


SelfAware Compute Introduces Predictive Software Optimization for Compute Efficiency

SelfAware's Optimization Engine helps software make better execution decisions before compute is spent in the wrong places

SCOTTSDALE, AZ, UNITED STATES, May 27, 2026 /EINPresswire.com/ -- As data centers expand across the country and communities raise harder questions about energy use, water consumption, and infrastructure strain, [SelfAware Compute](#) is introducing its approach to predictive software optimization: improving how code runs before organizations default to more compute power.



SelfAware Compute asks what happens when software can understand how it actually runs.

“

Every industry now has parts of the business that require serious compute. The question is how much more useful work we can get from the systems already in place.”

Kevin Howard, Chief Technologist, SelfAware Compute

Demand for compute is rising quickly, but the pressure is not limited to AI. Compute has powered critical industries for decades, hidden inside logistics systems, agricultural models, weather tools, factory software, medical systems, financial infrastructure, energy operations, and legacy enterprise code. As AI draws more attention to compute demand, SelfAware is focused on a broader issue: helping organizations make software execution more efficient before assuming the only answer is more compute.

“Every industry now has parts of the business that require serious compute,” said Kevin Howard, chief technologist, SelfAware Compute. “The question is not just how much more infrastructure can be built. The question is how

much more useful work we can get from the software and systems already in place.”

Most software is deployed as if one execution strategy should fit every run. SelfAware changes that. Its Optimization Engine breaks execution into analyzable pathways, giving software more

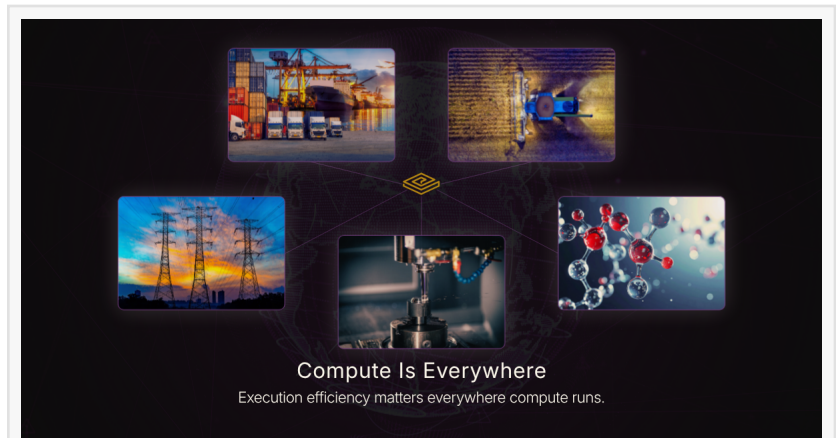
flexibility in how it runs. Instead of treating performance, memory use, and energy demand as results discovered after the fact, SelfAware models how software is likely to behave and uses that model to guide execution before compute is spent.

SelfAware's Optimization Engine analyzes code, identifies its possible execution pathways, predicts how those pathways behave under different inputs and hardware constraints, and selects an execution strategy based on the user's objective, whether that objective is runtime, memory behavior, energy use, performance, or a specific deployment requirement.

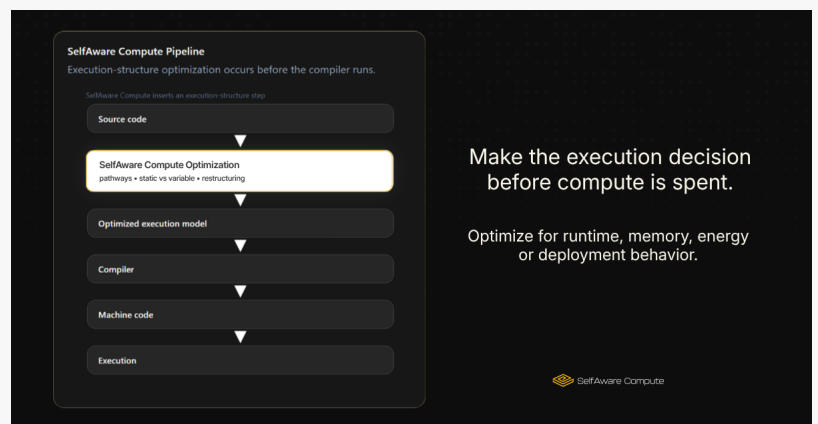
For software development teams, SelfAware is designed to improve execution without making another framework, architecture, parallelization pattern, programming model, or expensive hardware upgrade the starting point. The technology generates optimized execution as part of the run process, rather than asking engineers to hand-maintain separate original and optimized codebases. By working below the application workflow level and focusing on the execution structure of the code itself, SelfAware improves execution efficiency without forcing teams to rebuild their software stack around yet another tool or infrastructure cycle.

SelfAware Compute represents the next phase of Massively Parallel Technologies' work and reflects the company's sharper focus on [software execution pathways](#), prediction, and compute efficiency. The shift reflects a larger market reality: compute demand is increasing, data center expansion is becoming more controversial, and the physical footprint of digital infrastructure is becoming harder for communities and organizations to ignore.

Compute waste is not only a hardware problem. More servers, faster chips, and larger facilities increase capacity, but they do not necessarily make programs run more efficiently. If execution inefficiencies remain hidden, organizations can still spend power, memory, time, and infrastructure capacity in the wrong places. At data center scale, even small execution



Compute touches logistics, agriculture, weather, manufacturing, medicine, energy, and everyday business operations.



SelfAware's Optimization Engine helps software make better execution decisions before compute is spent in the wrong places.

improvements matter because the same inefficiencies repeat across large volumes of workloads.

“Most systems compensate for execution uncertainty with excess,” Howard said. “Extra hardware, extra memory, extra cooling, and extra budget all become safety margins. SelfAware is focused on reducing that uncertainty by helping software explain how it will behave before the cost is already committed.”

SelfAware is preparing multiple paths for making its Optimization Engine available, including enterprise engagements and service-based delivery models. Technical briefings and live demonstrations are available by request for organizations evaluating how predictive software optimization could apply to their own code, systems, and applications.

For enterprises, technical teams, and infrastructure operators, more predictable execution improves decisions about where work runs, how resources are allocated, and when additional capacity is truly required. The value is not limited to application teams. It applies anywhere compute is planned, scheduled, powered, cooled, or constrained.

SelfAware’s message is direct: compute demand will continue to grow, but advances in hardware alone will not solve the efficiency problem. Software execution itself must become self-aware.

To learn more or request a technical briefing, visit SelfAware Compute or contact Shannon Kendall at shannon.kendall@selfawarecompute.com.

About SelfAware Compute

SelfAware Compute is a software execution technology company focused on predictive optimization, execution-pathway analysis, and compute efficiency. SelfAware analyzes how software runs, models how inputs, hardware constraints, and runtime conditions affect behavior, and helps teams make better execution decisions using existing code and existing infrastructure where possible. SelfAware Compute represents the next phase of Massively Parallel Technologies’ work and is focused on software that finally understands itself.

Shannon Kendall
SelfAware Compute
shannon.kendall@selfawarecompute.com

Visit us on social media:

[LinkedIn](#)

[X](#)

[Bluesky](#)

[YouTube](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/915344477>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.