

How to ensure readiness of next-generation defense technology: Solving the issue of massive heat



Sponsored by



High-performance systems generate tremendous amounts of heat that can compromise missions and make high-tech systems less effective in tactical environments. But there is a solution to this modern-age challenge.

When the 2024 National Defense Authorization Act (NDAA) highlights the importance of a technology, it's usually an enabler of the highest order like artificial intelligence or quantum computing. Rarely, though, does it dive down to the microscopic level as it did when the 2024 NDAA House bill identified "aligned carbon nanotube technology and the predictable system interfaces it enables" as a key catalyst for Department of Defense mission critical systems and also cited Chinese competition and investment in carbon nanotubes.

The key word here is "aligned," as non-aligned carbon nanotubes have been around for two decades, finding applications from electron microscopes to drug delivery in biomedicine.

When carbon nanotubes are aligned, however, they become one of the most efficient conductors of heat ever developed – more so even than diamonds. And because the Joint All Domain Command and Control (JADC2) architecture will encompass all five warfighting domains, consuming power and needing to dissipate heat for an enormous number of systems, it's paramount that JADC2 architects focus on thermal system management.

Everything in the architecture – sensors at the edge, forward-deployed bases, dismounted troops, satellites in proliferated LEO constellations, naval vessels, aircraft,



Rafael Spears is general manager of Global Aerospace & Defense for Carbice.

ground mobile vehicles, and especially the computers to process and make sense of all the data being collected and turned into information to be shared throughout the battlespace – will require dissipation of more waste heat than ever before. It will put the "cooling" aspect of SWaP-C directly in the spotlight.

"Our JADC2 assets face a formidable adversary – heat – and it's a challenge whether you're on a satellite in space, a warfighter carrying a Manpack, or a UAV collecting sensor data," said Rafael Spears, general manager of Global Aerospace & Defense for [Carbice](#), a privately held tech-innovation company whose leadership team has received over \$60 million in funding for aligned carbon nanotube development from DARPA, Air Force and Army research labs, the National Science Foundation, and NASA.

"Power consumption and heat dissipation at the edge, especially for computational power at remote nodes and in forward deployed operations, significantly impact hardware performance, energy efficiency, and overall system reliability, as thermal management is a challenge across all domains," Spears noted. "Addressing this is integral to the future JADC2 architecture. Thermal challenges are not mere technical hurdles; they directly influence mission success and operational effectiveness."

Keeping edge devices cool

As JADC2 is envisioned for multi-domain operations, intensive computation and communication at the edge in forward-deployed facilities and other nodes in the architecture must take place in remote, diverse environments ranging from extreme cold to intense heat while maintaining optimal performance. This will introduce multiple challenges in the effort to provide efficient, reliable methods of managing the waste heat generated by these systems.

Hardware Reliability and Lifespan: Thermal stress and shock can significantly impact the performance, reliability, and longevity of electronic components. Semiconductor devices such as CPUs, GPUs, and AI chips are particularly vulnerable, and memory devices can be corrupted. Vibration is also common; many domain operations exasperate these effects and accelerate failure.

Performance Degradation: Unpredictable thermal systems management results in temperature extremes that affect processor clock speed, reducing computational power and increasing latency for real-time applications.

On the cover: U.S. Air Force Staff Sgt. John Engerraan, 3rd Air Support Operations Squadron sensors and effects team supervisor, and Airman 1st Class Cameron St Amand, 3 ASOS precision strike operator, survey the area during an arctic training in Delta Junction, Alaska, March 22, 2023. The training exercise pushed sensor and effects teams deep into simulated enemy territory and required SETs to be self-sustained for 72 hours. (U.S. Air Force photo by Airman 1st Class Ricardo Sandoval)

Sensor Accuracy and Calibration: Many JADC2 applications will rely on radar, lidar or EO/IR sensors having performance characteristics that will be at best altered and at worst substantially degraded by high temperatures or hot spots in the system.

In the context of military operations, power availability is also a crucial factor. Consider an Army Forward Operating Base (FOB) with a typical power constraint of around 250 kW. Much of this power is allocated for cooling electronic systems, which are essential for communication, surveillance, and other critical functions.

However, delivering this power poses a significant logistics burden for the DoD; the cost of logistics is measured not only in dollars but also, more importantly, in lives. A significant percentage of casualties in conflicts can be attributed to logistics operations, attacks on fuel convoys, supply routes, and vulnerable infrastructure. Studies undertaken during the Afghanistan and Iraq wars drew a correlation between fuel usage and casualties of soldiers and marines operating fuel and water convoys.


Said Spears: “That’s why energy efficiency matters. Better thermal efficiency saves both money and reduces risk to our service members.”

Computing at the Edge

One of the biggest challenges facing JADC2 architects will be supplying power to the processors that run the algorithms to make sense of the data collected, processed, and shared at the tactical edge. Almost all of that power goes into generating waste heat rather than flipping bits. Getting rid of that heat will be a challenge as formidable as providing the power in the first place.

Here’s why. The JADC2 vision will rely upon a significant amount of AI computation at the edge. DoD advocates of AI are still working on connecting a key element of their approach – the retraining of algorithms – with the enormous local compute requirements for that task.

Most significant AI applications occur today in the private sector, with ready access to gigantic server farms via large pipes. For military operations outside the continental U.S. (OCONUS), however, where the desired compute power cannot be displaced to a relatively secure, fiber-connected server farm, the need for very low latency and to operate through the assumed loss of connectivity in combat are critical requirements for JADC2 architects to consider.



U.S. Army graphic illustrating multi-domain operations.

Before 2012, compute requirements for computer vision algorithms were roughly following Moore's law – doubling every two years. In 2012 – training AlexNet, an early neural network, required 10^{18} Flops (Floating Point Operations Per Second – a measure of a computer's processing speed, specifically the rate at which it performs mathematical operations). Power required to support 10^{18} Flops can exceed 100 megawatts.

In a forward-deployed contested environment where a sensor or edge device has been disconnected from a server farm, retraining those edge sensors can require up to a Petaflop, 10^{15} Flops, of processing speed.

"Are JADC2 architects anticipating the need to retrain algorithms at the edge, because you must retrain or they will become useless over time?" said Spears. "The retraining of these edge devices will be the most challenging power use and thermal management application for DoD AI."

These concerns will only be amplified in a DoD Area of Responsibility like INDOPACOM, where OCONUS forces led by the U.S. Navy operate nothing but edge devices and nodes, whether they be ships, aircraft, radars, radios, processors, or storage devices. The complexity of multi-domain environments and the data they exchange will drive increased power consumption and cause them to run hotter than ever, and there won't be connections back to CONUS-based server farms, data centers, and cooling towers.

"You're going to be limited by the amount of power that you can either store or bring to remote locations, ships and submarines; that's half the equation," said Spears, "while the other half is, once you figure out how to provide the power you need, how are you going to predictably dissipate the waste heat that is generated?"

"This isn't just about removing heat; there are architectures that require a system to throttle down if it can't predict how it's going to move the heat throughout its range of operational modes and over its mission lifetime. When you can't predict how well the system is going to dissipate heat as a function of time and circumstance, you actually have to degrade its average performance to be sure of remaining within operational limits."

The cost of redundancy, the price of reliability

To combat this unpredictability, today's thermal system engineers build margin and redundancy into their designs. Spacecraft in particular are crafted with multiple layers of backup systems as insurance policies against on-orbit failure. Redundancy and margin aren't free, however; they come with increased mass and cost for most systems. Because of the unpredictability of how traditional thermal interface materials such as thermal grease, liquid silicone rubber, gap pads, or graphite degrade over time, thermal engineers must typically incorporate margins of 30 percent or more into their designs – what Spears calls the "belt and suspenders" approach. Traditional liquid and pad materials suffer from migration and plastic deformation and will always degrade due to basic physics – posing a high cost to the system.

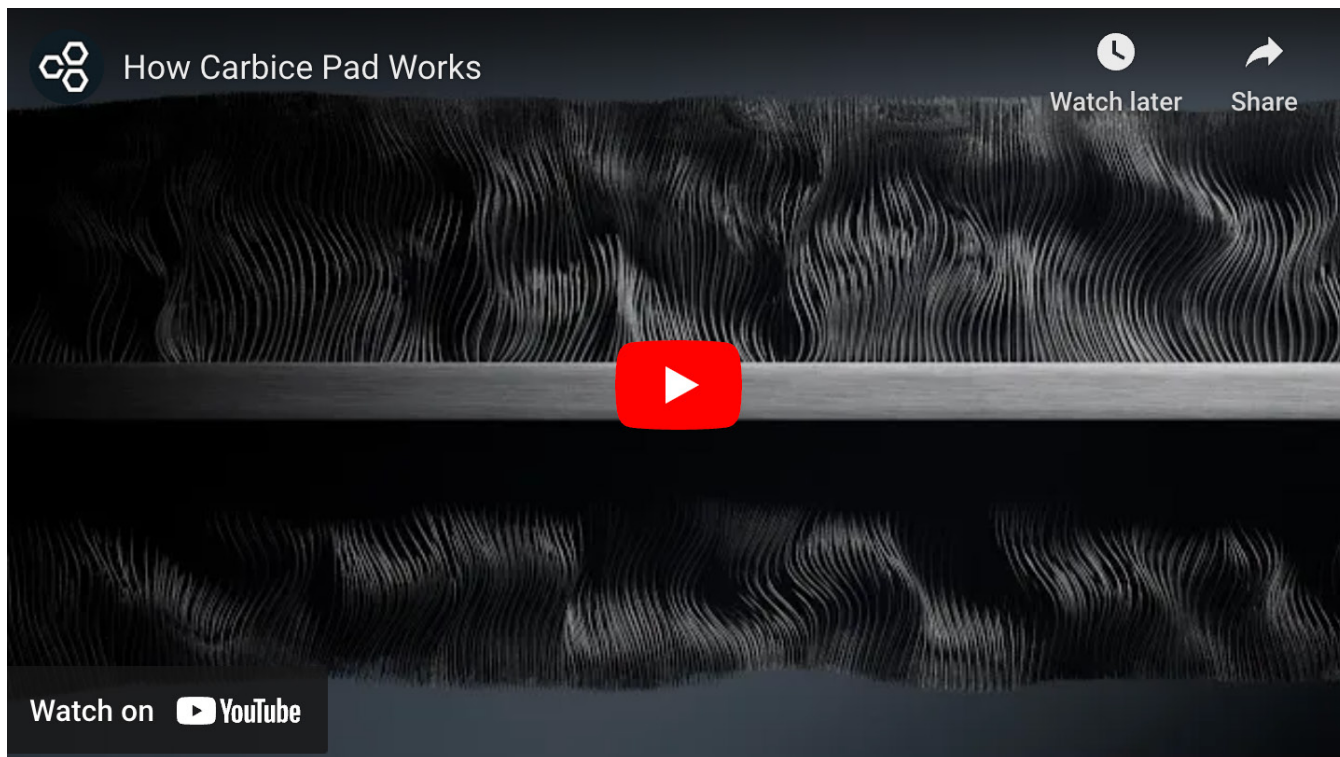
"We are disrupting that thought process with science and showing that you don't need all that margin if you have a predictable thermal system and materials that do not degrade or fail," said Spears. "When system architects use our product, the Carbice Pad, they can eliminate that 30-percent margin because our product provides predictable thermal performance for life. No spacecraft or mission-critical system will ever fly or operate with zero redundancy or design margin, but reducing it will provide increased system capability and reliable performance over the life of mission-critical systems."

The solution is aligned carbon nanotube interface technology

The Carbice thermal solution vertically aligns billions of carbon nanotubes to form a highway for heat to move from the hot side of systems to the cold. The result is unlike any solid, liquid, or phase-change material otherwise available.

Carbice nanotubes flex when compressed by temperature-induced expansion, mechanical stress from assembly, shock, or vibration. When the compressive load decreases, they rebound and maintain surface contact. This prevents air gaps that form with conventional thermal interface materials (TIMs); thermal performance is thus both more predictable and more consistent.

When carbon nanotubes are aligned they become one of the most efficient conductors of heat ever developed.



The Carbice Pad is simply manufactured with essentially three material inputs: waste carbon, aluminum foil, and hydrogen. With those elements, Carbice “grows” a vertically oriented, covalently anchored carbon nanotube forest on both sides of an aluminum substrate.

“What’s revolutionary about the accomplishment of Carbice CEO and Founder Dr. Baratunde Cola is that he not only created this carbon nanotube form factor, he created it at scale,” said Spears. “The Carbice product is produced in a 23,000 square-foot-manufacturing facility in Atlanta, GA. We have the current capacity to produce Carbice Pad for the entire U.S. space ecosystem. In developing the manufacturing process, Dr. Cola was very prescriptive and intentional about designing the Carbice thermal solution for mass production and scale. As a result, production and lead times are as fast as a few days.”

The Carbice Pad is a fully tailored, cost-neutral drop-in replacement thermal systems solution that does not require re-architecting current space, terrestrial, or computing systems.

The demand signal

The demand signal from the government is clear in roadmaps like the 2024 NDAA. JADC2 and all-domain operations linked by JADC2 will require industry to step forward with power consumption and heat dissipation solutions to enable the connected battlespace. Carbice has helped to address the heat dissipation challenge with the development of aligned carbon nanotube interface technology and the Carbice Pad.

“Carbice is not a project or a product development effort,” said Spears. “It’s TRL9 qualified, flying today, and is delivered from our AS9100D-certified manufacturing facility.”

The product is currently flying on orbit with numerous companies that support both commercial and national security space, including RTX Blue Canyon, Rocket Labs, Maxar, and Capella Space.