


# BMC HELIXGPT: AN EXPERT OF (YOUR) SYSTEMS



We have reached a historic point, where artificial intelligence (AI) is unleashing the next level of human efficiency. We haven't achieved artificial general intelligence (AGI), nowhere near close, but we will see increasing applications of AI in diverse verticals. Enterprise operations will be one of the first domains that will see significant impact with this shift.

Rapidly declining hardware costs, increasing availability of operational data, and recent modelling breakthroughs are enabling a tectonic transformation for the way enterprises will run their operations. We believe the future of enterprise IT and telecommunication networks will be more data driven, proactive, efficient, and less chaotic, which in turn will enable more innovation.

## **BMC's application of GPT technology**

Our mission here at BMC is to enable efficient digital enterprises. We believe the next level of efficiency in enterprise operations will require near real-time processing of operational data (observational and control plane data from infrastructure, applications, and end users) such that AI can establish an unadulterated view of reality, while ticket, incident, and change data (the activities of humans touching these systems) inform it on how human beings have been responding to issues native to the environment and business.

BMC HelixGPT, the working name for emerging GPT technology, integrates all these traditionally siloed data sources together to derive actionable insights for anomalies and autonomous resolution for mundane issues.

For this purpose, we have architected an AI platform that distills operational data into actionable insights (more on this later). It's a pipeline of AI models that automates L1/L2 IT analyst use cases

(anomaly detection, root cause analysis, and mitigation) by eyeballing operational data from IT or telecom services.

We ingest and integrate operational data from all observable layers. We distill anomalous patterns from metrics and logs into events and correlate them with all sorts of other events and change data in the system, with the help of topological data sources. The result of our operational data analysis is neatly packaged insights that we call Situations. Situations are causally ordered event chains that allow us to do root cause analysis and fingerprinting of repeat issues. They also give us a corpus of event patterns, which we then use to build predictive models around brewing problems. Our data pipeline looks roughly like the following diagram.

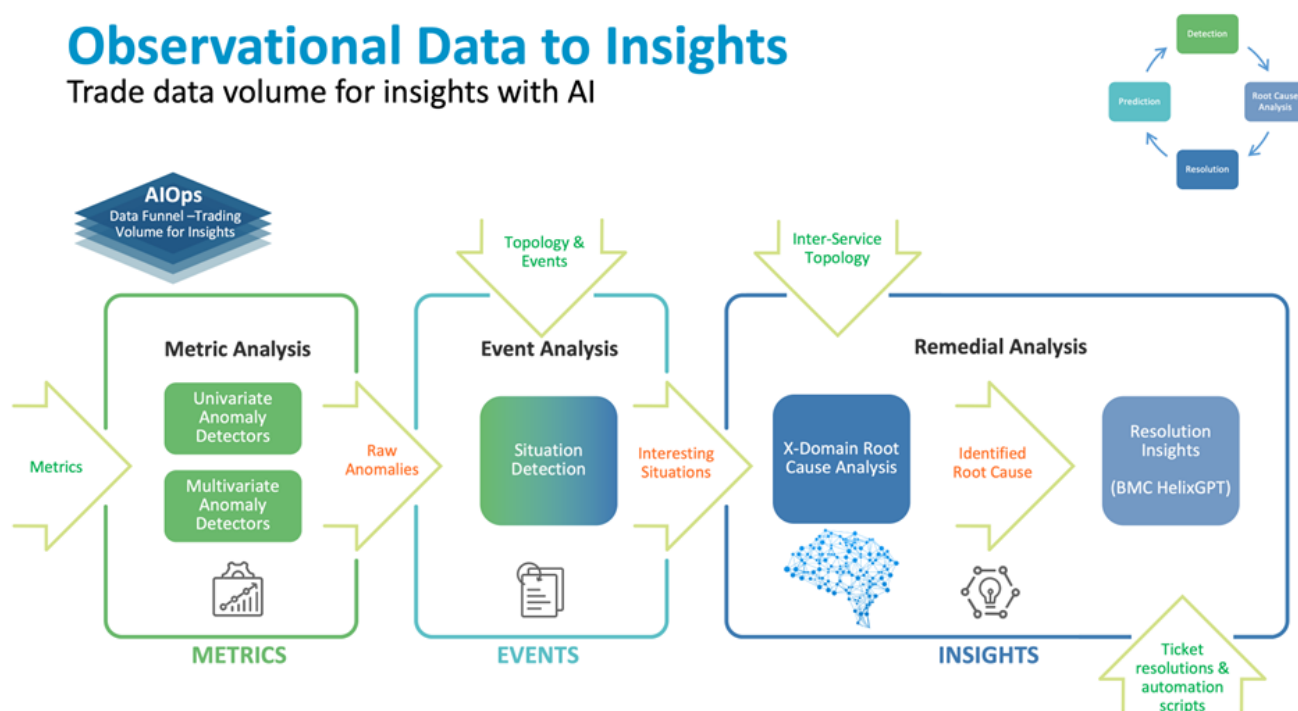


Figure 1. An AIOps data path where diverse data streams about operations are run through a series of analysis steps to identify root cause and resolution of preventable issues.

In our observation, telecom networks and sufficiently complex IT services often create tribal support groups around different layers of their architectures. Telecom businesses, for instance, span several related but distinct domains. These domains are vastly different from each other as they cover operational characteristics of facilities and buildings (power, cooling, etc.), transmission networks (aggregate data networks across vast geographies), core IP/WAN (high throughput networking for software definable circuits), network services (dynamic network definitions), and end user services (services rendered/metered for a customer). Operational knowledge about these domains is often recorded in ticket resolutions but gets trapped in different systems, authored by and catering towards different subject matter experts.

We leverage these tribal resolutions by tapping into ticket worklogs. We separate redundant entries from actual resolutions and use these resolutions to train a large neural network that complements a Generative Pretrained Transformer (GPT) model, as shown in the below diagram.

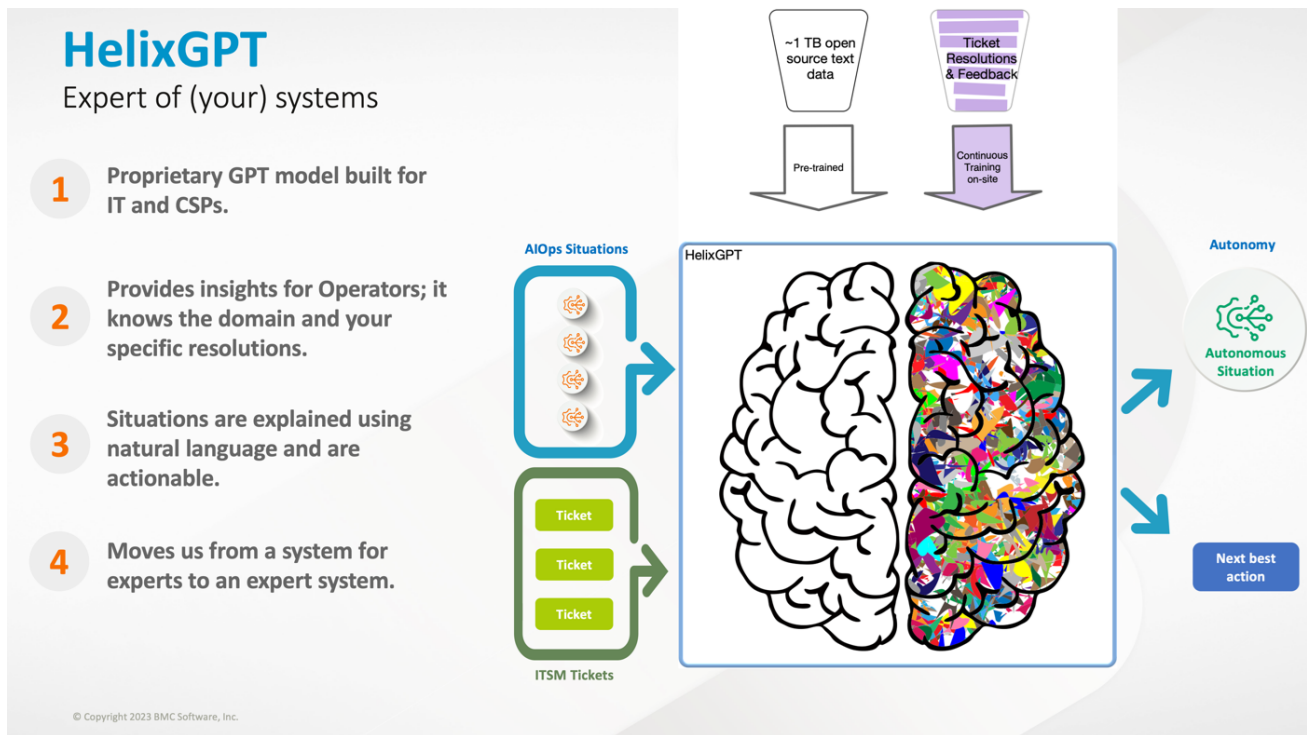


Figure 2. BMC HelixGPT analyzes AIOps and ITSM data to provide better insight into potential problems and automatically remediate low-risk fixes.

We are very excited about recent developments in GPT models, as they have validated our approach to solving full-cycle assurance. We have been using a GPT model to explain our insights to operators as a part of BMC HelixGPT, and we have now started to decorate our root-cause insights with ticket resolution-based prescriptions. We constrain our GPT with a domain and tenant-specific model that continuously learns resolutions from subject matter experts (SMEs) in the organization—hence the phrase “expert of (your) systems”—and prompt our GPT with root-cause analysis we conduct using observational data in prior steps in our AI architecture.

We developed this model to keep the insights provided by BMC Helix human-relatable and, yet, still specific to the enterprise so answers are based on your organization's expertise.