, and in that link sits a program called JUMP 2.0.

JUMP 2.0 stands for Joint University Microelectronics Program 2.0. The name is long and stiff, but the work it aims to do is clear. It ties research centers at several universities into one network so they can build the next generation of semiconductor technology together. DARPA places this program under the larger umbrella of the Electronics Resurgence Initiative. The word “resurgence” is there for a reason. It carries the intent to redraw the map of the electronics industry, an industry in which the United States once led.

This program did not appear overnight. It stands on an extension of 25 years of cooperation among DARPA, the Semiconductor Research Corporation, SRC, and industry. That long relationship matters. A structure in which the government sets the broad direction, industry supplies money and people, and universities take on risky basic research is not built in a day. It takes a long time just to find out whether each side can trust the others. Through this public-private partnership, DARPA helps university research centers take the lead in solving major problems in information and communications technology. The program’s goal is plain. Through research that shakes the field, it aims to keep the United States at the front of microelectronics.

Why universities, of all places? It may seem faster to pour money into one corporate laboratory, yet DARPA ties several universities into a single consortium. The reason lies in spreading risk. If all basic research is bet on one direction whose success is unknown, the loss is large when that path closes. If several universities explore different materials and structures at the same time, the chance that one of them opens a path rises. Universities are also a spring that keeps producing young researchers and new ideas. While corporate laboratories watch the products that can be sold now, universities have the room to hold on to questions that may matter only ten years later. DARPA tries to keep that spring from drying up while turning its watercourse toward defense and industry.

Then what problems have been handed to university researchers? The U.S. Department of Defense budget documents for fiscal year 2026 list five major technical challenges. How newly can

analog hardware be built? How can memory and storage be prepared for the explosion of data? How can the gap be narrowed between the volume of data pouring in and the communications capacity needed to carry it? As everything becomes connected through artificial intelligence, how can new security holes be blocked from the hardware level? How can the hard-to-manage rise in power demand for computation be controlled? The five questions are not separate. Solving one makes another grab the ankle again. Change a device to save power, and the communications method becomes a problem. Touch the communications system, and security starts to wobble.

To take on these challenges, university research teams divide seven basic technology themes among themselves: cognition, communications, sensing to action, computing and processing, memory and storage, integration and packaging, and devices that deliver high performance while using little power. Listed by name, they look like a table of contents from a textbook. Yet these themes are closer to an instruction to mix new materials and devices, structures and algorithms, design and integration technologies in one place. Inside them is a demand that physicists, materials engineers, and software researchers collide within one project and cross the borders of existing fields.

A closer look at the budget documents reveals one notable change. Starting in fiscal year 2026, the budget line for JUMP 2.0 moves. The program, which had been placed under “Electronics Science” within “Defense Research Sciences,” is reorganized under a new category called “Emerging Opportunities,” in a line named “Disruptive Microtechnologies Sciences.” If we look only at the fact that a budget box changed, it can seem like a small administrative matter. Yet it is worth remembering that the word “disruptive,” meaning something that overturns the field, has now been attached to the line’s name. Even the same research carries a different organizational weight depending on which drawer it is placed in. This change in name quietly shows where DARPA wants to place university semiconductor research.

There is something to keep in mind. Most of what I have written so far is the “goals” and “plans” contained in DARPA’s program description and budget documents. It is not a results table saying that the five challenges have been solved to a certain extent or that the seven themes have produced specific outcomes. A program page is a document that records promises about what will be done, and a budget document records a plan for where money will be placed. It is hard to say with confidence, based only on the public materials now available, what JUMP 2.0 has actually produced and how much. I do not want to color that blank space as if it were an achievement. I will keep separating plans from results to the end. Introducing DARPA as a well-built institution is different from boasting early about results that have not yet appeared.

Even so, the outline this program draws is clear. If universities are a spring of ideas, JUMP 2.0 is the channel that turns the water from that spring toward defense and industry. The materials and structures refined at universities lead to NGMM, the manufacturing program we will meet next in this chapter. They also extend to PICASSO, which handles light as a signal, and to RAAPTR, which

handles the radio-frequency environment. The trouble of chips growing hot is taken up by research such as THREADS and Thermonat. Basic research, manufacturing, and devices are designed to interlock like a chain rather than move separately. If you read this chapter to the end, you will see the route by which an idea that begins in a university laboratory flows outward.

The name of the office that houses JUMP 2.0 fits this picture as well. The program belongs to DARPA's Multi X Office, MXO, and its listed topics include electronics, microchips, and sensors. The office's character, which crosses several technologies rather than staying confined to one field, overlaps with the program's direction of broadly tying together the university research ecosystem. Even a position on an organization chart carries intent about how a program will be treated.

For Korean readers, this part leaves both envy and homework. A familiar model in which one government research institute takes charge of a core technology starts from a different place than a model that has tied several universities and companies into one consortium over 25 years. Before cutting off the question by saying which side is right, there is something we must ask first. Have we ever connected basic research at universities to the problems of defense and industry for this long and this tightly? If the answer does not come easily, that hesitation itself is a starting point. I will return to this question in the final part.

NGMM and Next-Generation Microelectronics Manufacturing

In a research complex in Austin, Texas, a new semiconductor production line is going up. Its goal feels a little unfamiliar. It is not a race to etch smaller transistors. It is the work of stacking chip pieces made in different factories, with different materials, into one whole. A communications chip is placed on top of a chip that handles computation, and another device is layered above it. It is close to picturing Lego blocks stacked vertically.

DARPA calls this method three-dimensional heterogeneous integration. Its abbreviation is 3DHI. The name is difficult, but the idea is clear. There is a physical limit to continuing to divide circuits ever more finely on a single silicon wafer. Once transistors have been reduced to the width of only a few atoms, there is no more room to shrink them. For more than half a century, the force that drove the semiconductor industry was the promise of etching more transistors into the same area. That promise has reached the wall of physics. So the answer is to build upward instead of shrinking sideways. It means not insisting on silicon alone, but attaching parts with different characters, such as compound semiconductors strong in communications or devices that carry signals through light, inside one package. DARPA sees future microelectronics innovation as turning on this kind of fusion. The contest is not only the performance of each part, but how tightly parts with different characters can be attached and made to move as one.

The program that moves this idea into an actual factory is Next-Generation Microelectronics Manufacturing, or NGMM. DARPA's stated goal is clear. It wants to establish a manufacturing center inside the United States that can support research and development and low-volume pilot production

using 3DHI. It also aims to make a self-sustaining center that continues operating after government support ends. Even if a researcher wants to produce a prototype, an idea stops on paper if there is no domestic line capable of stacking and bonding chips in three dimensions. NGMM is an attempt to place that line on American soil.

There is one point Korean readers may easily confuse. When the United States talks about building semiconductor factories at home, the CHIPS and Science Act of 2022 comes to mind first. NGMM is funded from a different pocket. It is a Department of Defense budget item. If the CHIPS and Science Act is directed at securing production capacity inside the United States now, NGMM puts defense money in advance into somewhat more distant technologies that have not yet settled into commercial markets. The two projects have different destinations. One fills today's production gap. The other plants seeds, in the name of defense, in technologies whose value may be priced only years later. Because it aims at both commercial and security uses, NGMM was designed from the start as a dual-use project.

The center is called TNC. It stands for TIE NGMM Center. DARPA is building it with the University of Texas at Austin and its Texas Institute for Electronics, TIE. TNC is not a factory belonging to one company. It is an open research line where defense companies, domestic foundries, startups, design firms, and universities can all enter. This character sets the direction of NGMM. If it were a private factory locked behind the doors of one large corporation, a graduate student's experiment and a small startup's prototype would both turn back at the entrance. Bringing equipment into one place for shared use has the same meaning. If manufacturing, assembly, measurement, and testing equipment are scattered across different buildings, making even one chip can take months. When they are gathered under one roof, when standards between machines are matched, and when processes are standardized and automated, the time needed to turn a new chip into a prototype shrinks. The distance from idea to physical object becomes shorter.

Building new hardware is not enough. Design software must follow. The electronic design automation tools used in the market today are made for drawing flat chips. They are not enough to handle the complexity of stacked three-dimensional structures. U.S. Department of Defense budget documents also state that commercial EDA tools lack the capability and flexibility to handle 3DHI systems. If the tools cannot keep up, designers cannot command the hardware no matter how far the hardware has advanced. That is why NGMM builds software as well. It includes a three-dimensional assembly design kit, 3D-ADK, containing design rules for assembling multiple chips into one; an EDA workflow for three-dimensional circuits; digital twin simulations that run performance virtually before making the physical object; and test standards for judging performance. These tools are useless if they do not match the processes of the actual factory. The program lays the ground for foundry, packaging, and software engineers to work in step in the same place. Rather than making the design kit chase the factory process after the fact, the structure ties the design kit to process changes as they happen.

Let me be clear about the posture here. It would be wrong to read this center as if it has already produced results. NGMM is a plan under construction and a worksite. The program is divided into phases. Phase 0, which drew the basic picture, ended in 2023. Participating research teams defined and analyzed what kind of 3DHI systems they would build, setting the framework for establishing the center. Phase 1, which builds the center's skeleton, was formally awarded in July 2024, and a kickoff event was held in Austin that September. Each build-out phase is set at two and a half years, and the work of producing actual prototypes belongs to the following Phase 2. A long stretch still lies between holding a kickoff event and printing chips.

What happens in Phase 1 gives a sense of the project's scale. It establishes a manufacturing execution system that oversees processes and equipment, data collection, and experimental schedules, and it develops a security plan that covers physical security and data protection. It buys, installs, and qualifies the core equipment needed to bond chips in three dimensions. Reference processes such as thermocompression bonding, which joins chips with heat and pressure; hybrid bonding, which directly joins surfaces; and patterning, which forms metal interconnects, are demonstrated in this phase. All of these are basic household provisions needed before prototypes can be made. It is not about producing a dazzling result first. It is about building the body that can produce results.

Here is the point that must be read honestly. If we look at the fiscal year 2026 task list in DARPA's budget documents, the items are not completed results but a row of things to be established going forward. At this point NGMM also moves one budget line over, from the existing PE 0603739E, MT-16, to a line under Advanced Disruptive Microtechnologies, PE 0603469E. The newly listed homework is as follows: a user access model that determines how industry and academia will share intellectual property rights, a self-sufficiency plan through which the center survives on members and customers after government money ends, and the actual production of reference test chips such as power amplifiers, low-noise amplifiers, and interposers. Whether a self-sustaining center can truly sustain itself has not yet been proven. It is only written as a goal. After government funding pours in the priming water, will member companies and customers pay enough to keep the line running? The answer to that question will come not from the budget document, but from invoices several years later.

So I recommend reading this section this way. Do not see NGMM as a weapon already in America's hand. See it as a large bet the United States is placing. It is an attempt to bring advanced chip production, concentrated in several Asian foundries, back inside the United States at least on the defense side. For Korean readers, this is not someone else's problem. Korean companies stand in the middle of the "Asian foundries" the United States wants to bring back. Why would the United States spend even defense money to plant a manufacturing ecosystem on its own soil? Because entrusting advanced chips to factories in other countries has come to be read as a weak link in security. Among the factories toward which that gaze is directed is Korea's name. It is an object of envy and, at the same time, a mirror that makes us look again at our own position. At the place

where the United States calls factories back in the name of defense, we place the position of Korean semiconductors beside it.

Reading PICASSO, RAAPTR, and the Heat Problem Together

Anyone who has put a hand behind a radar antenna knows it. The harder radio waves are pushed out, the hotter the equipment becomes. Heat rises along with signal strength. On the battlefield, this heat eats away at performance, shortens component life, and adds the extra burden of cooling equipment. Size, weight, power, and cost. The military bundles these four as SWaP-C. The more performance one wants, the more these four grow together. Raise performance by a handspan, and weight grows by a fistful. When weight grows, more cooling is needed. Add more cooling, and power demand and volume grow again. Heat sits at the center of this vicious cycle.

When DARPA deals with next-generation microelectronics and photonics, several programs with different names appear on the surface: PICASSO, RAAPTR, and THREADS and Thermonat, which aim directly at heat. By name, they seem separate. Step back one pace, and they are striking the same wall from different places. That wall is heat. Rather than separating the three paths, I will read them together as three faces of one problem.

First comes the side that meets heat head-on. In June 2026, DARPA released an article announcing a program called THREADS. Its title is interesting. It used the image of threading a needle through the thermal barrier created by high-frequency power. The program aims at how heat creates a ceiling for performance in devices that handle high-frequency, that is RF, power, and how that ceiling can be pierced. The phrase thermal management sounds bland. In the field, this one factor determines radar output and communications range. The temperature a device can withstand is the limit line of the power that equipment can produce. THREADS tries to push that line upward from the device and heat-path side.

Beside it is Thermonat. In May 2026, DARPA said this research had raised nanoscale thermal prediction to a level usable in actual chip design. The idea is to calculate in advance where and how much heat will build up inside a chip, down to extremely small scales, and act at the design stage. It reverses the order of making a chip first and then checking whether it runs hot. If hot spots can be identified on the design drawing in advance, the cost and time of remaking prototypes several times can be reduced. Yet this announcement is DARPA's own description. It is more honest to read it as a progress report on ongoing research than as a result backed by a thick pile of outside independent verification. If we blur that distinction, we mistake a plan for an achievement.

If what we have seen so far is the path of controlling heat through materials and prediction, PICASSO is an attempt to change the very cause of heat generation. PICASSO's full name is Photonic Integrated Circuit Architectures for Scalable System Objectives. The idea is this. Today's electronics convert light into electricity and electricity back into light when processing signals. Each time this optical-electrical conversion occurs, energy leaks and heat is produced. Each conversion

point is a small stove. Then what if signals were not changed into electricity, but calculated longer and deeper while remaining as light? That is PICASSO's question.

The metric to watch in PICASSO is circuit depth. DARPA defines it as the number of optically connected functional elements through which a light signal passes before ending in optical-electrical conversion. A circuit that turns light into electricity after one or two stages differs in heat generation and efficiency from a circuit that passes through many stages while remaining in the light state. If the number of conversions is reduced, the leaking heat falls by the same amount. So this metric does not ask how many optical parts were inserted. It asks how deeply the work was done with light.

To reach this goal, participating teams use commercial foundries and try new material combinations, such as placing thin-film lithium niobate on silicon. It is a way to attach and use the strengths of different materials on a single chip. To draw large-scale photonic designs in reality, electronic design automation tools are also needed. As circuits grow more complex, drawing interconnects by hand becomes impossible. DARPA maintains a design repository so related organizations can access it, and it allows participants to bring in intellectual property they previously created at their own expense. In return, the government secures data rights at the level of government purpose rights so the technology can move into other weapons systems. The structure puts a lock in place early so a good circuit does not remain trapped only in one company's vault. It is worth remembering that this document is not a results report boasting of achievements, but a question-and-answer document that defines how the ground will be laid.

If PICASSO changes computation inside a device into light, RAAPTR tries to change the radar and communications through which the device exchanges signals with the outside into light. RAAPTR's full name is Radio Frequency Architectures Applying Photonic Timing and Routing. Its core material is a hardened optical frequency comb. An optical frequency comb lines up the colors of light densely and evenly like a ruler and is used to measure time and frequency with extreme precision. For a long time it belonged to precise laboratory experiments. RAAPTR tries to pull this object into a battlefield that shakes, heats up, and fills with dust.

The real difficulty appears in this scene. An oscillator that runs well on a laboratory bench is shaken by vibration on a vehicle or aircraft. Vibration means signal jitter, and jitter means noise. RAAPTR demands analog subsystems that satisfy constraints on size and weight, power and cost, as well as harsh environmental standards. At the center is the design of a durable oscillator that can suppress and withstand vibration. DARPA has said it will select three to five teams for this program and put in 58 million dollars, hundreds of billions of won in Korean currency, through an Other Transaction agreement. The plan is to design architectures in Phase 1 and then expand in Phase 2 to a working system, reaching a demonstration at the level of light-supported radar. It is a staged structure: first make paper designs compete, then let the surviving side build the physical system.

This program does not open its door to everyone. Federally funded research and development centers, including national laboratories, and university-affiliated research centers may not participate

as performers or subcontractors. Only organizations inside the United States that have formally received the controlled unclassified information annex from DARPA may serve as prime contractors. The durable oscillator design itself is also classified as non-fundamental research protected under controlled unclassified information guidance. A threshold at this level tells us that the Department of Defense sees this technology as capable of determining superiority in electronic warfare and communications. The program stands somewhere between basic research that anyone can publish as a paper and applied research handled behind locked doors.

Set the three paths side by side, and the picture becomes clear. THREADS and Thermonat control heat that already exists through materials and prediction. PICASSO tries to reduce the cause of heat at the site of computation. RAAPTR aims at the same thing at the site of sensing and communications. The structure strikes one heat problem simultaneously from three points: materials, computation, and communications. It is no accident that most of these programs fall under the same Multi X Office. JUMP 2.0, which ties together the university research ecosystem, and NGMM, which targets next-generation microelectronics manufacturing, also sit under the same office. One office appears to be placing materials, manufacturing, computation, and communications on a single map and moving them together. When different programs sit under one roof, materials or design tools gained on one side can more easily cross to another.

Here lies the question for readers who deal with policy. The point to envy is not the flashiness of each program, but the way the problem is divided. Taking heat as the large wall, one office separately sets up teams that cool heat already produced, teams that reduce heat at its computational source, and teams that reduce it in communications, then runs them side by side. The budgets and goals are managed in different documents, but the hand overseeing them is one. Designed this way, even if one program falls, another branch survives, and a result from one branch becomes material for the neighboring branch.

Even so, it would be wrong to read this picture as a victory already won. PICASSO's question-and-answer materials, as well as RAAPTR's Q&A and proposers day presentation, are documents from before contracts are signed. Circuits that compute deeply with light and optical oscillators that withstand battlefield vibration are still goals and bets. They are not score sheets guaranteeing success, but more like tables showing how much will be wagered and where. The figure of 58 million dollars is not the size of the achievement. It is the size of the bet. The useful point for Korean readers is right here. The thing to learn from is not a completed photonic radar, but the design method: arrange programs so they strike one wall from several angles at once, and record each of them honestly as plans or results. The day that design actually breaks through the wall remains something each person will have to confirm with their own eyes at a demonstration site years from now.

Chapter 11 Space Robotic Servicing and Cybersecurity

The Meaning of Geostationary Satellite Servicing

There is a place 36,000 kilometers above Earth. In miles, that is about 22,000 miles. The number alone feels remote. Yet what happens there is unexpectedly orderly. Hundreds of satellites orbit at the same speed as Earth rotates. From the ground, they seem fixed at one point in the sky. Once a dish antenna is pointed in that direction, it can keep looking at the same satellite without being adjusted again. Because of this fixed quality, services used for military communications, weather observation, and national security continue without interruption at this altitude. This place is called geostationary orbit.

The distance that supports this convenience is also a curse. Once a satellite reaches geostationary orbit, human hands can never touch it again. If even one part starts behaving strangely, there is no way to go close and look inside. It cannot be diagnosed, repaired, or fitted with new equipment in place of old hardware. There is an even more painful case. The satellite body may still work perfectly, yet the mission must end only because the equipment inside has become outdated. DARPA described this situation as deeply frustrating for owners of assets worth hundreds of millions of dollars. An expensive machine is still alive, but it is abandoned because there is no way to act.

This frustration has distorted the way satellites are built. Since they cannot be repaired after a failure, satellites carry the same function in two or three layers. If one dies, another is supposed to take over. To extend service life even a little, they also carry as much fuel as possible. That makes satellites heavier and more complex. As weight rises, manufacturing costs rise, and launch costs rise with them. Because repair is impossible, they carry many spare parts, and because they carry so much, they become more expensive. A large share of the money spent on a satellite goes not to mission equipment, but to insurance against possible failure. Before the satellite even reaches orbit, fearful design consumes half its weight.

DARPA's question begins at this point. Could that unreachable altitude become a place that can be serviced in a reliable way? The program that gathered this question into one effort is Robotic Servicing of Geosynchronous Satellites, or RSGS. The name is long and stiff, but the work is easy to picture. It means building a vehicle that approaches a failed or aging satellite in space, inspects it with robotic arms, and services it. When that vehicle arrives, a satellite gains years of extra life in orbit. In DARPA's program description, RSGS is listed as a robotics, satellite, and space program led by the Tactical Technology Office. It is not an idea locked inside a concept drawing. It is classified as a plan to take real action in orbit.

The way it is being built deserves attention. DARPA is not building the whole satellite alone. DARPA is responsible for the robotic payload. This is the part that combines hardware, software, and tools that can be swapped depending on the situation. A private company brings and attaches

the spacecraft bus that serves as the body. Together, they form the robotic servicing vehicle, or RSV. The finished vehicle is owned and operated not by the government but by a commercial partner. The structure is that the state gets through the difficult technical threshold, and private industry turns it into a working business. Control is not left only to full automation. It mixes a high level of automation with remote operation from the ground. Humans keep final judgment on Earth, while machines handle fine movements beyond human reach.

The work itself sits on the edge of human technology. The vehicle must approach a satellite moving through vacuum faster than a bullet and align alongside it without collision. Then it must use robotic arms to touch it precisely. The hard problems DARPA has named show the difficulty clearly. They include making tools mounted at the ends of robotic arms meet requirements, moving the servicing vehicle between orbits while saving fuel, moving the arms with precision, automating certain actions, and building a foundation for coordinated control between the servicing vehicle's operating team and the customer's team that operates the satellite being serviced. The last item is easy to misunderstand. The moment one satellite is serviced, the satellite's original owner and the repair vehicle's operators issue commands related to the same object at the same time. If the two teams' signals fall out of sync, repair turns into an accident. How to divide and transfer control authority becomes a research topic in itself.

If that is the picture, the schedule is the plan that pulls the picture into reality. The program is entering the stage just before launch. It completes integrated testing of the robotic payload and spacecraft bus, trains the commercial partner, and runs rehearsals based on real mission scenarios. In 2026, the integrated robotic payload is to be launched into orbit, then checked and calibrated there. The goal is not to service once and stop. The goal is for the servicing spacecraft to carry enough propellant and equipment capacity to conduct dozens of operations over many years. One repair vehicle survives for a long time and takes care of multiple satellites in orbit one after another. If this picture works, satellite designers can remove some of the fearful insurance load and fill that space with better sensors and transponders. Weight falls, cost drops, and the cycle begins to turn the other way, toward better mission performance.

The United States is not only seeking the technology to launch robots. It is trying to grasp the rules of this new industry first. Through a consortium called CONFERS, DARPA is helping create standards for on-orbit servicing. It brings private and government experts together to study and publish consensus-based safety standards without forcing them on participants. The purpose is to document in advance the procedures for approaching another party's satellite safely. If technology comes first and rules follow late, accidents write the rules. This is an attempt to reverse that order and lay down rules first.

One uncomfortable truth must be stated honestly. The ability to approach another party's satellite and replace parts becomes, if turned in the other direction, the ability to damage that satellite. The distance between a repair robot and an attack satellite is shorter than it may seem. If a machine can

touch a solar panel in orbit, it can also fold or break that panel. The United States' rush to set standards and norms suggests an intent to restrain this double-edged quality through international rules. A country that sits in the seat where norms are written first also gains the basis to call another country's later conduct a violation of those norms. Getting ahead in technology and getting ahead in language become one task.

For Korean readers, this story leaves two questions. One is an industrial question. If a market opens for repairing satellites in orbit and extending their service life, the calculations of countries that build satellites and countries that supply parts will change. Once repair becomes possible, a satellite stops being a consumable and becomes a long-lived asset, and the industry's center of gravity moves from a single launch to many years of operation. The other is a question of norms. Whether Korea sits at the table where standards are made or accepts standards after they are made will make a large difference later. As Korea increases its satellite fleet, these two questions are not someone else's issue. They are items that will soon appear in our own budget documents.

I add one point. Most of what we have seen here is a schedule and a set of targets. The numbers, a 2026 launch and dozens of operations, are not achievements already completed. They are promises still to be kept. I will not present those promises as accomplishments. Whether a robotic arm actually revives an aging satellite in orbit can be independently confirmed only after launch. RSGS must therefore be read in two layers. If it succeeds, it will become a turning point that ends satellite design built around disposal like a consumable. If it is delayed or falls short, it will show how difficult space servicing still is. Either way, it will remain in the record as one of the first attempts to reach out to the once unreachable 36,000 kilometers.

RSGS Technology Transition and the 2026 Launch Plan

If the previous section asked why repair in orbit is needed, this section asks who takes on that need and how. Even if the technology sounds plausible, it ends as a press release unless the party that will build it through to the end and the party that will fund and operate it are defined. RSGS draws attention because of this division of roles.

Let us begin with the scene. About 36,000 kilometers above Earth, imagine a communications satellite in geostationary orbit that has almost exhausted its fuel. The body is intact. The camera works, and the antenna points in the right direction. But the moment the propellant needed to maintain attitude runs out, the satellite's mission ends. A machine launched at a cost of hundreds of millions of dollars becomes debris in orbit even though not one part has broken. For a long time, space satellites were treated this way. Once launched, they were one-time assets. Since they sat at an altitude beyond human reach, the answer to failure or fuel exhaustion was always the same. Abandon it and launch a new one.

DARPA's Robotic Servicing of Geosynchronous Satellites program targets this point. Its English name is Robotic Servicing of Geosynchronous Satellites, shortened to RSGS. The program's

question is brief. If a satellite is still sound but has run out of fuel, and if replacing one part could let it work for several more years, why should it be impossible to service it in orbit?

First we need to see why this program singled out geostationary orbit. In geostationary orbit, a satellite moves at the same speed as Earth, so from the ground it appears to stop at one place in the sky. Because the satellite is always in the same spot, a ground antenna can remain fixed in one direction. There is no need to rotate a dish to chase the satellite. That makes continuous communications and surveillance possible. According to DARPA's official 2026 announcement, hundreds of satellites responsible for national security and global communications still orbit there. They are expensive, and their roles are heavy. Military early warning, reconnaissance, and command communications depend on this altitude. If such an asset must be abandoned because of one tank of fuel, the loss cannot be counted only in money. During the years needed to build and launch a new satellite, the empty slot itself becomes a security gap.

The substance of RSGS is a mission robotic vehicle. In English, it is called a Mission Robotic Vehicle, or MRV. This spacecraft carries a set of robotic arms that move with the finesse of human hands. Its work falls into four broad categories. It services satellite performance in orbit, inspects spacecraft condition at close range, resolves unexpected failures, and, when needed, moves satellites to other positions. Moving a satellite may sound unfamiliar. The list includes returning a satellite that has drifted from its orbit to its proper position, or pushing a spent satellite to a location where it will not interfere with other satellites. The presence or absence of one robot determines whether order can be restored in orbit.

The words sound easy, but having a robot repair another satellite in the vacuum of geostationary orbit means crossing several barriers at once. First, approach and docking must work without collision. The two objects are moving in the same direction at thousands of kilometers per hour, so even if they appear to be stopped side by side, a small error can create a large impact. Then the robotic arm must grasp, turn, and insert parts with no meaningful deviation. It is difficult for a person on the ground to control every motion in real time. Because signals take time to travel, some portion must be designed so the spacecraft can judge for itself. One mistake can damage both the satellite being repaired and the satellite sent to repair it. There is no workshop margin here.

DARPA's chosen approach is the core of this section. The government is not taking the path of building and operating a space repair vehicle from start to finish by itself. DARPA, with the U.S. Naval Research Laboratory and NASA, develops the front end of the robot, the core hardware corresponding to the hand. That front end is the set of robotic arms that move like human hands. The task of carrying this robotic payload into space and integrating it into a final spacecraft for operation passes to the private sector. The private actor is SpaceLogistics, a subsidiary of Northrop Grumman. A commercial company mounts government-built robotic arms on its own spacecraft and runs them as its own business. In DARPA's own wording, this is technology transition. The government serves as the primer that breaks through the early hard problems no one else touches, then steps back at

the point where a market may form.

The importance of this division of roles becomes clear if we imagine the opposite failure. If the government owned and operated the entire repair vehicle, every satellite service would draw on the defense budget, and when the program ended, the technology might go into storage with it. If a private company owns the spacecraft instead, it can sell repair and refueling not only to the U.S. government but also to communications satellite operators around the world. Technology opened by defense needs flows into a commercial market, and if that market grows, the cost of the next satellite falls and defense costs fall with it. This is a dual-use design in which one technology serves defense and commerce at the same time.

James Shoemaker, the RSGS program manager, described the program as a government and private-sector partnership for next-generation satellite servicing. His stated goal is also clear. DARPA's robotics technology will be combined with a commercially proven spacecraft platform to show in practice that on-orbit servicing can work as a business in geostationary orbit. The word to watch is show. It does not mean that the business has already been established. It means the demonstration will test whether it can be established.

Then where does the 2026 launch fit? DARPA announced that it would launch this robotic servicing technology in the summer of 2026. Launch is not completion. On the ground, the robotic payload and spacecraft are integrated and tested, the commercial partner that will receive the technology is trained, and after launch the system is checked in orbit to confirm that the equipment works properly. Launch is one knot in a long preparation. It is not the finish line. The real test begins when the robotic arm reaches out to another satellite for the first time.

This must be stated honestly. RSGS is not a proven business. It is a demonstration. Shoemaker's statement is about proving commercial viability, not about viability already proved. Whether repair in orbit can actually make money, and whether the first demonstration runs as planned, are questions for launch and later operation to answer. Like other DARPA programs, this may go well or fall short of expectations. What we can confirm now is the plan and the announcements. I promised at the start to read plans as plans, and I will keep that line. I will not convert an impressive roadmap into an achievement.

From a policy perspective, the lesson of this section lies not in the list of technologies but in the casting sheet. How far does the state go, and where does it step away? RSGS draws that line at difficult hardware, the robotic arm. The government breaks through the place where no one moves first because the business case is uncertain, and leaves to the private sector the parts the market can already build, such as the spacecraft bus. This casting sheet is often missed when Korea talks about Korean-style research institutions. Unless the government first decides what it will carry and at what point it will transfer the work to the market, good technology ends as a report without transition.

Even so, the question this plan raises is not light. If a spacecraft with robotic arms can approach another satellite and remove and insert parts, the same hands can reach someone else's satellite.

An approach for repair and an approach for threat do not look very different from outside. The trajectory of a vehicle approaching to refuel an allied satellite and the trajectory of one closing in on an enemy satellite may look similar on radar. As the new market called on-orbit servicing opens, the question of who controls and protects those robots grows with it. If communications are cut or control signals are hijacked, a repair robot can become a dangerous object. The title of a DARPA document from April 2026 shows this tension directly. Victory in the new space race begins with cyber. Placing hardware in orbit and protecting that hardware are bound together. The next section continues that story.

The New Space Race and Cybersecurity

At 36,000 kilometers in geostationary orbit, one robotic arm approaches an aging satellite. This is RSGS, the robotic repair technology for geostationary satellites discussed in the previous two sections. Another machine approaches a satellite that has run out of fuel or whose parts no longer respond, grasps it, and services it. DARPA has said it will send this technology into orbit in 2026. It is a program led by the Tactical Technology Office. Up to this point, the story is about a repair worker in the sky.

Let us change only one question. What happens if someone intercepts the signal that commands that robotic arm? A machine that can repair another party's satellite can, in reverse, seize another party's satellite. The boundary where an orbital repair worker becomes an orbital threat rests on one line of software and one communications link. This is where the story of space robotic servicing becomes a story of cybersecurity. If the previous section discussed the robot's hand, this section discusses the nerves that send commands to that hand.

One short essay DARPA published in 2026 confronts this issue directly. Its title is the argument itself. Victory in the new space race begins with cyber. The author is a senior military officer assigned to DARPA, a so-called Service Chief Fellow. He looks back and says that the core of his mission was space cyber. In his view, it is an important but often neglected place on the space stage.

The essay begins with one calm fact. The cost of reaching space has fallen beyond comparison with the past. A threshold once open only to a tiny number of great powers with vast capital and technology has come down. Rocket launches are often no longer news, and satellites fill the sky by the hundreds. So the world has entered a new space race. Up to this point, anyone might say the same.

The essay moves one step further after that. In this competition, the country that gains lasting strategic advantage is not the one that launches the largest rockets. Advantage goes to the side that first figures out how to build and operate systems safely in this new environment. It is not launch vehicle thrust but the ability to protect assets once they are in orbit that decides the contest. The question is not how many satellites have been launched, but whether those satellites have been made unreachable by an adversary's hand. This shift of attention from the launchpad to orbit is the

backbone of the essay.

The problem is that this safety cannot simply be brought over from the ground. Most cybersecurity we know is a story of Earth. We install heavy antivirus tools, download patches when vulnerabilities appear, and let administrators log in to servers and fix them. This approach does not work well for satellites in orbit. The essay states clearly that the unique constraints of space systems must be understood. These are places human hands cannot reach, places with tight power and computing resources, and places where remote commands take time to travel. Security common sense built on the ground does not fit if it is placed on top unchanged.

So the essay presses one principle. Security for space assets must not be added later after the satellite has already been built. It must be planted inside from the first moment of design. It is the difference between bolting something on afterward and baking it in at the dough stage. The importance of this distinction becomes clear if we picture failure. Imagine each company launching satellites with its own hardware and software, then patching security on late and separately. Space assets built that way become a messy patchwork that is extremely hard to defend. No one can see at a glance where the weak point is or how far a breach in one place will spread.

To prevent this patchwork, the essay proposes a government reference framework. It is a common set of rules that can be applied across orbital spacecraft, ground control stations, and terminals. Think of it as a standard script. Rather than letting everyone build on their own and match systems later, it tries to make them build on the same baseline from the start. The officer who wrote the essay says he is turning experience gained from offensive cyber tools toward defense and applying it to make space infrastructure stronger. It also means that someone who has sharpened spears knows the weak joints of shields first.

We should pause and identify the nature of this material. This essay is not a completed performance report. It is not a record showing what a specific program has achieved. It is closer to an argument by an officer assigned to DARPA, persuading readers that this is the direction to take. The principle of not portraying DARPA only as a success story also holds in this section. The direction of space cyber is clear, but we do not yet have ample external material to verify from the outside how far that direction has taken root in real budgets and forces, or where money has been wasted. So I will introduce the direction and logic, but I will not speak as if they were confirmed victories. Claims and achievements must be weighed separately.

Even so, if we place the earlier sections of this chapter beside this one, one picture fits together. RSGS is a technology for physically touching satellites in orbit. If a robot can grab and repair another machine, orbit is no longer a sanctuary beyond reach. As physical access opens, protecting the communications and software that direct that access becomes heavier work. The repair robot's chain of command must stay clean for the repair worker to remain a repair worker. The moment that chain is contaminated, the same arm does the opposite work.

Let us unpack the reality of operation a little more. A satellite in geostationary orbit moves by receiving commands sent from the ground. Those commands travel through a communications link. For a servicing robot like RSGS, physical actions such as approach, stop, grasp, and manipulation ride on that link. If the link is breached, or if the ground system that creates commands is contaminated, the robot's fingertips point in the wrong direction. So the security planted from the beginning, as the essay describes it, is not an abstract slogan. It means enabling the satellite itself to distinguish which command truly came from the control station, and designing the system so that interrupted or unstable communications do not become dangerous malfunction. These matters must be decided at the first design drawing stage. After the satellite is built, replacing this skeleton is almost impossible.

From a policy perspective, the implication of this section reaches the order of budget allocation. The question of how many more new satellites to launch is easy to see. Launches leave photographs and numbers. By contrast, the question of whether a satellite's command chain was designed safely from the start is hard to see. Until an accident occurs, it looks as if nothing has happened. The essay's argument targets this asymmetry. It calls for money and people to be assigned first to work that gets pushed back because it is invisible. The reason it puts forward a standard script, that is, a government reference framework, is the same. If each project is left to do well on its own, each will leave different weaknesses on its own.

CONFERS from the previous section and the government reference framework in this section may look like separate stories, but they share the same grain. One is an attempt to set safety procedures for on-orbit servicing by agreement among countries. The other is an attempt to lay down common specifications for satellite cybersecurity inside one country. Both place rules ahead of technology. The ability to put robotic arms in orbit and the ability to write documents first so those arms do not operate in a lawless zone without rules. The United States is pushing both side by side. If technology runs ahead with no rules, accidents write the rules instead, and rules written that way are usually expensive.

For Korean readers, this scene will not sound like someone else's story. Korea is also increasing its satellites and relying more on orbit for communications and reconnaissance. The question raised by the DARPA essay becomes our question as it stands. More than how many additional satellites to launch, are we ready to protect those satellites' command chains from the first design drawing? If different agencies and companies are launching satellites under different specifications, the messy patchwork the United States worries about can appear in our sky as well. As the door to space widens, the side that first learns how to lock that door lasts longer. DARPA's short essay is closer to a warning than a boast, and that warning does not stop at national borders.

As this chapter closes, let us thread the three sections into one line. We began with why satellites in geostationary orbit had to be so heavy and expensive, moved to who builds and operates the robots that may lighten that burden, and arrived finally at how to protect the commands going to

those robots. The moment a hand reaches orbit, it opens more than the possibility of repair. It also opens the possibility that the same hand can become a threat. Space robotic servicing and cybersecurity are therefore the front and back of one story. In the summer of 2026, what the robotic arm actually does in orbit, and who protects the signals going to that arm, will write the next chapter of this story.

Chapter 12 Quantum: Experiments That Separate Hype from Reality

QBI and Quantum Computing Verification

In March 2026, a short but provocative headline appeared on DARPA's official news page. It said the Quantum Benchmarking Initiative was expanding its effort to separate hype from reality. Press releases from government research agencies usually lead with achievements. A new device was built, a record was set, a limit was passed. This headline began not with pride but with doubt. It was a declaration that DARPA would sort out how much of the talk around quantum computers was real. It was an attempt to cool down the heat around its own field, a rare stance for an agency that spends defense money.

Picture one scene. A company unveils a new chip with hundreds of qubits and announces that it will soon release a machine that will change the world. Investors applaud, and the stock price moves. Now imagine one government official sitting in that room. On what grounds can he decide whether the announcement is real or inflated? Quantum computing is now a field where money and expectations are rushing in at the same time. Announcements pour out about larger qubit counts and lower error rates, but whether the machine is actually worth its cost is a different question. There is no easy yardstick. This is where DARPA brought out a program called QBI.

QBI's question can be compressed into one sentence. Can a quantum computing approach that works at utility scale actually exist by 2033? The term utility scale may sound unfamiliar, but the meaning is not hard. It points to the moment when the value of the calculations a machine solves becomes greater than the cost of building and running it. It is the moment when the machine stops being a laboratory exhibit and becomes a tool worth paying for. QBI asks whether that moment can arrive within the next eight years. This program expanded an earlier effort called US2QC. Beneath it lies the aim of preventing a so-called strategic surprise, a situation in which an unexpected commercial breakthrough suddenly hits the United States from behind. The security calculation is that who builds it first matters, but so does when you realize it has been built.

There is one point that must not be missed when reading this program. QBI is far from a contest that chooses one winner. It is not a race in which a prize is offered and a trophy goes to the team that crosses the finish line first. DARPA intends to evaluate every workable approach side by side, as far as its budget allows. That is why verification is arranged in three phases. The demands grow heavier at each phase, so let us follow them one by one.

In Phase A, participating companies present the concept of a utility-scale quantum computer that could be realized in the near future, and they offer evidence that the concept is plausible. The demand to build the machine has not arrived yet. The demand is to draw a path that could be taken. In Phase B, the demand becomes specific. Companies must set out the research and development

plan for actually building the concept, the risks along the way, how they will reduce those risks, and what prototypes will be needed. DARPA looks less at how plausible the words sound and more at how tight the plan is. By Phase C, verification leaves the documents and moves to hardware. The team works directly with the U.S. government to show through demonstration whether the designed machine can be built and operate as intended. The weight shifts from concept to plan, and from plan to physical system. The sieve that filters out words becomes finer at each phase.

DARPA's self-assigned role is closer to a neutral third party than a referee. It stands one step away from the promotion of any company and the expectations of any camp, and serves as a channel that delivers verification results to other stakeholders inside the government. It is a device meant to keep national quantum investment from rushing in one direction on the strength of rumor and expectation alone. The picture of defense money measuring the hardware of commercial firms may feel unfamiliar. One may ask why money meant for weapons is being used to inspect corporate chips. But if the goal is to prevent resources from gathering in the wrong place, verification itself becomes defense work. Reducing wasted moves is also a form of strength.

The numbers show the scale of this verification. Since QBI began in mid-2024, it has evaluated about twenty commercial companies using different qubit architectures. Eleven organizations have passed Phase A and moved up to Phase B. That list shows the current map of the quantum computing market. IBM in superconducting systems, IonQ and Quantinuum in trapped ions, Atom Computing and QuEra in neutral atoms, Diraq and Quantum Motion in silicon spin, and photon-based Xanadu all appear side by side. It is a landscape in which several paths remain alive at once rather than converging on a single answer. Microsoft and PsiQuantum, which have been involved since the US2QC period, have already moved to Phase C, where system-level operation is checked, and Google has newly joined Phase A. Teams that are ahead and teams that have just entered are placed on the same board.

The person leading the program has changed as well. The manager who first set up QBI was Joe Altepeter, and Micah Stotimore, formerly the deputy manager, has taken over as the new managing director. Stotimore said the chance that someone will build a utility-scale quantum computer by 2033 has grown compared with before, but he added that it remains unclear which team will cross that finish line first. That is why the Phase A call for new companies was reopened through the second half of 2026. Abstracts are due on July 31, and full proposals are due on September 30. The very act of keeping the door open is an honest admission that the answer is not known. Even after selecting more than eleven promising candidates, the program does not close the door, and that tells us what kind of program this is. Organizations that believe they know the answer close the door. Only organizations that admit they do not know keep it open.

Why does DARPA look at several methods together instead of pushing everything toward one? A clue lies in the view that different qubits may be better together. Superconducting systems, trapped ions, neutral atoms, photons, and silicon spin each have different strengths and weaknesses. Some

approaches are fast, some last longer, and some are easier to connect. At this point, declaring one winner is close to gambling. So DARPA does not bind itself to one method. It asks only whether the method can reach the goal called utility scale. The next section will unfold this view again. For now, it is enough to carry forward one fact: QBI does not take the side of a specific technology.

There is material that lets us compare whether this stance belongs only to the United States. A report by the Electronics and Telecommunications Research Institute, or ETRI, on competition among quantum hardware platforms points out that the United States and Europe are taking different paths. The European Union has set a clear target of securing 100 error-correcting qubits in one system by 2030, and it is closer to a policy approach that selects a small number of platforms to support through regular evaluation. The government names likely winners from above. U.S. QBI stands on the other side. Rather than decide in advance which method is superior, it places weight on measuring each one against the absolute standard of whether it can reach utility scale. Whether the method is superconducting, silicon spin, or photonic, the purpose is to put a machine that truly works on American soil first, not to raise the hand of a particular technology. One path puts a target number at the front. The other measures the possibility of reaching the destination itself. The two paths aim at the same destination, but they differ in how far the government's hand reaches.

What Korean readers should take from this section is not a glittering success story. QBI has not proved anything yet. DARPA itself says it does not know whether a utility-scale quantum computer will stand in 2033, or who will be there if it does. That is why this program should be read as a verification device rather than a record of achievement. The headline about separating hype from reality shows plainly that what DARPA holds in its hand is not certainty but a question mark. Measuring technology without inflating it, and admitting that the answer is unknown while building the yardstick first. When Korea talks about quantum technology, the point worth envying may be less the machine's performance than this posture of verification. Performance may be caught up with over time, but the eye that separates inflated words from the real thing must be trained through institutions. Why that eye ends up looking at many qubits at once is what we will follow in the next section, holding each qubit in hand one by one.

Why Different Qubits Must Be Seen Together

One day in April 2026, an article appeared on DARPA's website. The title had a slightly crooked angle. It said that in quantum computing, different qubits are better together. Writing about technology competition usually asks who is ahead. This article moves in the opposite direction. It says not to choose one and push it, but to watch several methods side by side. Why would a defense agency funded by taxpayers say this? In the previous section, I said QBI does not take the side of a specific technology. The roots of that stance are in this article. To find the answer, we first need to hold the unfamiliar object called a qubit in our hands.

Ordinary computers calculate only with 0s and 1s. If a switch is off, it is 0. If it is on, it is 1. The switch in a quantum computer is strange. It can be spread between on and off. The longer this

spread-out state can be kept, and the more cleanly many such states can be linked, the greater the computational power becomes. The problem is that the state collapses easily. A slight rise in temperature or a tiny outside vibration can disturb it. That is why what you make a qubit from decides the fate of the machine. The material sets the performance, and the performance sets the use.

A 2026 report from the Electronics and Telecommunications Research Institute divides the quantum hardware now under development into six broad branches. It also points out that all six carry serious weaknesses. Let us follow each branch briefly.

IBM and Google stand in the superconducting branch. Because this path can bring in circuit technology already used in semiconductor factories, it is favorable for scaling up, and its operation speed is high. But the spread-out state does not last long. The system barely runs only when cooled close to absolute zero with a huge refrigeration unit. Imagine a computer that works only inside a refrigerator filling an entire room.

The trapped-ion method has a different texture. IonQ and Quantinuum are taking this path. Each atom is held by an electric field and used as a qubit, so the individual units are nearly identical, and the spread-out state lasts long. But the operations are slow. The number of atoms that can be held in one line is limited, so scaling up runs into a wall.

The neutral-atom method uses laser-made tweezers to hold atoms and freely rearrange their arrays. It is good for scaling up and the state lasts long, but the laser equipment is demanding, and operations are slow. The photonic method carries information on particles of light. It can run at room temperature without cooling and suits long-distance transmission. But operations that link two qubits are hard, and light leaks away little by little as it travels.

The two remaining branches are silicon spin and diamond defects. The first fits well with semiconductor processes and saves space, but wiring becomes the bottleneck. The second runs at room temperature, but it is hard to scale up and connect units to each other. Place all six in one room, and the picture becomes clear. No universal qubit yet satisfies speed, stability, scalability, and temperature at once. If one side is fast, another is weak. If one side lasts long, another is sluggish. If you hold on to only one, that method's weakness becomes the ceiling. The comparison table explains in the language of physics why DARPA did not pour everything into one path in the previous section.

If you look only at the market's surface, the match can seem settled. According to the June 2025 MIT Quantum Index report cited by the ETRI report, half of the 40 commercially released quantum processing units up to that point used the superconducting method. By numbers alone, superconducting systems look as if they have taken the lead. But when qubit count, error rate, and operation speed are assessed together, trapped ions and neutral atoms are close behind, and the diagnosis that it is hard to name a clear winner appears alongside the figures. A larger installed count and better performance are separate matters.

Experts who watch the market point in a similar direction. In a 2025 survey cited by the ETRI report, experts interviewed across different quantum platforms thought the methods were likely to fill one another's gaps. They also expected the field to specialize by use case rather than one method taking the whole board. The needs of places that urgently require fast operation are different from the needs of places that must carry a small device without a refrigerator. The diagnosis is that the very idea of filling every role with one machine is strained.

A more practical reason follows for why different qubits must be seen together. Companies have begun mixing completely different methods inside one machine. The key to scaling a quantum computer to a useful size lies in connecting many modules well, and there is analysis that light is advantageous for connecting one module to another. IonQ, which uses trapped ions as its main path, announced a plan to scale up to 2 million physical qubits by 2030 and said it would also use photonic interconnect technology it acquired in 2025. Quantum Brilliance, which backs the diamond approach, built a hybrid system that combines its room-temperature processor with ordinary computer central processing units and graphics processing units, and installed it at Oak Ridge National Laboratory in the United States. Internal operations may use one method, while information between modules moves by another. At this point, the statement that clinging to only one method makes it hard to reach the finish line begins to feel real. The contest may be decided not by one pure method but by a well-mixed combination.

Over this current, DARPA appears again. In March 2026, DARPA announced that it was expanding the Quantum Benchmarking Initiative. Its abbreviation is QBI. The article title shows the program's character plainly. It says the effort to separate hype from reality is being expanded. QBI aims at a sharp question. Is there truly a path to building a quantum computer at a useful scale by 2033? What will separate promotion from reality? The three-phase verification described in the previous section was the device built to answer that question.

The words of Micah Stotimore, who leads QBI, point to the core of the question. He says it is likely that someone will produce a quantum computer at a useful scale within 2033, but that it is not yet known which team will cross the finish line. That is why QBI draws a line and says it is not a contest to pick one winner. As long as the budget reaches, it will measure every plausible route one by one. The list of eleven organizations from the previous section was the result of that measurement.

It did not stop at words. The list of companies that moved to the next QBI phase in 2025 is evenly mixed across methods. Neutral atoms, silicon spin, superconducting systems, trapped ions, and photons, the branches discussed above all appear side by side. Instead of handing the prize to one method, DARPA laid out a board where different paths are measured by the same yardstick. Measuring teams with different methods by the same standard is not as easy as it sounds. Building that difficult yardstick is the substance of this program.

This stance becomes clearer when placed beside Europe's choice. As noted in the previous section, the ETRI report describes the European Union as setting a numerical target of 100

error-correcting qubits per system by 2030 and trying to choose from above which directions to support. The United States stands on the opposite side. Rather than rank specific methods, it keeps technical neutrality and places several branches on the scale together. It is too early to declare which side is right. But in a field where the winner is not yet known, the United States is answering through organizational design how risky it is to name a winner too early.

Why keep so many methods alive together? From the arithmetic of a defense agency, it makes sense. If a country pours its money into only one method and an unexpected breakthrough comes from a completely different one, that judgment becomes a hole in security. Seeing several qubits together is like taking out insurance by keeping a hand on several doors in advance, so that the United States is present wherever the wall is breached. The prevention of strategic surprise described in the previous section appears here as the act of placing hands evenly on six branches of qubits.

This point must be read honestly. QBI is not yet an achievement. It is an experiment. The year 2033 and the team that will cross the finish line are not fixed today. What DARPA has offered is not a declaration that we have won, but a method and a yardstick for sorting hype from reality. The sentence that different qubits are better together is not a record of victory. It is closer to a confession that the contest is still unknown. The value of this program lies in the place where that confession is turned into a budget and an evaluation system. And this experimental posture does not remain only with calculators. In the next section, the same question moves from computing to time and communication.

Operations Without GPS and the Debate Over Quantum Networks

An electronic warfare unit fires a strong radio signal into the sky. Its target is not the incoming missile itself, but the signal on which the missile depends. If the weak radio waves from satellites are covered by noise and disappear, an expensive guided weapon loses both where it is and what time it is. A weapon that has lost direction becomes close to scrap metal at that instant. A precision weapon that costs hundreds of millions of won per shot loses its use when a fingernail-sized signal is cut off. The previous sections dealt with the accuracy of computation. This section aims at the accuracy of time.

Modern militaries bundle position, navigation, and timing under the name PNT. DARPA's introduction to the ROCKN program, released on March 2, 2026, shows how tightly all three are bound to time. Advanced missiles, sensors, aircraft, ships, and artillery systems all take nanosecond-level timing from atomic clocks on GPS satellites. A nanosecond is one billionth of a second. To human sensation, it is almost no time at all. On the battlefield, this time that seems not to exist decides whether a target lives or dies.

The reason becomes clear if we calculate by the speed of light. Light travels about 30 centimeters in 1 nanosecond. If timing is off by only a few nanoseconds, the position calculation shifts by the

distance the signal traveled during that time. An error of a few nanoseconds can spread into a position error of more than 1 meter. A weapon that misses its target by 1 meter has already missed. The problem grows when several units aim at the same target at the same time. If each clock differs slightly, the order and spacing of a joint strike break down, and a well-planned operation collapses in the wrong rhythm. It is not one clock being wrong. It is the breathing of an entire unit falling out of step.

The weakness of GPS signals has long been known. Because the signal is so weak, it is easy to deceive. An adversary can jam it by covering it with strong noise, or lure a weapon toward false coordinates through spoofing that sends out a fake signal looking like the real one. DARPA calls this vulnerability a well-known weakness. It is not a newly discovered risk, but a hole that has been pointed out again and again. As weapons become more precise, the cost of this hole grows. Precision becomes dependence, and dependence becomes weakness. It is like a sharp knife becoming more dangerous when the handle breaks.

The idea behind ROCKN is to detach time from satellites and move it beside weapons. The name is short for Resilient Optical Clock Network. The plan is to build clocks that can hold accurate time on their own for a long period even when satellite signals are cut off, and to make those clocks small enough to carry on the battlefield. They must not be precision clocks sitting in the middle of a laboratory. They must be clocks a soldier can carry or a vehicle can load. That is why size and power become conditions as important as performance. No matter how accurate a clock is, if it fills a room and eats the output of a generator, it cannot be used on the battlefield.

DARPA's introduction speaks of two kinds of clocks. One is about the size of a shoebox. It runs on the small amount of power needed to light one bulb, and even without receiving any outside signal, it keeps sub-nanosecond precision on its own for two weeks. The other is a regional reference clock about the size of a washing machine. Without GPS synchronization, it shares the same time with multiple units spread across a wide area for more than six months. If the first clock is the wrist of an individual soldier, the second is the village clock tower that those wrists can check against. The wristwatch moves with the person, and the clock tower stays in one place to set the standard for many people. The two clocks have different roles, so their targets for size and duration differ as well.

This part must be read carefully. The figures above are not the performance table of weapons already completed and deployed. They are closer to program targets. In the official materials I secured for this chapter, I did not find a record showing that a shoebox-sized clock has operated in actual units for years. DARPA's article is more accurately read as what it is trying to build than as what it has already achieved. This is exactly why I promised earlier not to collect success stories alone. The moment target numbers and verified achievements are copied down with the same weight, a plan turns into a battle report. People who deal with policy must be sensitive to this disguise. If those who set budgets and decide adoption read targets as achievements, they count a capability that has not yet arrived as if they already possess it. Just as QBI tried to measure

corporate announcements, we must measure DARPA's own announcements with the same eye.

The story does not end just because time has been detached from satellites. Even if weapons and sensors scattered across the battlefield each carry accurate clocks, those clocks are useless if the communications network connecting them is pierced. If an adversary quietly inserts a fake node, secretly listens to timing signals in transit, or pushes false information into the communications path, the carefully aligned clocks shake again. Protecting clocks and protecting the network that shares time are separate assignments. Making accurate time and sharing that time safely are different problems. This is where the debate over quantum networks raises its head.

DARPA has a program called QuANET, or Quantum-Augmented Network. It belongs to the Information Processing Techniques Office, and its program introduction presents quantum, communications, security, and information as its themes. The root of the idea is one property of quantum mechanics. If someone secretly measures a quantum state, the act of measuring changes the state. Eavesdropping leaves a trace. If this property is placed onto a communications network, one can hope to build a line that detects listening attempts right away. The vision is a network in which, when units share well-aligned time and an enemy tries to intercept or manipulate that signal, the attempt itself is exposed. It is less like protecting a key and more like setting off an alarm the moment a door is opened.

But the distance between this hope and reality must be stated honestly. This distance is why the chapter is titled an experiment that separates hype from reality. The sentence that quantum communication is theoretically impossible to eavesdrop on is often quoted. But it assumes ideal conditions. Real equipment has gaps, and battlefield-scale quantum networks now sit at the stage of programs and debate. They are not a completed shield, but a question mark placed on a test stand to see whether they can become one. When readers encounter writing that erases this difference, they should pause. Between what theory promises and what parts can do, there usually lie many years of failure.

It is worth noting that DARPA itself acknowledges this temperature gap. On the quantum computing side, DARPA runs QBI, the Quantum Benchmarking Initiative. It belongs to the Microsystems Technology Office, and the title of the March 10, 2026 introduction states its purpose plainly. It says DARPA is expanding the journey to separate hype from reality. The point is to measure coolly whether quantum computers presented by various companies can really perform useful calculations. It treats the bubble in its own field as an object of verification, and this attitude becomes the scale for reading ROCKN and QuANET. The same agency is making declarations about what it will build and measuring whether what has been built really works. The declaring hand and the verifying hand are attached to the same body, and that is the character of this organization.

DARPA's view of qubits does not gather on one side either. In an April 14, 2026 article, DARPA says different qubits are better together. There is not only one way to build a quantum computer. Superconducting circuits, trapped ions, and several other paths compete, each carrying its own

strengths and weaknesses. Rather than place all the stakes on one in advance, DARPA judges that it is better to place multiple methods side by side and let them fill one another's empty spaces. The reason we took apart the six branches of qubits one by one in the previous section connects back to this sentence. In a field where the winner has not yet been decided, it is a posture of refusing to close judgment too early.

Korea is also watching this current. In April 2026, the Electronics and Telecommunications Research Institute, ETRI, released a report comparing quantum computing hardware platforms by type. The report treats DARPA's move to evaluate competition in quantum hardware technology as a point of reference. It notes that which qubit method will take the lead has not yet been decided, and that several methods therefore need to be placed on the scale together. The place where the United States is trying to clear away bubbles is being recorded by a Korean research institute with the same level gaze. This is not a place for envy and copying, but a place for looking at the same problem with the same caution. This point is also a clue for the Korean-style discussion in the final chapter.

The lesson this section leaves for policymakers is clear. When reading a technology description, they must separate the tense and character of each sentence. Targets like ROCKN's numbers for component size, power, and holdover time are verifiable promises. A few years later, one can ask whether those numbers were reached. By contrast, the phrase that a quantum network cannot be eavesdropped on swells in meaning if the conditions are not read alongside it. What Korea should learn when designing a defense research agency is not a flashy list of programs. It is restraint that places exaggeration in one's own field before a verification process, and a way to plant that restraint inside budgets and evaluation systems. Copying a few programs only requires money. Building an institutional eye for verification cannot be done with money alone.

The two axes covered in this section have different characters. ROCKN, which seeks to make time without GPS, has set numerical targets for component size, power, and holdover time as it moves toward a tangible goal. Quantum networks are still written in the language of debate and verification. When you encounter materials that bundle the two together as achievements of the same maturity, take one step back and ask: Is this deployed equipment, or a promise on the test stand? That one question is the first cut that separates hype from reality. And this cut is not a tool used only for quantum. Whatever DARPA program we read, the eye that separates a plan from a battle report runs through the chapters that remain after Chapter 12.

Chapter 13 Biological Computers and Optogenetics

Reading O-Circuit as a Plan, Not an Achievement

In the spring of 2026, a researcher opens SAM.gov, the official contracting site of the U.S. government. Into the search box goes a single program number: DARPA-PS-26-28. What appears on the screen is not a photograph of a finished device, nor is there a video of an experiment. It is a question-and-answer document packed with dozens of questions and replies. Which cells may be used, how much the budget is, what must not be built. This document is the current substance of O-Circuit. Not a photograph, but questions and answers. Not a device, but rules. Nothing shows the program's present state more accurately.

Stories are circulating. DARPA is said to be building computers out of living brains. Brains in glass jars are said to think on their own. The picture is provocative. Yet when the official materials are opened, the scene quiets down. O-Circuit is not a biological computer sitting somewhere. It is a basic research blueprint that is still selecting the teams that will study it. It is a plan that has only just raised its anchor. We should read this blueprint as a blueprint. Read it as an achievement and it becomes exaggeration. Read it as a plan and it finally becomes accurate.

We start with why such a program appeared. Artificial intelligence today consumes enormous amounts of electricity. A single data center can draw as much power as a city. The battlefield is different. There are no outlets, and generator fuel runs out when supply lines break. DARPA identified this point as the problem. Can a device learn and compute on its own with only a tiny amount of energy in places where electricity is scarce and resupply is difficult? O-Circuit begins with the idea of looking for that answer not in semiconductors, but in living cells. It reverses the direction of the question. It did not ask how to make a stronger chip. It asked how to compute without a chip.

Start with the program's official name. O-Circuit is an abbreviation. Written out, it means "Organoid Cellular Intelligence from Convergent Understanding and Information Transfer." According to DARPA's program description, the program belongs to the Biological Technologies Office. Its goal gathers into one point. It seeks to build a biological processing unit, or BPU, that can learn and compute on its own with only a tiny amount of energy even at the "tactical edge," where electricity is scarce and resupply can easily fail. This is a new device in which a mass of cells handles computation instead of a semiconductor chip. It is not yet a thing that has been made. It is a question about whether it can be made.

The fact that this program is in an early stage becomes clear from a few dates. According to the program description, O-Circuit began on March 10, 2026. On May 1, the solicitation was posted on SAM.gov. After questions from proposers poured in, a second amended version of the Q&A was issued on May 12. A program launched in March and then had its detailed Q&A revised twice in May. That means the program is now in the middle of recruitment and planning, gathering participating teams and aligning rules. This is not a venue for boasting about research results. It is closer to an

announcement seeking researchers. The trace of two Q&A revisions within two months also signals that the rules have not yet hardened and are still being shaped by the questions.

The scale of the money also tells us that this is a plan. In the Q&A document, DARPA said the total estimated budget set aside for Phases 1 and 2 is \$36 million. It did not place a hard ceiling that proposals cannot exceed, and it announced that multiple teams would be selected depending on the quality of the submitted proposals. University participation was officially allowed. Put these three points together and the picture becomes clear. This is not a project handed to one specific defense contractor. It is an early round that seeks to draw broadly on the basic research capacity of universities and industry. DARPA's way of working is visible as it is. It does not place an order after deciding the answer in advance. It puts up one clear goal and makes smart people in the world run toward it by different routes.

When reading a plan, the important point is not only what it will build, but what it will not build. In the Q&A, DARPA drew firm boundary lines so researchers would not spend money in the wrong places. One proposer asked whether the program was moving toward building an isolated brain inside a bioreactor. DARPA said no, plainly. O-Circuit is not a program to build a brain. The rumor of a brain in a glass jar is directly denied. Ideas for reducing energy with leading-edge semiconductors alone, so-called "pure silicon-based" solutions, were also pushed outside the scope. Existing silicon devices such as electrodes may be used to help put neural signals in and take them out, but proposals must foreground biological processing and reduce the share of silicon computation if they are to be selected. Drawing the boundary in advance reduces money leaking into futile directions. It is one face of DARPA's style of problem-setting, which starts by making the target clear.

The target metrics show how far this program is aiming. The energy target for a BPU is 10 milliwatt-hours per day. When proposers asked whether this figure would be measured only during computation, DARPA stated that it would be measured over 24 hours, including both standby state and active computation. As for where this energy value came from, the explanation was that it was derived from the energy use of the human brain scaled to a culture of about two million neurons, as discussed at the proposer workshop. DARPA also left room by saying that it had not imposed a strict ceiling on the number of neurons and that what matters is whether the whole system reaches the target. One more sense is added. DARPA recognized that research to identify receptors for detecting specific operationally relevant odorants can fall within scope if it is directly connected to making a BPU capable of olfactory sensing. Millions of cells calculating day and night on a small amount of nutrients, while smelling like an insect's antenna. That is the outline of the final design drawn by the Q&A document.

The honest evidence that this program is not yet at the achievement stage appears in DARPA's stance toward transition into real operations. One proposer asked whether there was a transition path to hand this technology to warfighters. DARPA's answer was calm. It said it is always interested in transition, but that there is no specific transition plan at this time, and that how the program unfolds

will determine the path. The same applied to mass production and scalability. Asked whether Phase 1 must examine the scalability or manufacturability of BPUs, DARPA said Phases 1 and 2 are research and development phases, while transition and scaling are deferred to Phase 3. The program is now unclassified and has no plan to assign another classification. A project building a weapon to plug directly into the battlefield would not leave things this open. Instead of rushing achievement, this is a document that writes down the parts that still have no answer as having no answer.

Add the fact that regulation is being handled in advance, and the nature of the program becomes clearer. Because animal cells or tissues may be used, ethical procedures are coordinated ahead of time. Asked whether approval from an animal care and use committee is needed when using neurons derived from commercially supplied induced pluripotent stem cells, DARPA explained that approval is needed only when the work qualifies as animal research. Asked whether government-run testing and evaluation would take place indoors or outdoors, DARPA answered that details would be released later as part of the test and evaluation plan. Even the testing environment has not yet been fixed.

Readers who deal with policy should keep two things in mind. First, develop the habit of weighing the force of official documents separately as plans and as achievements. Even within the same program, the sentence “\$36 million has been put on the table” and the sentence “computation with cells has succeeded” carry completely different weight. The first is a confirmed plan. The second appears nowhere in this document. Second, setting numerical targets and drawing prohibition lines first is itself an institution. Ten milliwatt-hours per day, measured over 24 hours. Do not build a brain. Reduce silicon computation. These few lines align dozens of research teams in one direction for the next several years. When Korea discusses a similar organization, what it should envy is not a dazzling product, but this skill of writing documents that attach budgets, rules, and deadlines to questions that have no answer and make people run.

That is how O-Circuit must be read to be accurate. It is not a completed biological computer. It is a stage for a challenge in which DARPA has laid down budgets, rules, and target numbers so that U.S. universities and research institutions can try to make computing devices out of masses of cells. One document draws a line saying it will not build a brain, pushes real-world transition to later, and sets the harsh threshold of 10 milliwatt-hours per day. The promise to read a planning document as a planning document is tested in passages like this. A device in which cells compute sounds impressive. For now, the substance of that impressiveness is a \$36 million invitation to experiments that have not yet begun.

The GO Program and Optogenetics

Light of a specific wavelength falls onto a small culture plate containing cells. There is no needle, and no pipette dropping reagents is visible. What changes is only the intensity of the light and the rhythm of its blinking. Taking that light as a signal, a new genetic sequence begins to be made one

letter at a time inside living cells. No tangible material has moved in or out, yet the cells now carry information that did not exist moments before.

This scene is not a laboratory success story. It is the target picture drawn by the Q&A document for the GO program, the program FAQ that DARPA released in 2026. GO is short for “Generative Optogenetics.” What we hold in our hands is not a result, but a plan. The FAQ is a document in which DARPA answers questions researchers might have before submitting proposals. It contains not what has been accomplished, but what DARPA intends to ask for. This distinction must be held from the start. Just as O-Circuit in the previous section is a program pushed by the Biological Technologies Office, GO also stands on the same current of trying to implant computation in living systems. Both are experiments with stakes placed on them, not achievements to brag about.

The term optogenetics may sound unfamiliar, but its roots are old. It is a technology for turning cellular switches on and off with light. Neuroscientists have used this method to shine light on specific neurons and control their activity. They implant light-gated channels in cell membranes, and when illumination is turned on, the cells respond. GO aims one step deeper. It does not stop at turning something on and off with light. It seeks to produce desired nucleic acid sequences directly inside cells using only light. It is a move from a switch that turns on and off to a pen that takes dictation letter by letter.

At the center of this goal is a device called a “nucleic acid compiler.” Its English abbreviation is NAC. According to the FAQ, NAC is a protein complex designed to be expressed inside living cells. Existing biotechnology has usually inserted genes made outside the cell into the cell, or copied DNA from templates already inside the cell. A template is the original strand used as the model when copying a sequence. NAC reverses that common sense. It operates in a template-free manner and combines optogenetic control responsive to light, the ability to bind substrates that serve as materials, and enzyme activity that actually links sequences. In response to light signals supplied from outside, it synthesizes DNA or RNA sequences on its own inside the cell. DARPA wrote that it hopes this direction will become the basis for “mass-free genetic programming.” That means transmitting genetic information with only the signal of light, without transporting material.

One strict rule is attached. The information encoding the sequence must be only light. A researcher asked whether nonchemical signals such as electricity, sound, or magnetic fields could be used as inputs. DARPA’s answer was limited. Such signals may be used only for auxiliary control, such as limiting the number of bases added at one time. The information that determines which sequence will be made must be carried by light. Other signals cannot take that place. One rule makes the program’s character clear. It means light must become more than illumination. It must become language. The task is to load letters onto intensity, rhythm, and wavelength, and make cells read those letters and transcribe them into sequences.

DARPA also drew the experimental boundaries clearly. In the GO document, “in vivo” refers to the inside of living cells. Whether floating in culture medium or attached to a culture plate, a living cell

culture counts as in vivo. Systems without cells are classified as in vitro. The final gate of Phase 2 must be passed by showing light-driven programming and sequence synthesis inside living cells. Destroying cells to read the sequence that has been made is allowed. What works well in a test tube and what works inside living cells are difficulties of entirely different order. That is why DARPA fixed the final gate inside the cell.

The exclusions are also clear. Human-subject research and animal-subject research are outside the program's scope. The use of embryonic stem cells as experimental material was also barred. Proposals may seek to make the system operate only in a specific organelle inside the cell rather than throughout the whole cell, but DARPA views that as carrying higher technical risk and requires a justification for why such precision is needed. It also requires that the demonstration be possible with commercially available optical equipment, not expensive specialized equipment. These conditions are not meant to lower the ceiling of the technology. They are devices to make results verifiable by others.

The document also states what will count as success. One ability DARPA watches closely is the ability to create sequences that are difficult to make with existing methods. These are complex sequences with high GC content. In DNA's four letters, regions with many Gs and Cs bind tightly to each other and often stall during synthesis. In Phase 2, DARPA may intentionally assign difficult sequences as tasks to test the proposing teams' limits. The nucleic acids produced must do more than exist. They must visibly change cellular function. Showing activation of a reporter such as green fluorescent protein or a luciferase enzyme is preferred. A reporter is a marker protein that signals success with light or color. Larger plans that address cellular metabolism, movement, or differentiation are also welcomed. The program also leaves open the use of generative AI in designs such as genetic circuits, or AI support for sequence design, as means to accelerate the goal.

The way money is handled shows the nature of this program well. In Phase 1, a proposer may request up to \$1.7 million if participating in one research area, and up to \$1.99 million if participating in both areas. The initial proposal states only the cost and work for Phase 1. The scope and budget of Phase 2 are not fixed from the beginning. They are negotiated separately with DARPA around the conceptual design review in month nine of Phase 1. It is a structure that opens the next door and budget only to teams that have shown performance. This is the opposite of handing over all research funds in advance and waiting for results. Several doors are set up, and only teams that pass the earlier door are placed in front of the next one.

The idea of building genes inside cells with light alone could become a blessing or a threat. DARPA pointed this out within the document itself. To address the question of how new generative optogenetics changes biosecurity and cybersecurity risks, DARPA issued a separate announcement that includes three dedicated groups, including a biosecurity working group, apart from the technical development. It also included a device for fairness. As a rule, DARPA is wary of allowing the same institution to serve both as a technology developer and a working group operator. If an institution

insists on doing both, it must establish a complete firewall between the two teams. The same person is not allowed to appear on both sides. An exception allows a faculty member from the same institution to join a working group if that person has no involvement in the technical task. From the start, DARPA intends to separate the hand that builds the technology from the hand that watches its risks.

At a proposer workshop held in Washington on January 7, 2026, researchers from academia and industry introduced themselves and formed teams through posters and brief ten-minute presentations. Foreign nationals and permanent residents may participate, but they must go through a separate security process to confirm whether they can handle controlled unclassified information, and if needed, they may face access restrictions. The door is open to international researchers, but it narrows again in front of sensitive information.

Seen through a policy lens, GO is both a technology plan and a control plan. What Korean readers should take from this document is not a feeling that success is near. No milestone has yet been passed. It would be wrong to read GO as an achievement of U.S. biotechnology. The passage to watch lies elsewhere. The more dangerous the technology, the more the organization that monitors the risk is established on the same day as the development announcement. The budget is broken into stages and tied to performance. Success criteria are fixed in forms that others can verify again. There is the ambition to make cells write genes with light, and the caution that writes down in advance the path by which that ambition could turn into a weapon. GO is an experiment not yet resolved, with those two things placed side by side in one document.

Where Biological Computation Meets the Power Problem

Imagine a single reconnaissance drone. Its propellers spin, its camera scans the ground, and an artificial intelligence inside that video keeps calculating without rest to distinguish people from vehicles. Algorithm performance improves every year. The reason that drone in the sky usually cannot endure for long is not the algorithm, but the battery. High-performance computing devices consume a lot of electricity. In the rear, where power plants run and data centers stand in rows, that thirst can be satisfied. At the front line, where fuel and batteries are hard to transport, the story changes completely. No matter how smart a machine's judgments are, the moment its power goes out it becomes a heavy lump of scrap metal.

There is a document that begins from this predicament. It is the introductory material for DARPA's O-Circuit program, released in 2026. The program's official page opens by saying that modern artificial intelligence is powerful but requires significant electrical power to operate. It then states that at the tactical edge, where power supply is tight and resupply is difficult, that energy demand itself becomes a wall blocking operations. This sentence shifts the center of gravity of the problem. Making smarter AI and providing the electricity to keep that AI on are different problems.

The approach DARPA brought out in front of this wall is somewhat unfamiliar. It is not to make better semiconductors, but to use semiconductors less. The name O-Circuit is an acronym drawn from a long official title. In Korean, it is close to “organoid cellular intelligence from convergent understanding and information transfer.” The words are difficult, but the structure is clear. This program belongs to DARPA’s Biological Technologies Office, or BTO, and brings biology, high-performance computing, and neuroscience into one place. The goal is to make a computing device from living cell tissue, a biological processing unit, or BPU. It aims for a device that uses minimal energy while learning and computing on its own. The idea is to entrust computation to cell cultures instead of silicon chips.

Why cells? This is the first point where biology and the power problem meet. In O-Circuit’s Q&A document, one proposer asks, “What about ways to reduce energy on a silicon basis?” DARPA’s answer is firm. A solution handled purely inside silicon is outside the scope of this program. The use of silicon hardware such as microelectrode arrays as auxiliary devices to read signals from cells and send signals into cells is allowed. Yet the actual computation should be pushed as much as possible toward biological processing, and the share of silicon computation should be reduced. An old observation lies underneath. In today’s computers, the place where data is stored and the place where calculation occurs are separated, so information is endlessly transported between them. A large amount of electricity leaks as heat during that round trip. In neurons, by contrast, computation and memory occur together within the network itself. The distance traveled is shorter, and the expectation is that less energy will leak away.

The size of that expectation becomes clear in a single number. O-Circuit’s energy target is 10 milliwatt-hours per day. It is not an intuitive unit. If we recall that fully charging one smartphone takes roughly around 10 watt-hours of electricity, this is about one-thousandth of that. A device that computes all day must run on one-thousandth of a single smartphone charge. The document does not allow this number to be used loosely. When another proposer asked whether periods of rest taken to save electricity could be excluded from measurement, DARPA answered no. The full 24 hours, including both rest time and actual computation time, must be the basis. This power is not measured for the whole device, but as the direct metabolic consumption of cells burning nutrients. The electricity used by an incubator, gas exchanger, or electrode array that keeps the cells alive must be reported, but no separate target is placed on it. The evaluative eye is trained on the consumption of the cells that perform the computation itself.

The document also records where this destination of 10 milliwatt-hours came from. According to an explanation from the proposer workshop, it is an estimate obtained by scaling down the energy used by the human brain to the size of a culture of about two million neurons. The next point is interesting. DARPA adds that researchers need not be bound to this figure of two million. Even if they use more neurons, they can do so as long as they meet the energy standard of 10 milliwatt-hours per day. This means the target is not the number of cells, but energy efficiency itself. The message to researchers is not “use this many,” but “cross this threshold.” It is a typical DARPA

gesture in posing a problem. It does not specify the method. It fixes the destination.

The planning document also states where this biological computer might actually be carried. In later stages of the program, it imagines a completed BPU loaded onto a drone. The document estimates the drone payload limit at 10 pounds, about 4.5 kilograms, while explaining that this value is only an estimate based on current drone performance and may vary depending on whether the platform is aerial or ground-based. Within that limited weight, the picture is of a cell computer that senses the smell of chemical substances on its own and produces control signals to move through terrain toward a target without help from a central server. It imagines a small machine that judges alone, without a battery charging station or a communications relay.

The lines must be kept clear. The 10 milliwatt-hours, two million neurons, and 4.5-kilogram drone described so far are not achievements already accomplished. They are all targets and conditions DARPA set out to gather proposers, numbers contained in the solicitation and its Q&A. There is not yet any verified result showing that a mass of cells has met these criteria and actually found its way. That is why this section should not be read as a success story. O-Circuit is now a program in the planning stage, and what we hold in our hands is the size of the ambition set out in that plan. If ambition and achievement are mixed together, readers may misunderstand that the United States is already flying cell-computer drones. The document does not support such a misunderstanding.

Still, the point this ambition aims at is clear. How can artificial intelligence be kept running where power is scarce? Most defense research answers this question with more efficient chips and lighter batteries. O-Circuit knocks on a different door. It asks whether computation can be moved from silicon to cells and made to run on nutrients like sugar water instead of electricity. Viewed through policy, this question goes beyond the success or failure of one technology. At the front line, where the energy supply line itself becomes combat power, an attempt to reduce the energy cost of computation by several orders of magnitude also rewrites the old military problem of supply. As the optogenetics program GO discussed in the same chapter seeks to handle biological neural circuits with light, these experiments by the Biological Technologies Office move toward blurring the boundary between computation and life.

The question this scene poses to Korean readers is not envy. It is whether we have the eye to read such planning documents as plans. The cool-headedness not to mistake dazzling numbers for achievements, the habit of separating an open attempt from a verified result. What DARPA shows through O-Circuit is not a result, but a coordinate. It is a coordinate marking how far, and in what unfamiliar direction, U.S. defense research intends to push into the physical limit called power. Whether the attempt will reach the targeted efficiency or end as expensive learning is something no one can promise for now.

Chapter 14 Battlefield Medical Autonomy: Triage, VITAL, RAPIID

The Field Problem Behind the DARPA Triage Challenge

Imagine that a building has collapsed. Chunks of concrete are piled up like a mountain, and people are trapped beneath them. Smoke fills the air, and because it is night, visibility is poor. One medic arrives. There are dozens of casualties. Who should be seen first? Who can survive if treated now, and who can wait a little? That decision must be made within minutes. This is the scene DARPA is trying to capture through a competition called the DARPA Triage Challenge, or DTC.

Triage is a medical term that means sorting casualties. It is the work of deciding who should receive treatment first. The word originally comes from French, meaning 'to sort,' and the concept took hold in battlefield field hospitals. When resources are scarce, the cold calculation is to treat first those who can be saved. The difficulty is that current sorting methods do not fit well when mass casualty incidents, or MCIs, send large numbers of victims into care at once. DARPA defines this program as a field problem for the U.S. Department of Defense, noting that current triage approaches do not fit mass casualty situations and have limited power to foresee when lifesaving intervention will be needed. The program is run by DARPA's Biological Technologies Office.

Why is the current method not enough? Primary triage requires a medic to approach each casualty directly and check the person's condition by hand. When every person must be touched and examined one by one, the method breaks down as the number of casualties rises. Judgment also varies by observer. An experienced medic and a new medic can make different decisions about the same patient. In secondary triage, where evacuation and resuscitation are handled, the ability to predict how a casualty's condition will worsen later remains moderate or lower. A patient who looks stable on the outside can collapse sharply a few minutes later. This gap delays treatment, and delayed treatment can become death. Someone who could have survived at first is lost to time. That is the real enemy DTC is aimed at.

DARPA wants to send machines in before people. The picture is one in which unmanned aerial vehicles, drones, and unmanned ground vehicles enter dangerous sites first and find casualties. Before a human medic's hands can reach them, robots scan the site and draw a map of who is where and in what condition. The actual field is nothing like a laboratory. The competition rules spell out the terrain the robots must cross. They must pass through mud, asphalt, and roads covered with ash left by fire. They must find people not only in bright daylight but also at night with almost no light, in thick smoke and dust, and amid flames. Multistory buildings and large debris piles block the way, and there are no markers the robots can use to judge their own location. To imitate this chaos, DARPA uses the Guardian Centers in Georgia and Disaster City in Texas as real competition courses. Algorithms that ran well on experimental data are tested there to see whether they fail in front of real rubble.

Time is brutally short as well. Competing teams are dropped into the site with only vague information, such as 'this is a combat situation' or 'a building collapsed in an earthquake.' They are not told where people are, how many there are, or what injuries they have. No scenario lasts more than 30 minutes. Within those 30 minutes, the robots must search an unfamiliar site and find casualties. That number, 30 minutes, silently shows how tightly initial response must be organized in a real disaster.

One more strict principle is attached. Robots must never touch the casualties. They cannot contact any part of a casualty's body or clear debris at the site. The injuries the robots must detect are not minor. They must identify head injuries, penetrating wounds, burns, respiratory distress, severe bleeding, multiple simultaneous injuries, and even amputations. They must read all of this in real time, from a distance, using only sensors and without touching the person. The recognition challenge is severe. What a human would confirm by checking a pulse and pressing on a wound, a machine must replace with vision and thermal sensing.

The rules also treat false steps by robots harshly. To imitate a battlefield with poor communications, DARPA command posts forcibly limit the number of casualty reports a team can submit. This is meant to stop teams from sending guesses at random and to preserve communications capacity. The same principle applies in scoring. If a team wrongly reports a casualty where no casualty exists, it is counted as a false positive and receives a weighted penalty. One machine error can send a medic to the wrong place, which wastes people and time. The scorecard reflects not only what the robot found, but also what it saw incorrectly.

The target rises over time. At first, remote human operation is allowed, but the final goal of Phase 3 assumes a situation where communications are cut and there are too few people to operate the machines. At that point, ten robots must move with full autonomy and without human intervention. The robots must find casualties, judge their condition, and assign triage categories on their own. The image is not one operator holding onto one robot, but a formation that spreads out by itself and divides the site.

The next stage after finding casualties has a different character. In secondary triage, when a medic arrives and attaches biosensors to a patient, the fight becomes a fight with data. Participants in the data competition work with real trauma patient records collected by groups such as the U.S. Army Institute of Surgical Research. They use 4,230 cases in Phase 1, 6,061 cases in Phase 2, and in the end more than 8,000 real cases. The algorithm must look only at the first few minutes of data after a patient arrives and predict whether a lifesaving intervention will be needed up to four hours later. It must report changes through sliding 15-minute windows and adapt to environments with different resources, such as a civilian hospital in the rear and a poor medical tent at the front. Even for the same patient, what can be done differs depending on which hospital that patient is in.

The program also meets the wall of ethics and regulation. Research involving human subjects is governed by strict rules. Self-funded teams that do not receive direct DARPA funding cannot be

supervised by DARPA, so they are forbidden to collect or use new human-subject data outside the officially provided materials. They are also barred from trying to reidentify patients from the military trauma data they receive or from sending that data outside. To get past this wall, teams use simulators built by connecting Unreal Engine 5, a game engine, with a physiology engine. They create detailed virtual people of different ages, sexes, and skin colors, then train algorithms inside that world by imitating blood pressure falling, heart rate rising, and capillaries changing because of burns. Since real casualties cannot be brought in carelessly, virtual patients become the training ground.

Seen through the eyes of policy, every rule of the competition exposes a U.S. Department of Defense concern. The 30-minute limit reflects the speed of initial response. The no-contact rule reflects the desire to send fewer people into dangerous zones. The penalty for false positives reflects the calculation that resources must not be wasted. The goal of a ten-robot autonomous formation already assumes a battlefield where communications fail. This is not a stage for showing off technology. It is closer to an exam sheet that writes down the worst conditions in advance and forces teams to compete within them.

One point needs to be marked clearly. DTC is an ongoing competition, not an achievement that has already changed the battlefield. The rulebook and program summary describe goals and test conditions, not verified results. A system in which ten robots sort casualties in rubble within 30 minutes without contact, look ahead four hours, and support lifesaving unmanned evacuation is the problem this program says it will solve. It is not proof that the problem has been solved. It is more honest to read it as homework still to be done. For Korea's disaster response or military medicine as well, the thing to envy may not be a finished robot, but the method of breaking a problem into small pieces and fixing those pieces as conditions.

The Design of VITAL and Medical Autonomy

Suppose one soldier collapses on the battlefield. The person beside him is usually not an emergency physician. He may be a trained medic, but he comes into the field with limited clinical experience from hospitals. He also has few grounds for judgment in hand. He estimates the patient's danger from a few thresholds, such as whether the heart rate has passed 100 beats per minute or whether the respiratory rate has fallen below 15 per minute. These numbers are averages drawn from many bodies. They do not tell him which way this one patient in front of him is now falling.

The problem is time. A patient who looks stable can suddenly sink a few minutes later. Once bleeding passes a certain threshold, it is hard to reverse. DARPA calls this point 'delayed lethality.' It means a condition in which, by the time the danger is noticed, it is already too late and the few minutes when action was possible have already passed. On past battlefields, where evacuation could happen at once, losing those few minutes may have been less fatal. DARPA's future battlefield is different. It assumes a front line where enemy interference blocks helicopter evacuation, communications and supply are cut, and wounded soldiers must hold out for a long time. The farther

the road to the hospital becomes, the heavier those few minutes in the field become.

VITAL is a program built to regain those few minutes. Its name means Virtual Integrated Twin for Autonomous Lifesaving. The core idea is to build a patient's cardiovascular system inside a computer. By receiving biosignals from the actual body, it creates that person's own 'digital twin' on the screen. Then it asks that twin about the near future. What will happen to this patient 15 minutes from now? If a tourniquet is applied, if transfusion begins, if CPR is performed, how will the results diverge? VITAL tries to run those branches virtually in advance and then give the medic the better choice to act on now. The aim is to move judgment away from paper thresholds and toward predictions fitted to the actual movement of this patient's body.

The idea is hard because performance and weight pull in opposite directions. To imitate the human body faithfully in physical terms, the model becomes heavy. It must take in rich data from MRI, CT, ultrasound, and chemical sensors, and calculate the movements of blood vessels and the heart one by one. DARPA calls this heavy version a high-fidelity model, or HF. HF is precise, but it demands large computing equipment. It cannot be run inside a trench where communications and power have failed. A model that gives an accurate answer in a laboratory and a model that must give an answer within five seconds on dirt have different requirements from the start.

VITAL prepares a second version. It is a reduced-order model, or ROM. This is a lighter version that selects only the truly important movements from HF and greatly reduces the amount of computation. One sentence in the program's own definition stands out. VITAL says the work is not making a model, but 'learning' the method of transferring a heavy HF into a light ROM. The tool used here is 'mechanistic artificial intelligence.' It takes body states such as bleeding rate or vascular tone, and interventions such as hemostasis and transfusion, as inputs, then predicts the waveforms that sensors would detect. Unlike a statistical model that memorizes answers directly from observed data, this approach tries to hold onto why the body reacts that way through physical laws.

The targets VITAL sets for itself are harsh. The target conditions are divided into two. The acute condition is hemorrhagic shock, and the chronic condition is heart failure. A proposing research team must choose one of the two when submitting an abstract. For hemorrhagic shock, the risk must be identified at least 15 minutes before the event occurs. For heart failure, deterioration from Stage 3 to Stage 4 must be predicted before 80 percent of the time needed for that worsening has passed. By Phase 2 of the program, that prediction time must be increased by more than 30 percent compared with the first phase.

The model must become lighter without becoming dull. When HF is reduced to ROM, performance loss must be held to 15 percent or less, and the model must not invent strange waveforms that could not come from a real body. At the same time, it must run more than 1,000 times faster than a well-tuned HF. The time needed to refit the model to the patient's real-time condition must not exceed five seconds. When it meets an unfamiliar situation it has never learned, it must recognize that fact with a probability above 90 percent and report its own chance of error. This last condition is

striking. The program fixes honesty in the face of the unknown as a performance metric, rather than only the ability to get familiar things right. It writes into the design that a machine that is confidently wrong is the most dangerous machine when lives are at stake. It is also a condition that the system must not become a machine that only throws out numbers. VITAL must produce expert-level reasoning more than ten times richer than standard evaluation methods. It must not stop at listing heart rate and blood pressure values. It must also carry the grounds for a judgment that makes clinical sense.

Verification mechanisms are attached so that these goals do not end as talk. An independent evaluation team separately supported by DARPA checks the research results from the outside. If animal or human research is included, the required ethics review must be completed before the work begins. This is a safeguard against a structure in which a team grades its own performance. From the view of a policy practitioner, this design matters. The more ambitious the target, the stronger the temptation to inflate the result. Giving evaluation authority to an independent team rather than to the researchers shows that the dull procedure called performance management is, in fact, the frame of trust.

VITAL is not aimed only at the battlefield. Program materials also present calculations that lead into civilian medicine. In the United States, cardiovascular disease accounts for one quarter of all deaths. The estimate is that reducing hospitalizations from heart failure by only 10 percent would save 7 billion dollars each year. The picture leaves open the possibility that predictive technology sharpened on the battlefield can move to hospital emergency rooms and intensive care units. Research begun with defense money while also aiming at civilian value is a calculation DARPA has used for a long time.

What has been written so far is not an achievement, but a design. A digital twin that looks 15 minutes ahead and a lightweight model 1,000 times faster are still targets. VITAL is a program run by DARPA's biotechnology office, and the numbers just listed are conditions the program sets for research teams, not records already achieved. Moving a heavy model into a light one without loss and detecting unfamiliar patient states by itself are not easy even in a laboratory. The higher the target, the greater the risk of failing to reach it. The official documents now available do not tell us how this program will cross that risk. A promise not to draw DARPA only as a success myth is kept in passages like this.

When VITAL is viewed in a slightly wider frame, the picture becomes clearer. DARPA's biotechnology office is dividing battlefield medicine into several programs and pushing them at the same time. The DARPA Triage Challenge discussed in this chapter tries to find casualties remotely inside collapsed buildings and conduct primary triage, while RAPIID tries to move blood substitutes that can be stored for a long time at room temperature onto the battlefield. Program materials also include a scenario in which these pieces fit together into one future. A drone finds a casualty, VITAL predicts the body's near future, and hemostasis and blood administration follow. That scenario,

however, is not a fact verified in the field. It is a sketch of the direction these programs point to together. When each piece will be completed, and whether they will actually connect, remains unsettled. Within that sketch, VITAL occupies the place of judgment, telling other equipment when and what to do.

The phrase medical autonomy can sound frightening. It can be read as handing human life and death to machine prediction. Looking closely at VITAL's design, however, it is closer to having machines alert humans to the few minutes they used to miss than to an ambition to replace people. It is not designed to push away the medic's hands, but to make those hands move a little earlier. The lesson Korean policymakers can take from this program lies in the same place. What deserves attention may not be flashy predictive performance, but the attitude of fixing goals in numbers, handing scoring to outsiders, and stating the possibility of failure in advance. Whether that design moves beyond paper targets and becomes battlefield reality will be answered by the test results still to come.

RAPIID and the Transition to Battlefield Blood Substitutes

First imagine a wounded soldier lying on the ground. Whether the injury came from a gunshot or a fragment, what brings that person down within minutes is not the wound itself. It is the blood leaving the body through the wound. When a large vessel tears, blood runs out quickly. Then a condition begins in which oxygen cannot reach every corner of the body. The muscles of the limbs, the kidneys, and the brain begin to starve for oxygen. Medicine has long known the answer to this scene. Give whole human blood, meaning complete blood from a person. The problem is that carrying that answer to the front line is almost impossible.

Official materials DARPA released in 2026 state this limit plainly. Human blood must be kept cold. If the refrigerated distribution system, the cold chain, breaks even once, the blood cannot be used. Its storage life is short. And the supply depends entirely on how much someone has donated. In a forward unit with no hospital facilities, or in a trench where no one knows when an evacuation helicopter will arrive, it is difficult to meet these three conditions at the same time. Even if bleeding is controlled well, a wounded soldier slowly collapses from oxygen shortage if there is no blood to give. DARPA calls this state, in which tissue lacks oxygen and function breaks down, ischemia. The program's name comes from that word.

RAPIID is an English acronym for Resuscitation and Prevention of Ischemic Injury and Dysfunction. It is run by DARPA's biotechnology office, the Biological Technologies Office, or BTO. This program has a preceding story. It is a program called FSHARP. FSHARP first tested whether biosynthetic components could imitate the core functions of human blood. Those functions are carrying oxygen, stopping leaks to achieve hemostasis, and restoring reduced blood volume and clotting factors. DARPA summarizes FSHARP as having reduced the risks in the process of making these components. It was the stage of confirming inside the laboratory that the principle works.

What works in the laboratory and what can actually be used on the battlefield are entirely different matters. Crossing that gap was the next gate. RAPIID was launched to target precisely that gap. The program manager, Major Robert Murray, Ph.D., explains its character this way. If FSHARP proved that making a blood-like material that can be stored at room temperature is scientifically possible, RAPIID focuses on turning that possibility into a 'deployable reality' that saves lives at the point of injury. The wording is worth watching. This is not invention, but transition. It is not a story of making something new, but a story of taking what has been made and putting it into a medic's pack.

The target timing is fixed. DARPA aims to push clinical trials, regulatory approval, manufacturing scale-up, and field deployment technology in sequence, and as early as fiscal year 2029, deliver to the battlefield a complete deployable system approved by the U.S. Food and Drug Administration, or FDA. The full schedule is 36 months. It is divided into two broad phases. During the first 12 months, teams prove safety and efficacy before use in humans and design manufacturing plans for mass production. Those plans must meet Good Manufacturing Practice, or GMP, standards. During the next 24 months, the work moves into early human clinical trials and confirms a commercialization strategy to move the technology into civilian medical markets or defense procurement. The structure is to build laboratory evidence within one year, then spend the next two years passing through the two gates of the human body and the market.

One striking part of RAPIID's design is that it does not place every burden on one research team. The program is divided into four technical areas that run at the same time. The first area creates the room-temperature-storable components themselves. These are oxygen carriers, platelet-derived products, and dried plasma. Each of the three is developed independently and pushed through clinical, regulatory, and scaled manufacturing stages. The second area handles the hardware that lets those liquids be used on the battlefield. This includes rugged packaging that can survive harsh environments, such as prefilled syringes or packs. Point-of-care diagnostic devices also belong to this area. They are judgment tools that tell a medic this wounded soldier needs transfusion now. The third and fourth areas support clinical work, regulation, commercialization, and sustainment. To cross the section where research results often disappear because they cannot move into the market, the point commonly called the 'valley of death,' DARPA runs this commercialization area alongside the first two areas. The sequence is not to finish the components first and then think about how to sell them. It means carrying regulation and market paths from day one.

It is worth pausing on why this structure is interesting from a policy view. Defense research in many countries often reaches the stage of making a good prototype. The problem comes next. A finished prototype may lose years in regulatory review, or it may end as one or two prototypes because no one prepared a mass-production process. RAPIID pulls that failure point into the program in advance. Regulatory response and manufacturing planning are put on the same timetable as research. It is an example of how DARPA's way of defining problems, which puts deployable results ahead of good science alone, appears in medicine.

A point needs to be stated honestly. RAPIID is not yet an achievement. It is a plan. The year 2029 is not an arrived result, but a target. Nowhere in these materials is there a record that room-temperature blood has actually been supplied to the battlefield. DARPA itself lowers its posture before regulation. The program's Q&A materials state clearly that the FDA is an independent government agency and that DARPA will respect and follow the FDA's authority on regulatory matters. They also say DARPA will work to obtain FDA emergency use authorization or formal approval for rapid transition. That is not a sentence guaranteeing approval. It is a sentence saying they will try to obtain approval. If this difference is blurred, a plan is misread as a performance report. Clinical trials that put artificial blood into the human body are an extremely careful field. Oxygen-carrier substitutes have been tried in several countries over the past decades, but many stopped after failing to cross the thresholds of side effects and safety. Future clinical results will have to tell how RAPIID will overcome that history.

This section therefore asks that RAPIID be read as a roadmap. It is a plan to break the 20th-century blood supply chain of refrigerators and donors, and to create blood that can be stored at room temperature and administered immediately at the point of injury. The ambition is large. The gates that must be verified are just as many. Human clinical trials remain, GMP-level mass production remains, and the FDA's judgment remains. Even if the casualty sorting of the Triage Challenge and VITAL's prediction of patient status discussed earlier become faster, those judgments lose their destination if there is no blood to give. If a robot finds a wounded soldier and artificial intelligence warns of danger 15 minutes ahead, prediction remains a number on paper without a treatment tool in hand. RAPIID is trying to fill that blank. That blank, however, is not filled yet. It is being filled toward 2029. Anyone thinking about Korea's military medicine should hold both admiration and caution at this scene. What deserves admiration is the design that calculates failure points and puts them on the timetable. What deserves caution is copying a target that has not yet arrived as though it had already arrived.

Chapter 15 Strategic Materials and Bio-Based Manufacturing

A Map on Industry Day

On March 12, 2026, a map appeared on the screen at an Industry Day presentation hosted by DARPA's Biological Technologies Office. It showed the Gulf Coast of Texas. Red lines crowded into one narrow stretch of coast. The image showed that refineries and product pipelines were concentrated in the area the materials called PADD3. It meant that almost all the materials the United States uses, such as plastics, adhesives, and composites, come from oil, and that the factories that handle that oil sit together in one region.

The presentation placed other images beside the map. They were news photos of Russian refineries burning. The reports described Ukraine's so-called Operation Spiderweb, in which drones struck refineries deep in Russia's rear, putting half of the fuel supply to Moscow at risk. On one side stood America's chemical infrastructure gathered in one place. Opposite it were enemy refineries collapsing under a swarm of drones. The meaning of placing the two images together was clear. DARPA first wanted to show how brittle America's current mode of chemical manufacturing could become in wartime.

The person leading the presentation was Dr. William Moundfield. The message he delivered from the stage was about chemistry, but also geography and security. DARPA summed up the picture in two phrases. One was 'Peer Adversary in Conflict.' The other was 'Vulnerable Infrastructure and Supply Chains.' If one missile or drone strike can stop a nation's chemical production, that is no longer only an industrial problem. It is a security problem. The plan born from that awareness is the subject of this chapter, the Fleetwood program.

For Korean readers, one point of background matters. The concentration of refining and petrochemicals along the Gulf Coast inside the United States is not an accident. It is history. Ports were nearby, crude oil flowed, and pipelines were already in the ground, so factories kept being built on top of that base. In peacetime it is efficiency. In wartime it is a vital point. When resources gather in one place, costs fall. When that one place is hit, everything stops. Fleetwood aims straight at the back side of efficiency, the weakness created by concentration.

Fleetwood is not yet an achievement. It is a plan. The materials this chapter rests on are the Industry Day slide deck (PDF-034), the FAQ document that explains the proposal instructions (PDF-012), and two official DARPA news releases (NEWS-020, NEWS-011). Nowhere is there a results report saying lignin has been decomposed 100 percent, or that aviation fuel has been made from palm leaves on a battlefield. For that reason, this chapter reads Fleetwood not as 'they did it,' but as 'they demanded that it be done this way.' If that distinction blurs, a plan turns into a report card. The promise not to write about DARPA only as a success myth is kept in places like this.

What It Means to Read a Plan Document

There is a good way to understand a DARPA program. Look at what it forces researchers to do. Not a few papers, but the lines they must cross to pass a proposal review. When you see which teams are accepted and which teams are rejected, you see what the organization truly wants. Fleetwood's FAQ document draws those lines tightly.

First comes money. The program set expected support at 4 million to 8 million dollars per team across Phase 1 and Phase 2. In Korean currency, that is roughly 5.5 billion to 11 billion won. Each phase lasts twelve months. Short and fast. This is not research that leisurely digs into principles inside a laboratory. It is designed to demand proof in numbers within a fixed period. A 24-month clock is attached to the proposal from the start.

The document also fixes who can lead. The prime institution that leads the research cannot remain only a paper administrator. The FAQ explicitly rejects a prime that serves only as a 'program manager with minimal technical contribution.' The prime must possess the core expertise needed for the program goals and take on a substantial technical role directly. A consortium cannot put a famous name on the front and push the work down to subcontractors. One line in the document cuts away the habit of taking money on reputation alone.

Another striking point appears. Government entities, including federally funded research and development centers (FFRDCs), university affiliated research centers (UARCs), and the U.S. Army Research Laboratory (ARL), cannot participate either as primes or subcontractors. National laboratories are excluded wholesale. Why would a defense program block government laboratories? The materials themselves do not give a reason. [Inference: This can be read as an intent to avoid locking the technology inside a military secret cabinet and to move it quickly into the capital and commercial ecosystem of private industry and universities. Technologies involving government laboratories often move more slowly into commercial use. Still, this interpretation is not a stated fact in the document. It is an inference drawn from the exclusion rule.]

The door is also open to foreign institutions. But the gate they must pass through is heavy. They must go through a process called Fundamental Research Risk-Based Security (FRRBS), and foreign researchers must explain not only their CVs and assigned duties, but why a U.S. person could not be used instead. They must submit a security plan to prevent access to controlled unclassified information (CUI), and they can join only after DARPA approves it. This clause is not someone else's story for Korean readers. It shows plainly what kind of wall stands before Korean universities or companies when they try to enter U.S. defense research. The door is open, but a mountain of paperwork sits on the threshold. Before speaking of international cooperation, we must understand the weight of those documents.

The proposal process itself is built like a survival game. Researchers do not submit a thick proposal at the beginning. They first submit an abstract of no more than five pages and one summary slide. A DARPA program manager screens those first. Only selected teams earn the right

to submit an oral presentation proposal (OPP), and only after presenting in person in Washington do they become candidates for an Other Transaction (OT) agreement. The design does not ask teams to win with a long document. It makes them pass several short gates. Instead of forcing everything into one submission, it filters teams at each stage. It is the opposite of spreading research money broadly. The program sets gates and opens the next door only to teams that survive.

The Trouble Called Lignin

Fleetwood's feedstock is not oil. It is plant residue. More precisely, it is what is called lignocellulosic biomass, residue from agriculture and forestry. DARPA's news material (NEWS-020) says this residue is discarded worldwide each year at gigaton scale. A gigaton means one billion tons. The point is that an enormous resource lies underfoot, yet remains unused.

The difficult component in this biomass is lignin. It gives plants shape and firmness. It is rich in carbon and has high energy density, but its chemical structure is tangled, making it hard to process. You can think of it as the glue-like component that lets trees stand and keeps stems from snapping. For that reason, lignin has long been burned as cheap fuel or thrown away. Burning the lignin left after papermaking in boilers has been a long-standing practice. Fleetwood targets this troublesome material head-on. It aims to break down lignin with atomic-level precision and turn it into materials that can be made inside the United States immediately when needed.

The program divides the technology into two areas. The first area (FA1) is to create catalysts that efficiently separate and break down lignin from real biomass. The second area (FA2) is to create catalysts that convert the resulting lignin fragments into expensive final chemicals. The target materials are specialty high-value chemicals such as pharmaceuticals, defense materials, and lubricants. The FAQ document gives no easy passes. Proposal teams must offer a solution that addresses both FA1 and FA2 to be evaluated. A team that only handles feedstock breakdown well and leaves the back end to someone else will not even be accepted. The program demands a complete connected process, a structure in which one team takes responsibility from raw material to final product.

The clause allowing almost any method is also interesting. Teams may mix any catalytic approach, including heat, electricity, light, cell-free biocatalysts, plasma, and living microbes. The method is open. DARPA does not care whether teams use microbes, run electricity, or shine light. But the destination is frighteningly high. The road is open, but the destination is not negotiable. This is the posture the program demands from researchers.

The Number 100 Percent

Fleetwood's required performance metrics leave almost no room for compromise. Participating teams cannot take a detour by selecting only easier carbohydrate components such as cellulose or hemicellulose. That strategy is immediately pushed outside the program scope and rejected. They must target lignin alone. The program blocks at the door the familiar easy path of processing only the

easy parts and burning the difficult lignin.

The numbers rise like this. By the end of the twelve-month Phase 1, teams must show more than 50 percent in delignification and aromatic monomer yield. By the end of Phase 2, they must prove more than 90 percent delignification and 100 percent separation and conversion of lignin into monomers by mass. In a chemical process, 100 percent is close to a demand for no waste. In the laboratory, a 100 percent yield is almost mythical. Reactions always leave residue, byproducts form, and losses occur. Knowing that, DARPA still wrote 100 percent into the target line. This reads as a way to see how far teams can push in that direction, rather than only whether they arrive there. The steepness of the target is contained in that one number.

The experimental conditions are also controlled. To prevent participating teams from bringing in feedstock that suits their own preferences and showing only good results, all experimental biomass must be unprocessed feedstock procured and provided by the government test and evaluation (T&E) team. Once teams touch the feedstock, the objectivity of the result can wobble, so the government holds that handle from the start. Anyone can produce good numbers with good feedstock. Real ability is judged by rough feedstock handed to you by someone else.

The feedstock list draws attention. DARPA's example biomass includes rice residue, banana waste, coconut shell, coffee grounds, nut shells, palm leaves, corn stover, sugarcane bagasse, soybean residue, macadamia shells, sargassum seaweed, switchgrass, elephant grass, forest waste, and wheat and barley residues. There are more than seventeen types. Why list so many different things? Because the battlefield is not one place. In a jungle, palm leaves lie everywhere. On a coast, seaweed washes ashore. In farmland, corn stalks remain. Since no one knows where the fight will occur, the idea is to make materials from whatever the land provides. That thinking is embedded in this list.

Old paths are blocked too. Existing methods that break everything apart through high heat, such as pyrolysis, gasification, and hydrothermal liquefaction, are outside the program scope. Strategies that aim at low-value C1 to C2 compounds such as carbon dioxide, methanol, and ethanol, or hydrogen, as final targets are also barred. DARPA wants a new path such as 'reductive catalytic fractionation,' not a rough burning down of everything. The goal is not to make piles of cheap output. It is to preserve the complex molecule called lignin and turn it into valuable fragments. The order is to overturn a century of treating lignin as something to burn away, and to revive lignin as a raw material.

Why Money Calculations Are Forced

Up to this point, the program may read like a difficult chemistry assignment. But Fleetwood does not demand chemistry alone. It demands money calculations as well. Every participating team must submit techno-economic analysis (TEA) and life-cycle assessment (LCA) models. They must translate not only chemical reaction equations, but capital expenditures, operating costs,

maintenance costs, market prices of catalyst reagents, and energy consumption by reaction temperature and pressure into money. They must write in the calculation sheet how hot the reactor must be, how much electricity is needed to produce that heat, and how much that costs.

The demands continue. Teams must show how the system can scale across three production levels, from small modular production deployed forward to large chemical plants. The government T&E team uses these materials to verify the minimum selling price (MSP) of the resulting materials. DARPA fixes in place from the proposal stage the fact that something that works in the laboratory and something that survives in the market are different. No matter how elegant the chemistry is, no one will use it if it costs more than oil. Researchers are made to face that cold arithmetic from the start.

At the end of each phase, teams that prove outstanding economic value receive a 50,000-dollar cash performance incentive. It is not a large amount of money. In an 8-million-dollar program, 50,000 dollars is closer to a symbol. But the prize aims not at the amount, but at the direction of competition. It makes teams compete on economic value, not chemical performance alone. Teams are pushed to compete not over who breaks lignin better, but over who can make it at a price that can survive in the market. It is a device that moves laboratory pride onto the market's calculator.

Even Waste from Units

Biomass is not the only resource here. Waste produced by U.S. forces during operations is also a burden at the front line. The U.S. Department of Defense Fiscal Year 2026 budget document (PDF-001) includes a program called 'Turning Upcycled Waste into Novel, Sustainable Materials.' In the budget, it is listed as the WP-01 task under PE 0602024E and the MBT-02 task under PE 0602715E. The fact that the budget document assigns item numbers means this concept has moved beyond the idea stage and become a funded project.

In expeditionary operations, the military depends on products made from petroleum, rubber, and wood. But after use, they pile up as waste around units. They bring no value and add only logistics and environmental burdens. Picture piles of scrap tires, broken wood, and discarded paper rising on one side of a forward base. To handle them, trucks must come and go again, and those trucks become targets. This program studies the feasibility of turning common defense waste such as scrap tires, waste wood, and paper back into durable materials. If it succeeds, the military could recycle waste generated by units on site and turn it into emergency construction materials or basic chemicals, instead of hauling expensive materials from the homeland. It is worth remembering that this picture is still at the feasibility study stage. The scene of trash becoming material is attractive, but the budget document says only that the possibility will be examined.

Transition Is Harder Than Invention

Another axis in the story of strategic materials is rare earths and critical minerals. From fighter jets to smartphones, advanced devices always contain rare earth elements and critical minerals. Electric

vehicle motors, missile guidance devices, and data center servers cannot run without these materials. The problem is that the United States and other Western countries depend overwhelmingly on foreign sources for them. DARPA called this a 'critical vulnerability in the supply chain' (NEWS-011). The description of the 'OPEN' program in the budget document gives the same warning. It says critical mineral markets and supply chains are brittle, and when an international supply shock comes, prices jump quickly and sharply, sending immediate ripples through the economy.

Shashall Manning, director of DARPA's Commercial Strategy Office, left a sharp line: 'The challenge in this space is not invention, but transition.' It is short, but heavy. Even if a brilliant technology for extracting rare earths is created in the laboratory, supply chain independence remains out of reach unless it grows to real factory scale and withstands the market. Between good science and usable capacity lies a deep gap often called the 'valley of death.' Papers exist, but factories do not. Patents exist, but products do not. Countless promising technologies disappear in that gap. DARPA's question is how to cross it.

The answer that emerged is SMART. The name stands for 'Strategic Materials Accelerator & Research Test Bed.' DARPA and the state of Utah agreed to establish it through the University of Utah. The goal is to increase domestic U.S. supply of critical minerals and rare earths needed for national and economic security. This facility gives innovators a physical place to test the performance of their technologies directly at meaningful scale beyond the laboratory. Jakob Jensen, the University of Utah's vice president for research, described what makes SMART special this way: the ability to connect world-class discovery with equally strong benchmarking and scale-up infrastructure. In other words, the place of discovery and the place of growth have been put side by side. It provides a factory front yard where results from the laboratory can immediately be tested at larger scale.

Rare Earths Mined by Microbes

The first team to test SMART's scale-up capability will be researchers from DARPA's EMBER program. EMBER stands for 'Environmental Microbes as a BioEngineering Resource.' They are creating biotechnology methods to separate and refine rare earths from resources inside the United States that have not been properly used. The target feedstocks are striking. They are not overseas mines, but phosphate mine waste left inside the U.S. mainland, acid mine drainage, and residue from electronics recycling processes. The idea is to mine what has been discarded. Red water flowing out of abandoned mines, piles of old electronics. Today they are costly problems to manage, but the story changes if rare earths are inside them.

The idea that microbes can mine minerals may sound unfamiliar. The principle is this. Some microbes have properties that let them bind to particular metals or dissolve them. If those properties are refined, minerals can be selected using living microbes instead of hot acids or harsh solvents. Dr. Tiffany Prest, EMBER program manager, says performers are developing new biotechnology for

mineral processing, and SMART will provide an important opportunity to test and refine this innovation. She expects the test bed to help teams understand the real bottlenecks and challenges they face as they scale the technology, and to mature laboratory-level microbial technology into usable domestic resources. Her words carry the same concern Manning pointed out earlier. The fact that microbes can extract rare earths is not enough. The process must be repeatable at scale.

The budget document contains other pieces that broaden this picture. The SMASH program (PDF-030) sets an extreme goal: to individually separate 80 elements of the periodic table from government-provided solid feedstock at 100 percent. It is an ambition to split the periodic table apart, element by element. The NASCENT program (PDF-040) links economic mechanism design and computer science under the name national security economic theory. Beyond chemistry and biology, economics too is drawn into the material of defense programs.

When these pieces are placed side by side, one rough concept begins to appear. [Inference: One can imagine a state-led coordination that works like gears: mine with microbes (EMBER), separate elements with precision (SMASH), scale and verify it at physical size (SMART), and wrap an economic defense net around it so market prices are not shaken by hostile dumping (OPEN, NASCENT). Still, this assembly is an interpretation drawn from several documents, and these materials do not show that DARPA has formalized it as one integrated blueprint.] Why is an economic defense net needed? Because no matter how much rare earth is extracted inside the United States, it is useless if the price does not work. If a country cuts prices sharply and floods the market, U.S. mineral startups that have just begun to walk cannot withstand it and collapse. They die by price, not technology. To prevent that death, there must be eyes watching prices and devices that help firms endure when shocks arrive. This idea itself shows a view of security not as chemistry or mining alone, but as a market problem.

Do Not Confuse Plans with Results

The thing I was careful about while writing this chapter was not reading plans as results. Fleetwood's presentation slides and FAQ contain powerful scenes. Drone aviation fuel drawn from palm leaves in a jungle, precursors for explosives and pharmaceuticals stamped out from seaweed on a beach, isolated units continuing the fight endlessly with only the plants underfoot. A small refining unit set up beside a tent, and aviation fuel coming out when pieces of wood gathered nearby are fed into it. It is a narrative that makes the reader's heart beat faster.

That dramatic battlefield self-sufficiency scenario appears in the materials as [inference], not as fact. Phrases such as a self-sufficient kill web or ultimate material autonomy are interpretations in which commentators push the ambition drawn in the presentation materials further. The actual documents speak much more plainly. Meet these metrics. Follow these rules. Prove this economic case. There is not yet a results report saying performance has reached the target line. These materials contain neither experimental results showing that lignin has actually been decomposed 90 percent nor reports that fuel has actually been produced on a battlefield.

This book's attitude of refusing to read DARPA only as a success myth is tested at this very scene. A dazzling future image is the purpose of a program, not its report card. The 100 percent written in the plan is not a number reached. It is a number demanded. If the two are mixed, work that has not even begun is mistaken for work already finished. That kind of mistake is dangerous above all in the making of policy. Respect DARPA's ambition, but do not say it has been realized. That balance is the principle running through this entire chapter.

A Mirror for Korean Policy

The story in this chapter reflects directly back on Korea. Heavy dependence on certain countries for rare earths and critical minerals, and a structure in which chemical and materials supply chains gather at a few points, are problems Korea faces with even greater urgency. If we place the discussion of Korean defense AI from my earlier books only as background, the points worth envying and the points that should not be copied as they are begin to separate.

What is worth envying is the view that defines the problem as security. When America saw chemical plants concentrated in one place, it saw a vital point rather than only efficiency. The design that does not stop at invention but ties transition and scaling into one program is also worth learning from. It is a structure that does not end with one paper, but pushes the result toward the market. The cold insistence on money calculations and minimum selling price from the proposal stage, the method of preserving objectivity by having the government control the feedstock, and the rule that filters out reputation-only teams with a single line in the document. In such details lies the operational wisdom DARPA has built over many years.

There are also things Korea should not copy as they are. A short phase worth 8 million dollars and a steep metric of 100 percent work only when a thick industrial ecosystem and a budget that can absorb failure stand behind them. In the United States, layers of chemical companies and universities can attempt those metrics, and even if several teams collapse, the program keeps moving. There is thickness that can bear failure. If Korea copies only the metrics without that support, it can easily create a program whose plan is dazzling but whose transition is absent. If the target line is set high but only two or three teams can attempt it, the high number becomes not a record of achievement, but a record of frustration. This is where Korea often falls when it speaks of a Korean-style DARPA. Flashy goals are easy to copy, but the ecosystem and failure-tolerant budget that support them are not easy to build.

The questions suited to Korea are different. Which materials do we depend on certain countries for? Which industries stop first if they are cut off? Who will test the alternatives at factory scale, not in a laboratory? Does Korea have a validation space like the test bed America built in Utah? These questions must come before copying the number 100 percent. Envy makes us imitate goals. Honest self-examination makes us build the support first.

DARPA's real question was not what percentage of lignin can be broken down. It was how to move laboratory marvels into real capacity that can survive enemy resource weaponization and market disruption, and who will cross that gap and how. This question is not America's alone. No one yet knows when the scene of making aviation fuel from palm leaves will become real. Perhaps it never will. But an organization exists that writes that scene into a plan, puts the number 100 percent and money calculations into that plan, and prepares a factory front yard to verify it. That fact, at least, is confirmed now. And in front of that fact, Korea's question should not be how much it envies America, but how honestly it is looking at its own vital points.

Chapter 16. Small Businesses and Startup Transition

So far, this manuscript has followed large programs. Artificial intelligence that repairs code on its own, vehicles that drive over dirt roads without roads, robots that repair satellites in orbit. When you look at pictures like these, it is easy to see DARPA only as a stage for enormous experiments. Behind that stage are much smaller scenes. A company with thirty employees. An engineer revising documents late into the night. A five-page business plan. This chapter looks at how those small scenes walk into the inside of defense infrastructure. It also looks at how narrow that path really is.

SimVentions and a Case of Successful Transition

Picture an officer sitting in front of a command post screen. The situation in front of him changes fast. He has to make a judgment within minutes. What he needs now is not a new weapon. He needs to find, among hundreds of past simulations, the one that fits the present phase. Those materials are scattered across units and agencies. Their formats differ. People prepare the lists by hand, so items go missing and fall out of order. When the moment becomes urgent, he cannot find a scenario worth using. It is close to a warehouse overflowing with goods but no inventory.

The U.S. military calls this kind of virtual experiment modeling and simulation. In short, M&S. It means running countless trials inside a computer before deciding where to place forces or how the enemy may move. As weapons systems have grown complex and operations across several domains have increased, commanders have come to rely on it before making decisions. The problem was not the experiment itself. The path to taking out the results and using them again was blocked. Data from different agencies did not work together, and errors entered when people manually listed assets. Even if valuable scenarios already made for an emergency were sleeping somewhere in a warehouse, nobody could pull them out. DARPA's challenge began from this scene. Gather scattered M&S assets in one place, reduce the time needed to find them, and build a repository that units could share. It was not a request for a good paper. It was a request to remove the inconvenience in front of them.

The company that answered this challenge was SimVentions. The solicitation topic number was HR001121S0007-03. DARPA supported the company with \$1.5 million through the Small Business Innovation Research program, or SBIR. With that money, SimVentions built a secure web application called SmartKM. The first SmartKM was not a grand object. It was closer to a warehouse that safely gathered M&S items such as datasets, scenarios, and reports dealing with real battlefield problems. The focus was on how to store, organize, and find knowledge again. It did not invent a new weapon. It made existing materials usable.

It is worth pausing over the shape of SBIR as a system. The U.S. Small Business Innovation Research program gives money in stages. First, it checks on a small scale whether an idea may work. Then the company builds a prototype. At the end, it moves to a contract that a real customer buys. DARPA's SBIR and STTR proposal instructions put this flow into documents and set the

conditions for moving to the next stage. This means the thresholds are placed one after another. For that reason, most small companies stop at the prototype stage. More often, something that looked plausible in the lab disappears before it crosses into real military infrastructure. The SimVentions case remained in the record because it crossed that last threshold. If the companies that cross it stand out, the reverse also tells us that many companies did not.

In December 2024, DARPA signed a three-year contract with SimVentions through its Adaptive Capabilities Office. The amount was \$14 million. In the SBIR system, this stage is called Phase III. It is where early research passes validation and moves toward real deployment. Research that began with \$1.5 million led to a main contract nearly ten times larger. This scene shows how seed funding grows into a much larger adoption contract. The important point is not the tenfold multiple itself. It is the order: the government first put in early money while bearing the risk, and only after that risk receded did the larger money follow.

The core of this contract was to add artificial intelligence and machine learning to the existing SmartKM. The request was to change it from a warehouse that only collected materials into a repository that learns and reasons on its own. DARPA wanted SmartKM redesigned so it could raise the speed of commanders' decisions and widen the range of usable solutions. The company described three results. Lowering the information walls that had stood between the Department of Defense and the services, making real-time cooperation possible. Protecting security by strictly matching each user's credentials against sensitive items inside the repository. Reducing the time needed for search and evaluation, so the system could point first to operational options that human eyes might have missed in a busy situation. The third point catches the eye. It is a story about a machine picking up answers that people missed. A warehouse organizing tool had begun to look toward the seat of a staff officer.

The direction of development also appears. SimVentions set a goal of formally transferring SmartKM to the Office of the Under Secretary of Defense for Research and Engineering in 2026. The company said it was not stopping at use inside the U.S. military, but was also continuing talks about adoption with U.S. Indo-Pacific Command and the Department of Homeland Security. It goes one step further. The plan is to expand the technology into civilian industries that use a great deal of M&S, such as aviation, automobiles, logistics, medicine, and machinery. This is the familiar picture of dual use, in which a tool grown with defense money flows into civilian markets. The route by which an object used in the military spreads to hospitals and airlines is a structure that lets technology made with defense spending earn its value once more inside the U.S. economy.

The background that allowed SimVentions to handle a contract like this is also worth noting. The company opened in 2000 and now has more than 450 employees. In 2020, it became an employee-owned holding company. It had long provided systems engineering, software development, and M&S services to the federal government. It supplied cybersecurity work to the Tomahawk missile program at the Naval Surface Warfare Center Dahlgren Division and had also

handled Navy electronic warfare tasks. This was not a newborn company that fell out of the sky. It already knew the grammar of defense contracts. Even such a company needed five years to push one new repository into field use. Tram Chase, vice president of the company's Technology Solutions Group, said SimVentions had spent the past five years building SmartKM with DARPA, that it had become a point of pride for the company, and that it had benefited both DARPA and the company. The fact that even a company that already knew the defense market needed five years gives a sense of how long this road must feel to a startup knocking on the door for the first time.

I want to draw a clear line for the reader. The story so far comes from DARPA's own transition success story document. That means it is a piece in which the contracting agency introduces the fruits of its own program. Facts such as contract amounts, dates, and the company's history can be accepted as stated. But independent evaluation of how effective the technology was in actual field use, or what made it better than other repository attempts, cannot be filled in by this one document alone. A success story should be read as a success story, not immediately treated as the whole of verified performance. If this manuscript is not to become a collection of applause for success stories, it has to keep making that distinction.

Even so, the picture this case gives Korean readers is clear. A small seed fund of \$1.5 million. A path by which that money grew over five years into a \$14 million main contract. An attempt to spread the tool again into civilian industry. The work of taking something already in the hands of a small company, adding artificial intelligence to it, and pushing it into defense infrastructure, no less important than inventing a new technology from the start. A structure that places thresholds layer by layer and filters out most firms, but attaches a contract ten times larger to the one that crosses them. This later part is what DARPA does well. Transition rather than invention. Settlement rather than invention. Behind the name SimVentions, that sentence stands. The same sentence repeats once more with the next company, with the digits changed.

Jarjet Technologies and the SBIR Path

If the previous story was software that handles materials, this time it is a component you can hold in your hand. It walks the same SBIR path, but the scale of money changes by one digit.

There are radio waves that the eye cannot see. GPS signals, military radio, satellite communications, and mobile phone calls all pass over them. On the battlefield, the side that captures these waves first and cleanly sees earlier than the other side. It must see early to judge early. The U.S. Department of Defense gave this capability a name: spectral dominance. It means a state in which access to the spectrum is not disrupted. The term sounds grand, but when you take it apart, it eventually comes down to components.

The problem is that these waves are analog. For a computer to read the analog signal received by an antenna, it must first be converted into digital form. The component that handles this conversion is an analog-to-digital converter, or ADC. If this conversion is slow or consumes too much power,

even a good computing device behind it cannot run at full speed. However smart a brain may be, it is useless if the eyes are cloudy. The bottleneck choking many defense electronic systems was right here. So DARPA asked researchers a question. Can you make a more efficient and cheaper ADC? If the question in the previous section was how to gather scattered materials, the question here was how to brighten blurred eyes.

The small company that answered this question was Jariet Technologies. DARPA's transition case document, PDF-021, places the company's path on the SBIR route and explains it there. SBIR stands for Small Business Innovation Research. It is seed funding designed so that small firms with limited capital but advanced technology can offer bold ideas for national security problems. It is not a method of pouring large sums into large companies. It is a method of first placing seeds in small hands that have ideas.

SBIR is divided into three stages. Phase I looks at feasibility, whether the idea can become technology. Only seeds that seem likely move to the next square. Phase II develops the technology and builds a prototype. This is the section where a tangible object is made. Phase III meets large government or private capital and moves into fielding and commercialization. Phase III is the gate among gates. DARPA's proposal instructions, PDF-036, require two things in this process. One is a commercialization strategy. The other is dual-use capability aimed at both military and civilian markets. Good technology is not enough. It means the company must write from the start a plan for selling that technology and a use that works in both markets. If SimVentions in the previous section showed the result of dual use, these instructions show that dual use was not an accident but a condition the system required from the beginning.

There is a reason for this demand. The industry calls the section where technology fails to move from the lab to the market the 'valley of death.' Companies often make it as far as a prototype, then disappear because they cannot find the money and partners needed for mass production and real sales. SBIR is a bridge that helps them cross that valley. Requiring a commercialization strategy and dual use is close to telling companies to build the muscles for crossing the bridge in advance. The fact that DARPA provides separate guidance for small businesses also appears in the 2026 small business FAQ, PDF-016. It separately organizes frequently asked questions so that startups unfamiliar with the procedure do not turn back at the door. The hand that builds a high gate and the hand that places steps in front of that gate exist together inside the same agency.

Now look at Jariet's numbers. The topic number it received was SB153-004. According to PDF-021, the company received \$3 million across two Phase II SBIR awards, and \$1.5 million in Phase III government matching funds was added. Together, that makes \$4.5 million. In the eyes of a defense budget, this is not even the price of one major weapon. With this money, the company made an ADC chiplet with high sampling speed.

The word chiplet may feel unfamiliar. Let us draw the scene first. Older semiconductors printed the needed functions as one large mass. Everything was carved into a single large plate. A chiplet is

different. It connects small functional pieces like toy blocks to make the desired shape. Jariet's chiplet performs faster analog-to-digital conversion while greatly reducing power consumption. When these pieces are laid like tiles on a larger chip, several communication channels can be installed at once. That means it becomes easier to increase channels. Chiplets made this way go inside digital phased-array antennas. 5G base stations, radio-frequency transceivers, satellites, and radar all use these antennas. The fact that one component can be used across many boards is the seed of the story that follows. If a component that goes into military radar also goes into civilian base stations, the condition called dual use is not paperwork forced into shape. It becomes an inborn property of the component.

That was the story of the technology. Next comes the story of money and markets. After DARPA's early investment ended, Jariet attracted more than \$150 million in customer investment. The money called in by \$4.5 million was more than thirty times larger. Most of this revenue came from NRE, non-recurring engineering, meaning one-time custom contracts to design and make new products. Companies whose names are widely known joined in. Lockheed Martin, L3Harris, and Altera put capital into adopting the chiplet. On the military side, Jariet signed a contract with the Naval Surface Warfare Center's 'STEAM PIPE 24' initiative. A single chip from a small company walked into the inside of U.S. military weapons systems. After the seed money dried up, the market took over so the waterway would keep flowing. If the government attached ten times the money in the SimVentions case, here the government planted only the seed and private and military heavyweights filled in the other thirtyfold.

The company's size grew with it. Jariet had 30 employees when it was founded in 2015. It now has 80. In 2025, it established Jariet International, a Swiss subsidiary, with 23 people. That was a step aimed at Europe and international markets. This chiplet does not remain only in defense. The booming civilian commercial satellite network industry needs the same component. Dual use by the military and civilians, the condition SBIR required from the beginning, returned as real revenue. David Clarke, the company's senior vice president for program and contract management, said DARPA's investment played a decisive role in developing cutting-edge integrated circuit technology that supports spectral dominance.

This point needs to be handled honestly. The numbers and quotation just introduced all came from a 'success story' document that DARPA published itself. As the name says, it is a favorable account that puts the successful parts first. The document on another company introduced in the same success story format, SimVentions, PDF-010, has the same character. Such documents do not contain failed proposals or the stories of other small businesses that failed to cross the valley and disappeared. It would be wrong to read the path from \$4.5 million to \$150 million as the average for every company that passes through SBIR. Jariet is one case that ran well, and we leave it as a case. How many technologies SBIR actually transitions and how much money it moves cannot be known from this one document alone. You cannot look at one successful tree and say the whole forest is dense.

Even so, the path leaves a clear picture for Korean readers. The state did not pour in a large sum all at once. It first planted a small seed of \$4.5 million in a small company with technology. It required a commercialization strategy and dual use so that the company could cross the valley of death. On the other side of the valley, Lockheed Martin and the Navy were waiting. The seed grew not with defense spending but with market money. The sequence by which small seed funding calls in large capital, together with the institutional texture of proposal instructions and small-business guidance that supports that sequence, must be placed together for the picture to be complete. The point Korea should watch is not the name of the successful company. It lies in the design of the system that got that company across. Names can be swapped out, but design remains.

Between the Proposal System and Real Transition

When the stories of the two companies are placed side by side, a common frame appears. The first square in that frame is the proposal. Success is completed in the contract amount, but the beginning was always in front of a few pages of documents.

Imagine the office of a small company. About thirty employees. A few engineers stay late into the night, revising a bundle of documents. On the desk are circuit diagrams and graphs of test results. The document now being revised is not a technical manual. It is a five-page document called a 'transition and commercialization strategy' that will go inside a Phase II proposal to DARPA. Within those five pages, the company must prove that its technology can be sold to the military and to civilians, and show who has already sent a letter saying, "We will use this." However good the technology is, trouble begins if this scene is empty. A company stays up all night not because of the circuit but because of the sales route. That scene is the starting point of this section.

This demand is written into DARPA's SBIR and STTR Phase II proposal instructions, PDF-036. Let us unpack the terms briefly. SBIR refers to Small Business Innovation Research, and STTR refers to Small Business Technology Transfer, both long-standing U.S. support programs. They are channels made so that small companies, not large defense contractors, can receive defense research funds. This channel is usually divided into three stages. Phase I is a short stage that checks whether an idea may work. Phase II is the main stage of building a prototype. Phase III moves into actual procurement and the market. The problem is that the program already requires Phase III preparation during Phase II. It asks a company that has not even finished its prototype to write down first whom it will sell that prototype to.

Follow the instructions as they stand. For a small business to receive Phase II funding, it cannot only write how outstanding its technology is. It must include a transition and commercialization strategy within no more than five pages, and show that the technology has dual-use capability, meaning it can be used in both defense and civilian fields. Letters of intent from the Department of Defense or defense contractors provide strong support. At the proposal stage, the question is already, 'Show us where you will sell it.' Before research even begins, the structure makes the company write where that research will go after it ends. The five-page limit also matters. It signals

that DARPA wants not a long list of aspirations but a compressed execution plan. If a company cannot draw its sales route within five pages, the system sees it as not yet ready. We have already seen how this demand bore fruit in the two companies above. SimVentions' move into civilian markets and Jariet's entry into the satellite market both began from one line in the first proposal: dual use.

The title of this section comes from this scene. Between the proposal system and real transition. Between those two points lies a wide valley. Researchers call it the 'valley of death.' It is the point where technology has clearly been proven in the lab, but it cannot cross into real equipment or the market and quietly disappears in a pile of papers. The ability to write a good proposal and the ability to move that technology into an actual battlefield or market are different things. The former is close to writing skill. The latter is a matter of money, contracts, and customers. One good paper does not turn by itself into one good product. Between them, production facilities are needed. Buyers are needed. Time to pass validation is needed. The stories of the two companies that DARPA published under the name success stories are records of crossing this valley.

First, look again at SimVentions. The U.S. Department of Defense had an old headache. Modeling and simulation assets, meaning tools that imitate and test weapons and operations by computer, were scattered separately across departments and did not connect with one another. The same thing was made several times, and waste accumulated as offices bought new tools without knowing what the office next door already had. SimVentions received \$1.5 million in SBIR funding from DARPA and built a web program called SmartKM to bind these scattered assets together. At first, it was close to a warehouse organizing tool. It was closer to an inventory sheet that showed what was where.

The story after that is the heart of this section. In December 2024, this technology moved into a \$14 million Phase III contract. One and a half million became fourteen million. As the contract grew, artificial intelligence entered the system, and its character changed into a tool that first finds operational alternatives a commander might have missed. A program that once showed a list changed into one that speaks first and says, "There is also this option." The company now has more than 450 employees, and in 2026 it is expected to make an official transition to the Office of the Under Secretary of Defense for Research and Engineering. It is also entering civilian markets such as aviation and medicine. The promise of dual use written into the proposal did not end as a phrase on paper. It led to real revenue, a rare case.

The record of Jariet Technologies is one digit larger in scale. The company received a total of \$4.5 million from DARPA: \$3 million in Phase II and \$1.5 million at the start of Phase III. What it made with that money was an ultra-high-speed analog-to-digital converter, a chip that turns radio-wave signals into numbers. You can think of it as a device that moves wave-like signals captured by an antenna into numbers a computer can handle. The faster this conversion is, the wider the frequency band it can handle, and in electronic warfare it can read and respond to signals before the other side.

Jariet designed this chip in chiplet form. Instead of printing one huge chip as a whole, it connects small pieces like Lego blocks. Because the assembly can be changed as needed, it is far more flexible. If one part fails, only that piece can be replaced, and the configuration can also be rebuilt to fit another weapons system.

The result was larger than a few proposal pages could have imagined. This chiplet drew investment from major defense contractors such as Lockheed Martin and L3Harris and from the Naval Surface Warfare Center, and customer funding attached in Phase III exceeded \$150 million. Compared with the first \$4.5 million, the digit changed. The seed money supplied by DARPA became a small portion of the total. The government took on the risky early section first, and when the technology proved useful, large money from private firms and the military followed. The staff grew from thirty to eighty, and the company even established a subsidiary in Switzerland to enter the commercial satellite communications market. A chip that began in defense research extended into civilian satellite communications. Once again, it was a case of dual use.

Small businesses entering this path for the first time have many questions. What documents are needed, how the procedure works, and when to submit what. The requirement to write a five-page strategy document is itself a threshold. For a company trying for the first time, this threshold feels high. DARPA issues separate guidance documents for small businesses to answer those questions. The BARK-related FAQ from 2026, PDF-016, is one such document. It is not a flashy performance record. It is closer to practical guidance for lowering the doorway. If the proposal is a high gate, this places one step in front of it. Behind major success stories, these quiet devices that help entry also exist. If you look only at the two successful companies, the system appears dazzling, but underneath it lie documents like these that guide companies coming for the first time. Success is not made only by contracts. It is also made by the small repairs that prepare the road leading to those contracts.

Step back, and there is something to state honestly. The stories of SimVentions and Jariet came from documents that DARPA itself published under the title success stories. That means they are stories selected from successful companies. These documents do not say how many proposals were rejected, how many technologies disappeared in the valley of death, or how much money was wasted in the process. What this section can show is a map of those who crossed the valley. It does not contain the traces of those who did not. We cannot know from these materials alone what the success rate is, or whether the two companies are close to the average or exceptions. If we blur this point, DARPA begins to look like an agency without failure. That is not fact. It is a limit of the materials. To draw the map properly, one has to count not only those who crossed but also those who fell, and the materials in our hands now show only the side that crossed.

So before envying these two companies, Korean policymakers should first ask about the invisible side. A system that makes firms write a commercialization plan and likely buyers at the proposal stage. A channel that keeps Phase III open so seed money can lead to a truly large contract. The few companies that survive on that channel. What Korea may need to bring over is not this success

story, but the frame of the system that made the success possible. Not stopping at a review that distributes research funds, but asking about sales routes and dual use from the proposal stage, and making a place in advance for a large contract to attach after seed money. The many failures filtered out by that frame remain absent from these documents.

As I close this chapter, I want to leave one last point. What the United States gained from this route was not the names of two companies. It was the sequence itself: putting small money first and attaching large money later. Invention happens anywhere. The bridge that carries invention to the market and the battlefield, the rule that demands that bridge from the first page of a proposal, and the guidance for small companies hesitating in front of the bridge. If Korea talks about a Korean-style DARPA and needs something to copy, it is not the portrait of a successful company but the blueprint of this bridge. It must also remember that the names of those who failed to cross are not written on that blueprint.

Chapter 17 Academic Links and DARPAConnect

Twenty Years of the Young Faculty Award

A young professor who had been sitting in a university lab wrestling with equations climbs into a military vehicle a few months later instead of handling lab equipment. DARPA has arranged a field visit. The professor sees with his own eyes what kind of heat, vibration, and shortage of power the circuit diagram he had been drawing on paper must endure in practice. That day, he learns in his body how wide the gap is between the clean conditions of the lab and the rough conditions of the field. The familiar equations remain the same, but the place where those equations must live has changed.

This scene shows the character of the Young Faculty Award, or YFA. From the name alone, it may sound like a scholarship program that adds research money for new university professors. The way DARPA has run this award for 20 years is more tightly designed than that. It is less a program that hands over money and stops there than a channel that draws young researchers into the ecosystem of defense research.

First, we need to ask why this award targets young people. A researcher's interests harden early in a career. The problem that will become a life's field is usually set within a few years after the doctorate. DARPA aims at that period. Instead of bringing in established masters who already have names, it chooses to catch people who will run for the next 20 or 30 years at an early stage and turn their direction together. It reaches in before the seed has decided which way to grow.

DARPA's official material released in June 2026 states that YFA has supported more than 500 rising early-career researchers at over 60 universities and research institutions in the United States over the past 20 years. Judged by numbers alone, it is a sizable research funding program. The old aim of the program is clear. It draws young faculty at universities, or peer researchers at nonprofit science and technology research institutions, into sponsored research that builds the capabilities of future defense systems. It is a design that turns the eyes of the next generation of scientists, engineers, and mathematicians toward the U.S. Department of Defense and national security problems early in their careers.

A researcher selected for YFA does not receive only a research bank account. The researcher sits close to DARPA program managers, existing research teams already carrying out DARPA projects, and military users who will eventually use the technology. They receive mentoring and gain lines of contact with industry and defense. In that process, they refine their research ideas again, not in the language of pure scholarship but in the context of what national security needs. That is where the program's inner meaning lies. The demand is not to write one good paper, but to frame the problem again in the language of defense. Even for the same research, the direction that wins recognition at a conference and the direction that works on a battlefield can diverge. YFA shines its light toward the latter at that fork.

One striking rule in this program is the military site visit. A YFA principal investigator must visit a military site at least once to better understand the actual needs and environment of the Department of Defense. It is a device that shows the field to someone who has held a pen. The scene described earlier comes from this rule. It means that the researcher should experience at least once the place where theory works and the place where theory collapses. One visit is not some grand training program. Still, there are times when that one visit is enough to change the baseline inside a researcher's mind. It tries to plant the habit of designing while imagining not the paper reviewer, but the soldier who will actually carry the equipment.

The range of research topics is wide as well. From machine learning and many-body physics, to the meeting point of broadband transmitters and antennas, to models that examine the spread of infectious disease across multiple scales, the materials alone list 15 topic areas. The focus is on advanced fields meant to lift microsystems, biotechnology, and defense science. Rather than pour everything into one field, the program spreads seeds across several fields that may matter on future battlefields. Since no one can know in advance which seed will bear fruit, it sets a wide range. That breadth itself is a strategy. It presses down the risk of betting big in one direction and missing, by making smaller bets in several directions.

Within the 20-year path, there is also a device one level higher. It is the Director's Fellowship. DARPA selects a small number of YFA awardees who have shown outstanding results against their initial project goals and gives them this fellowship. On top of the basic two years of research, it adds up to \$500,000 for third-year research funding. The message is to build on early success, refine the technology, and fit it to the Department of Defense's needs. It gives those who have done well one more chance, while tying that chance clearly to the direction of defense. It does not pour in a large sum from the start, but adds money later after seeing results. This order, with risk placed up front and reward placed later, resembles the way DARPA works in many programs.

In the 2024 YFA list, seven people received this fellowship. A glance at their research titles shows the interests of future defense technology. Ultrathin nanophotonic thermal emitters for adaptive high-power heat rejection, gas-surface scattering dynamics in very low Earth orbit, and end-to-end modeling and mitigation of fault injection attacks. How should heat be discarded? How should air resistance be handled for objects moving in low orbit? How should physical attacks aimed at hardware be blocked? These real problems sit behind unfamiliar laboratory titles. The titles alone may look unrelated, but the heat of high-power weapons, air resistance for low-orbit satellites, and hardware security are walls that next-generation weapons systems will hit.

DARPA has recently adjusted this program. The 2026 YFA announcement added a new item. It allows up to \$350,000 more to support U.S. graduate students who participate in a role that helps a YFA project. DARPA said this measure aligns with its interest in expanding a pool of people who can receive security clearances and actually be deployed in future programs. The idea is not to stop at supporting one young researcher, but to embrace the graduate students growing beside that

researcher as a reservoir of security talent. It has tied the work of building technology and the work of raising people who can handle that technology to the same budget line. The calculation is not just to make one weapon, but to prepare in advance the people who will handle that weapon in the next generation.

The size of the money is also worth noting. The Department of Defense's fiscal year 2026 budget materials show that YFA had a little over \$17 million assigned under basic research in fiscal year 2024. Compared with the budgets of massive weapons programs, it is not a large amount. Given the nature of a program that divides early-stage seed funding among many rising researchers, the scale does not feel strange. It spreads small amounts widely, then pours more later only where sprouts appear. The face of this program is not the size of the budget, but the order in which the budget flows.

Most of what I have written so far rests on material DARPA released itself. Numbers such as 500 people over 20 years, 60 institutions, and a \$500,000 fellowship can be verified. The problem comes next. I was not able to gather enough independent material this time to judge from the outside which weapons and systems this program has actually led to, or which projects spent money and ended without results. Defense research sometimes buries results quietly and does not readily reveal failures. So it is more accurate to call YFA a well-designed program than to fix it as a successful one. Design and results are different matters. The fact that the blueprint of a program is excellent and the fact that fruit actually grew according to that blueprint require different evidence.

For Korean readers, the question this program raises is not the size of the budget. Many countries give money to young researchers. The point worth watching is the way direction is attached to money. A mandatory field visit, a seat across from users, one more conditional investment for those who do well, and even the graduate students beside them counted as part of the talent pool. A country that spreads research funds and a country that digs a defense waterway into those funds may look similar on the surface, but they differ inside. How far should we follow this connection, and where should we stop? Separating what deserves admiration from what would be dangerous to copy as is. That is the homework left before moving to the next section.

Deep Thoughts Industry Day and Researcher Entry

On the morning of February 6, 2026, people begin to gather one by one in front of 675 North Randolph Street in Arlington, Virginia. The time is 7:45 a.m. The event has not yet started, but attendees are already standing in line, waiting for identity checks. The mobile phones in their hands cannot cross this threshold. Nor can recorders or cameras enter. Small electronic devices must be left at the storage area by the entrance. The cafe inside the building opens at 7 a.m. Over a cup of coffee, researchers meeting for the first time exchange business cards. This is the opening scene of DARPA's Deep Thoughts Industry Day.

The office in charge of this event is DARPA's Tactical Technology Office. The program lead is Kevin Sloan. The program's name is Deep Thoughts. The technology at issue is small autonomous unmanned underwater vehicles. The aim is to build small vessels that move through the sea on their own without a person aboard. Building this kind of technology requires more than the strength of one defense contractor. People who write software, people who handle batteries, and people who study navigation underwater must gather in one place. Industry Day is the place that first calls these people into one room.

Here we should pause to look at how DARPA runs a program. DARPA is not an agency that scatters research funds and waits for good results. A program manager decides on the problem he or she wants to solve, puts that problem before the world, and gathers people who may be able to answer it. Industry Day is the first button. It is a place to disclose the problem, see who is interested, and induce people to form teams. If we take apart the structure of the day, we can see how DARPA filters researchers and calls them inside.

For a researcher to enter this program, the first step is to obtain one document. The notice first appeared on January 20, 2026. Nine days later, on January 29, DARPA issued Amendment 01. The document number is DARPA-SN-26-17. The changed point was how to receive the technical appendix. This appendix is controlled unclassified information, or CUI. It is material at an intermediate level, not open to anyone but not top secret either. At first, the notice said only that guidance would be given later through the procurement system. The amendment changed this, directing organizations capable of handling CUI to request the document directly through a designated link. The deadline for this request was 4 p.m. Eastern Time on February 4, 2026. Registration for the event itself had to be completed on the dedicated website by 4 p.m. on February 2, two days earlier. Only those who registered in advance could pass through the door. This small revision, made nine days after the first notice, redefined from the start whose hands could receive the document.

The first barrier to researcher entry appears in this scene. It is nationality. This program deals with information controlled under the International Traffic in Arms Regulations, or ITAR. For that reason, attendance is limited to U.S. citizens representing U.S. companies. Foreign national researchers and overseas institutions cannot enter this room. A Korean researcher would have to stop at this point. The inner seats of U.S. defense research are divided first by passport color, no matter how strong the ability may be.

The event starts at 8:30 a.m. and ends at 3:30 p.m. The place is the DARPA Conference Center. The day's schedule is divided into two parts. In the morning, the overall program overview and Technical Area 1 are explained at the CUI level. Those who enter this session must agree to comply with Department of Defense Instruction 5200.48 and National Institute of Standards and Technology Special Publication 800-171. If security officer Eric Bafour asks, they must present a government-issued photo ID. In the afternoon, the classification level rises by one step. Technical

Area 2 is handled at the top secret level. To sit in that room, a person must already hold a security clearance, and the facility security officer at the person's organization must send a visit request to DARPA through the Defense Information System for Security. The system code is designated as DDAAUS4. Fax and encrypted email are also available as alternative channels. With one door between them, some of the people who sat together in the morning cannot enter the afternoon room.

Even after setting up these layered doors, DARPA beckons toward unfamiliar names. The notice strongly encourages participation by nontraditional proposers. This includes small businesses, universities and research institutions, and new companies contracting with the government for the first time. It means this is not a gathering only for giant defense contractors. DARPA lists five goals for this day. To inform industry, academia, and government of the program's vision and goals. To let people whose interests and abilities match meet one another. To explain to new entrants what principles the program follows and what interim goals it will pass through. To help organizations with needed expertise and facilities form teams. To collect questions from the site and later return them in a question-and-answer document. If we look quietly at these five lines, they gather in one direction. DARPA keeps the threshold high, but tries to make the range of people who can cross that threshold as wide as possible.

It is also worth noting that the event does not end as a one-way presentation. During lunch and in the afternoon, separate time is set aside for attendees to discuss cooperation with one another. It is a place where a software-strong startup and an established company that understands hardware can fill each other's gaps and form a team. After the event, contact information for those who agree is distributed as a list. The point is to let them continue conversations they could not finish that day. For a researcher, the real benefit is a one-on-one meeting with the program lead. Individual meetings with Kevin Sloan do not take place on the crowded event day, but continue online on Monday, February 9, and Tuesday, February 10, after the weekend. To get such a slot, a person must apply separately to the designated email address before the registration deadline. The structure brings the person who posed the problem and the person who may answer it face to face, not during a crowded day but in front of a quiet screen.

There is a point to mark. DARPA attached disclaimers at the end of the notice. Attendance at Industry Day is not a requirement for submitting a later proposal. The event itself is not a formal request for proposals or request for information. Nor will the government reimburse the costs of participation. It may look like an impressive invitation, but legally it is no more than a preliminary meeting for sharing information. This attitude connects with the way we should read it. It would be difficult to look at one Industry Day and say DARPA has thrown the doors wide open. Open doors and closed doors sit side by side in the same notice.

When this day is read together with the previous section, the picture becomes clearer. If the Young Faculty Award is a device that keeps the door open for young professors over a long time, Industry Day is a device that calls people inside when a particular program opens its door. Security is

suffocatingly tight, yet the range of invitation is unexpectedly wide. The point where these two forces collide is how DARPA filters researchers. The thresholds of nationality and clearance remain, but for those who can cross them, DARPA does not reject a small company or someone doing government work for the first time.

If Korea were to imagine a similar system, this single day shows both what is worth admiring and what is hard to carry over as is. The method of having one program manager put forward a problem and draw in unfamiliar startups and universities to build teams is worth learning from. The ITAR wall that cuts eligibility by nationality is, inside the United States, a device for protecting domestic industry, but from outside it is a wall that cannot be crossed. When discussing a Korean model, what may matter is not the impressive stage, but the fine design that writes the rules for opening the door and the rules for closing it in the same document.

The Threshold DARPAConnect Tries to Lower

Imagine a young assistant professor at a university. In hand is a technology made through nights in the laboratory. Three papers have been published, and it has won an award at a conference. One day, the professor feels the urge to test this technology against defense research. At that moment, a wall appears. No one explains how to contact DARPA, what papers must be submitted, or what a contract even is. Companies such as Lockheed Martin or Boeing have departments dedicated to this procedure. The assistant professor has no such department. All that exists is one good idea.

DARPAConnect is the passage opened for this assistant professor.

Start with the name. DARPAConnect is an online guidance portal run by DARPA. It targets people who are trying to set foot in defense research for the first time, not existing defense contractors. The targets are university researchers, startups that have just been founded, and small and medium-sized businesses that have never done business with the government. The idea is to place a map and signposts in their hands in advance so they do not get lost in the unfamiliar world called DARPA. Underneath lies the recognition that leaving a door open is not enough. For someone who does not know where the door is or what happens after entering it, an open door is no different from a closed one.

Let us look specifically at what it gives them. The Industry Day document for the Deep Thoughts program, released in late January 2026, introduces three free courses offered by DARPAConnect. They are on demand, so people do not have to connect at a fixed time and can watch them whenever they want.

The first course explains DARPA's award system and requests for proposals. The Department of Defense's contracting methods differ greatly from transactions between ordinary companies. It helps people encountering unfamiliar systems for the first time, such as other transaction agreements or Small Business Innovation Research funding mentioned in earlier chapters, understand the frame. The second course explains how to make good use of Proposers Day. DARPA almost always holds

a Proposers Day event before starting a new program. This course teaches people in advance how to introduce their technology and what kinds of companies they should team with. The third course gives guidance on getting a proposal accepted. The performance metrics DARPA demands are severe by academic standards. It helps a researcher who has only written papers orient a document toward national security needs.

The three courses share one point. They all target the threshold of procedure and information. They are guidance for people who failed to enter not because they lacked technology, but because they did not know how to enter. This distinction matters more than it may seem. The threshold of ability and the threshold of information are different problems. If the threshold of ability is lowered, research quality may fall. If the threshold of information is lowered, quality does not fall. Instead, it brings onto the stage people who had good technology but turned away because they did not know the path. DARPAConnect touches this second threshold.

Lowering the threshold does not end with online courses. DARPA's program notices themselves gesture in the same direction. Look again at the Deep Thoughts Industry Day document just mentioned. This official notice strongly encourages participation by nontraditional proposers. The phrase nontraditional proposer is not mere rhetoric. The document spells out the range clearly. It includes small businesses, academic and research institutions, and new companies entering into government contracts for the first time. The assistant professor imagined earlier falls exactly within this range. A window that once faced only giant defense contractors has turned toward university laboratories and garage startups through one sentence in a notice.

We should stop for a moment here. When a notice calls new companies, the reverse meaning is that it had not been calling them before. The defense research market has long been occupied by a small number of large companies. There are reasons. The regulations are thick, verification is long, and a culture that does not readily accept failure has taken root. In that environment, only organizations with dedicated personnel able to handle the procedure remain. DARPAConnect and the nontraditional proposer clause can be read as attempts to shake that hardened board by hand.

There is another route toward academia. In June 2026, DARPA announced new Director's Fellows while noting that the Young Faculty Award had turned 20. The Young Faculty Award, examined in the previous section of this chapter, is an old ladder that draws early-career researchers into defense research. Twenty years means this ladder is not a passing experiment but a settled institution. If DARPAConnect's online guidance hands a map to new entrants, the Young Faculty Award provides first funding to researchers who actually enter holding that map. The two operate on different layers and point in the same direction. One shows the way at the door, and the other takes a hand inside the door.

Other kinds of doors are open as well. DARPA's Multi-domain X Office held a Spark Tank Pitch Day in June 2026. The name is playful. The event, reminiscent of a startup audition program, puts technologies that can be used by both civilians and the military on stage under the name of dual-use

design. The target fields include electronics, hardware, materials, microsystems, and robotics. Participants stand on stage, give short presentations, and receive interest on the spot. It is a structure that lets them knock on the door with a few minutes of presentation instead of a thick proposal. It can be read as an attempt to make the form of entry itself lighter. For people who could not even start under the weight of paperwork, it opened another route of a few minutes before a microphone.

This is the picture the documents show. Now comes the point where we must ask honestly.

The materials in this section clearly show that DARPA has opened doors. There are free courses, notices calling new companies, a Young Faculty Award that has passed 20 years, and a pitch day. Opening a door and knowing how many people actually walked through it, won contracts, and transferred technology to units are different questions. The former is confirmed by these documents. The latter is not confirmed by these documents alone. This section does not have outside statistics showing how many people took the courses, how many of them passed proposals, or how much the success rate rose compared with the past. So I will not fill the size of the results with guesswork.

There is a claim that it takes an average of 14 years to field a U.S. weapon. I have treated this figure in my earlier writing as background for explaining the rigidity of the defense acquisition system. It can serve as a reference line for judging how thick the wall was that DARPAConnect seeks to cross. But I will not set this number up as independent evidence for this manuscript. Background belongs in the place of background, and I do not want to mix what is confirmed with what is not. I hold to this principle not because I am timid. If DARPA is portrayed only as a success myth, the design wisdom we actually need to learn gets buried under applause.

So the judgment of this section is cautious. DARPAConnect is an attempt to redesign the entry road into defense research. It fixed its target as new entrants, explained procedures through courses, attached invitations through the language of official notices, and even prepared a presentation stage through pitch days. The direction is clear. The size of the results remains open.

The point Korea should watch when imagining a Korean-style research and development window lies here as well. What is worth learning is the design that opens the door. It is the sense to separate the threshold of information from the threshold of ability and lower only the side that can be lowered. Explaining procedures through online guidance, writing the invited targets clearly in the notice, and reducing the weight of paperwork through a presentation stage are methods worth trying. One more habit must be kept before copying them. Before boasting that the door has been opened, count whether people actually entered through that door. If the number of entrants, the share that led to contracts, and the list of technologies transferred to units are not quietly counted afterward, the window remains a publicity phrase and the board stays hardened. We can admire DARPA's skill in opening doors, but we must also bring with it the habit of verifying the results inside those doors ourselves. Without that eye for verification, the window we build may end as another closed door with only an impressive name.

Chapter 18. The Starting Point for a Korean DARPA

A Door Opened by One Signature in Brussels

On July 17, 2025, a quiet signing ceremony took place at the headquarters of the European Commission in Brussels, Belgium. Camera flashes went off a few times, a brief report passed through the news cycle, and by that evening other stories had pushed it aside. The country that placed its name on the document was the Republic of Korea. It was the procedure for joining Horizon Europe, Europe's vast research program, as an associated country. It was the first case of its kind for an Asian country.

It may seem abrupt for a book about DARPA to end with Europe and a signing ceremony in Brussels. Open the weekly briefing published on August 21, 2025, by the Korea-EU Research Centre, or KEREC, and the scene looks different. That signing ceremony was a window into how the name DARPA crossed to policy documents on the other side of the world and lodged there. One organization created by the United States became a model written into Europe's budget design, and Korea had just stepped onto that board. If we are going to talk about a Korean DARPA, this is the honest starting point. I use the word honest because I want to hold down, from the beginning, the impatience that jumps straight from envy to imitation.

According to the KEREC briefing, Korea became able to take part immediately in Pillar 2 of Horizon Europe as an associated country from January 1, 2025, under transitional measures. Pillar 2 is the axis that deals with global challenges and industrial competitiveness. Over the next three years, Korea's contribution will be 22.5 million euros, roughly 33 billion Korean won. If we look only at the amount, it is modest in the scale of a national budget. The door that money opens is not modest. We need to know what lies inside that door before we can see the real weight of the number.

Europe Changes Direction

Why did Europe open the door now? If we read the KEREC briefing slowly, it becomes clear that Europe's research landscape is changing direction as a whole. In the past, climate change and pure basic science stood at the center of Europe's research budget. Now security and defense have moved to the front. This change first appears in slogans.

Poland, which held the rotating presidency in the first half of 2025, put the slogan "Security, Europe!" at the forefront. It declared that it would treat defense, the economy, food, and cybersecurity as one bundle, and said it would expand the use of artificial intelligence. Denmark, which took over in the second half, also put security independence and competitiveness first under the banner "A Strong Europe in a Changing World." Even when the country changed, the center of gravity in the slogans pointed to the same place. When slogans change in Europe, money moves behind them.

The flow of money can be checked in numbers. The Commission presented a proposal for the next Multiannual Financial Framework, worth 2 trillion euros. Inside it sits the newly created European Competitiveness Fund. Its scale is 409 billion euros, roughly 600 trillion Korean won. Among the four areas into which this fund will pour money, "defense and space" appears plainly. Alongside clean transition and decarbonization, digital transition, and health and biotechnology, defense was named as a core investment target. This means Europe has started to connect the purpose of technology development directly to defense and the capacity to survive. Put plainly for Korean readers, Europe long hesitated to fund military research directly. The caution of a continent that had lived through two world wars was written into its budget structure. That caution is now collapsing.

The Moment DARPA Appears in European Documents

The protagonist of this book reveals its name inside European documents. The Commission said it is changing the European Innovation Council, or EIC, by taking the U.S. Defense Advanced Research Projects Agency, DARPA, as a model so that it can focus support on defense and dual-use startups. It also declared that in the next Horizon Europe, dual-use projects would be supported by default across all fields.

Let me pause to explain the term dual-use. Dual-use refers to technology developed for civilian commercial purposes that can move directly into military and security purposes as well. AI, unmanned drones, and quantum computing are typical examples. Technology for delivery drones can move into reconnaissance drones, and logistics optimization AI can move into target-processing AI. Europe, which had funded only purely civilian research, is erasing that line by its own hand.

The weight of saying that DARPA is the model is not light. [Inference] It reads as a willingness to fund technologies that carry high risk of failure but, if they succeed, can overturn the battlefield. European research, long tied to complex bureaucracy and academic output, appears to be declaring that it will move like a pipeline combining American-style venture capital and rapid acquisition. This is not something the KERK briefing says directly. It is my interpretation of that sentence. I mark it so that the facts in the source and my judgment do not get mixed.

The case of Israel shows how sensitive a weapon dual-use money can become. In July 2025, as the humanitarian crisis deepened because of the food blockade in Gaza, the Commission proposed restricting Israeli startups from taking part in the EIC Accelerator. The KERK report assessed this measure as a targeted sanction aimed at fields where Israeli technology had a high possibility of military and security dual-use. The hand that gives money also becomes the hand that cuts it off. Dual-use funding has begun to move beyond technical support and become a lever of diplomacy and military power. Korea has one point to keep in mind. Entering the door of European research funding also means being tied to Europe's political judgments.

A Budget Doubled, a Fight for People

Supporting all these currents is the next Horizon Europe, the program called FP10 in the documents. The Commission proposed raising this budget, which will run from 2028 to 2034, to 175 billion euros, roughly 260 trillion Korean won, about twice its current size. It is organized around four axes: excellent science, competitiveness and society, innovation, and the European Research Area. The second axis, which touches the real-world deployment of technology, works together with the 409 billion euro Competitiveness Fund mentioned earlier.

Growing the money was not the whole story. Commissioner Zaharieva made clear that before increasing the budget, she would first pursue reform to remove wasteful procedures and run the system with leaner operations. President von der Leyen also confirmed that FP10 would keep its independence while being firmly connected to the Competitiveness Fund. They are holding two goals at once: expanding the money while trimming the fat. Anyone who knows Korea's research and development administration will find the scene familiar. A budget grows, but procedures stay the same, so money cannot flow at its own speed. That concern looks similar in Brussels and in Seoul.

The fight to capture people is taking place in parallel. The Commission released a package called "Choose Europe for Science." Its core is to give early-career researchers financial stability and draw them to Europe. The 22.5 million euro Marie Curie pilot ties support for researchers for two to three years to an obligation by the employing institution to provide two more years of employment and a path to a permanent position. It uses 制度 to hold onto the problem of young researchers losing their place when a three-year contract ends. France opened "Choose France for Science," centered on space and digital fields, while Germany is running the "1,000 Heads" program to bring in 1,000 international talents, including American scholars. The European Research Council has announced a super grant from 2026 that can concentrate up to 7 million euros on one person over seven years. WIDERA, which aims to close regional gaps, receives 29 million euros. It is an all-out money fight for talent.

Korea has walked into the middle of this vast current. As the first Asian associated country, Korea can now reach directly into the funding pools of the 409 billion euro Competitiveness Fund and FP10 Pillar 2's 175 billion euros. Korea and the EU are widening the stride of cyber cooperation by completing negotiations on a digital trade agreement. KERCA builds bridges by summarizing the core points of cluster work programs so that Korean researchers can connect more easily with this European capital. At the moment Europe imitates DARPA, Korea has gained a place on that board. That single signature in Brussels was more than a modest contribution.

What AI Does on the Battlefield

While Europe was turning toward defense, AI was already changing the board on the battlefield. The report "U.S. Policy Trends and Implications for Securing AI-Cyber Superiority," published in 2026 by the Software Policy & Research Institute, or SPRI, traces that reality in close detail.

This report should be read with care. The operation names and numbers it contains are statements compiled and conveyed by SPRI, not facts that we have confirmed directly in the field. From the start, I promised to read planning documents as planning documents and reported accounts as reported accounts. I keep that promise here as well. Please read the scenes below within the frame of "the report says this."

The SPRI report shows the moment when AI and cyber converge through several scenes. It says that in the 2022 Russia-Ukraine war, satellite communications disruption and AI precision strikes appeared, and that in the 2024 Iran-Israel conflict, AI became a deterrent force that held battlefield initiative. It also reports a case from one U.S.-Israeli operation in which Palantir's targeting platform processed more than 1,000 targets within 24 hours. A case of air defense intercepting more than 99 percent of launched missiles and drones appears alongside it. We cannot verify the truth of these numbers ourselves. The direction the report is trying to show is clear. AI has now become both spear and shield.

DARPA's path explains how this change grew in cyberspace. In 2014, DARPA opened the Cyber Grand Challenge and began an experiment in which machines found and fixed code vulnerabilities on their own. Work once examined one by one by human hackers was handed to machines. In the recently completed AI Cyber Challenge, or AlxCC, a cyber reasoning system created by Team Atlanta automatically found vulnerabilities in large-scale open-source code and produced patches at the low average cost of 152 dollars. Machines fixed code that humans could not fully inspect.

The same technology flows to the other side. The report notes that an AI model released by Anthropic in April 2026 showed advanced capability in vulnerability analysis and attack path generation, pointing to AI's two faces as both defensive tool and offensive weapon. The hand that fixes code and the hand that breaks into code come from the same technology. This duality is the heavy burden carried by defense cyber.

Eyes That Changed with Each Administration

U.S. cyber strategy has changed its view of AI each time the administration changed. The SPRI report divides that transition into four periods.

During the Obama period, the focus was on improving federal infrastructure and defending intellectual property. AI was treated as a future engine of innovation, but still mainly at the level of blocking technology theft by adversaries. Trump's first term raised the idea of preemptive defense under the banner of "peace through strength" and began to recognize AI as a technical foundation for cyber defense capability. During the Biden period, resilience and corporate responsibility came to the front, with demands for "security by design" and "responsible AI." AI was encouraged as a defensive tool, but also seen as a risk factor that made systems more complex.

The national cyber strategy released by Trump's second administration in March 2026 redefined AI from a passive defense measure into a means of offensive advantage. With "technological

superiority" and "technological dominance" at the front, it fixed "agentic AI," which can carry out preemptive detection and disruption operations on its own, as a core execution tool for cyber operations. Agentic AI means AI that can judge and continue acting on its own without a human issuing commands each time. AI, which had remained in the seat of defense, was pushed to the front line of attack.

The policy framework appeared in the "AI Action Plan" announced by the United States in July 2025. The plan divides into two paths. One is "AI for cybersecurity." This path places AI into defense and offense as a working tool to raise response speed. The other is "cybersecurity for AI." Since AI systems and data themselves become targets of hacking and contamination, this path protects that stack of technology. Sharpening the spear and protecting the hand that holds it were placed side by side. For this, the United States gathers university talent for hackathons, builds high-security data centers that handle raw intelligence, and creates an "AI Information Sharing and Analysis Center" to share threat information among critical infrastructure sectors.

One DARPA program fits the policy exactly. [Inference] The line of protecting the AI technology stack appears to overlap with DARPA's SABER program. SABER is an attempt to block adversarial attacks in which an enemy subtly disturbs AI inputs and to maintain a sustainable "AI red team" inside the Department of Defense. An adversarial attack is the kind of attack that mixes tiny noise into an image that looks fine to human eyes, causing AI to mistake a tank for a bus. The sentences of the action plan have gained a body through a DARPA program, but this connection is my interpretation based on the SPRi report.

The Shadow of a Strategy Chasing Superiority

A strategy that chases superiority carries a shadow. The SPRi report points to three contradictions the United States must resolve.

One is the clash between loosening regulation for innovation and demanding security at the same time. If companies are given freer hands, speed rises, but if heavy security requirements are placed on top, that speed slows again. Speed and safety pull against each other. Another is the security dilemma in which pushing superiority stimulates the other side's fear and produces an arms race. If cutting-edge AI is put forward as an offensive weapon, anxiety grows in competitors such as China and Russia, and they in turn expand AI cyber threats, creating a vicious cycle. The last is friction with allies. The United States seeks to lead multilateral norms, yet also tries to pass the responsibility and cost of cyber defense to allies. A demand that says, let us make norms together but share the costs separately, cracks alliance cohesion more as time passes.

These three contradictions quietly tell us one thing. No matter how outstanding the technology DARPA creates may be, it can creak badly at the policy level. Technological success is not the same as strategic success. The promise not to portray DARPA only as a success myth is kept in passages like this.

The warning sound is also loud. In a March 2026 report, the RAND Corporation warned of a "cyber surprise," in which an adversary armed with artificial general intelligence launches a large-scale cyberattack at an unexpected moment. By contrast, the report notes that Korea's "AI Action Plan," released in February 2026, was assessed as disappointing in strategic depth because it mostly listed tasks to be done. Europe changed direction with money, and the United States pushed weaponization through doctrine, while Korea's document still remains at the stage of listing items. Korea's starting point appears when we take this contrast seriously.

Reading Earlier Books and New Sources Separately

Before moving to the Korean story, I will divide the grounds into two layers. I use the discussion of Korean defense AI in my earlier book, "The AI Defense Revolution," only as background. This means I will not set that book's account up as an independent basis for new factual claims. Why is this distinction needed? If I pull in an earlier book by the same author as if it were evidence, unverified judgments can disguise themselves as verified facts. So I place the earlier book only as background for recalling Korea's context, and I tie new facts to the source_id in this source ledger.

As background, the current Korean situation described by the earlier book is as follows. Facing a "troop cliff" caused by population decline, Korea is pursuing "Defense Innovation 4.0" and aiming to build an intelligent three-axis system and kill web. The Agency for Defense Development, ADD, is at the center, and its Defense AI Center grew in January 2025 into the Defense Artificial Intelligence Technology Research Institute, with a budget of 15 billion Korean won and roughly 1,000 people. Hanwha Systems, LIG Nex1, Korea Aerospace Industries, and Hyundai Rotem combine from the top down to build weapons systems. To stress it again, this account is the diagnosis offered by the earlier book, and I do not count it as an independently confirmed factual claim.

The painful point identified by that book was not technology but institutions. It diagnosed that the average time to introduce a weapon in Korea is 14 years. That speed does not fit AI software, which can become outdated in one or two years. By the time a newly purchased AI system is deployed, the technology may already be seven or eight generations old. A bureaucratic culture that does not tolerate failure and R&D budgets crowded into short-term projects also hold things back. So the book presented as tasks a rapid acquisition track suited to software, a high-risk "Korean DARPA," and a talent pipeline connected by unconventional treatment. It identified the problem in theory.

An Honest Confession About Sources Filled with Success Stories

The new sources show how that theory actually works in the United States. Let me stop for a moment and state something honestly. Most of the sources we hold are DARPA program introductions and transition success cases. We have almost no material that externally verifies DARPA's failures and waste.

This is not a small problem. Anyone who knows this organization can guess that among the many programs DARPA has pursued, only some moved into the field, and far more quietly folded. The

materials left in our hands lean toward surviving success stories. When reading the transition success stories below, it would be good to remember that these may be scenes selected for success alone. We can learn the success formula, but countless eliminations hidden by that formula sit inside the silence of the sources. If we pretend not to see that silence, the discussion of a Korean DARPA begins by seeing only half the picture.

With that premise in place, let us look at transition cases. Jarjet Technologies is a company that makes signal conversion chiplets for ultra-high-speed communication networks. These are small semiconductor pieces that move between analog and digital signals, and in the fight over radio spectrum, the performance of this component leads directly to superiority. The company began with 4.5 million dollars in SBIR basic funding from DARPA and won deployment contracts worth 150 million dollars from places such as Lockheed Martin. Seed funding of 4.5 million dollars grew into a contract more than thirty times larger.

SimVentions is a company that built a modeling and simulation platform. It started with 1.5 million dollars in funding and grew into a 14 million dollar Phase III contract with artificial intelligence added. The United States built this bridge from small seed funding to large contracts through SBIR and Other Transaction authority. SBIR is the Small Business Innovation Research program, and Other Transaction authority is a flexible contracting method that bypasses traditional procurement procedures. Where the earlier book cried out that a rapid acquisition track was needed, the United States had already built a physical bridge and was sending vehicles across it.

Reading Planning Documents as Planning Documents

There are also plans that shake the foundations of hardware and logistics. These should be read as plans, not achievements. If we miss this distinction, our judgment goes badly astray.

DARPA's O-Circuit is a plan to culture 2 million neurons, remove silicon semiconductors, and build a biological processing device that computes for 24 hours using the low power of 10 milliwatt-hours per day. It is an idea close to a laboratory concept: growing living nerve cells and making them compute. Fleetwood is a presentation aimed at supply-chain independence, fully breaking down lignin from biomass waste abandoned on the battlefield and producing fuel and strategic materials on-site even under an enemy naval blockade. The idea is to extract and use fuel from wood and grass lying around the battlefield even when the enemy blocks sea lanes and cuts off supply.

Both programs are not accomplishments already achieved, but plans announced as things to be attempted. They are images within the range of FAQ documents and industry briefing materials. If we read planning documents as achievements, exaggerations arise, such as "the United States is already running computers with living brains." That exaggeration soon leads to impatience: "we must do that right now too." The promise of this book is to keep the discipline of reading plans as plans and achievements as achievements. Under the same discipline, programs such as RAPIID in battlefield medicine and SMART for strategic materials are useful as signals because they build field

deployment and regulatory approval into the plan from the beginning, but I will not accept them as completed achievements.

Then What Is Korea's Task?

When the institutional diagnosis in the earlier book is placed over the reality shown by the new sources, the tasks Korea must now build become a little clearer. The four recommendations below are all my interpretations based on the sources. I will keep this mark to the end so that fact and judgment do not get mixed.

[Inference] The first path is that Korea cannot remain only with physical strike systems. The intelligent three-axis system described in the earlier book carries weight toward physical force such as missiles and robots. As the SPRi report warns, if an adversary uses artificial general intelligence to launch a "cyber surprise," Korea's command-and-control system itself can stop. No matter how precise a missile may be, it cannot be fired if the command system is paralyzed. This leads to the recommendation that the Defense Artificial Intelligence Technology Research Institute should grow, at the same strategic level as physical weapons, both defenses that prevent data contamination of friendly AI models, as in America's SABER, and offensive capabilities that disturb an enemy's neural networks.

[Inference] The second path is to build a real bridge that breaks the 14-year acquisition period. This means moving the Jarier and SimVentions formula onto Korean soil. Instead of staying in a top-down structure that concentrates budgets only on ADD and large defense companies, Korea should spread seed funding widely to ventures and university labs, attach large-company capital to the technologies that succeed, and create by law a transition highway that carries them directly into operational weapons. The United States has made this bridge through 制度. Korea still has it only as a slogan. Bridges are not built with words. They are built with law and budget.

[Inference] The third path is that Korea must not just pass by Europe's open door. Korea's newly gained status as an associated country in Horizon Europe is an admission ticket into the 409 billion euro Competitiveness Fund and the Pillar 2 funding pool. If Korean defense companies and AI startups form consortia with European universities and research institutes and win dual-use projects, a path opens for Korean weapons systems to seep naturally into European standards and civilian markets. It is a chance to grasp a third axis called Europe, beyond the domestic market and the framework of the Korea-U.S. alliance. As we saw in the Israel case, we must also remember that European money comes with European politics.

[Inference] The fourth path is to prepare independence in the foundations that run AI, that is, power and materials. The AI discussion in the earlier book leaned toward algorithms, data, and platforms. Plans such as O-Circuit and Fleetwood seek to change the power and logistics that run AI itself. If we think of a battlefield short on batteries and electricity, Korea also needs to allocate bold funding to low-power computation chips and technologies that can produce military supplies from

domestic waste when rare earth and fuel supplies are cut off. Here too, we must remember that this is a direction learned from foreign programs still at the planning stage. Preparation begins when we do not mistake someone else's plan for our own achievement.

Start with Envy, But Do Not Stop There

When the four paths are laid out, I worry that the phrase Korean DARPA could harden into a slogan for copying the United States as it is. It is true that DARPA is enviable. Even Europe has written the name into its own budget documents. The scenes in which America's seed-funding system grew small companies into giant contracts are worth learning from.

The attitude kept from the first chapter was not to applaud only success stories. The sources in our hands lean toward DARPA's bright side, and we lack enough grounds to verify the dark side. Folded programs, wasted budgets, and countless experiments that never moved into the field are hidden behind the silence of the sources. In that condition, the phrase "let us build a DARPA too" can easily become a slogan thrown after seeing only half the picture. There is a reason we must ask first what should not be copied as is, before asking what to copy.

Let us return to the signing ceremony in Brussels. What Korea gained that day was an admission ticket to European funding, but also a mirror. Europe is weighing which parts of DARPA to choose for its own circumstances and which parts to discard. A culture that endures failure, a system that spreads seed funding widely, an attitude that honestly separates plans from achievements. Among these, what Korea truly needs to transplant may not be the name of an organization, but the habits beneath it. The starting point for a Korean DARPA may not lie in drawing an organization chart, but in honestly counting what we do not know.

What Korea should learn from this scene is less an American-style label than the work of seeing in one line, before changing the signboard of a research institute, how failed projects are recorded in the language of the next budget review, how many years a program manager can stake personal judgment, which direction a researcher is rewarded to run between papers and delivery, whether military users come in early and fix the problem together, who takes responsibility until the fine sentences in a proposal reach the actual test range, procurement process, and the hands of soldiers, whether a technology that begins to look like an achievement survives inside which procedure of which unit, where small and midsize companies meet the next stage of funding, test ranges, and buyers after receiving their first contract, whether university researchers gain time inside the system to learn the language of field users, and whether records remain that make the next attempt more accurate than reports that hide failure.

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.



PayPal

<https://paypal.me/kimkjkj>

Scan the QR code or open the link above.