DECISION MAKER'S GUIDE TO CHOOSING AN OBSERVABILITY PIPELINE

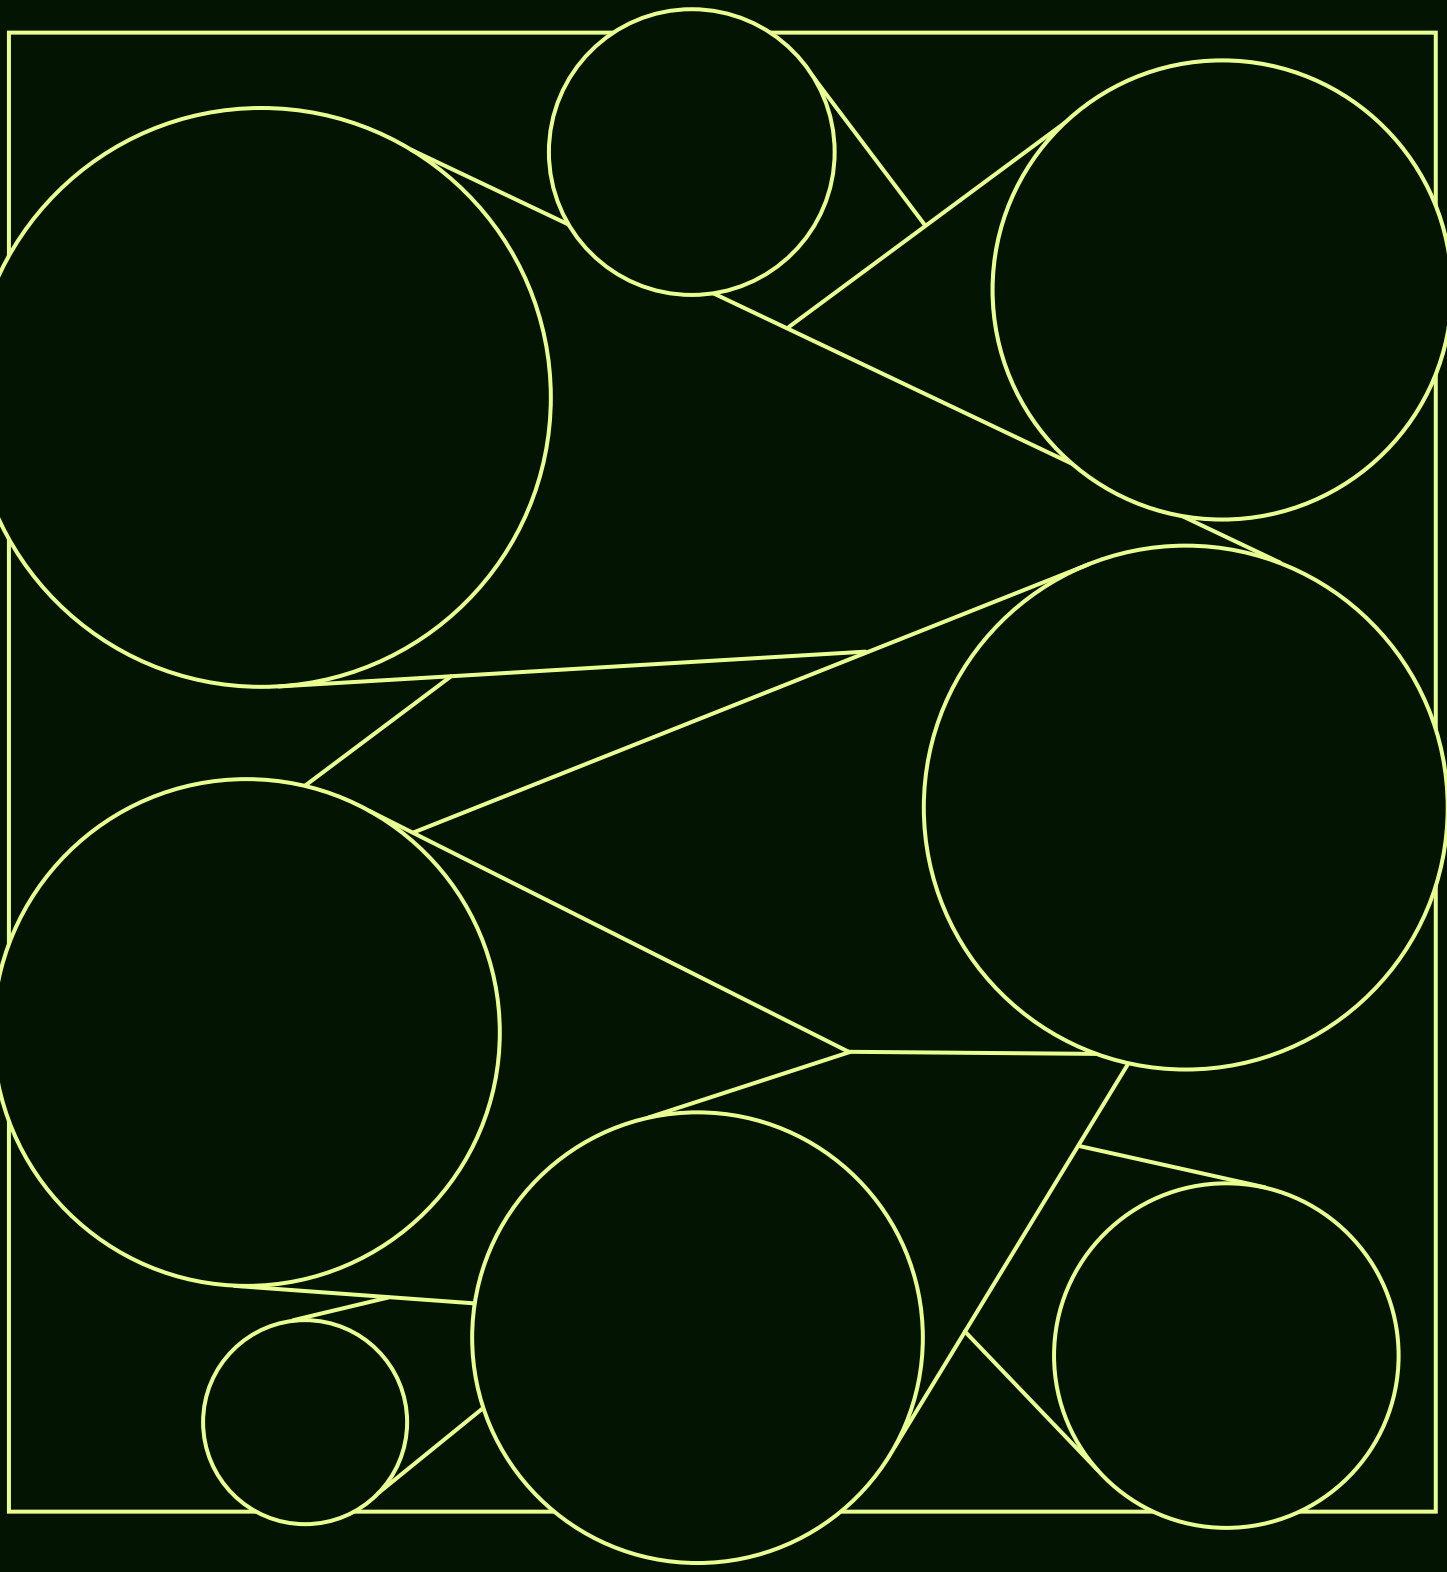

mezmo

WHITE PAPER

WHEN IT COMES TO ENSURING EFFICIENT, EFFECTIVE DATA-DRIVEN DECISIONS, RECOGNIZING THE VALUE THAT OBSERVABILITY PIPELINES CAN DELIVER TO YOUR BUSINESS IS ONLY HALF THE BATTLE.

The other half involves choosing an observability pipeline solution – a task that is not always simple given the many options available today. As a growing number of vendors and developers offer observability pipeline platforms – including both paid services and open source solutions – it can be challenging to understand the differences between the various options on the market and determine which best align with your business's needs.

We've prepared this white paper to help decision-makers conquer this challenge. In the following pages, you'll find actionable guidance about what to look for when choosing an observability pipeline solution in terms of features, flexibility, cost, and more. You'll also gain perspective on how observability pipeline functionality fits within your business so that you can make informed decisions about which pipeline features offer actual value and which amount to "fluff."

We offer this guidance in the form of five critical questions leaders should ask themselves as they evaluate observability pipeline solutions. The answer to each question plays a central role in determining the best observability pipeline solution for your needs.

WHAT IS AN OBSERVABILITY PIPELINE?

Before diving into guidance on assessing observability pipeline options, let's briefly define what observability pipelines do and how they provide business value.

An observability pipeline is a tool or process that centralizes the ingestion, transformation, correlation, and routing of data from across a business. The purpose of an observability pipeline, in other words, is to provide an efficient, cost-effective means of working with data from across a wide variety of sources, as well as transforming data to make it more actionable. Observability pipelines also ensure that data is made available quickly to whichever stakeholders need to consume it, in whichever format is most actionable for their varying use cases.

By implementing an observability pipeline, businesses gain the advantage of faster, more effective decision-making. It also becomes much easier to gain critical context on information by correlating multiple types of data with each other. They can even take action on data while it is still in motion. Beyond this, observability pipelines can reduce the total cost of ingesting, processing, and analyzing data through features that make it possible to optimize the way it is stored and interpreted.

Check out our white paper on the business value of observability pipelines for a deeper dive into their essentials and role in modern business.

HOW TO CHOOSE AN OBSERVABILITY PIPELINE: FIVE QUESTIONS TO CONSIDER

Again, there is no shortage of observability pipeline solutions on the market today. But not all offerings are made the same. Beyond apparent differences like whether solutions are commercial or open source, there are a variety of small but important distinctions between various offerings that determine how flexible they are, how easy they are to use, and more.

That's why it's vital to ask yourself the following questions as you evaluate observability pipeline solutions.

HOW OPEN IS THE PLATFORM?

Most observability pipelines promise to be able to collect data from a wide variety of sources and route it to many different destinations. In this sense, they claim to be flexible enough to support diverse environments and use cases.

However, what ultimately determines the flexibility of an observability pipeline solution is whether the technology that underpins it is open – meaning it is based on transparent, community-defined standards. Open technology ensures that observability pipelines can be adapted or extended to fit virtually any type of data or use case, resulting in businesses never being constrained to working with only certain types of data. Without an open approach to data, pipelines have a fundamental limit concerning which data sources and use cases they can support.

That's why it's wise to look for observability pipelines that leverage open source agents for ingesting data and that fully support open instrumentation frameworks like OpenTelemetry. These solutions offer maximum flexibility for enabling the data types and use cases your business requires today and those that may arise in the future.

HOW RAPIDLY DOES THE PLATFORM DRIVE ACTIONABILITY?

The integrations and policies that drive observability pipeline functionality are complex. However, that doesn't mean that teams should have to spend copious amounts of time setting up and managing pipelines before they can begin deriving value from them.

Instead, observability pipelines should be capable of driving actionable operations from the start. This is possible when the pipelines offer features for easy set up and integration, as well as the ability to control costs and analyze pipeline data in motion through search, alerting and visualizations.

This functionality is important not just because it saves time and effort for IT organizations, but also because it makes it easy to expand observability pipelines quickly across the organization in order to accommodate new use cases and business requirements.

WHICH DATA ENRICHMENTS ARE AVAILABLE?

In many cases, the raw data your observability pipeline ingests is of limited value on its own. To provide maximum insight, the data must be enriched.

There are many ways to enrich data. Enrichment could involve adding critical context to a data set by, for example, translating IP addresses within a log file to their associated host names to make it easier to interpret the data. Or, you could enrich data by removing redundant entries, a process known as deduplication. Sampling data to reduce total data set size (and, by extension, reduce the time and cost required to analyze data) is another example of data enrichment.

It's always possible to enrich data manually before analyzing it. However, that approach demands time and effort, making it impossible to analyze data in real time.

That's why the best observability pipelines enrich data automatically. Using algorithmic decision-making, or by providing easy-to-use configuration capabilities, they determine which additional information to add to a data set, or how to improve data quality to make it more actionable. In turn, they allow businesses to make decisions based on fully enriched data in real time without waiting on manual data enrichment before taking action.

WHICH CORRELATIONS DOES THE PIPELINE ENABLE?

On their own, individual data sets tend to deliver limited value. By correlating and comparing corresponding data across types and domains, teams can gain the deepest visibility into relevant trends and anomalies.

As with data enrichment, correlation shouldn't be a complex, tedious operation within the context of an observability pipeline. Observability pipelines should instead make it easy to set up correlations across data sets – and correlation processes should scale seamlessly as data volume increases. Simple, scalable correlation becomes possible when a data store exists alongside the pipeline in which correlations can take place.

In practice, the ability to establish correlations means that businesses should be able, for example, to correlate logs from multiple microservices to highlight for IT and SRE teams how the performance of one impacts the others. As another example, they should be able to link application performance data to customer engagement data so that CRM teams can determine which performance events trigger application abandonment.

DOES THE PIPELINE ENABLE DECISIONS IN TRUE REAL TIME?

It's easy to promise real-time decision-making, and most observability pipelines do. But in practice, a user's ability to make real-time decisions tends to vary. It hinges on the following factors:

PLATFORM OPENNESS AND USABILITY:
An open and intuitive platform makes it possible to enrich and correlate data quickly, while it is in motion. In turn, decision-makers can derive insights rapidly and in real time.

PIPELINE DATA SEARCHES:
The ability to search for data while it is in motion within a pipeline is critical for translating data into insights as quickly as possible.

ALERTING:
Configuring alerts allows teams to glean actionable insights and respond accordingly while data is moving through the pipeline.

REAL-TIME VISUALIZATIONS:
Visualizations also play a central role in helping decision-makers to take real-time action by making it easy for them to interpret data, even when it is still in motion.

Making decisions in real time is critical because delays of even a few minutes can prove problematic in many situations. For example, if you rely on an observability pipeline to detect cyberattacks, waiting minutes to identify a breach could mean that attackers will have already compromised mission-critical resources by the time you react. Or, failing to detect an application performance problem in real time during peak usage could translate to substantial losses in revenue if users abandon your app because of failed transactions or slow response rates.

CONCLUSION

In short, the best observability pipelines are those that deliver:

- **Data ingestion based on open technology, providing maximum flexibility for working with any type of data and supporting any and all use cases.**
- **An intuitive, user-friendly experience that allows any user in the business to make data-driven decisions.**
- **Automated data enrichments and correlations, ensuring that businesses can optimize the value of their data and the insights it provides.**
- **The ability to make real-time decisions by searching, alerting on and visualizing data in motion.**

Mezmo was designed from the start to provide these benefits. Using open source agents and automated data enrichment and correlation processes, Mezmo ingests, processes, and routes data from virtually any source to virtually any consumer. And when combined with native log analysis features to drive actionability while data is in motion, Mezmo is truly designed to help you set your data free. With more than 5 petabytes of data processed every month and over 28,000 active users, Mezmo empowers businesses across the globe with intuitive, real-time observability pipelines so that they can make fully informed decisions based on all data available to them.

To see Mezmo in action for yourself, request a demo with one of our solutions experts.

SALES CONTACT: OUTREACH@MEZMO.COM
SUPPORT CONTACT: SUPPORT@MEZMO.COM
MEDIA INQUIRIES: PRESS@MEZMO.COM

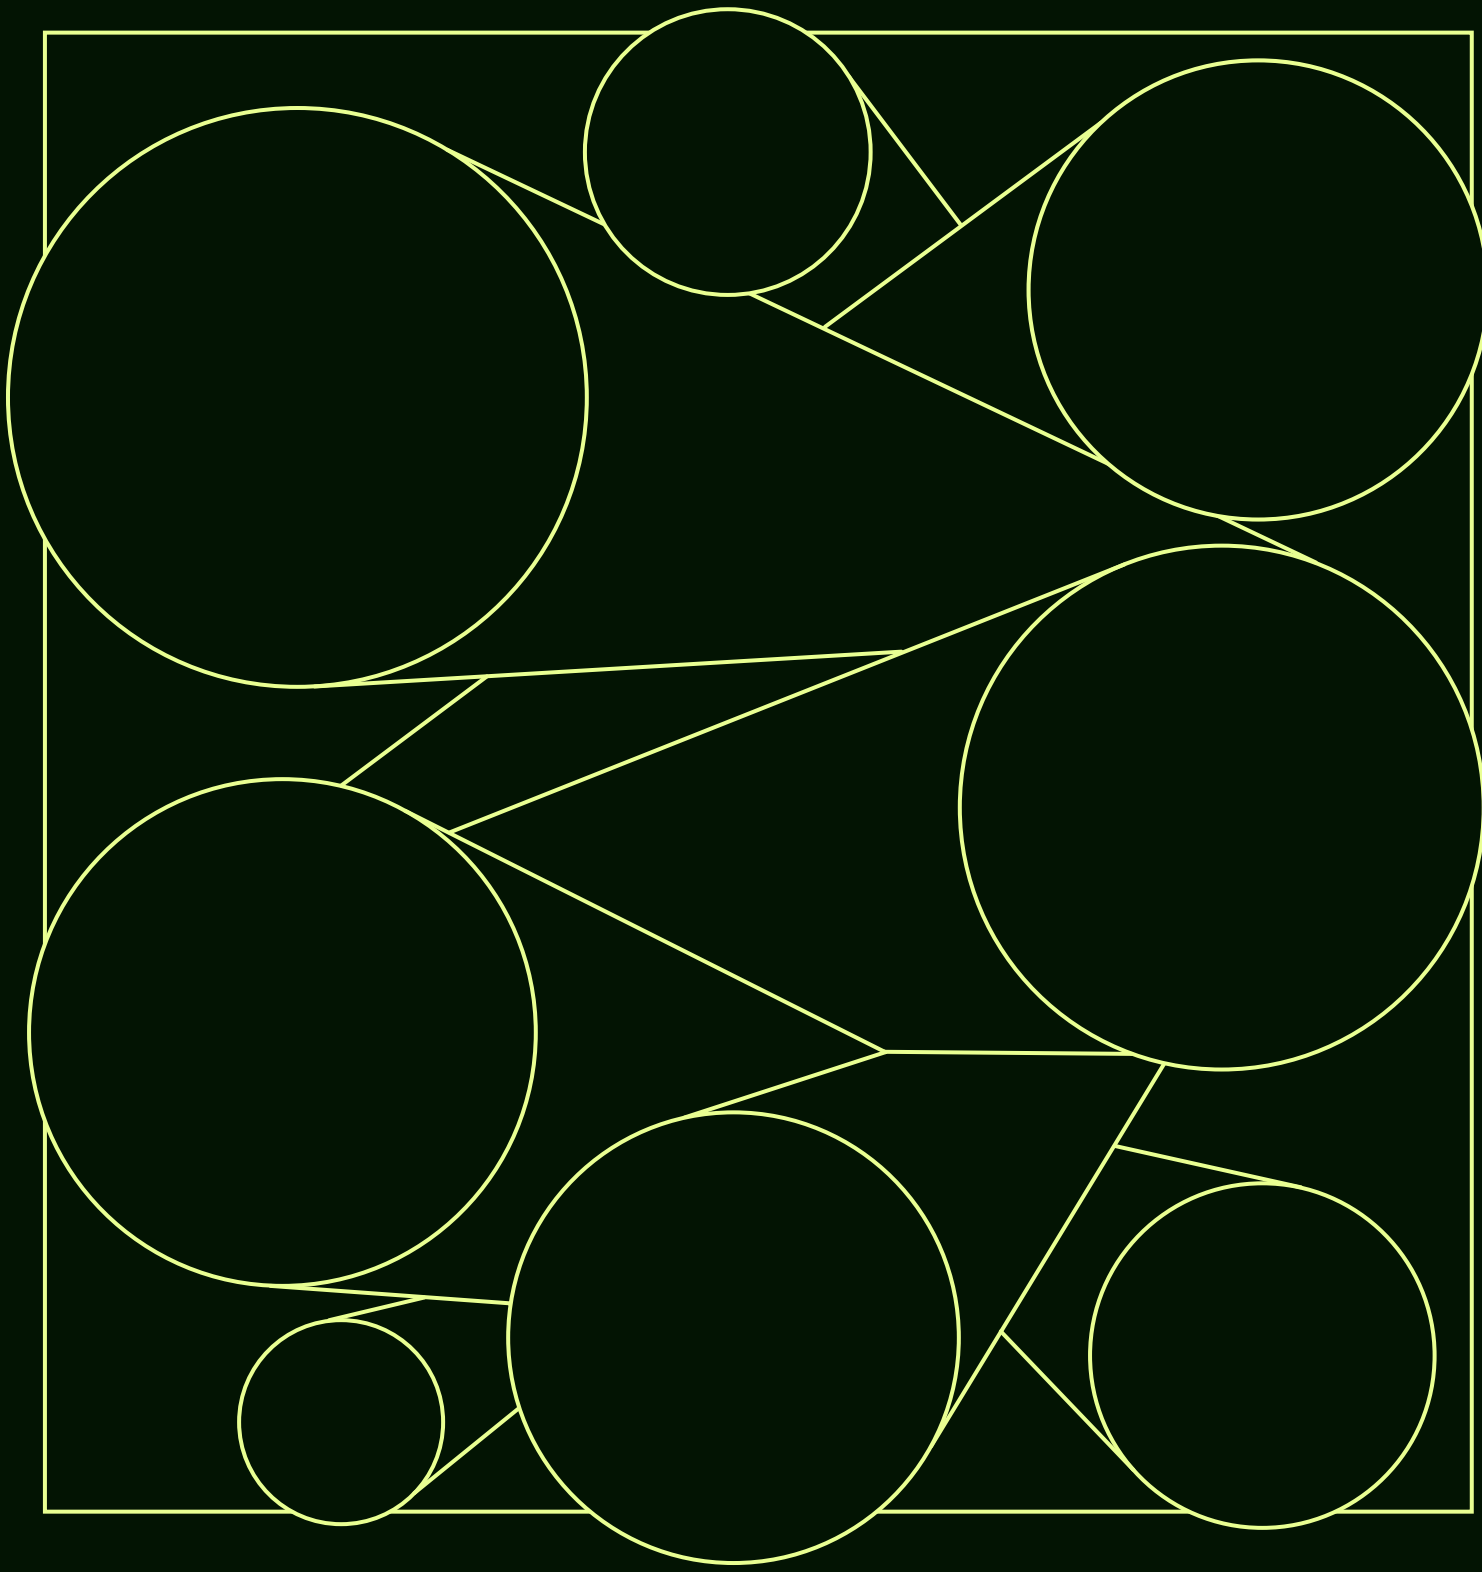

mezmo