Mphasis
The Next Applied

Coverage Report

Mphasis urges CTOs to rethink legacy core as HFS research highlights ontology as a foundation for trusted agentic AI

March 23, 2026

Online

CXO Today

Mphasis urges CTOs to rethink legacy core as HFS research highlights ontology as a foundation for trusted agentic AI

Mphasis believes enterprise CTOs must rethink legacy core architectures as agentic AI adoption accelerates. This perspective is reinforced by HFS Research, which identifies ontology-based enterprise knowledge graphs as a critical foundation for safe and scalable AI. HFS Research notes that while agentic systems can accelerate execution, without an explicit semantic control layer, they risk propagating incorrect intent and scaling flaws embedded in legacy logic.

At the heart of the NeoIPTM platform is Mphasis OntosphereTM, an ontology-based intelligence layer within the broader Mphasis NeoIPTM suite that captures business rules, code paths, documentation, and operational artifacts into a structured, continuously updated knowledge graph. By defining what core enterprise concepts mean within a specific organisational context and how they interrelate, Ontosphere creates a persistent semantic layer that can govern and guide agentic systems safely.

Across financial services, insurance, and public sector organizations, modernization has typically been incremental. Many have adopted cloud infrastructure and digital interfaces while leaving core business logic largely unchanged. As agentic AI systems begin to ingest applications and automate decision flows, this embedded logic becomes mission-critical.

Mphasis cited measurable results from client engagements applying its ontology and knowledge graph framework. In a global insurance ITOps and observability programme, the client achieved:

- 67% accuracy in major incident prediction;
- 3-5 hours of early warning capability;
- A 50% reduction in mean time to detect, acknowledge, and resolve incidents.

“At Mphasis, we see AI as an architectural inflection point, not just a technology overlay. Enterprises are entering an AI era where the core can no longer be treated as untouchable. Layering intelligence on top of fragmented systems only scales complexity. The shift now is to make intelligence part of the architecture itself – and NeoIP is helping in doing so,” said Nitin Rakesh, Chief Executive Officer and Managing Director, Mphasis.

“Our research shows that ontologies and knowledge graphs are increasingly important for enterprises adopting agentic AI. Mphasis’ NeoIP platform reflects these principles, embedding semantic intelligence at the core of enterprise architecture to support continuous modernisation and safer AI deployment. The challenge isn’t the absence of intelligence, rather its fragmentation. Enterprises already have the intelligence embedded in code, workflows, and documents – but what they lack is a system of truth for intelligence. The task is to extract, structure, and make it reusable, so AI can drive transformation without repeating legacy mistakes,” said David Cushman, Executive Research Leader, HFS Research.

As agentic AI matures, enterprises will require far greater architectural discipline. Without embedded intelligence grounded in an enterprise context, AI remains merely artificial. The organizations that succeed will be those that can

codify their own institutional meaning and leverage it to enable continuous, secure, and well-governed transformation.

NeoIP is designed to do so to reduce repeated relearning cycles that can consume significant portions of traditional or legacy transformation effort. HFS Research describes this approach as a centralised meaning layer that enables enterprises to build and control AI agents safely, while supporting continuous modernisation rather than episodic transformation.

Global Security Magazine

Mphasis urges CTOs to rethink legacy core as HFS research highlights ontology as a foundation for trusted agentic AI

Mphasis believes enterprise CTOs must rethink legacy core architectures as agentic AI adoption accelerates. This perspective is reinforced by HFS Research, which identifies ontology-based enterprise knowledge graphs as a critical foundation for safe and scalable AI. HFS Research notes that while agentic systems can accelerate execution, without an explicit semantic control layer, they risk propagating incorrect intent and scaling flaws embedded in legacy logic.

At the heart of the NeoIPTM platform is Mphasis OntosphereTM, an ontology-based intelligence layer within the broader Mphasis NeoIPTM suite that captures business rules, code paths, documentation, and operational artifacts into a structured, continuously updated knowledge graph. By defining what core enterprise concepts mean within a specific organisational context and how they interrelate, Ontosphere creates a persistent semantic layer that can govern and guide agentic systems safely.

Across financial services, insurance, and public sector organisations, modernisation has typically been incremental. Many have adopted cloud infrastructure and digital interfaces while leaving core business logic largely unchanged. As agentic AI systems begin to ingest applications and automate decision flows, this embedded logic becomes mission-critical.

Mphasis cited measurable results from client engagements applying its ontology and knowledge graph framework. In a global insurance ITOps and observability programme, the client achieved:

- 67% accuracy in major incident prediction;
- 3-5 hours of early warning capability;
- A 50% reduction in mean time to detect, acknowledge, and resolve incidents.

“At Mphasis, we see AI as an architectural inflection point, not just a technology overlay. Enterprises are entering an AI era where the core can no longer be treated as untouchable. Layering intelligence on top of fragmented systems only scales complexity. The shift now is to make intelligence part of the architecture itself – and NeoIP is helping in doing so,” said Nitin Rakesh, Chief Executive Officer and Managing Director, Mphasis.

“Our research shows that ontologies and knowledge graphs are increasingly important for enterprises adopting agentic AI. Mphasis’ NeoIP platform reflects these principles, embedding semantic intelligence at the core of enterprise architecture to support continuous modernisation and safer AI deployment. The challenge isn’t the absence of intelligence, rather its fragmentation. Enterprises already have the intelligence embedded in code, workflows, and documents – but what they lack is a system of truth for intelligence. The task is to extract, structure, and make it reusable, so AI can drive transformation without repeating legacy mistakes,” said David Cushman, Executive Research Leader, HFS Research.

As agentic AI matures, enterprises will require far greater architectural discipline. Without embedded intelligence grounded in an enterprise context, AI remains merely artificial. The organisations that succeed will be those that

can codify their own institutional meaning and leverage it to enable continuous, secure, and well-governed transformation.

NeoIP is designed to do so to reduce repeated relearning cycles that can consume significant portions of traditional or legacy transformation effort. HFS Research describes this approach as a centralised meaning layer that enables enterprises to build and control AI agents safely, while supporting continuous modernisation rather than episodic transformation.

IT Brief UK

Mphasis urges CTOs to fix legacy cores for agentic AI

Mphasis has urged chief technology officers to revisit legacy core architectures as companies accelerate the adoption of agentic artificial intelligence, warning that automation can amplify flaws embedded in older systems.

HFS Research has pointed to ontology-based enterprise knowledge graphs as a foundation for "safe and scalable" AI. It warns that agentic systems can propagate incorrect intent when organisations lack an explicit semantic control layer across applications, processes and data.

The warning comes as many large organisations push AI into production while leaving core platforms largely intact. In financial services, insurance and the public sector, modernisation programmes often focus on cloud infrastructure and digital channels. Core business logic frequently remains unchanged, even as new automation layers depend on it.

Legacy exposure

Agentic AI refers to software agents that can interpret goals, make decisions and execute steps across systems with limited human intervention. In large enterprises, that can include agents reading documentation, tracing workflows, calling internal services and initiating operational actions.

This model increases the importance of underlying rules and dependencies. When logic is fragmented or undocumented, an agent can scale errors faster than a human operator. It can also entrench inconsistent definitions of key business concepts across functions, such as what constitutes a customer, a claim, a risk event or a service outage.

HFS Research argues that ontologies and knowledge graphs provide a practical control mechanism. An ontology formalises the meaning of concepts and the relationships between them. A knowledge graph links those concepts to real assets such as code, workflows, data sources and documents. Together, they create a shared semantic layer that guides how agents interpret information and what actions they take.

NeoIP focus

Mphasis has positioned its NeoIP platform around that approach. Its NeoIP suite includes Mphasis Ontosphere, which it describes as an ontology-based intelligence layer. Ontosphere captures business rules, code paths, documentation and operational artefacts in a structured knowledge graph that is continuously updated.

In practice, Ontosphere aims to standardise definitions across an organisation and map how they connect to systems and processes. This creates a reference layer above legacy platforms and newer applications, and a way to govern what AI agents can infer and execute across a complex technology estate.

Mphasis says it is seeing measurable operational outcomes in client work where it has applied this ontology and knowledge graph framework. In one example from a global insurance IT operations and observability programme, it reported 67% accuracy in major-incident prediction and three to five hours of early warning. It also reported a 50% reduction in mean time to detect, acknowledge and resolve incidents.

Such metrics matter for CTOs under pressure to industrialise AI quickly. Faster incident prediction and resolution can reduce downtime and operational risk, and improve service management where modern and legacy systems interact in unpredictable ways.

Architectural shift

Mphasis frames the current moment as an architectural shift rather than another technology rollout. It argues that adding AI on top of fragmented systems increases complexity and amplifies hidden weaknesses.

"At Mphasis, we see AI as an architectural inflection point, not just a technology overlay. Enterprises are entering an AI era where the core can no longer be treated as untouchable. Layering intelligence on top of fragmented systems only scales complexity. The shift now is to make intelligence part of the architecture itself - and NeoIP is helping in doing so," said Nitin Rakesh, Chief Executive Officer and Managing Director, Mphasis.

HFS Research makes a similar point about fragmented enterprise knowledge. It says organisations often already have the necessary logic and context across code, workflows and documents, but lack a single system that can express it consistently for humans and machines.

"Our research shows that ontologies and knowledge graphs are increasingly important for enterprises adopting agentic AI. Mphasis' NeoIP platform reflects these principles, embedding semantic intelligence at the core of enterprise architecture to support continuous modernisation and safer AI deployment. The challenge isn't the absence of intelligence, rather its fragmentation. Enterprises already have the intelligence embedded in code, workflows, and documents - but what they lack is a system of truth for intelligence. The task is to extract, structure, and make it reusable, so AI can drive transformation without repeating legacy mistakes," said David Cushman, Executive Research Leader, HFS Research.

For technology leaders, the debate extends beyond tools and platforms to governance, risk management and accountability as agents gain autonomy. A semantic layer that defines intent and meaning can also support auditability by providing an explicit model of how an organisation understands its own processes and terms.

Mphasis expects the need for "architectural discipline" to grow as agentic AI matures and more work moves from pilots to production. It also expects ontology-driven knowledge graphs to become a key building block for continuous modernisation programmes, particularly in organisations where legacy cores still run essential operations.

Online		
Date	**Publication/ Portal**	**Headline**
March 19, 2026	CXO Today	Mphasis urges CTOs to rethink legacy core as HFS research highlights ontology as a foundation for trusted agentic AI
March 19, 2026	Global Security Magazine	Mphasis urges CTOs to rethink legacy core as HFS research highlights ontology as a foundation for trusted agentic AI
March 19, 2026	IT Brief UK	Mphasis urges CTOs to fix legacy cores for agentic AI