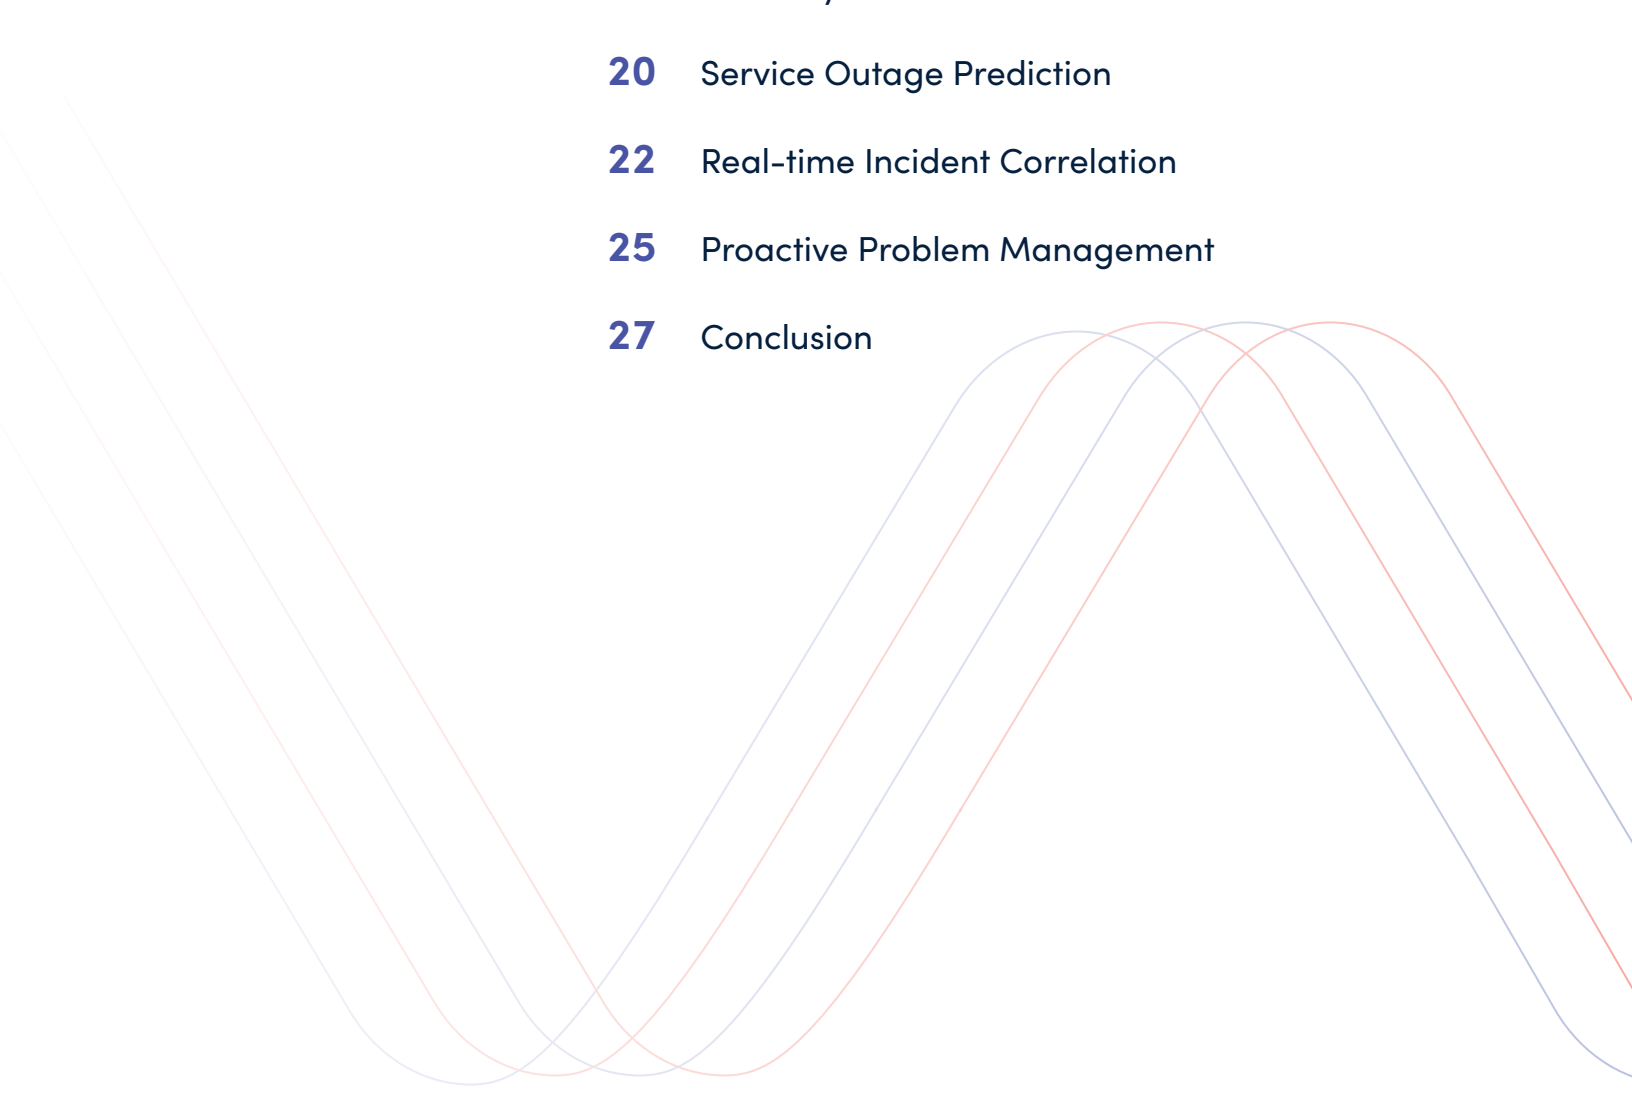


Harnessing the Power of AI Across Service and Operations Management

Implement IT best practices with BMC Helix AI-enabled ITSM and ITOps solutions to deliver value-driven outcomes through a single platform.

Table of Contents

- 03** Executive Summary
 - 04** Challenges Addressed by BMC AISM and AIOps Solutions
 - 05** Dynamic Service Modeling
 - 08** Service-centric Noise Reduction and Probable Root Cause Analysis Using Situations
 - 14** Anomaly Detection
 - 20** Service Outage Prediction
 - 22** Real-time Incident Correlation
 - 25** Proactive Problem Management
 - 27** Conclusion
- 

Executive Summary

Defense organizations can leverage the power of generative AI and observability to identify and prevent IT issues before they arise.

Operations teams face many difficult challenges. They are overwhelmed with increasingly large amounts of complex data from multiple sources, and event noise hides problems that need to be solved quickly.

Using artificial intelligence for IT service management (AISM) and artificial intelligence for IT operations (AIOps) is a paradigm shift that enable machines to solve IT issues without human assistance.

BMC Helix Operations Management with AIOps uses monitoring, advanced anomaly detection, artificial intelligence, and machine learning (AI/ML)-based event management and root cause isolation, open integrations, and intelligent automation to reduce MTTR and maximize service performance and availability.

The machine learning (ML)-based, AI-powered solutions provided by BMC analyze, predict, spot, and help IT organizations proactively resolve issues in real time.

Why would your organization require AISM and/or AIOps solutions?

These solutions help you address complex military IT challenges and manage the exponential growth of data that many organizations are experiencing by offering the following capabilities:

- Automate the entire operations process across multi-cloud and hybrid environments

- Apply AI and ML to detect patterns and reduce noise, save time and labor, and lower mean time to repair (MTTR)
- Identify the likelihood of future outcomes using predictive analysis and, monitor business services and visualize status using heat maps and tile views
- Automate the identification of frequently recurring incidents by applying natural language processing (NLP) and AI
- Leverage AI to accurately detect emerging situations such as major incidents in real time
- Proactively remediate issues before any impact on service level agreements (SLAs) occurs
- Enable rapid, SaaS-based deployments with containerized, microservices architecture
- Use out-of-the-box adapters and REST APIs for policy-driven data collection and ingestion of topologies from third-party solutions to improve visibility and business service management
- Use rules to quickly determine the most likely root cause of an event

This paper discusses the challenges posed by the need for rapid changes in IT and how they can be addressed with AISM and AIOps solutions from BMC.

Challenges Addressed by BMC AISM and AIOps Solutions

Digital transformation is essential for a program's competitiveness and growth. If your program has not already introduced the new digital services and mission models that the warfighter demands, they're working quickly to do so. This transformation places unprecedented pressure on IT to play a more strategic role to the mission and comes with its own set of challenges, which BMC's AISM and AIOps solutions can help address, as follows.

Break down data silos

The inability to manage large chunks of data is a key reason many organizations are not able to monitor events and systems effectively in their environment. With BMC AIOps solutions, data is ingested in the form of topology fragments, events, metrics, and logs, and is processed through a set of algorithms that select specific data points. After those data points are selected, a correlation or set of patterns is identified and inferences are drawn and then passed to a collaborative work environment. Similarly, the application of AI in service management helps with creating and linking problems to a recurring set of incidents—the set of data that traditionally resides in silos.

Eliminate IT operational noise

If you are part of an IT operations (ITOps) team, operational noise is your primary concern. This noise creates several problems for the business, including higher operating costs, degraded performance and availability, and risks to the enterprise's digital initiatives.

BMC AIOps solutions make a tangible difference across industries because they don't just reduce the noise; they also eliminate it by creating correlated incidents that point to a probable common root cause.

Deliver a seamless customer experience

Ensuring a seamless customer experience with predictive analysis is an important business objective. BMC's AISM and AIOps solutions can automate complex decisions by collecting, integrating, and analyzing data. By leveraging this data, the solutions can reduce manual efforts and predict future events before they become an issue and start affecting availability and performance.

With BMC's AISM and AIOps solutions, you can:

- Reduce efforts required for incident analysis and problem creation
- Easily automate decision-making
- Speed up problem-solving
- Continuously integrate and deploy automated solutions

Overcome monitoring and analytics challenges

Data collection is the primary step for enabling AIOps; you must collect and correlate data from multiple sources to analyze it effectively. BMC Helix Operations Management, one of BMC's AIOps-powered solutions, provides a single pane of glass to analyze the data collected from across all the underlying domains and services, including data from 3rd party sources, reducing the need to use multiple tools for analysis.

Dynamic Service Modeling

BMC Helix Operations Management

leverages and extends the dynamic service models (DSMs) that define the boundaries of assets and relationships in your IT environment. DSMs created by BMC Helix streamline service model management and updates.

Using those models, BMC Helix:

- Enables integration and normalization of data from third parties
- Creates rich visualizations
- Leverages BMC Helix capabilities to combine inventory and event data into a [unified topology view](#)

For example, Becky is a service designer with Apex Global. As an operator, she needs to monitor and manage the organization's various business services, such as its data center operations.

She also needs to be able to model the needs of her various internal customers. These customers will care about different boundaries around the assets, and they will have different goals. While creating the service models to map these services, Becky realized that there are a few common processes that are part of each of these service models. For example, both the Kubernetes cluster service and virtual applications management service are shared services used by many other services, such as Kubernetes deployment, network availability monitoring, virtual data center operations, and so on.

In this example, the common basic services—such as Kubernetes deployment, network availability monitoring, and virtual data center operations—can be defined as service blueprints.

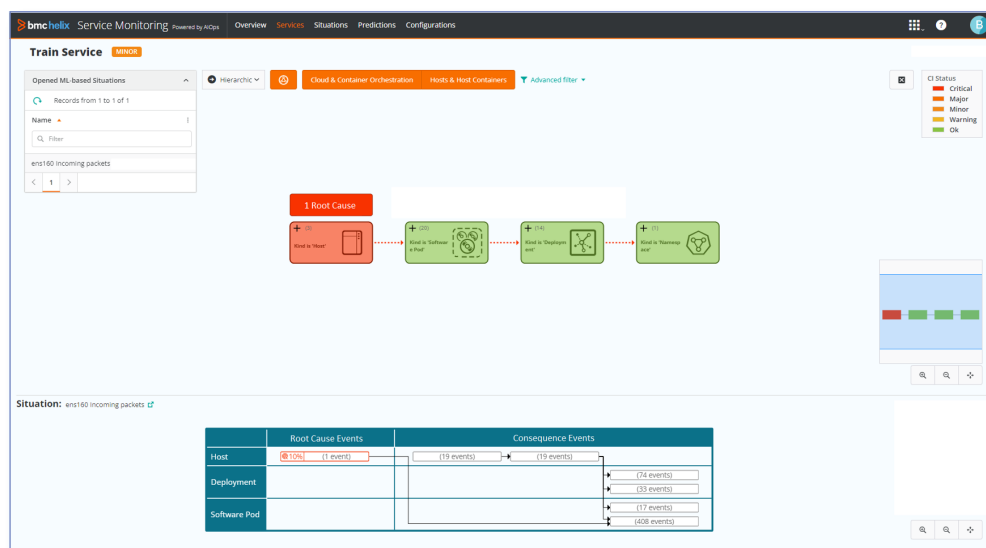


Figure 1: Service topology view with a detected Situation that points to the root cause and associated impact

To manage commonly deployed services automatically, BMC Helix provides service blueprints, which capture the architecture of services that are repeatedly deployed within the organization or across industries.

A service blueprint is a parameterized set of rules that identifies the elements of a business service. It contains conditions for matching some starting nodes that should belong to a service and a set of graph traversals that are used to find related nodes.

Building blocks of blueprints are nodes that represent CI queries and links that represent traversals for inclusion. You can set node and link filters to define the criteria for inclusion.

BMC offers an ever-growing catalog of predefined service blueprints for commonly available architectures. In addition, you can create and edit your own blueprints for your organization to use.

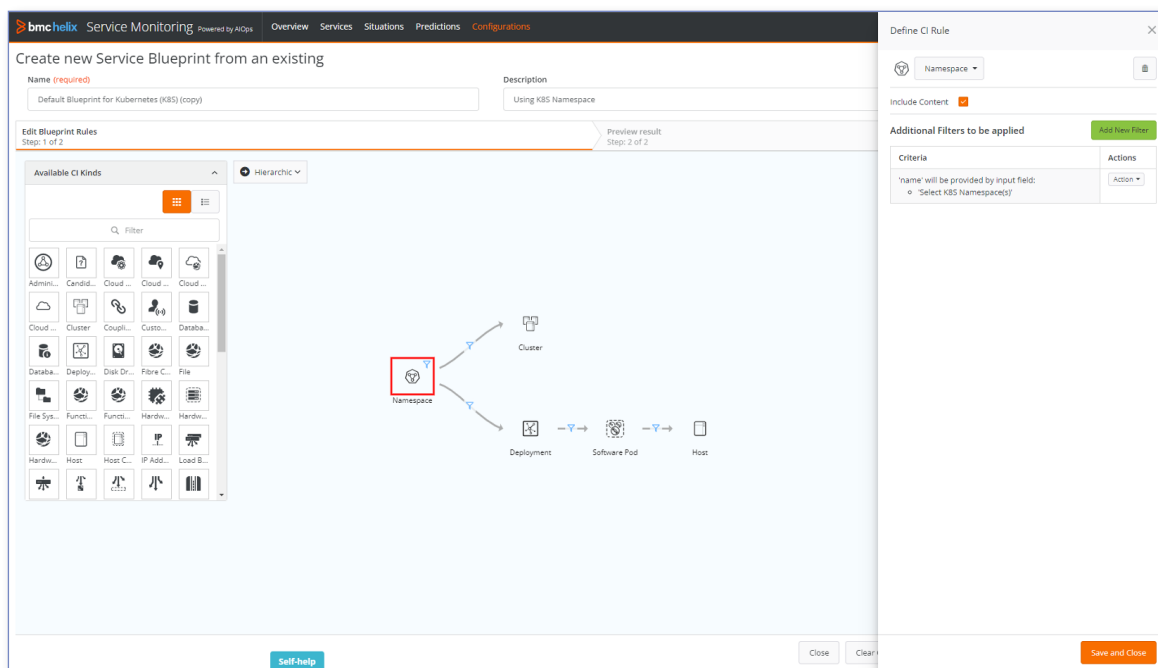


Figure 2: The blueprint editor provides a list of node types the blueprint can use (left pane), shows a set of graph traversals (middle pane), and provides the ability to define filters (right pane)

You can create services by using blueprints. First, you need to specify the “variable” required by the blueprint definition. Once defined, the service is maintained and kept up to date in case of any CI or blueprint change.

To use a blueprint to define a business service, you must create an instance of the blueprint by providing the input parameters to match the required nodes. The system can then use the traversal rules to find the matching related nodes and include them in the service. As the data changes, the system automatically reapplies the rules so that the service model stays up to date.

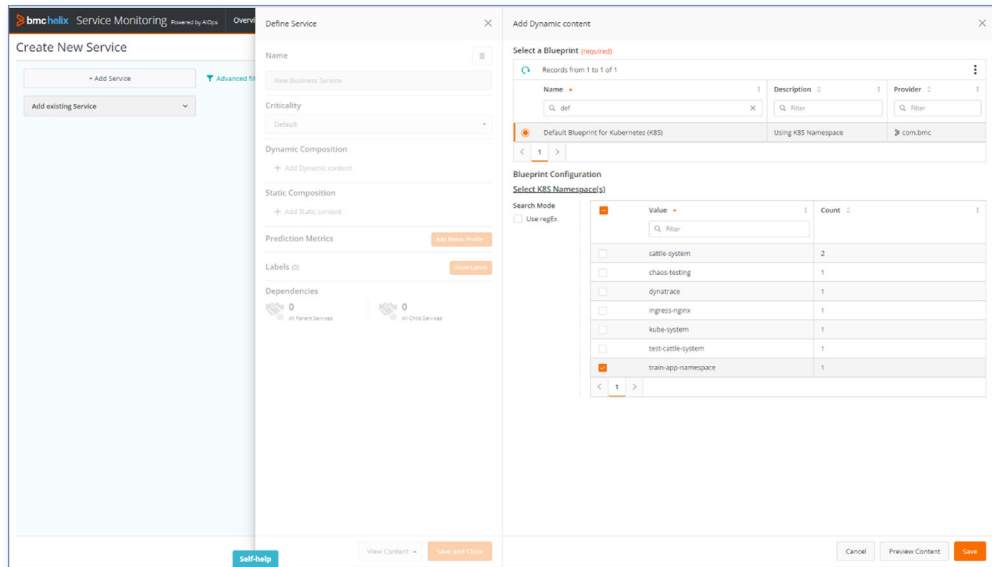


Figure 3: The service editor prompts you to select values that should be used as "variables" for the selected blueprint

In addition to the CIs defined in the blueprints, you can manually select a set of CIs to be added statically to a service.

Using BMC Helix, you can model dependencies between services. Therefore, large organizations or structures can be divided into smaller, more manageable sets of services. In the following example, each service can contain its own set of CIs.

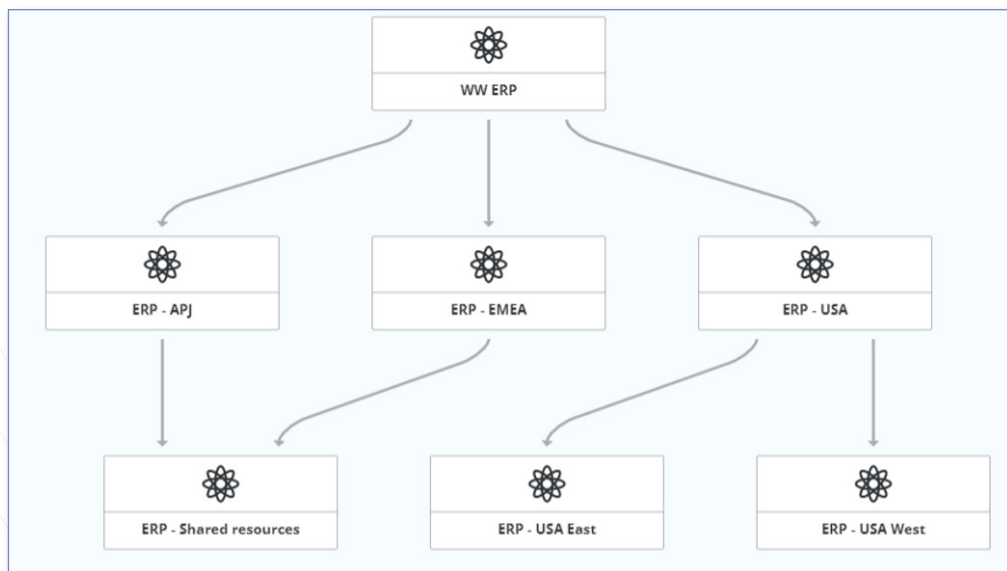


Figure 4: You can manage dependencies between services by dragging and dropping them in the service editor hierarchical view

Service-centric Noise Reduction and Probable Root Cause Analysis Using Situations

Not every anomaly, outlier, or even failure is necessarily a problem that will lead to business interruptions. In fact, in a complex environment, there is always some background event load, almost like a constant, cosmic background noise.

Some of these events, however, are leading indicators that later turn into problems with real business impact. In the past, various attempts have been made to reduce noise by using online or offline event processing algorithms that relied on event similarities to calculate correlations.

However, such techniques often fell short in achieving noise reduction because correlation doesn't necessarily mean causation. Also, most of the proposed techniques involve heavy hyperparameter tuning that simply pushes the model accuracy responsibility to the end users.

We believe that maximal event noise reduction is only possible after events are sorted by causality. After the events are sorted by causality, it is trivial to separate the root cause events from the symptomatic events that have transpired on the same

dependency chain. Furthermore, when the causal chains are grouped by the root cause event, we can correlate seemingly unrelated but causally related incidents. Hence, we postulate that identification of the root cause offers the best possible reduction of noise.

For this reason, we use a novel clustering algorithm that leverages observed topologies to measure the causal distances of event pairs. We rely on observed and ontological distances between the sources of these events. Resulting clusters, once mature, become "Situations" that identify the root cause and the impact.

We make no assumption about the size and number of these clusters because we rely on the dynamic topology to guide us to auto-tune our clustering.

Hence, the resulting algorithm requires zero maintenance from the end user in terms of data modeling and hyperparameter tuning or training.

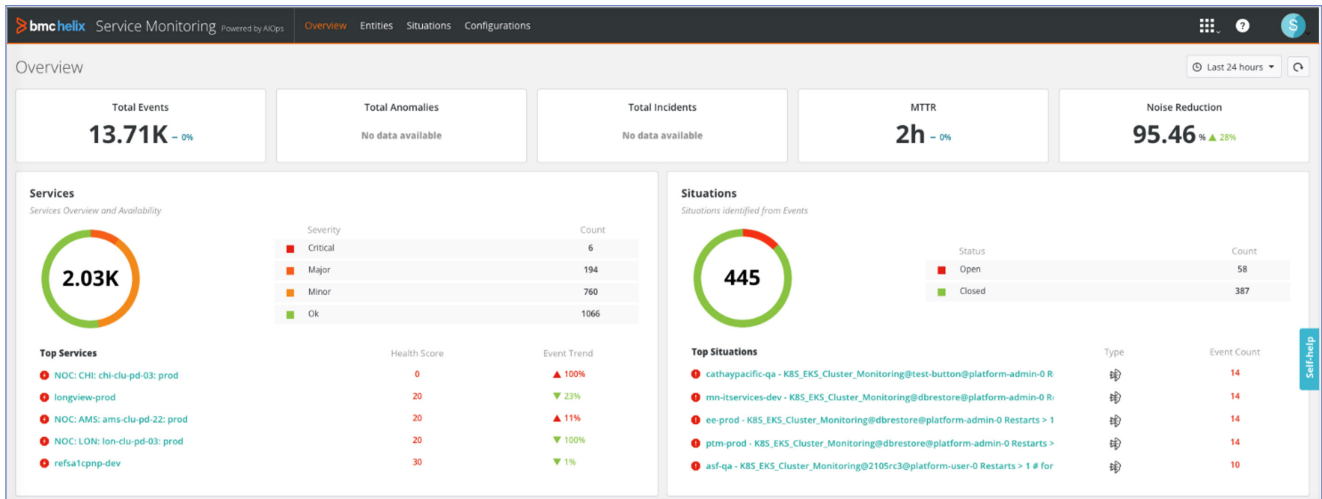


Figure 5: This view shows all events observed in an environment, the Situations created out of them, and the noise reduction percentage

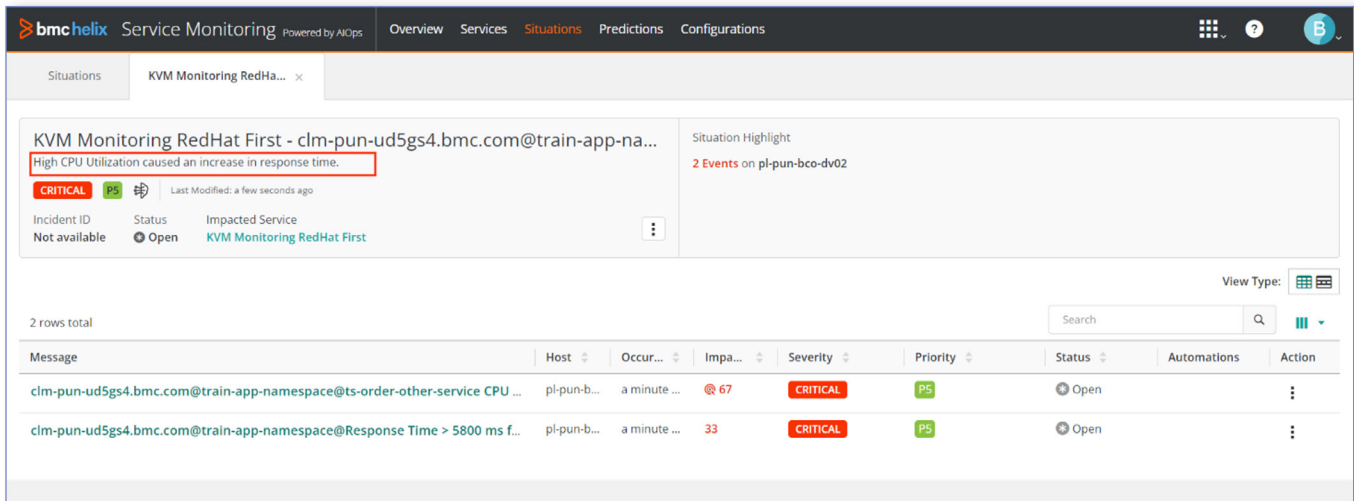


Figure 6: This Situation lists all 51 causally related events in ranked order, along with the probable root cause score and related symptoms; the highlighted text shows the AI-generated summary of the Situation

Algorithm used for event noise reduction

BMC Helix AIOps archives algorithmic noise reduction by identifying and filtering low entropy events and causally clustering filtered events as Situations using temporal, topological, and ontological dimensions. The events used for generating Situations can originate from different layers of the architecture, such as application, network, or infrastructure.

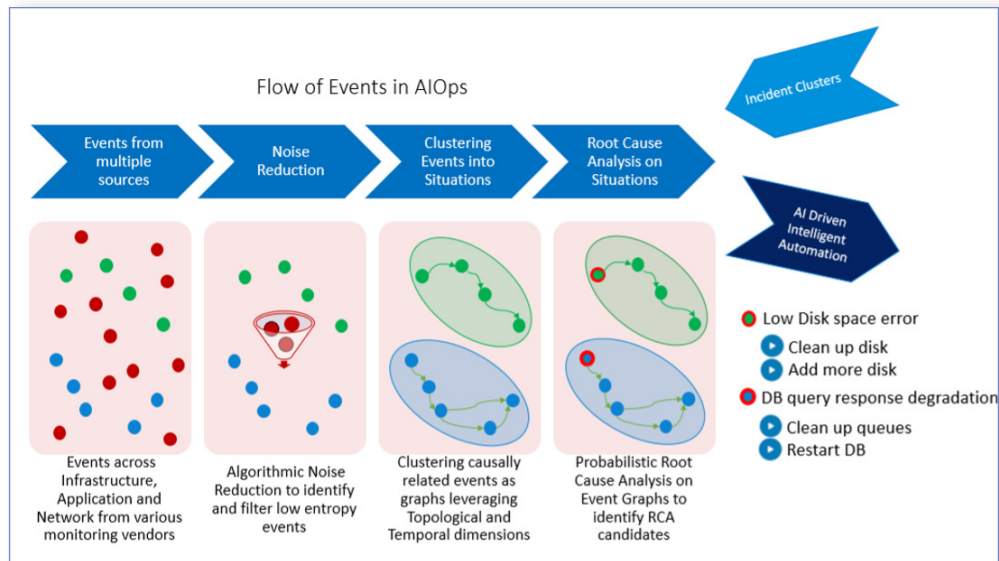


Figure 7: BMC Helix AIOps identifies, filters, and causally clusters events from multiple sources as Situations for a more streamlined and actionable view

The algorithm works in three phases:

1. Incremental differential clustering for causality:

Causality: Our clustering relies on temporal context, textual context, topology, and an internal knowledge graph. The knowledge graph maintains various domain-specific bits of knowledge so we can evaluate the encountered events with ontological inferences. This information enables us to generalize what we know about anomalies in new technologies and frameworks into cause-and-effect transmission conduits, which point us to the root cause.

2. Situation title generation: We use custom generative AI models to generate human-readable and actionable text from each Situation.

3. Situation lifecycle management: As events change state or get added to a Situation, or time elapses, the state and severity of the Situation changes. For example, a Situation gets closed when all events in the Situation are closed. Both the time and state of the events influence the state and severity of a Situation.

As shown in Figure 8 below, the following events have been generated for a service:

- "Response time degradation" on device A of type "Micro-service"
- "High memory utilization" on device B of type "Physical server"

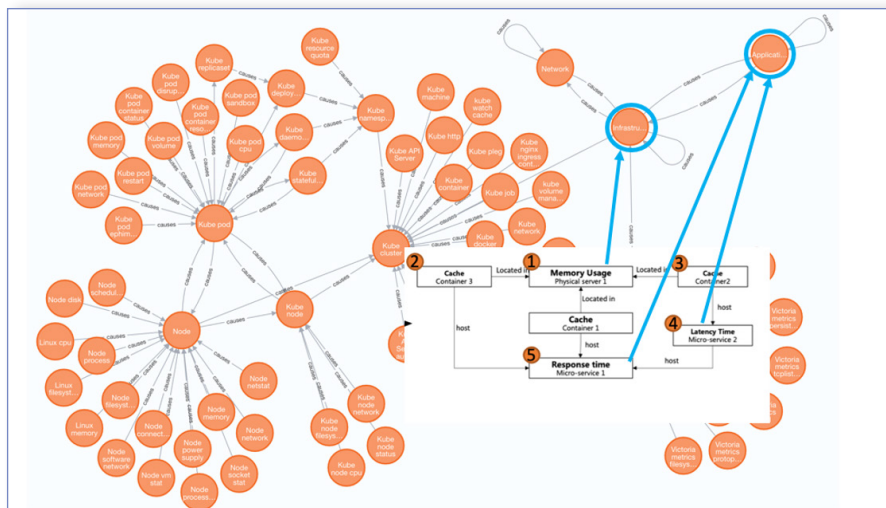


Figure 8: The figure above shows the knowledge graph of the “Response time degradation” and “High memory utilization” problems and how they are related to each other topologically and temporally

Using the ML-based algorithms, we perform the following tasks:

1. Extract the temporal similarity between these events.
2. Check how these events are topologically related in the observed topology.
3. From the knowledge graph, hypothesize whether “High memory utilization” on device B of type “Physical server” can cause “Response time degradation” on device A of type “Micro-service.”
4. Based on the above information, create a directed causal relationship between these events.
5. After identifying the directed causal relationships between all the events in the service, based on our patented clustering technique, identify and group events as Situation clusters that are evolved from the same root cause. Therefore, each Situation cluster can be seen as a directed event graph involving the root cause and its related symptoms.

6. If the service has different evolving problems, generate multiple Situations related to a problem.

Service-centric probable root cause analysis

Service-centric probable root cause analysis (RCA) is a key differentiator for BMC and one of the primary benefits of the BMC Helix open AIOps solution.

Probable RCA enables ITOps teams to:

- Identify services that are impacted
- View the top-scoring Situations for any impacted service
- Determine the most likely root cause of a potential problem or issue in each Situation after factoring in time, metrics, events, and topology
- Change the time window to view how the Situations adjust over time
- Drill down to the details of the Situation’s RCA score, including events, metrics, and topology

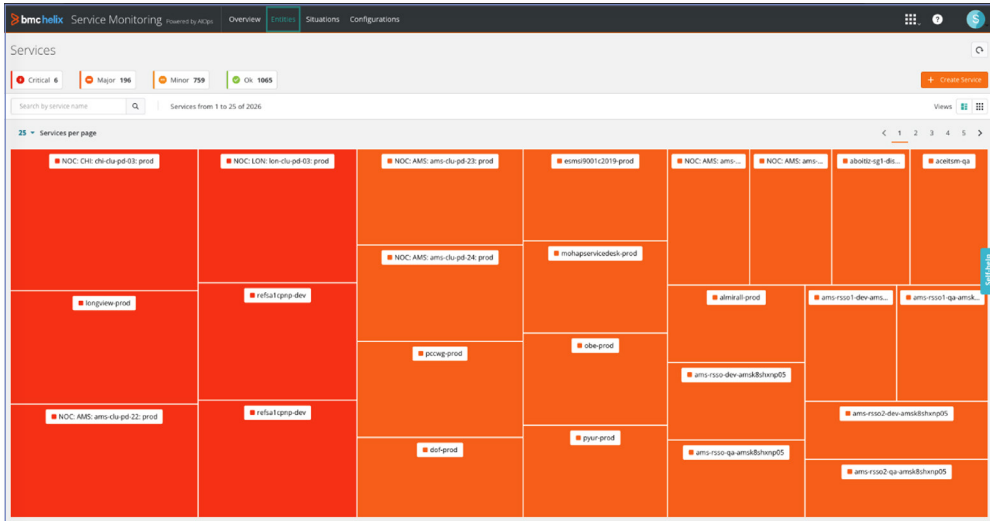


Figure 9: The heat map above shows critical and minor services from a total of 2,000 services

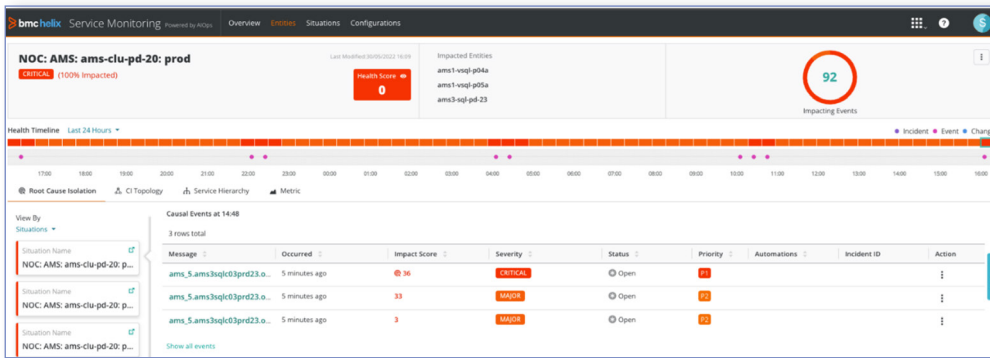


Figure 10: The figure above shows an impacted service that has multiple Situations, and the root cause events with the impact score

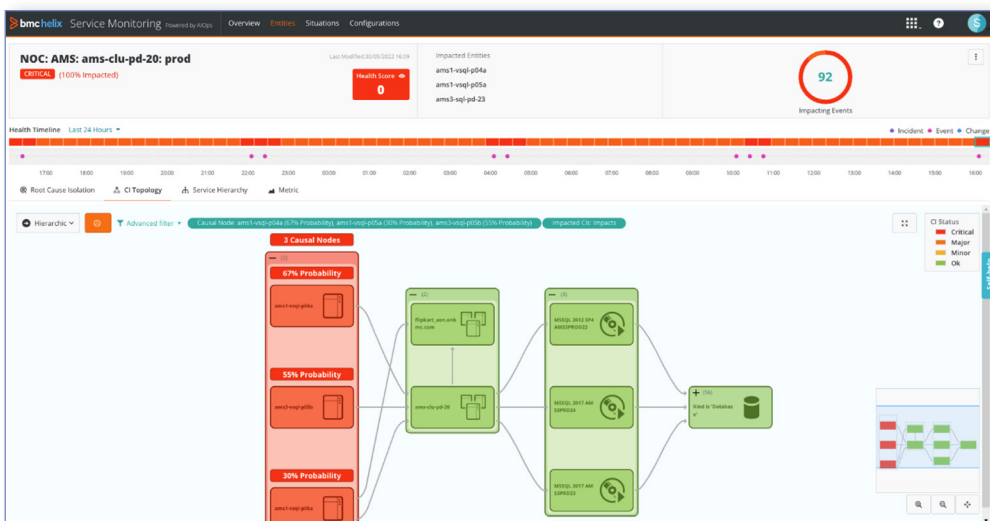


Figure 11: The Situations overlay on the service topology indicates the impacted devices and the devices that are root cause candidates

Algorithm used for RCA

The algorithm used for RCA is based on causal analysis of data (events, incidents, change requests, and so on) collected from the service and operations management systems across the dimensions of topology, time, criticality, text, and service model. It uses ML to reduce noise and determine the probable causes of a Situation that affects an IT service.

The algorithm works in three phases:

- 1. Collect and cluster:** It collects Situations formed from all the relevant data—such as events, metrics, logs, incidents, and change requests—across the topology of a model in the DSM, and change integration with ITSM is built into the solution.
- 2. Rank:** It runs a custom ranking algorithm that uses the event graphs generated from Situations and criticality, topology, knowledge graph, and time proximity to determine the causality rank order of Situations and root causes within Situations.
- 3. Learn:** It uses the supervised deep learning models that are built based on the feedback and historical Situations to learn and improve new causal relationships between existing and new types of entities and problems. This feature provides improved accuracy for both the noise reduction and RCA using Situations. This learned knowledge graph also provides an out-of-the-box RCA for new services, which can then be transferred across different services and tenants.

For example, the following events have been generated for a service:

- “Response time degradation” on device A of type “Deployment Service”
- “High CPU utilization” on device B of type “Kubernetes Pod”

Using the algorithm, we perform the following tasks:

1. Identify the Situation from the events.
2. Based on the page rank, try to identify the events causing more causal events with much higher impact scores and tag them as probable root cause candidates. In this example, we assign a high-impact score to the second event since it is causing the first event and tag it as probable root cause.

If the users believe the identified causal relationship is inaccurate, they can provide feedback around that relationship and feed it to the ML model. Root causes and Situations are continuously improved based on this feedback.

Anomaly Detection

Anomaly detection is one of the most sought-after technologies in IT today, and BMC Helix leverages ML techniques to provide a robust anomaly detection service.

BMC Helix proactively remediates issues before any impact on SLAs, optimizes the customer experience, increases productivity, and reduces the number of incidents generated from events by:


- Triggering events and notifications based on a single metric or group of metrics behaving abnormally
- Reducing event noise by using sensitivity controls to manage detection
- Displaying detailed graphs showing the anomalous metric or metrics for events

Univariate anomaly detection

The univariate anomaly detection service streams the incoming time-series data and detects any unusual occurrences that could be potential anomalies. After detecting the occurrences, it can trigger an alert, allowing you to identify underlying issues at an early stage and proactively avert any upcoming issues or outages.

This service has the following features:

- Support for real-time streaming to facilitate early alert generation
- Flexibility (relying on as few assumptions as possible)
- High efficiency (both in space and runtime)
- Interpretable output
- Customizability (for example, ability to add seasonality)
- Support for millions of metric data that have a low memory footprint
- Scalability



Algorithm used for univariate anomaly detection

The algorithm leveraged in BMC Helix uses a new data structure for an accurate online accumulation of rank-based statistics such as quantiles. It uses a variant of the one-dimensional K-means clustering algorithm to create a very compact data structure that allows accurate estimation of quantiles.

This algorithm works by capturing the distribution of data. The time-series data is captured into an hourly digest. For example, the observations for each Sunday from 3–4 p.m. are sent to the same digest each week. For each digest, the data is first sorted and then clusters are created based on their distance. It forces the clusters at the tail end to be smaller, thus providing a proper balance between a very accurate quantile estimate in the tail of a distribution and a reasonably accurate quantile estimate near the median, while keeping the number of clusters as small as possible.

When a new observation streams in, the algorithm measures its distance from the centroids of the existing clusters.

If the distance is far enough, the observation is merged with the existing clusters or, if the distance is greater, a new cluster is created on the tail. If this value breaches the lower or higher threshold of the digest, it is identified as an anomaly.

Baseline calculations

The anomaly detection service considers a weekly seasonality for hourly data. This consideration means that a digest is created for each hour of the week (for example, Monday 8–9 a.m.).

After the service starts, to create a baseline, data for the first six hours is captured, and the percentile value that is calculated based on these observations is used as a baseline for the next 24 hours. After data is available for 24 hours, its percentile values are used as the baseline for the next six days and the next week.

Finally, when we have data for an entire week, it is used to calculate the baseline for the same hour of the same day the following week. For example, Monday 9–10 a.m. data for the first week will eventually be the baseline for Monday 9–10 a.m. during the second week, and so on.




Figure 12: The blue line represents the actual data points, and the grey area represents the normalcy band (calculated dynamically); when the data breaches the normalcy band, an alarm is generated (red icon), when it falls in the normalcy band again, the alarm gets closed (green icon)

Capturing seasonality

One of the most important features of the univariate anomaly detection service is the ability to distinguish between an anomaly and seasonality. The underlying algorithm can identify weekly seasonal patterns for hourly data and exclude them from its anomaly alerts. For example, if there is a spike in the data on Monday between 9 and 10 a.m. and it occurs on the same day and hour in subsequent weeks, it is identified as a seasonality and not an anomaly.

Configuring alarm policies

You can configure additional alarm policies based on your requirements on top of the ML-based anomaly alerts. For example, if an alarm is generated when the anomaly detection baseline threshold is violated, you can add an additional check to generate a critical alarm if the metric value exceeds 65 percent for one minute and an anomaly detection baseline threshold is violated.

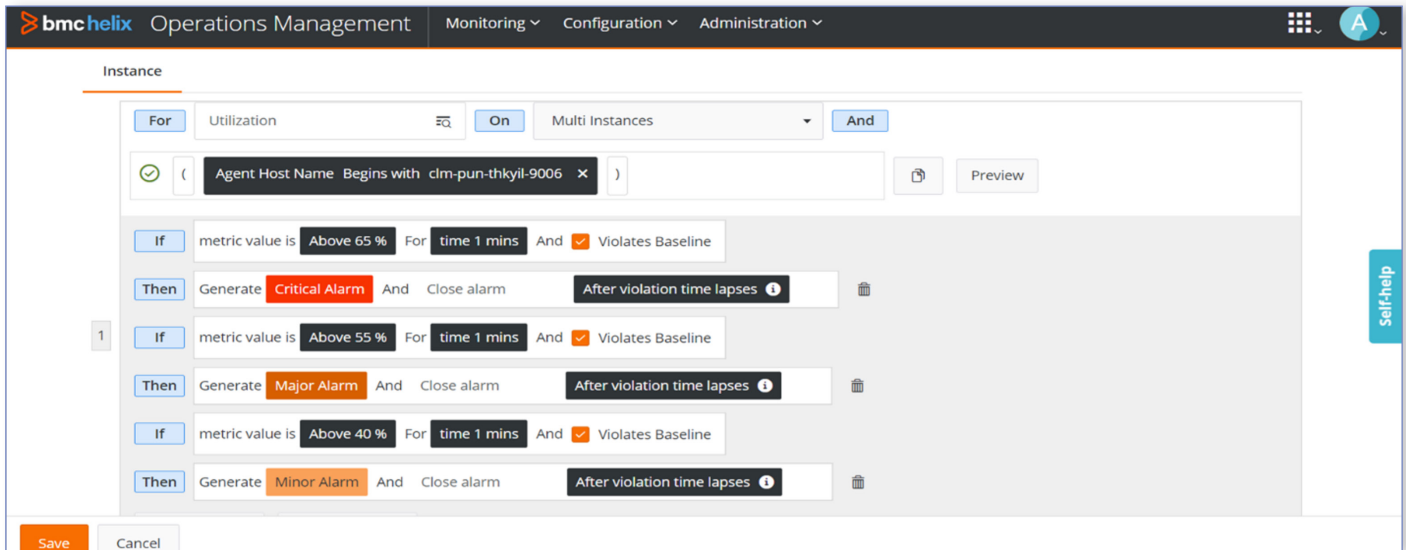


Figure 13: A simple alarm policy configuration

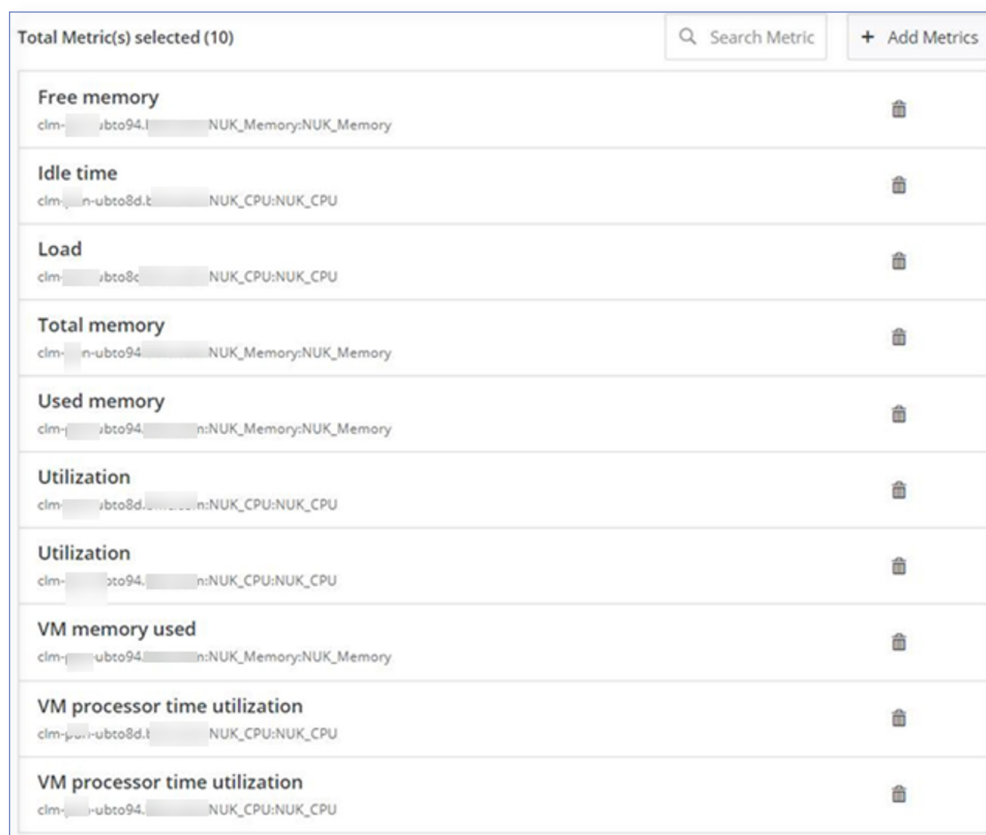
Multivariate anomaly detection

Multivariate anomaly detection is computationally more complex because it models a group of metrics (for example, all key service metrics or entity metrics). However, it can surface interesting interactions among different metrics.

Algorithm used for multivariate anomaly detection

For multivariate anomaly detection, BMC Helix enables you to define systems of related metrics as key performance indicator (KPI) groups and detect anomalies that capture the relationships and interactions among these metrics.

BMC Helix uses the Amazon SageMaker Random Cut Forest (RCF) anomaly detection algorithm to enable this capability. RCF is an unsupervised algorithm for detecting anomalous data points within a data set. All the metrics configured in a KPI group are analyzed in combined form and a single anomaly is detected, along with the contribution score of each of the metrics related to the anomaly.



| Total Metric(s) selected (10) | | Search Metric | + Add Metrics |
|--------------------------------------|---|---------------|---------------|
| Free memory | clm-...-ubto94.l...:NUK_Memory:NUK_Memory | | |
| Idle time | clm-...-ubto8d.l...:NUK_CPU:NUK_CPU | | |
| Load | clm-...-ubto8c.l...:NUK_CPU:NUK_CPU | | |
| Total memory | clm-...-ubto94.l...:NUK_Memory:NUK_Memory | | |
| Used memory | clm-...-ubto94.l...:NUK_Memory:NUK_Memory | | |
| Utilization | clm-...-ubto8d.l...:NUK_CPU:NUK_CPU | | |
| Utilization | clm-...-ubto94.l...:NUK_CPU:NUK_CPU | | |
| VM memory used | clm-...-ubto94.l...:NUK_Memory:NUK_Memory | | |
| VM processor time utilization | clm-...-ubto8d.l...:NUK_CPU:NUK_CPU | | |
| VM processor time utilization | clm-...-ubto94.l...:NUK_CPU:NUK_CPU | | |

Figure 14: If you have configured a policy with 10 metrics, all of them are analyzed together and a single anomaly event is generated

The algorithm combines and analyzes all 10 metrics together, and provides a single anomaly score, along with how much each metric contributes to the anomaly score. The same metric from different machines is considered to be two different metrics. For this sample anomaly event, the anomaly score returned by the algorithm is 1.298457.



Figure 15: The table above lists how each metric has contributed to the anomaly score in order of most impactful to least impactful

In the above table:

- The last metric (Total Memory) without any change in its value is not contributing; hence, its contribution score is zero
- The second metric (Used Memory) has a spike of ~10.45, but at the same time all the other metrics are behaving normally, hence an anomaly is not raised
- The fifth through ninth metrics in the table above are contributing minimally to the overall abnormality score (if they had been configured as univariate, it might have been anomalous)

In addition, BMC Helix has a set of rules and hyperparameters that allow you to reduce false positives such as sustained anomaly detection and variability range configuration of the anomaly score. The RCF algorithm associates an anomaly score with each data point. A low score value indicates the data point is normal, and a high score value indicates the presence of an anomaly in the data. The definitions of “low” and “high” depend on the application that sends the data. However, in common practice, scores beyond three standard deviations from the mean score are considered anomalous.



Service Outage Prediction

Predictive analysis is the use of data, statistical algorithms, and ML techniques to identify the likelihood of future outcomes based on historical data. It helps in understanding what will happen in the future given the current state of a business service, if the infrastructure remains the same.

BMC Helix computes the predicted value of a KPI based on the current values of metrics. If the predicted value is in the abnormal range, alerts can be generated. This multivariate analysis is where multiple metrics are used to predict a KPI. Identifying the metrics that have a larger impact on KPIs enables you to prioritize your investigation.

The service outage prediction feature offered by BMC Helix has the following advantages:

- 1. Near-term prediction:** Such as for the next 15 minutes, 30 minutes, or few hours, including a sudden abnormality.
- 2. Cost-effective performance:** Streaming data that is driven by multiple infrastructure and business objects leads to a large number of metrics. Despite this large number, the prediction is near real-time, and processing is cost-effective.
- 3. Explainable:** The prediction model can be explained.

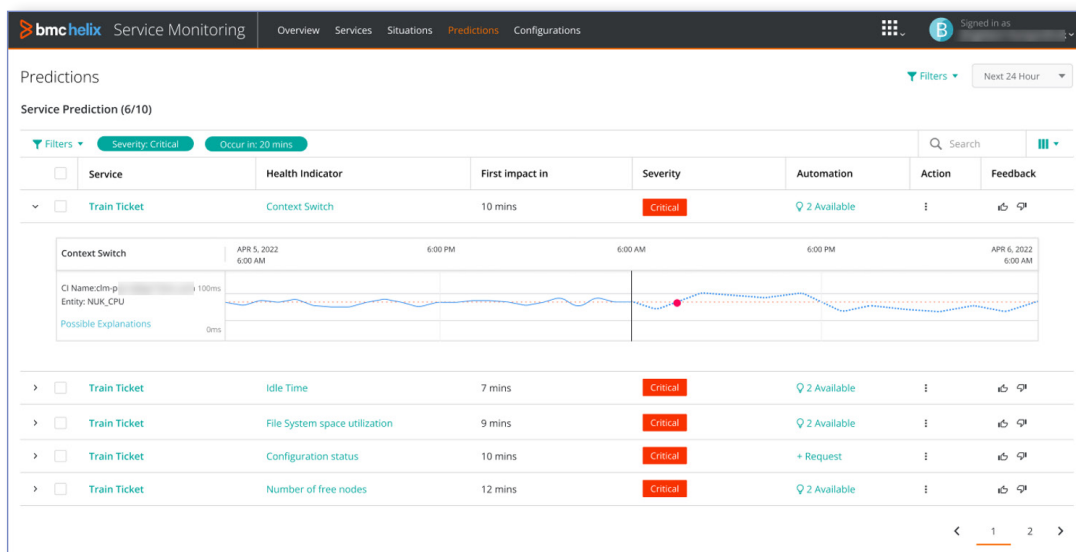


Figure 16: The figure above shows service outage predictions for a train ticket service. A red dot indicates the outage prediction time

Paradigm used for service outage prediction

For service prediction, BMC Helix uses the online, incremental ML paradigm, which helps with creating a model that learns incrementally with every new data point.

Algorithm used for service outage prediction

BMC Helix uses the SGDRegressor module based on the Stochastic Gradient Descent (SGD) algorithm to optimize linear regression. This module is quick and efficient and helps with achieving incremental learning; it also supports multivariate learning and analysis. The predicted data from SGDRegressor is extrapolated or used by the state-space ML model for forecasting future intervals.

For example, consider “response time” as the KPI for a service that we need to forecast. We take network metrics like packets in and packets out that could impact the response time into consideration. For every row of metrics data, the SGD algorithm first predicts and then learns, using gradient descent with a defined learning rate. This process continues for all the incoming data. The predicted value is stored as an array.

After every defined interval (e.g., 15 minutes), the predicted values are used to forecast the next 24 hours. The forecasting algorithm is the state-space algorithm that requires seasonality as a hyperparameter to forecast.



Real-time Incident Correlation

In organizations with a large inflow of incidents to the service desk from multiple channels, the major incident managers on call or service desk supervisors do not possess the analytical means to analyze an emerging Situation such as a major incident. They often need to rely on a report, a dashboard, or word of mouth to detect whether a set of incidents reflects the same (emerging) condition of an impacted IT service. A significant effort is required to detect and manage major incidents when time is of the essence and end users are affected.

The AISM capabilities of BMC Helix ITSM automate the identification of clusters of incoming incidents that are related to the

same underlying Situation in real-time. The automatic identification helps major incident managers or service desk supervisors identify major incidents or other significant disruptions quickly, enabling them to rapidly reduce service downtime and minimize any negative business impact.

In addition, the solution's ability to manage duplicate incidents reduces work for service desk supervisors. Teams can focus on resolving the original incident, which, when resolved and closed, leads to closure of all the child incidents as well. This frees up resources that can be used for other higher-value work and restoring services as quickly as possible.

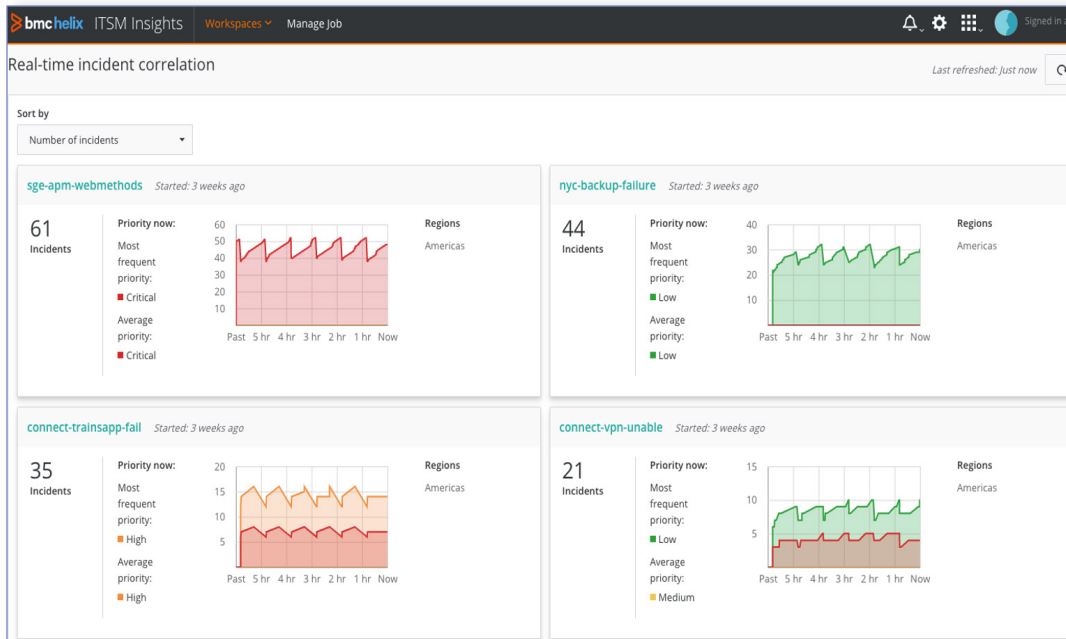


Figure 17: The Real-time incident correlation dashboard displays new, incoming incidents as clusters based on user-configurable grouping fields, text fields, and sliding time windows; each tile on the dashboard represents a cluster

Algorithm used for real-time incident correlation

BMC Helix ITSM uses a state-of-the-art ML algorithm, BERT, which is a language model. This algorithm is bidirectionally trained (as compared to the unidirectional models that read the text input only in one direction). Therefore, it has a deeper sense of language context to group incidents based on the semantic similarity of incidents and generate a meaningful caption. If you have two incidents with similar text (e.g., “I cannot connect to Skype” and “Fail to connect with Skype”), they are placed together into a group or cluster because the intent of these two incidents is the same. If you have configured the Proactive Service Resolution (PSR) feature to create incidents from events, BMC Helix ITSM can create event-based clusters.

For example, on a Monday morning, multiple users are facing network issues when connecting to a video conferencing system. These users start creating tickets because they are not able to connect properly.

Each user describes the problem differently (e.g., “I cannot connect to video conferencing system” or “Video conference system—connect failed,” etc.). However, all these tickets have a similar semantic representation, indicating an issue related to users not being able to connect to the video conferencing system. The BERT algorithm allows for learning the context of a word based on its surroundings (both to the left and right of the word) to automatically detect textual similarity between these tickets and start grouping them together in near real-time as they happen.

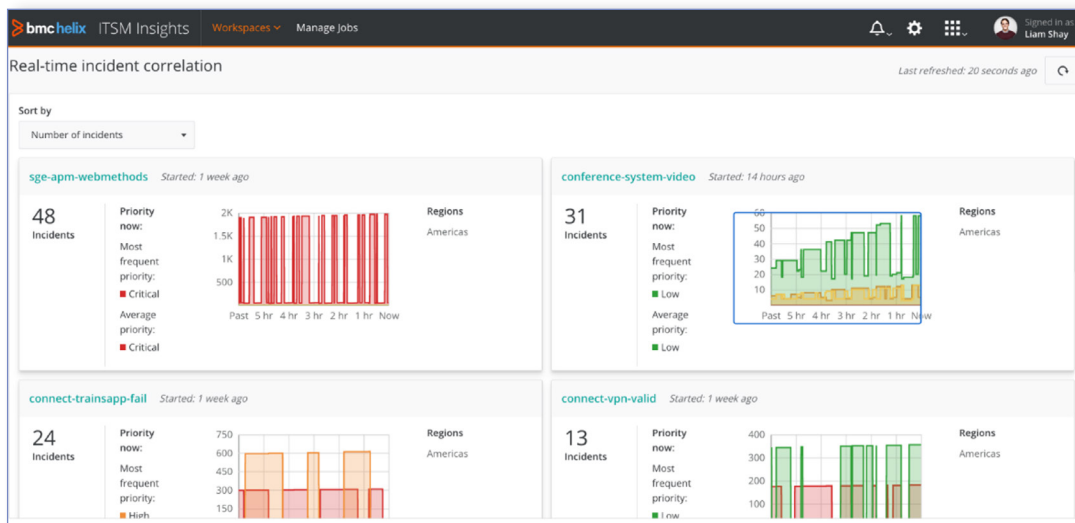



Figure 18: After a certain configurable threshold (usually five to seven tickets), the algorithm creates a new cluster (conference-system-video) on the dashboard

As new tickets are created, this cluster count increases; if the tickets are resolved, the cluster count decreases. By looking at these clusters, an incident manager immediately can conclude that the issue started four to five hours ago, and there is a steady increase in the number of tickets, from 30 to 60 in the past few hours. The incident manager can double-click on this cluster to drill down for a detailed view and take action, such as relating all these tickets to a parent-child relationship, in BMC Helix ITSM.

In addition to clustering, the algorithm continuously evaluates whether any of these clusters should be tagged as “Candidate Major Incidents” based on certain criteria, such as a rate of increase of tickets in a prescribed time or the number of tickets in a cluster. This evaluation helps rank the clusters to quickly identify fast-trending clusters.



Proactive Problem Management

Problem investigation is a costly process and therefore must be prioritized diligently. It is not possible for problem coordinators to look at every closed incident and launch problem investigations into them because of the associated cost. First, they need to determine the set of frequently recurring incidents. When determined manually using Microsoft Excel or any other tool, it becomes a very cumbersome process that might take days or even weeks depending on the volume of data. The other challenge is locating key information in free text fields. Often, the whole process suffers due to a lack of accuracy and reliability. For these reasons, the problems remain unmitigated.

BMC Helix ITSM automates the identification of recurring incidents and problem investigation recommendations by using proactive problem management, which improves efficiency and leads to simplification. Proactive problem management not only provides capabilities to find recurring patterns of incidents and visualize them in heat-map and drill-down views, but also to take action by quickly creating a problem investigation in BMC Helix ITSM.

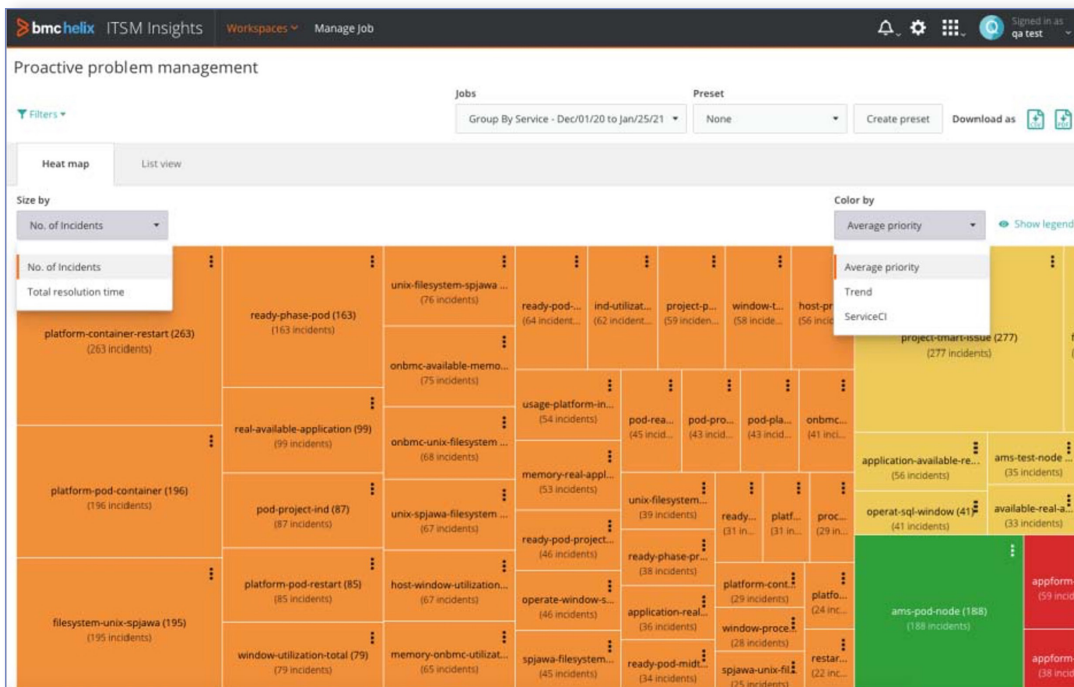


Figure 19: The proactive problem management dashboard displays recurring incidents as clusters in the form of a heat map

Algorithm used for proactive problem management

The K-means clustering algorithm is applied to the resolved and closed incidents to identify similar incidents and tightly group them into clusters in a way that is coherent. In minimizing the sum of distances between the incident data points and the corresponding clusters, it uses various techniques such as specifying initial clusters, initialization of centroids, and continuous iterations. The clustering algorithm also identifies interesting clusters that are candidates for “problem investigations.”

For example, proactive problem management can help problem managers understand recurring, similar incidents based on historical ticket data, such as data for the past week or past few months.

The first step in this process is to create a problem management job by configuring the lookback period (for example, three months), filters, the group-by-field, and text fields to apply ML. The group-by-field is optional but can be used to provide hard boundaries for clusters such as group-by “ServiceCI.” The text field used for ML is usually the “Description” field. Once these fields are configured, the algorithm fetches the historical data from BMC Helix ITSM according to the lookback period, with standard pre-processing and data-cleaning steps.

Next, the algorithm does the first-level grouping by looking at the group-by-field and then applies the K-means algorithm on the text field for natural language-based similarity.



Figure 20: The algorithm evaluates the quality and generates tiles for each of the clusters produced

This heat map shows the highest-volume critical priority cluster to be “prt-webmethod-alarm,” with 1,704 incidents in the past three months. Another “certificate-valid-issue” cluster points to a certificate issue in VPN causing a high volume of 201 critical incidents. The problem manager can click any of these clusters, drill down for a detailed view, and take action (e.g., create a problem investigation).

Conclusion

AI/ML technologies are key in cloud-native technology environments that have many more components, such as microservices and containers, than traditional applications and emit much larger volumes of operational data.

IT teams must adopt a holistic service and operations management strategy driven by intelligence and AI/ML across their entire hybrid IT environment. They must build monitoring into digital and cloud transformation processes and strive to eliminate gaps in visibility and control.

Additional benefits from an AISM and AIOps platform approach include the ability to:

- Identify and correlate Situations faster
- Restore service and remediate problems more efficiently
- Enable effective and seamless collaboration
- Improve technical resilience

It is imperative that IT teams deploy ML and analytics as part of an open AISM and AIOps strategy to manage the increasing volume, variety, and velocity of data, the management of which has grown beyond human scale across an increasingly hybrid, complex, and fast-moving IT landscape.

Additionally, IT needs a flexible solution that is easy to deploy and upgrade and supports fast-value realization to meet the constantly changing needs of a dynamic organization. As such, a SaaS deployment model is ideal for rapid onboarding and cost optimization across any environment.

Leverage AIOps-powered solutions provided by BMC for a comprehensive, end-to-end service and operations management solution.

BMC Helix ITSM, BMC Helix Operations Management, and BMC Helix Discovery are SaaS solutions that use a microservices-based containerized architecture to provide fast and scalable deployment in an easy-to-use environment. They also contain a full suite of capabilities including anomaly detection (univariate and multivariate), dynamic service modeling, next-generation probable root cause analysis, service-centric monitoring, open integrations, log analytics, and advanced, automated event management with event correlation and noise reduction.



For more information

To learn more about BMC's AISM and AIOps solutions, please visit bmc.com/it-solutions/serviceops.html.

About BMC

BMC works with 86% of the Forbes Global 50 and customers and partners around the world to create their future. With our history of innovation, industry-leading automation, operations, and service management solutions, combined with unmatched flexibility, we help organizations free up time and space to become an Autonomous Digital Enterprise that conquers the opportunities ahead.

www.bmc.com



BMC, the BMC logo, and BMC's other product names are the exclusive properties of BMC Software, Inc. or its affiliates, are registered or pending registration with the U.S. Patent and Trademark Office, and may be registered or pending registration in other countries. All other trademarks or registered trademarks are the property of their respective owners. ©Copyright 2024 BMC Software, Inc.



5 4 4 9 8 8