

# BMC HELIX COMPOSITE APPROACH TO ARTIFICIAL INTELLIGENCE IN THE GENAI ERA



Several months ago, I wrote [a blog](#) outlining BMC's application of Generative AI (GenAI) technology through [BMC HelixGPT](#). Since then, GenAI has demonstrated its potential for creating diverse content (text, images, audio, video), computer code, configuration, meaningful conversations, and even entire novels already developed – likely authored with just a bit of prompt engineering.

Our mission at BMC is to provide actionable insights to operations teams. Any assistive AI technology targeting service and operations teams must gain trust, and the bar to clear while assisting operations is really high. These teams are overworked and under stress most of the time. Their attention span is limited, so actions must be focused. They work in an environment where availability and performance reign supreme. **'Actionability' is the key KPI.** Correct & plausible don't make the cut in our efficacy benchmarks, especially in the operations management environment.

## Understanding composite AI

Composite AI integrates multiple AI models to create a more comprehensive and robust set of capabilities that complement each other. The advantage of Composite AI is that it leverages the strengths of various AI components, each specialized in different domains, to create a more versatile approach with more accurate, actionable outcomes.

Think of Composite AI as an analogy to the human brain, where researchers observe similar specialization and work breakdown (cite: <https://www.nature.com/articles/nature18914>). While the cortex is uniform under a microscope, various imaging techniques suggest different parts of the brain specialize to handle different tasks. These regions of the brain come together to gather and

process information, maintain context, make decisions, recommend actions, recall knowledge, and then communicate these recommended 'next step' actions to various motor subsystems. Each lobe is assigned to perform specific tasks within the human brain. The Frontal Lobe is responsible for thought, memory, and behavior. The Parietal Lobe regulates language and touch. The Temporal Lobe manages hearing, learning, and emotions. The Occipital Lobe performs visual processing. A human brain can recommend the next best actions only when all of the lobes and functions within the brain come together.

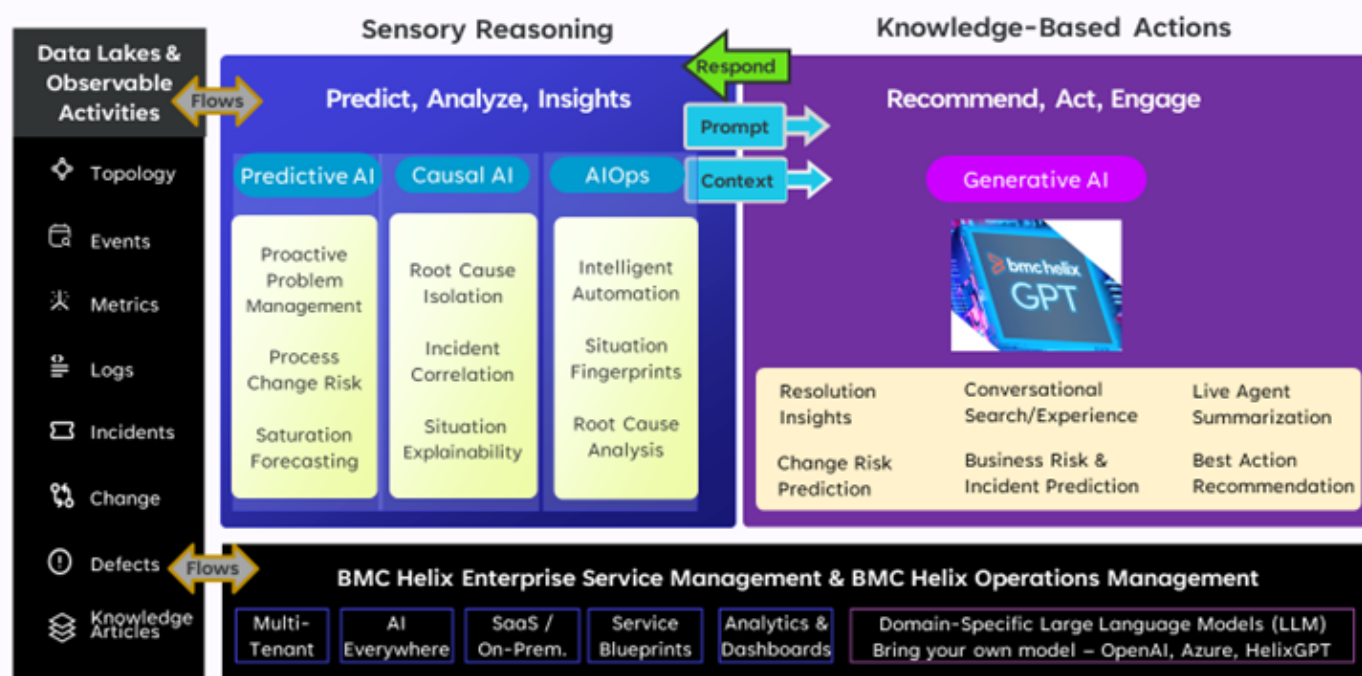
Composite AI within the context of enterprise ServiceOps, similar to the functions of the human brain, integrates and automates different types of intelligence to determine the best possible actions. However, Composite AI completes these functions on a massive enterprise scale across billions of data points in real-time by utilizing purpose-built processing pipelines for telemetry data to distill raw observations into facts that build up the context of a problem as it transpires.

With the help of Composite AI, we get to cast the monitoring products of the past as our eyes and ears and ticketing systems rich with domain and environment specific knowledge as our recallable memory.

## BMC Helix composite AI approach for improved actionability

The BMC Helix Composite AI approach consists of two main parts: sensory reasoning and knowledge-based action planning. The diagram below maps these two main parts in greater detail.

### BMC Helix Composite AI Approach



What you see on the far right-hand side of the diagram is data and a lot of it! BMC Helix captures data about all observable activities constantly flowing within your organization. Observable reality manifests itself on streams of topology, events, metrics, logs, incidents, change activities, defects, and even knowledge articles someone scribbled in a forgotten SharePoint folder somewhere. These traditionally siloed data lakes are often populated with information created automatically, user-generated information, and information through third-party integrations. Helix integrates all of that

data into a comprehensive model of your organization that is indexed by service topologies, as the structure and architecture of the service tends to help reasoning about all sorts of diagnostic and remedial automation functions down the line.

Sensory reasoning synthesizes and processes all of the incoming data to figure out what's going on in reality. Metric and event data from infrastructure, applications, networks (IP, Transport, Radio Access), and end users gets interpreted to detect anomalies. Here, various BMC Helix AI models are applied to detect anomalies such as unexpected traffic/load, resource utilization/saturation as patterns. BMC Helix then applies its proprietary AI algorithms to perform sensory reasoning to further process these anomalies into qualified situational explanations that capture what went wrong, what the root cause is what the impact seems to be. These BMC Helix AI algorithms include:

- **Predictive AI** applies AI techniques to predict future events or outcomes based on historical data and patterns. Components of predictive AI span machine learning (ML), training data, pre-trained models, regression, and time-series analysis. BMC Helix use case examples of predictive AI include proactive problem management, process change risk, and saturation forecasting.
- **Causal AI** integrates Knowledge Graph and Transformer-based AI techniques to understand and model relationships across observability data variables. It also determines the cause-and-effect relationships between events that unfold during a problem. Components of causal AI include reasoning about causal relations or patterns using topological data and a Knowledge Graph-based causality analysis, counterfactual 'what if' scenario analysis, graph modeling, and variability analysis assessing how causal relationships change depending on how the variables influence one another. BMC Helix use case examples of causal AI includes root cause isolation, incident correlation, and situation explainability.
- **BMC Helix for AIOps** leverages AI and ML to enhance enterprise operations by automating and optimizing tasks. BMC Helix for AIOps use cases include intelligent automation (such as for event management), root cause analysis, automated orchestration of routine tasks or workflows, automated integration with Enterprise Service Management, and third-party applications.

Through our Composite AI approach, the BMC Helix **platform** performs sensory reasoning across the entire IT stack: applications, containers, infrastructure, network, and even (if you have it) mainframe.

Now let's dive into the second area of the BMC Helix Composite AI Approach, operations-informed, Knowledge-Based Actions. Here, all of the distilled observability insights about Situations from the sensory and reasoning AI algorithms are used to build context for the generative AI –specifically BMC HelixGPT. **BMC HelixGPT** then produces, in human-style language, the situation explanations with recommended 'next best' actions.

The entire BMC Helix platform, across our Composite AI approach, is based on topology aware custom low rank adaptors that allow us **fine-tune models** for very specific tasks and based on your determined enterprise domains. We also use **retrieval augmented** generation to result in more contextual, detailed responses about realtime data sources such as transaction traces, live metric data, etc. These capabilities vastly improve the accuracy of AI insights, leading to improved actionability, which is the main KPI we track as discussed in the beginning.

# Applying the BMC Helix Composite AI to Operations Management

BMC Helix was built from the ground up to be a platform to process Observability and ITSM data at the telco scale. BMC Helix performs sensory reasoning based on observable reality – it provides the eyes and ears for the brain as it constantly processes vast amounts of monitoring data and formulates diagnostical reasoning as anomalies arise. BMC Helix harnesses all information flows specific to your enterprise data lakes, processing across time series and event streams. We employ Transformers and Knowledge Graph-based framework to achieve this data capture. In a future blog post, I will share a deep dive behind BMC Helix reasoning techniques involved.

We harvest and integrate monitoring data from existing tools into a unifying, comprehensive model that represents the structure and performance of targeted applications and IT services (modelled as a [property-graph](#)). BMC Helix does this dynamically without requiring any maintenance. As the architecture of the service changes with time, our AI discovers new boundaries/components, thanks to our BMC Helix for AIOps Service [Blueprints](#).

We employ a pipeline of AI&ML modules to convert near-real-time monitoring data into aggregations about emerging and impending anomalies likely to degrade service KPIs. We collect all the available ticket data to generalize resolutions people discuss in chat streams or work logs.

To gain credibility with operational teams, we have built explainability at the foundation of Helix. Any insight we derive from monitoring and/or ticket data can be mapped back to raw data sources or sometimes more advanced reasoning and feedback components that allow the experts to review how AI reasons. Explainability also serves as a conduit to harvest domain expertise from humans. Expert feedback is our source for learning new heuristics and domain-specific knowledge, which we then generalize so that they can be applied to future problems using GenerativeAI.

HelixGPT learns domain and environment-specific knowledge about resolutions from existing ticket/issue databases. It acts like the part of our brain that learns and generalizes new concepts. We collect all the available ticket data to generalize resolutions people discuss in chat streams or work logs. We have a propriety GPT-based neural network architecture that knows to pay 'attention' to actionable bits of these resolutions, so we can offer the operators remedial next best action even before the problem manifests at scale.

This necessitates the underlying GPT model to pay attention to vast graphs that describe the environment and the architecture of the target service, so we introduced graph-aware adapters that readily work on graph embeddings, as such vast data can't really be expressed in natural language in context. HelixGPT learns domain and environment specific knowledge about resolutions from existing ticket/issue databases. It acts like the part of our brain that learns and generalizes new concepts. These graph-aware adapters (an industry first, patent pending) sway the network's generation towards relational facts that matter in the environment (such as service dependencies, support-team memberships, et cetera), making us less prone to hallucination while keeping our generated insights highly actionable and specific to our users' environment.

Together, BMC's Composite AI approach with BMC Helix for ServiceOps, offers enterprises an integrated AI stack that sees/hears and learns/reasons about complex IT system issues – that's how operations teams can solve problems through clear actionability.