

Your Data Center is No Place for a Space Odyssey

The first movie I remember seeing in a theater was *2001: A Space Odyssey*. If you saw it, I am guessing that you remember it, too. At the core of the story is HAL, a sophisticated computer that controls everything on a space ship en route to Jupiter. The movie is ultimately a story of artificial intelligence gone awry.

When the astronauts realize that HAL has become dangerous due to a malfunction, they decide they need to turn HAL off. I still recall the chill I experienced when one of the astronauts issues the command, "Open the pod bay doors please, HAL." And HAL responds with, "I'm sorry, Dave. I'm afraid I can't do that."

Artificial Intelligence is Real Today, but not Perfect

Today, we are finally experiencing voice interaction with a computer that feels as sophisticated as what that movie depicted more than 50 years ago. But sometimes with unintended or unexpected consequences.

Artificial intelligence (AI) is great, except when it is not. My sister recently purchased a vehicle with collision avoidance technology built in. Surprisingly, it engaged the emergency stop procedure on a rural highway when no traffic was approaching. Fortunately, there was no vehicle following close behind or this safety feature might have actually *caused* an accident. (The dealer eventually accepted the return of the vehicle.)

Artificial Intelligence in Data Center Infrastructure Products

Artificial intelligence and machine learning technologies are being incorporated into data center infrastructure products. Some of these implementations are [delivering measurable value](#) to the customers who use these products. AI/ML enabled capabilities may include:

- AI/ML enabled by default... Yay!
- Cloud-based analytics...Yay!
- Proactive fault remediation... Yay!
- Recommendations... Yay!
- Totally autonomous operations... I'm not sure about that.

Examples of Artificial Intelligence and Machine Learning Done Right

- HPE [InfoSight](#) – all the “Yay!” items above. For example, HPE claims that with InfoSight, 86% of problems are predicted and automatically resolved before customers even realize there is an issue.
- HPE [Memory-Driven Flash](#) is now shipping for HPE 3PAR arrays. It is implemented as an 750 GB NVMe Intel Optane SSD add-in card that provides an extremely low-latency read cache. The read cache uses sophisticated caching algorithms to complete nearly all I/O operations in under 300 microseconds. Yet, system administrators can enable this cache per volume, giving humans the opportunity to specify which workloads are of the highest value to the business.
- Pivot3 [Dynamic QoS](#) provides policy-based quality of service management based on the business value of workloads. The system automatically applies a set of default policies, and dynamically enforces those policies. But administrators can change the policies and change which workloads are assigned to each policy on-

the-fly.

When evaluating the AI/ML capabilities of data center infrastructure products, enterprises should look for products that enable AI/ML by default, yet which humans can override based on site-specific priorities, preferably on a granular basis.

After all, when a critical line of business application is not getting the priority it deserves, the last thing you want to hear from your infrastructure is, "I'm sorry, Dave. I'm afraid I can't do that."