

The Future of IT Operations

Towards a high-quality and impact free IT Landscape

This document describes StackState's vision for preventing IT-related incidents. We start with current challenges in the market, followed by strategies to overcome these challenges. Finally, we explain how these strategies fit in the roadmap towards preventing IT incidents.

Increased complexity

Welcome to the era of digital transformation, where there's been a complete paradigm shift to dynamic environments. Here, software drives corporations to grow and compete, and agility rules the day. But, be forewarned—in this new fast-moving digital world, any problem in IT or application infrastructure can seriously disrupt business. Corporations risk losing millions of dollars and (potential) customers when these unforeseen incidents occur. That's why predicting and preventing IT incidents that could potentially impact business has become the "holy grail" of today's modern organization.

No control over chaos

Organizations have started to introduce new principles like DevOps and SRE to deliver better and more reliable software faster. This accelerated innovation in hybrid and multi-cloud environments and the dynamics around the changes in infrastructure and networks create the urgent need to stay in control more than ever. As a result, IT organizations have been adopting many different tools to monitor and manage every aspect of their IT landscape. More tools also mean more data generated by IT. When serious incidents threaten or strike, the infrastructure & operations (I&O) leader must make critical decisions. However, quickly finding the right data to make these decisions remains a challenge.

Lack of context despite all tools implemented

For data to be relevant and actionable, it must be placed in context. Knowing how all your data relates is imperative for making better, faster, holistic (business) decisions. However, current monitoring tools in the market do not cut across the multiple data types required for extracting holistic insights.

As Lodewijk Bogaards, CTO at StackState, explained, "These monitoring systems are not equipped to handle today's complex, scalable and dynamic IT landscapes. Too much work is left to be done manually, and the tools used by many IT organizations only address specific aspects of certain domains, such as applications, infrastructure or networks. These 'domain-centric monitoring tools' limit the IT operation's ability to detect, diagnose and address performance issues and to predict potential impact on other critical parts inside the organization. I&O leaders should look at the entire system, rather than specific technologies and infrastructure layers."

Beyond human capabilities: analyze the entire system

Domain-centric monitoring tools will continue to exist, providing data capture, analysis, and visualization for the DevOps or SRE specialist. However, to create holistic insights, they should forward their data streams to an Artificial Intelligence for IT Operations (AIOps) platform. An AIOps platform combines big data and machine learning functionality to support all primary IT operations functions. This is made possible through the scalable ingestion and automated analysis of the ever-increasing data generated by IT. The platform enables (1) the real-time use and collection of multiple data sources and (2) analytical and presentation technologies.

Start ingesting multiple data sources

Building or adopting a data-agnostic database to collect multiple data types is an essential strategy in any road map towards predicting and preventing IT-related incidents. To accelerate this journey for I&O leaders worldwide, StackState developed a versioned graph database as the fundament of its AIOps platform. This versioned graph database can collect and store any historical or streaming data type, including the characteristics of all nodes and the relationship between them. This includes tracing, telemetry and topology data. It also stores an incremental time slice of every change that took place inside your IT landscape, which makes it possible to travel back in time. All combined, this is known as StackState's 4T Data Model®.



Image: four essential dimensions for building rich context

Context as the fuel for AI and Machine learning

The data agnostic nature of this data model supports IT organizations in building the broadest context of any IT landscape ever seen in the market. As a result, StackState's AIOps platform acts as a lens where domain-centric data sources can be focused into a single, coherent, cross-domain analysis. The more data sources ingested into StackState, the more context can be delivered to the algorithms to analyze (1) why specific changes took place, (2) the impact the change had on the IT landscape and (3) how this data can be used to predict future incidents. StackState believes that context multiplies the value generated by artificial intelligence and machine learning. Without context, the capabilities of AI and machine learning for DevOps and SRE practices cannot be used to its full potential.

Embrace automation for accurate root cause analysis

When data generated by IT is too complex to process by hand, AI and machine learning are needed to extract useful insights. In this process, reducing false positives is extremely important for delivering relevant insights to the users. Within StackState, the algorithms and analytical capabilities are fueled by its versioned graph database, which contains context built up by multiple data types. This is a significant advantage, as the platform already understands where events occurred over time, including their up and downstream dependencies using graph- and bottleneck analysis. The AI and machine learning capabilities can therefore focus on extracting relevant information from the historical and streaming data collected, instead of using costly computing power to analyze how the data is related to each other, which on itself can result in false positives.

Moreover, StackState's algorithms can rely on more than one source. Drawing from multiple sources means events can be analyzed from different angles, which dramatically reduces the chance of delivering false positives to the users. These technologies combined deliver accurate root cause analysis and impact analysis to the StackState users. These analyses are visualized and presented through StackState's unique topology visualization interface and provide great insight as to where to focus remediation efforts.



Image: automated root cause analysis visualized through StackState's topology interface

Context for predicting and preventing IT incidents

StackState enables organizations to collect historical and streaming data to predict future incidents, an essential strategy for the future of IT operations. With its unique versioned graph database, StackState combines multiple data types in real-time and generates a full "post-mortem" style report—before an incident impacts the business. Powered by Artificial Intelligence, the report contains information on (1) the probable impact of the incident, (2) the likely root cause, (3) the probable trigger of the incident and the (4) predicted timeline. This is introduced as the 'Automated Pre-Mortem Analysis'.

With this technology, StackState helps I&O leaders make better decisions faster and avoid high-severity outages. Industry leaders like IBM, NS International, KPN Telecom, NN Bank and Schuberg Philis already trust StackState in taking this journey.