From: SCSP
To: John Wetzel
Subject: The Department of War Should Establish a Strategic Compute Reserve
Sent: 7/30/2026 1:30:14 PM

Forwarded this email? Subscribe here for more

SPECIAL COMPETITIVE STUDIES PROJECT

The Department of War Should Establish a Strategic Compute Reserve

SCSP
JUL 30

READ IN APP

Hi, I'm Ylli Bajraktari, President of the Special Competitive Studies Project. Today's newsletter is authored by Andreas Theodoropoulos, a Research Assistant on SCSP's AI Adoption team and a recent graduate of Tufts' Fletcher School of Law and Diplomacy, with contributions from James Ryseff, a Director on our AI Adoption team. Andreas argues that the recent conflict with Iran offered a preview of modern war — one in which the Department of War's need for classified compute surges precisely when new capacity is hardest to build, certify, and secure. The good news is that America already possesses the largest and most dynamic commercial data center base in the world. The task now is to ensure that this remarkable national asset can be mobilized when it matters most. To that end, Andreas proposes that the Department of War establish a Strategic Compute Reserve.

Andreas Theodoropoulos

Compute Is the New Currency of Military Power

Modern warfare is increasingly constrained not by the availability of weapons, platforms, or personnel, but by access to the computing power needed to connect and employ them at speed. The recent conflict with Iran made this vividly clear. During Operation Epic Fury, American forces leveraged AI tools such as the Maven Smart System to identify, prioritize, and strike 13,000 targets in the first 38 days of combat — an unprecedented degree of speed and scale in the history of warfare. This achievement echoes what the Ukrainian Armed Forces have demonstrated in their war against Russia, and what the Israel Defense Forces have shown in Gaza, Lebanon, and against Iran. Across the board, software and compute have become central to the way modern wars are fought — and won.

This is, at its core, an American advantage story. The United States leads the world in frontier AI, in cloud infrastructure, and in the data center buildout now reshaping the global economy. But that very dynamism means compute is a resource in fierce demand. Of the more than 35 gigawatts of data center capacity currently under construction in the United States, over 90 percent has already been pre-committed to a corporate customer. Companies like JPMorgan Chase now plan their compute needs five to ten years in advance, while firms such as Uber, Coinbase, and Walmart have had to ration compute expenditures as increasingly capable frontier models drive up costs. The lesson from the commercial sector is unmistakable: organizations can no longer assume compute will be available when they need it, in the quantities they need it. Those who prepare today will hold the advantage tomorrow.

The Canary in the Coal Mine

The compute constraints facing the commercial sector should sharpen the Department of War's focus, given its own rapidly growing appetite for data-driven military systems. The Pentagon has confirmed that Maven Smart System usage during Epic Fury surged month-to-month by 38 percent on unclassified workloads and 89 percent on classified workloads. By token volume, top daily usage surged by 4,425 percent, hitting a peak of 20 billion tokens a day.

Surging demand is only half the picture; the Department's current classified compute base is strikingly small by commercial standards. Consider U.S. Central Command, one of the military's most AI-forward organizations. As of early 2025, CENTCOM's classified network ran on just four NVIDIA H100 GPUs, with a partnership with the Air Force Research Laboratory set to add 24 more — a total the command's chief data officer described as more classified compute capability than anything else he was tracking in the Department. For perspective, a typical commercial server node today holds four to eight of those same chips, meaning one of America's premier combatant commands has been navigating the AI transition with roughly the classified computing power of a single commercial server rack. And a single mid-sized AI company, Inflection AI, assembled a cluster of some 22,000 H100s back in 2023 — hundreds of times the Department's most ambitious classified deployment. The Department itself recognizes the gap: its fiscal year 2027 budget request seeks nearly $30 billion to transition from

“scattered clusters” of GPUs to an integrated AI infrastructure portfolio. That candor is welcome — and it underscores both how far the Department has to travel and how decisive America’s commercial compute base will be in closing the distance.

As dramatic as these figures are, they would pale in comparison to the scale required to prevail in a major war against a peer competitor like China. A conflict in the Indo-Pacific would demand the analysis and fusion of far more sensor data, the prioritization of a far greater range of targets, and the coordination of operations and logistics across significantly more aircraft, warships, and personnel. These workloads would not be episodic. They would need to run continuously, at the speed of modern combat.

Here the Department faces a familiar dilemma — and history offers a proven answer. In peacetime, building and maintaining dedicated data centers that sit idle or underutilized is prohibitively expensive. In wartime, lead times for new facilities stretch between one and two years. But America has solved this problem before. For decades, civil reserve models have delivered critical resources and services in moments of need, coordinating supply and demand across military and civilian sectors alike. Just as the Civil Reserve Air Fleet and the Voluntary Intermodal Sealift Agreement allow the Department to surge transport capacity in a national crisis, the Department should now build the mechanisms to mobilize America’s massive commercial data center capacity in the event of war. Three lines of effort would put a Strategic Compute Reserve within reach.

Reevaluate the Requirements for Classified Compute

As the Federal Data Center Enhancement Act, which set minimum standards for federal data center on uptime, security, and energy use, lapses this September with no replacement in sight, the Administration has an opportunity to build a deliberate compute strategy for national security. This gap reflects a larger pattern: Department data infrastructure and cybersecurity policies have not fully kept pace with the shift to cloud and AI-driven computing. The Department’s leadership has long recognized that its cybersecurity requirements are excessively compliance-oriented rather than focused on delivering effective security outcomes for the warfighter. Many of these requirements pre-date the development of cloud computing, and they may no longer add value in the new operational paradigm.

The Department should reimagine its regulatory frameworks around risk-based accreditation, identifying which rules still make sense in a cloud computing environment and which are obsolete. In a full-scale war, the Department cannot wait months or years while additional classified compute resources are certified and inspected. The moment to reduce these unnecessary burdens is now, while there is still time to make deliberate reforms rather than deferring difficult decisions until the imminent approach of war forces the issue. By finally delivering on years of promised reforms, the Department can ensure that warfighters have the compute they need if called upon to “fight tonight.”

Pre-Build Security Into Commercial Data Centers

To prepare data centers for mobilization, the Department should work with cloud service providers to upgrade and maintain physical and cyber security environments, scaled to the sensitivity of the workload and the current capabilities of each facility. American industry is already well ahead here: Amazon Web Services, Microsoft, and Google have all developed the technology to virtually protect and separate commercial and sensitive military data below the Secret level, and the Department can work with providers to integrate this logical separation for classified workloads.

Enabling commercial data centers to run Secret and above workloads will require deeper collaboration and significant infrastructure investment — investment the Department can subsidize to incentivize hyperscalers to meet these stringent requirements. Financial packages could be structured along the lines of the Maritime Security Program, which provides financial backing to ship operators in return for enrollment in the sealift program. Some upgrades are straightforward: replacing magnetic strip key card systems with GSA-approved security locks, hardening air ducts larger than 96 square inches, and preinstalling physically separated SIPRNet wiring. Others — certifying classified supply chains and issuing security clearances to personnel — will take longer and demand significant lead times, which is precisely why the work should begin now.

Fortify Against Major Computing Security Threats

Finally, the Department must recognize that not every cybersecurity risk associated with commercial data centers can be remediated in advance. Facilities that operate as commercial or unclassified environments will contain hardware, networking equipment, or other components that cannot be fully trusted if repurposed for classified military workloads. For these cases, the Department should develop a rapid conversion model that accepts some equipment may need to be removed rather than converted. This approach allows the Department to distinguish between security requirements that can be pre-built into commercial data centers during peacetime and those that must be satisfied through wartime reconstitution before a formerly commercial facility connects to classified networks. The goal is not to pretend all commercial infrastructure can be made trusted in advance, but to build the plans, inventories, supply chains, workforce, and contracting mechanisms needed to transform that infrastructure into trusted military compute at speed.

The Work Ahead

In a world moving at the speed of data, slower queues and greater downtime translate directly into strategic disadvantage — and speed of preparation translates directly into strategic advantage. America enters this competition with assets no rival can match: the world's leading AI companies, the world's largest cloud infrastructure, and a proven tradition of harnessing commercial strength for national defense. What remains is to connect them. A Strategic Compute Reserve would do exactly that, ensuring that the extraordinary computing power America builds in peacetime stands ready to serve the

nation in war. It is the insurance policy the Department of War needs — and one the United States is uniquely positioned to create.

Subscribed

LIKE COMMENT RESTACK

© 2026 Special Competitive Studies Project
Crystal City, Virginia, 22202
Unsubscribe

Get the app Start writing