

Technology: How AI's Breakneck Pace Is Outrunning Enterprise Strategy and Your IT Refresh Cycle

Published: 19 February 2026 | Author: David Daoud

```
{ "channelId" : 20673, "language": "en-US", "commlId" : 663015, "displayMode" :  
"standalone", "height" : "auto" }
```

Transcript:

Anthropic - one of the leading AI companies in the world, backed by over \$30 billion in fresh funding at a \$380 billion valuation - just released its second major AI model in less than two weeks. Twelve days. Two flagship model releases. And the thing they said about it? They said the new mid-tier model - Sonnet 4.6 - now performs at a level that previously only their top-of-the-line model could reach. Let that sink in for a second. The capability ceiling they hit two months ago is now the midpoint. The premium is now the standard. And they'll do it again in another two weeks.

Now I know what some of you are thinking. "That's an AI company doing AI company things. What does that have to do with me managing software licenses, hardware assets, and end-of-life equipment?" And the answer is - everything. Because what's happening at companies like Anthropic, signals about the speed at which your environment is about to change, and whether your processes, your governance, and your business model are built to keep up with it.

Let's start with CIOs, because the pressure landing on your desk right now is unlike anything in recent memory - and I don't say that lightly.

You are managing a technology refresh cycle that was probably designed for a world where major capability shifts happened every 18 to 24 months. Enterprise software vendors would release a new version, you'd evaluate it, maybe pilot it, run it through procurement, and deploy it over the course of a year or two. That's the rhythm most IT organizations were built around. That rhythm is gone. It has been replaced by something closer to a sprint that



never ends.

When Anthropic releases two major models in twelve days, and when those models are being embedded into productivity tools, coding environments, design platforms, and knowledge work applications that your employees are either already using or actively asking to use - your evaluation cadence has to change. Your procurement policies have to change. Your software licensing assumptions absolutely have to change.

And here's the harder conversation nobody wants to have yet: some of the software you're currently paying for - tools your organization renewed last year, contracts you're locked into for another 18 months - may be delivering significantly less value than they were when you signed the deal. Not because the software got worse. But because AI just leapfrogged it. The market is already signaling this. The iShares Tech Software ETF is down more than 20% year to date. Investors are voting with their dollars that a meaningful chunk of traditional software value is at risk. As a CIO, you need to be asking your vendors the hard questions at renewal time. What does this tool do that an AI model can't replicate? And if the answer is unclear, that's your answer.

Now let me shift to the ITAM community, because you sit in a genuinely fascinating and somewhat precarious position right now.

On one hand, your core value proposition - helping organizations understand their assets, that is what they own, what they're paying for, and whether they're in compliance - has never been more important. Because here's what's actually happening on the ground in most enterprises: people are spinning up AI tools faster than procurement can track them. Someone in marketing is using one model. The engineering team is using another. Finance is experimenting with a third. And nobody has a complete picture of what's being used, what data is being fed into these systems, what the cost exposure looks like, or whether any of it complies with the organization's data governance policies.

That is an ITAM problem. A massive one. And the organizations that are coming to you for help managing their Oracle licenses or their Microsoft EA - they're going to need help managing their AI footprint too. The question is whether you're positioned to offer that. AI asset management is happening right now, and the window to establish yourself as the expert in that space is open today. It will not stay open forever.

But here's the uncomfortable flip side of that conversation. The AI models that are driving



this complexity are also capable of automating a significant portion of what ITAM has traditionally done manually. Think of such things as Software reconciliation. Entitlement mapping. Discovery and normalization. Contract analysis. These are tasks that have historically required skilled analysts doing painstaking, time-consuming work. AI agents are increasingly capable of doing versions of that work faster and at lower cost. So the question for every ITAM vendor and every ITAM professional is not “will AI affect my industry?” It already is. The question is whether you’re going to use it, or wait for someone else to use it against you.

Now let’s get to the ITAD side of this conversation, because in my opinion, this might be where the most immediate and tangible opportunity lives.

The pace of AI model development that we saw this week doesn’t happen in a vacuum. It happens because companies are pouring billions of dollars into AI infrastructure. And AI infrastructure means hardware. Specifically, it means GPUs. High-end, specialized, extraordinarily expensive GPUs - and the servers, the networking equipment, and the storage systems that surround them. And here’s the thing about hardware in a market that’s moving this fast: it gets old quickly. Not because it stops working. But because something better came out and the organizations running at the frontier need to upgrade to stay competitive.

That creates a disposition wave. And it’s not a small one. We’re talking about data centers full of hardware that was state-of-the-art 18 months ago and is now being rotated out in favor of the next generation. If you are in the ITAD sector that have the expertise to handle this equipment - proper data sanitization, responsible remarketing, environmental compliance - this is a significant and growing revenue opportunity. As you very well know, the secondary market for AI-capable hardware is genuinely hot right now. There are buyers. There is demand. And the volume of material coming out of hyperscalers and large enterprises is only going to increase as upgrade cycles shorten.

But - and this is important - not every ITAD operation is equipped to handle this category of assets. High-value GPU disposition requires different logistics, different security protocols, different remarketing channels, and different expertise than a standard PC refresh. If you’re in ITAD and you haven’t already built out that capability, this is the moment to do it, because the window between when this opportunity peaks and when it gets crowded out by competition is not going to be a long one.



So let me bring this back to the headline that started this talk. Two major AI model releases in twelve days. It sounds like a tech industry story. But what it actually is, is a clock. And that clock is telling every CIO, every ITAM professional, and every ITAD operator the same thing: the pace of change in your environment just went up another gear, and so you must treat this as a wake-up call to move your governance, services, and business models to meet this moment.

Subscribe to Compliance Standards' Sector Pulse Research

Contact us at 754.229.0095 or inquiries@Compliance-Standards.com

Our Offices: Boston, MA | 2361 Walnut Court, Pembroke Pines, FL 33026