IBM Newsroom

Open Source, After Mythos

By Rob Thomas | Senior Vice President, IBM Software and Chief Commercial Officer

ARMONK, N.Y., April 9, 2026 /PRNewswire/ -- There is a pattern we keep seeing in technology. When software moves from product to platform and then from platform to infrastructure, the rules change.

At the product stage, control can feel like an advantage. Closed systems move fast. They tightly manage the experience. They concentrate value inside a single company. That approach can work at the product stage. But once a technology becomes foundational, once other systems, institutions, and markets begin to rely on it, the bar shifts. At infrastructure scale, openness stops being ideological and starts being practical.

AI is crossing that threshold now.

Anthropic's limited preview of Claude Mythos brings that reality into sharper focus. The company says the model can discover and exploit software vulnerabilities at a level that few human experts can match, and it has launched a gated initiative, Project Glasswing, to put those capabilities into defenders' hands first. What matters most is not just the model itself, but what it signals: AI is no longer just a tool people experiment with. It is becoming embedded in how organizations secure systems, write code, make decisions, and create value.

At that point, the core question changes. The issue is no longer only what these models can do. It becomes how they are built, governed, inspected, and improved over time.

History is consistent on this point. As systems grow in importance and complexity, closed development becomes harder to defend. No single company can anticipate every failure mode, every adversarial use, or every operational requirement. Restricting access to powerful systems is an understandable instinct. It can look like caution. But at infrastructure scale, security improves more often through scrutiny than through concealment.

That is the enduring lesson of open source software.

Open source does not eliminate risk. It changes how risk is managed. It allows more researchers, developers, and defenders to examine systems, test assumptions, surface weaknesses, and harden code under real-world conditions. In security, visibility is not the enemy of resilience. It is often a prerequisite for it.

This matters even more in the age of AI. If frontier models are increasingly capable of finding vulnerabilities, writing exploits, and reshaping the security landscape, we should be cautious about concentrating understanding of those systems inside a small number of companies. Critical technologies tend to be safer when more people can inspect them, challenge them, and improve them.

This is also why open systems do not destroy value. They move it. One of the oldest misconceptions about open source is that it commoditizes innovation. In practice, it usually pushes competition up the stack. As common foundations mature, value shifts toward implementation, reliability, orchestration, trust, and domain expertise. The winners are not the companies that merely own the base layer, but those that know how to apply it best.

We have seen this before with operating systems, cloud infrastructure, and developer tooling. Open foundations expanded participation, accelerated improvement, and created larger markets on top. AI is likely to follow the same path. Even in enterprise technology, leaders increasingly view open source as strategically important, especially for infrastructure modernization and emerging capabilities like AI.

There is another reason openness matters. Who gets to participate shapes what gets built. Narrow access leads to narrow perspectives. Broad access enables more researchers, startups, governments, and institutions to influence how technology evolves and where it is applied. That does not just drive innovation. It builds legitimacy and adaptability.

The Mythos moment should give us pause, but not for the most obvious reason. Yes, the capabilities are striking. Yes, the risks are real. The deeper implication is structural. Once AI becomes critical infrastructure, opacity can no longer be the organizing principle for safety.

For decades, the most reliable model for secure software has been open foundations combined with serious governance, active maintenance, and broad scrutiny. As AI enters its infrastructure phase, that same logic increasingly applies to the models themselves. The more critical the technology, the stronger the case for openness.

If AI is becoming foundational, then openness is no longer a debate. It is a design requirement.

Media contact:
IBM Press Room
ibmpress@us.ibm.com

SOURCE IBM
