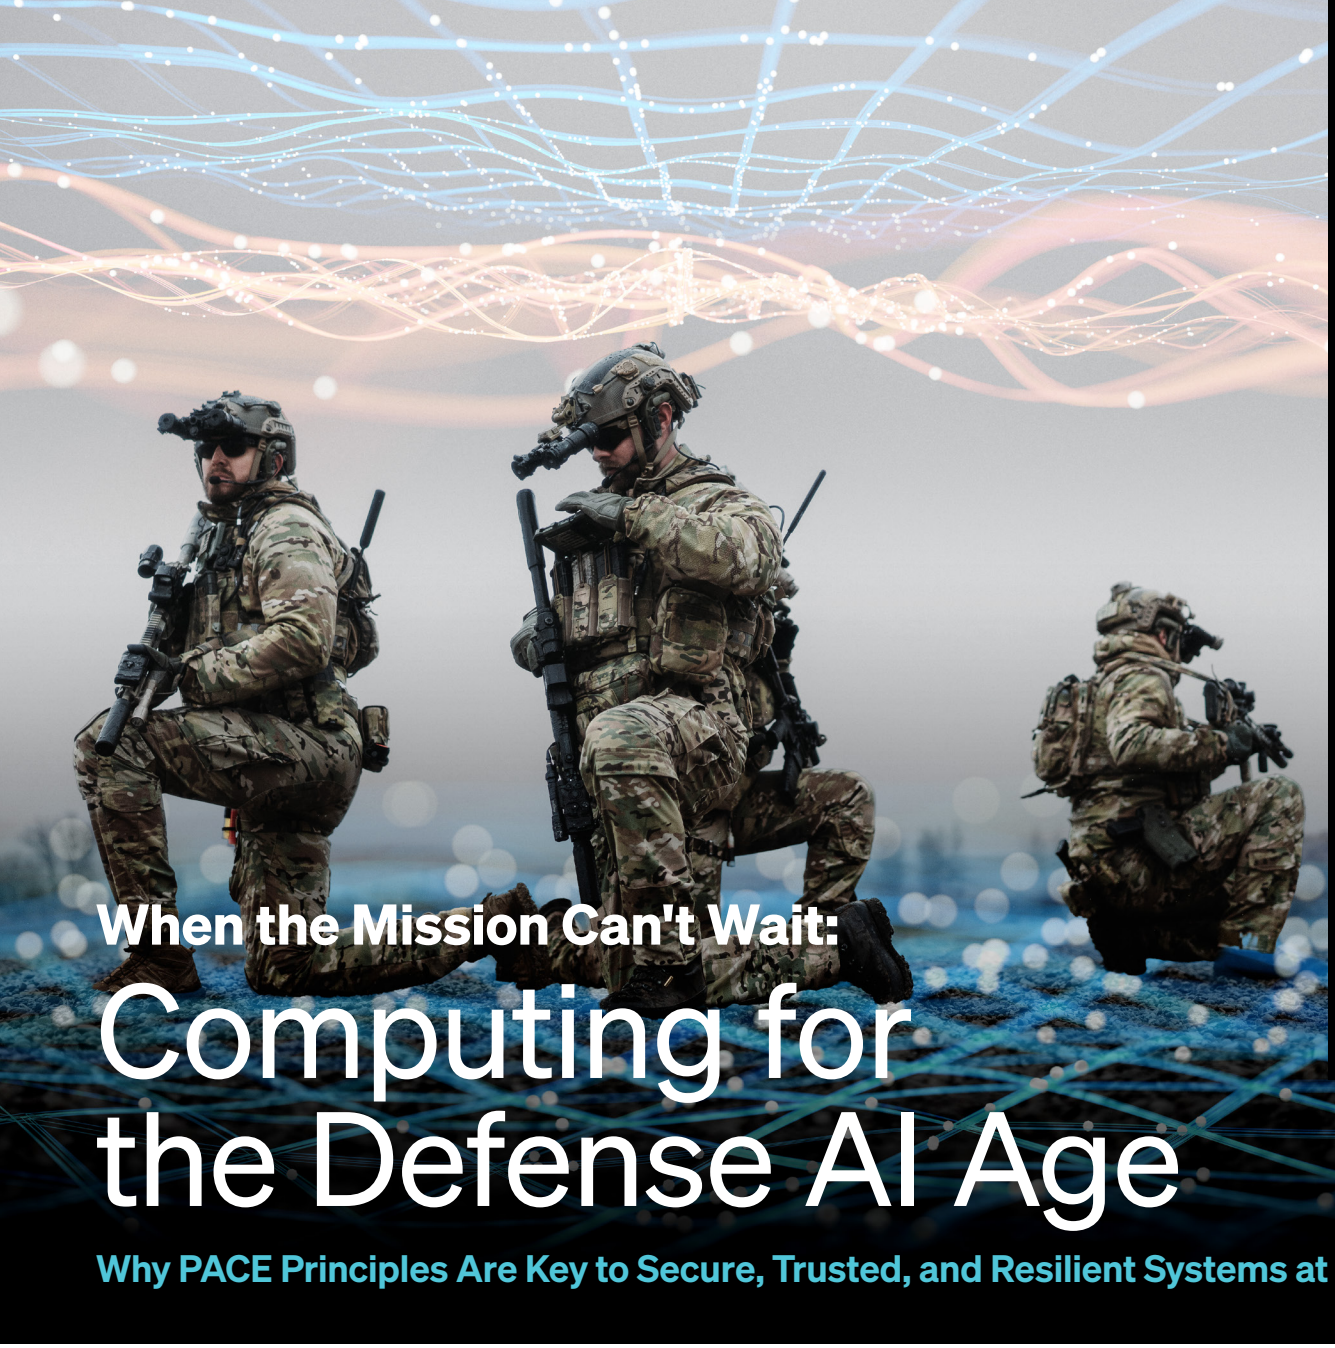

Booz Allen

Velocity

When the Mission Can't Wait:

Computing for the Defense AI Age

Why PACE Principles Are Key to Secure, Trusted, and Resilient Systems at Velocity

By Andrea Inserra

There's a moment in every mission planning cycle when the conversation shifts from what we intend to accomplish to how we keep accomplishing it when conditions turn against us. Throughout my career supporting defense missions, I've observed how good operators plan for that moment from the start. They don't treat disruption as an exception. They treat it as a design requirement.

Years ago, during a field engagement with one of our customer teams, I learned how special operations forces use a planning approach called PACE—Primary, Alternate, Contingency, Emergency. While it started as a communications tool, PACE can be understood more broadly as a flexible framework for building resilience. Primary and alternate options are actively maintained and updated as requirements shift, ensuring they stay ready for use. Contingency and emergency options remain in reserve, each governed by clear triggers that help keep operations moving even under stress.

This approach and redundancy reflect a simple truth: every critical pathway can eventually fail. The solution is to build multiple validated ways to achieve the objective—and to practice moving between them when pressure is high. Just as important, PACE isn't a strict step-by-step model; it is designed to balance performance with resilience across changing conditions.

As an engineer by training, I appreciate the logic of PACE as a design philosophy for defense computing in the AI age. AI has done to our systems what contested environments do to communications: compressed timelines, exposed brittle assumptions, and eliminated the margin for complacency that human-speed architectures quietly depended on. In my years overseeing complex defense systems, the underlying mission principles aren't new. What has changed is the speed, scale, and autonomy with which systems now operate.

That is why this issue of Velocity focuses on a single overarching challenge: how to build computing systems that remain secure, trusted, and resilient when AI is operating at AI speed under contested conditions.

What AI and Autonomy Have Actually Changed

The core disruption isn't simply that AI makes systems faster. It is that AI compresses the time available for humans to secure, validate, and recover those systems once they are operating at AI speed. Security models that depended on analyst review, trust models that assumed humans were the primary actors, and resilience models built around recovery after disruption are all under strain simultaneously.

AI-enabled autonomy—autonomous systems, agentic software, and AI-speed decision making—is accelerating both the threat and the operational tempo required to defend against it. The core mission hasn't changed: defense forces have always been built to operate amid friction, disruption, and active opposition. What's different now is the tempo and complexity of those pressures. The environment evolves faster, scales wider, and increasingly relies on autonomous systems that challenge our decision cycles in entirely new ways.

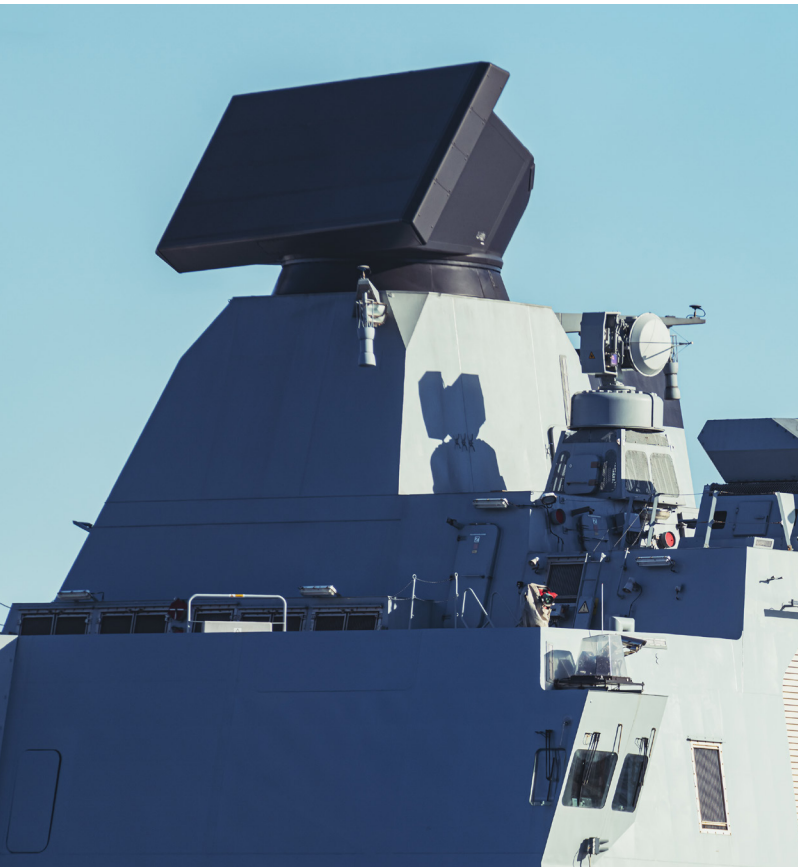

On the cyber threat side, automated reconnaissance, AI-assisted exploit generation, and AI-speed lateral movement have compressed timelines that once stretched across days into minutes. The detect-analyze-respond workflow that anchored security operations for two decades assumed analysts had time to work. In many scenarios, they no longer do.

Our survey of federal leaders found that 79% were extremely or very concerned about adversaries using AI to accelerate cyber-attacks against their agency. Furthermore, 84% reported high concern about adversarial AI compressing the window between vulnerability discovery and exploitation.

On the mission side, autonomous agents will increasingly execute complex tasks, interact with sensitive data, and coordinate with other agents—often without human review. The trust and authorization models built for a world where humans were the primary decision makers are straining under that shift. Only 28% of federal leaders surveyed reported high confidence in their ability to deploy agentic AI securely.

The result isn't simply a faster version of the old environment. It is a challenge to the assumptions underlying how we design secure, trusted, and resilient systems in the first place.

Secure: Matching the Adversary's Tempo

I've had this conversation with our government cyber leader Brad Medairy on multiple occasions, and his framing has stayed with me: "If defense can't operate at AI speed, it will be outpaced by those who can." That's not a warning about the future. It's a description of where we are.

What that means in practice differs by threat. Automated attacks—AI-generated exploits, AI-speed lateral movement, autonomous reconnaissance—compress attack cycles from weeks to minutes and demand equally automated defenses. But the harder problem may be the human-directed campaign: a deliberate, patient adversary who uses AI not to overwhelm but to persist—pre-positioning inside networks months or years before acting. Countering that threat requires more than automated defenses. It requires the cognitive countermeasures that only humans provide: recognizing adaptive behavior, reading intent, and disrupting campaigns that are designed to evade the signatures automated systems are trained to catch.

What strikes me most in our Velocity research on cyber autonomy and zero trust for operational technology (OT) is how clearly that second threat is playing out in OT environments—systems designed for isolation that are now fully networked. Adversaries exploit the gap between digital integration and the security controls that should have come with it, and they do it quietly, over time.

Closing that gap requires cyber autonomy: AI-supported detection and response that operates at the speed of the threat. It requires limiting the blast radius when a breach occurs—and it will occur. In PACE terms, this is the discipline of keeping your alternate path viable. A segmented, autonomy-enabled architecture doesn't just defend better; it ensures that when one path is compromised, the mission continues.

Trusted: Thinking Outside the Perimeter

Here's the harder conversation: the challenge to trust in AI-era systems isn't just technical. It's categorical. When autonomous agents are making real-time operational decisions, the old verification question—did this person have authorization?—gives way to something much harder: can we prove this system is behaving correctly, under all conditions, including the ones we didn't anticipate?

Answering that requires a fundamentally different mindset—one that moves from assumed trust to provable trust. The Velocity research points toward what that looks like in practice. As one example, automated reasoning can provide mathematical assurance that software and policy logic behave correctly under every condition, not just the scenarios a test team anticipated. Zero-knowledge proofs allow verification across organizational and classification boundaries without exposing the underlying data. And governing agentic AI—behavioral constraints, identity binding, continuous authorization, human checkpoints—is what makes autonomous systems governable rather than merely powerful.

In PACE terms, contingency mode requires knowing exactly what your systems will and won't do when primary conditions fail. In practice, that means commanders and operators must be able to verify not just who is accessing a system, but how autonomous systems will behave before those systems are trusted with mission-critical decisions. Trust built on assumption breaks precisely at that moment. Trust built on verifiable properties holds.

Resilient: Engineering to Survive

In my time both as an engineer and working with defense leaders on resilient architectures, one theme comes up again and again. Booz Allen CTO Bill Vass, whose work on distributed systems shapes the resilience thinking in this issue, frames it simply: "Everything fails. The only question is how well you survive it." In defense, surviving means the mission continues—under degraded bandwidth, compromised nodes, disrupted communications, or unexpected autonomous behavior.

That kind of resilience is an architectural commitment made before systems are built. It means distributing workloads across cloud, tactical edge, and on-premises enclaves with the ability to shift among them without catastrophic capability loss. It means deep observability—the instrumentation that keeps operators oriented and gives AI what it needs to diagnose and recover. And it

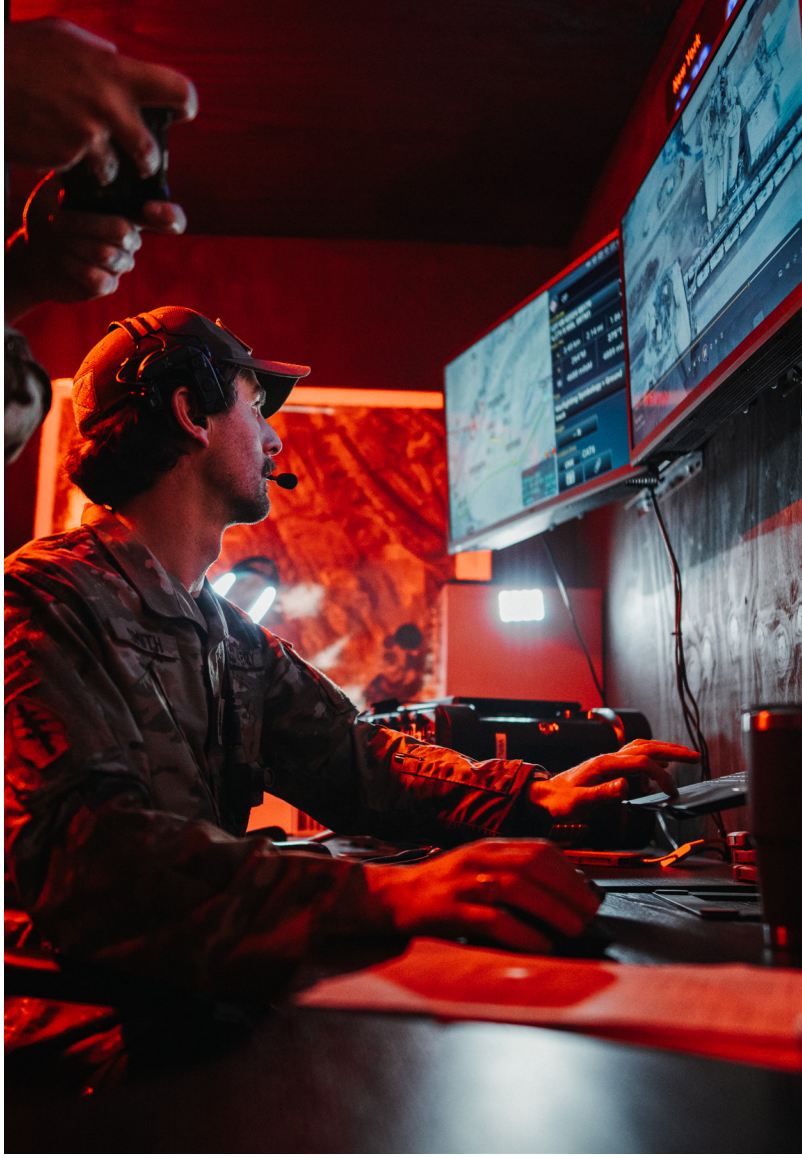

means designing graceful degradation explicitly: knowing what emergency mode looks like before you're in it.

The emergency path in PACE is not a retreat. It's proof the mission was engineered to survive.

Why These Three Are One Problem

Security, trust, and resilience cannot be pursued on parallel tracks. AI eliminates that luxury—and PACE illustrates why. You don't design primary and forget about emergency. The framework only holds because every path is validated, every transition is practiced, and the whole system is built for conditions that will actually occur.

A system that is secure but not resilient fails the moment conditions degrade. A system that is resilient but not provably trustworthy cannot be governed when autonomous behavior matters most. A system that earns trust but can't match adversarial tempo will be exploited before it can respond. All three, integrated from the start, is the only architecture that holds under pressure.

The Work Ahead

Defense has always found a way when the environment changed faster than the existing playbook. The operators who developed PACE didn't do it because communications are reliable. They did it because communications are often the thing that fails—and the mission still has to get done.

In the AI era, mission advantage will belong not to the systems that operate fastest under ideal conditions, but to the ones engineered to continue operating when conditions inevitably deteriorate.

PACE works because operators refuse to be surprised by failure. It's time our systems shared that refusal.

Andrea Inserra is president of Booz Allen's Global Defense business.

Booz Allen®

About Booz Allen

Booz Allen is an advanced technology company. We build commercial-grade products and solutions for America's most critical defense, civil, and national security priorities.

For more information, visit www.boozallen.com.

(NYSE: BAH)