greymatter.io

Use Case:
Mesh Networking Behavioral Anomaly Detection

Introduction

Service-based architectures and decentralized cloud-native computing delivers on the promise of rapid software release, risk reduction and business flexibility. They also generate an exponential growth in network communications and signal volume. These signals contain a wealth of operational and business intelligence if successfully captured and analyzed. However, doing so in a timely fashion requires vast resource expenditure and advanced technology. Greymatter.io's AI employs this data to provide human operators indicators, recommendations and actionable outcomes with respect to your enterprise's cloud-native mesh networks.

Current Solutions Miss the Mark

Today's conventional analysis solutions require significant human intervention, lack insight, and deliver little actual value. More often than not, they are relegated to shelfware status. At best these tools simply waste resources. At worst, they create more problems than they solve.

Bringing on additional human observers is not a viable solution for these data volumes, and would introduce a dual inefficiency: First, since enterprise systems perform without catastrophic failure most of the time, a full-time army of human monitors would spend most days observing a perceived healthy system. Second, human reliance is a gamble. Incidents occur quickly, and may be too subtle for a human to detect, particularly one conditioned to prolonged inactivity. Of course, even an alert watchman will have trouble drawing inferences from reams of numbers and statistical associations.

Even with more eyes on the problem, troubleshooting remains a challenge. When a problem occurs, the decentralized nature of today's networks means there are simply too many places to check before the problem can be effectively framed, much less resolved. Teams inevitably find it daunting merely to narrow the problem source down to a manageable set of possibilities, let alone remedy the cause, leading to scenarios such as the following:.

> VP Engineering: "Why isn't this service working?"
>
> DevOps: "Perhaps it's a bug in the service itself, and the logs might show something, but... it logs a lot, so we'll need to restart and catch it in the act. Or it's a bug in one of its five dependencies. We'll check their resource charts and logs. Or perhaps they're all working fine, but it's the service mesh configuration, and traffic isn't making it there. Or perhaps its working in one cloud on specific platforms and not another. We'll need to make some test queries while watching logs to confirm. Or maybe something else corrupted the database. We'll check those logs too. Or maybe it's a denial of service attack."

Simply stated, the cause could be many things. That statement has never been more true than today in an IT world of disparate systems, services, data stores, clouds, platforms and applications, all dependent on one another.

The Greymatter.io Solution

Today's enterprise requires a tireless, intelligent, and omnipresent capability that learns your mesh network's patterns and behavior, operating at machine speed. Greymatter.io's unobtrusive AI uniquely fills this gap with live-learning pattern detection powered by deep neural networks. Greymatter.io ignores noise and normal variation, sidestepping the problems inherent in rulebased systems and conventional AI. It learns the characteristic qualities of each service, application, node, traffic pattern, and uses them to identify outliers. When a potential issue is detected, Greymatter.io immediately delivers the context and clues an operator needs to take action.

Once deployed, Greymatter.io receives signals (logs, observables, and statistics) from services throughout a deployment, automatically training a neural network for each service to recognize its usual "look and feel". Greymatter.io then actively monitors the firehose of signals for any that deviate too much from the characteristic behavior it has observed. In further detail, Greymatter.io's deep semantic anomaly detection capability is powered by recurrent neural autoencoders. It requires no user configuration, and rarely requires tuning. Machine learning establishes the normal pattern of behavior for each service, and automatically detects deviations as they occur. The platform creates a central repository of abnormal behavior, eliminating the

need to hunt for TM & © 2020 Decipher Technology Studios 2 of 6 “something strange” during troubleshooting. Instead, evidence of such behavior is delivered to the user for direct review

Incidents are identified in the same way a human accomplishes the task. The platform:

1. Observes and establishes what normal signals look like,
2. Watches the stream for anything significantly abnormal given that understanding, and
3. Reports its findings continuously to the human operator.

Unlike a human, Greymatter.io can consume and analyze massive volumes of raw data. Such volumes are common in wide-spread, decentralized service and application networks. The following provides a schematic breakdown and walk-through of the Greymatter.io AI process.

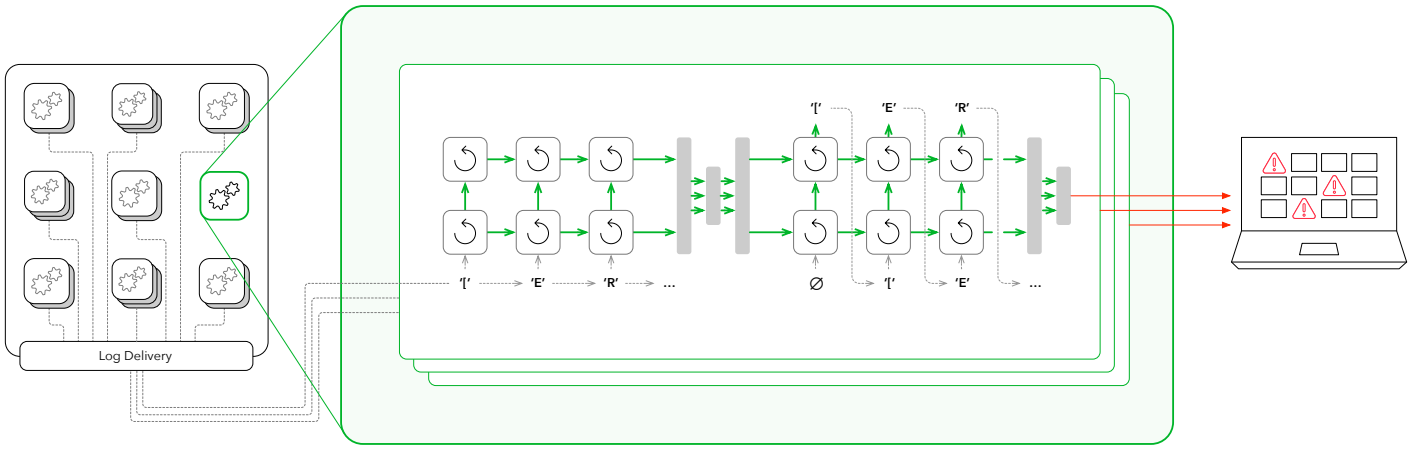

Mesh Networking Behavioral Anomaly Detection

1. All services, Greymatter.io sidecars, and other auxiliary software in a deployment produce signals in the usual way. Metrics, Statistics, Audits, and Logs are delivered to the Greymatter.io AI services, hosted within the mesh network.
2. Each new signal stream is treated separately and automatically, and no new configuration is required to begin accepting signals from a new configured source within the Greymatter.io mesh network.
3. Once training is complete, every signal stream is run through the fully-trained system, and any that fall too far outside of the training distribution (above a statistically computed threshold) are flagged as anomalies.
4. Greymatter.io presents anomalies to the user in conjunction with other system events as part of the larger operational picture.
5. Collaborative input from human users is incorporated into the future state of the behavioral models, further increasing precision and recall while uncovering potential issues throughout your environments.

Unmatched Enterprise Value

Greymatter.io is the *only* universal mesh networking platform on the market today to integrate behavioral anomaly detection. The platform is uniquely designed to capture and analyze the otherwise overwhelming firehose of mesh-emitted signals typically left untouched by similar platforms. Greymatter.io anomaly detection delivers:

Force multiplication: Greymatter.io keeps critical business operations performant with minimal human interaction. This helps keep teams streamlined, unifying and combining critical analysis functions typically fulfilled by multiple add-on 3rd party tools within a single platform.

An unblinking eye: Greymatter.io is ever present throughout the network, working at finely granular levels to provide continuous overwatch and warning on critical systems. Early indicators enable operators to fix issues before they impact users, ensuring maximum uptime.

Deep insight: Without Greymatter.io, the operator has much shallower health checks (often falsely reporting that all is well), user feedback (subject to human failure and limitation), manual checks (how fast can you type?), and mere optimism that "*silence is golden.*" None of these scale to thousands of services and none of these bring peace of mind.

Machine Speed: Greymatter.io conducts data capture and analysis of volumes of data at speeds no human can match. This means faster detection and mitigation of errors and malicious intrusion attempts, including those undetectable to the human eye, in a fraction of a second.

Adaptability: Greymatter.io anomaly surfacing requires no configuration, incorporating new signals as they're received. Per-service sensitivity is set automatically using the AI's own training statistics.

Adjustability: Sensitivity and other post-training parameters are exposed to the user for optional fine-tuning. If new versions of a service behave differently enough to warrant it, Greymatter.io can also be signaled to restart the training process on a fresh stream.

Course Correction: Greymatter.io integrates operator feedback incrementally, with training status and many other statistics available for situational awareness and troubleshooting by human operators.

Greymatter.io Case Studies

The following case studies demonstrate the real world applications and benefits of Greymatter.io's deep semantic anomaly detection capability.

Early Warning

1. During a canary deployment, a new version of Service A is released, and a portion of user traffic diverted to it. Greymatter.io begins detecting anomalous logs four hours into the test, indicating increasing response times on successful requests. This is a warning of unexpected resource contention and subsequent backup *before* any requests have failed, triggering an investigation well before the problem has actually impacted any users.

2. During normal operation mid-week, Service B begins throwing HostNotFound exceptions in its logs. This is only caught by anomaly detection. Further investigation reveals that a misspelled host name was added to a static list, subtly degrading performance. All services depending on Service B slow down, and buffers begin to back up. Without Greymatter.io's early warning, these services would eventually fail.

Troubleshooting

1. Service C, with which the human operator has little working experience, fails at unpredictable times with a segmentation fault. Nothing seems unusual in its own logs, truncated as they are, though the line "Segmentation Fault" is anomalous and shows up in the Greymatter.io anomaly detection application. Each time this happens, one second immediately prior to the segmentation fault, an anomaly is detected in a *dependent* of this service, and the line indicates a rare request to Service C. Greymatter.io's depth of data capture provides the operator with important information to share with the developer about an otherwise opaque problem.

2. Service D is returning corrupted responses in production, which the developer cannot reproduce locally. The service's copious logs would be overwhelming to the operator, normally requiring the developer to travel on-site. However, in this case, Greymatter.io analysis indicates the timing of corrupted responses correlates with anomalous logs from the database itself. The database log lines, which would otherwise have seemed irrelevant to the operator, are enough for the developer to reproduce the bug.

Security

1. During normal operation, Service E begins producing anomalies. At first the activity seems innocuous: the debug endpoint, which isn't frequently used, is now receiving occasional requests. When this persists the operator checks with the developer, and discovers that it wasn't meant to be accessible anymore. They check the user identity in the observables provided by Greymatter.io, which indicate the requesting user isn't associated with the project, and immediately alert security to a likely intrusion that would otherwise have passed unnoticed.

2. Service F, with a narrow purpose, is only ever used by an automated system, begins receiving very regular requests at particular times of day. Since there is little variation, Greymatter.io has learned to expect this regularity. When Service F begins receiving requests at random times of day, they are flagged as anomalies. The operator is disturbed to find that they correlate with otherwise normal sidecar egress anomalies for another sensitive service, indicating the possible compromise of the automated system credentials. The user alerts security while confirming the configured scope of the non-person credentials.

Summary

Service-based architectures and cloud-based capabilities are incredibly valuable advancements, but present challenges at-scale due to the speed of operations they enable and sheer volume of signals they emit. Greymatter.io's AI provides a cost effective means of continuous oversight at machine speed, taking advantage of mesh network operations unlike any other capability available today. Greymatter.io delivers on the promise of AI Operations, providing the enterprise an unblinking eye capable of monitoring operations with human-like insight at nearly imperceptible levels of detail, and extending the mental reach of human operators for optimal performance and uptime.

Connect With Us

+1 (887) 356-3011
info@greymatter.io
https://greymatter.io

106 N. Lee Street, Floor 2
Alexandria, Va, 22314
United States

GreyMatter.io @GreymatterIO https://www.linkedin.com/company/greymatterio/

For sales or further information, contact info@greymatter.io