RELIAQUEST®

The Ultimate Guide to **Threat Detection**

Everything You Need to Implement an **Effective Detection Strategy**

2024

Gold-Standard Detection for the Modern SOC

As adversaries continually evolve their attack techniques, maintaining an effective threat detection strategy is crucial for safeguarding organizational assets and data.

This design guide provides insights on detection orchestration, the most effective threat detection strategy that aims to address an organization's specific needs and challenges. By leveraging detection orchestration, security operations teams can stay ahead of threats with rapid, up-to-date, and comprehensive detections that are centrally managed.

By following this guide, you will be better equipped to design, implement, and maintain a modern security operations center (SOC) that not only meets today's security challenges but also anticipates and adapts to tomorrow's threats.

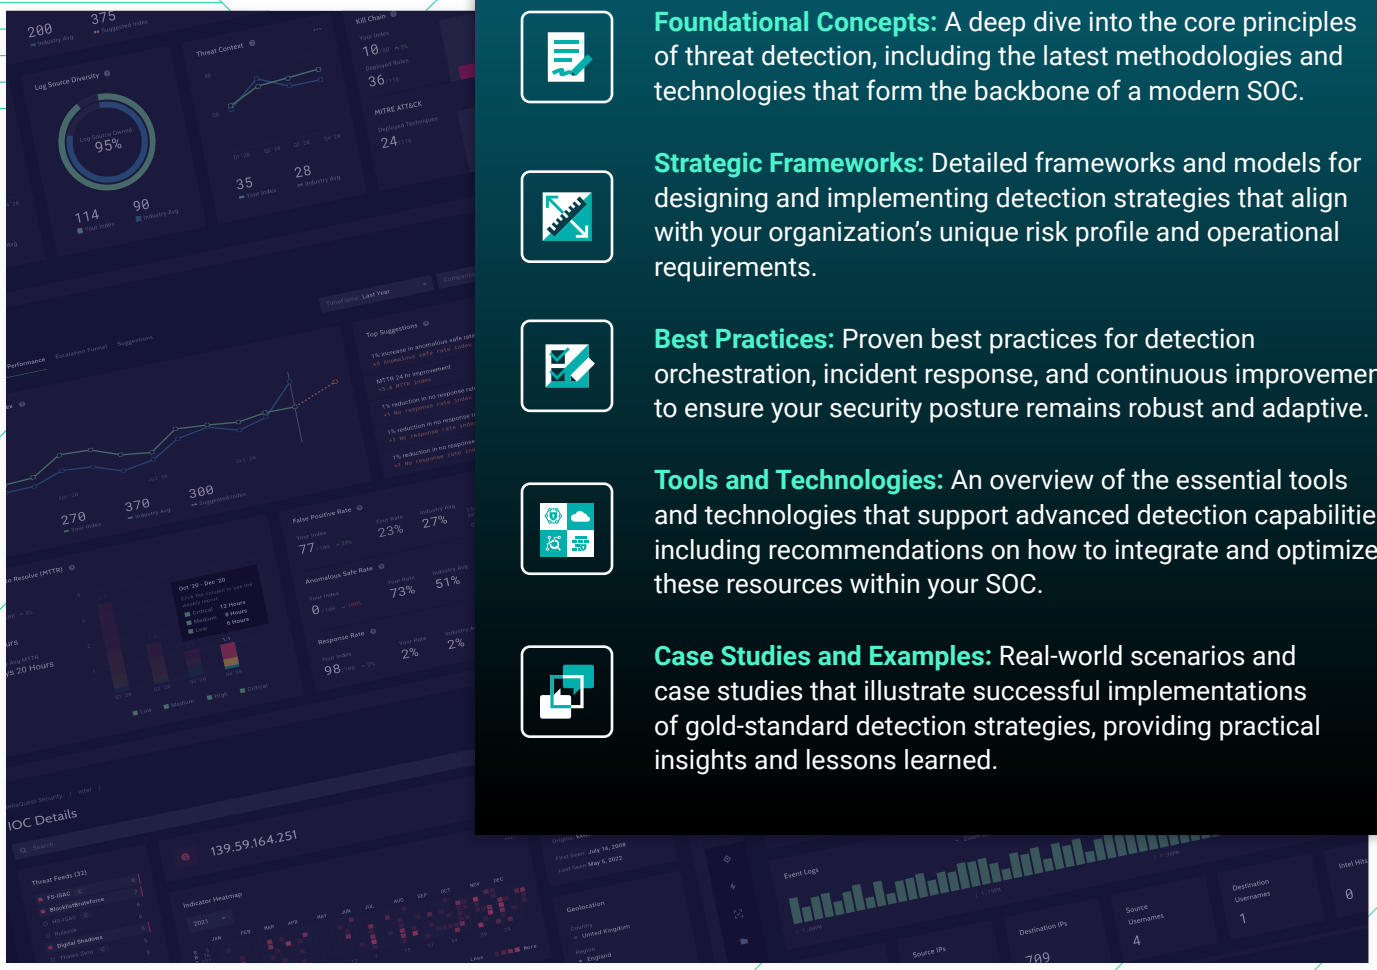

Within this guide, you will find:

Foundational Concepts: A deep dive into the core principles of threat detection, including the latest methodologies and technologies that form the backbone of a modern SOC.

Strategic Frameworks: Detailed frameworks and models for designing and implementing detection strategies that align with your organization's unique risk profile and operational requirements.

Best Practices: Proven best practices for detection orchestration, incident response, and continuous improvement to ensure your security posture remains robust and adaptive.

Tools and Technologies: An overview of the essential tools and technologies that support advanced detection capabilities, including recommendations on how to integrate and optimize these resources within your SOC.

Case Studies and Examples: Real-world scenarios and case studies that illustrate successful implementations of gold-standard detection strategies, providing practical insights and lessons learned.

Shifting Your Approach to Detection

The traditional approach to detection development typically involves crafting all the necessary logic directly within each detecting technology (SIEM, EDR, Cloud, etc.). However, this isolated approach quickly grows less effective in scalable architectures that have multiple detection technologies. Additionally, security operations teams that want the flexibility to swap out technologies will find this method limiting. Developing detections technology-by-technology can create extended detection development times that may result in coverage gaps, increasing the dwell time of threats within the network. This, in turn, leads to a higher mean time to resolve (MTTR). It also demands a distinct level of expertise for each type of technology, leading to resource restraint and excessive complexity for the security operations team.

A better approach is "detection orchestration." With detection orchestration, detections are designed once and then seamlessly deployed across a variety of technologies. Any updates to the logic or tuning adjustments are handled through the detection library that feeds the orchestrator. **This centralized management offers several benefits:**

- Easy detection building without various technology expertise
- Speed of deployment for immediate coverage
- Single source of truth for detection logic that reduces tool hopping
- Greater flexibility and scale when designing a detection architecture using multiple technologies
- Independence from the specific capabilities or limitations of the technology, enabling more sophisticated and tailored detection strategies

What Is Detection Orchestration?

Detection orchestration is the centralized coordination and management of the various components and systems involved in threat detection within an organization's cybersecurity infrastructure. It integrates and automates different security tools, technologies, and processes for streamlined threat detection, investigation, and response (TDIR) while measuring the effectiveness of detections across the platform. It incorporates the key elements of Detection-as-Code, detailed below, for a more streamlined approach to detection handling.

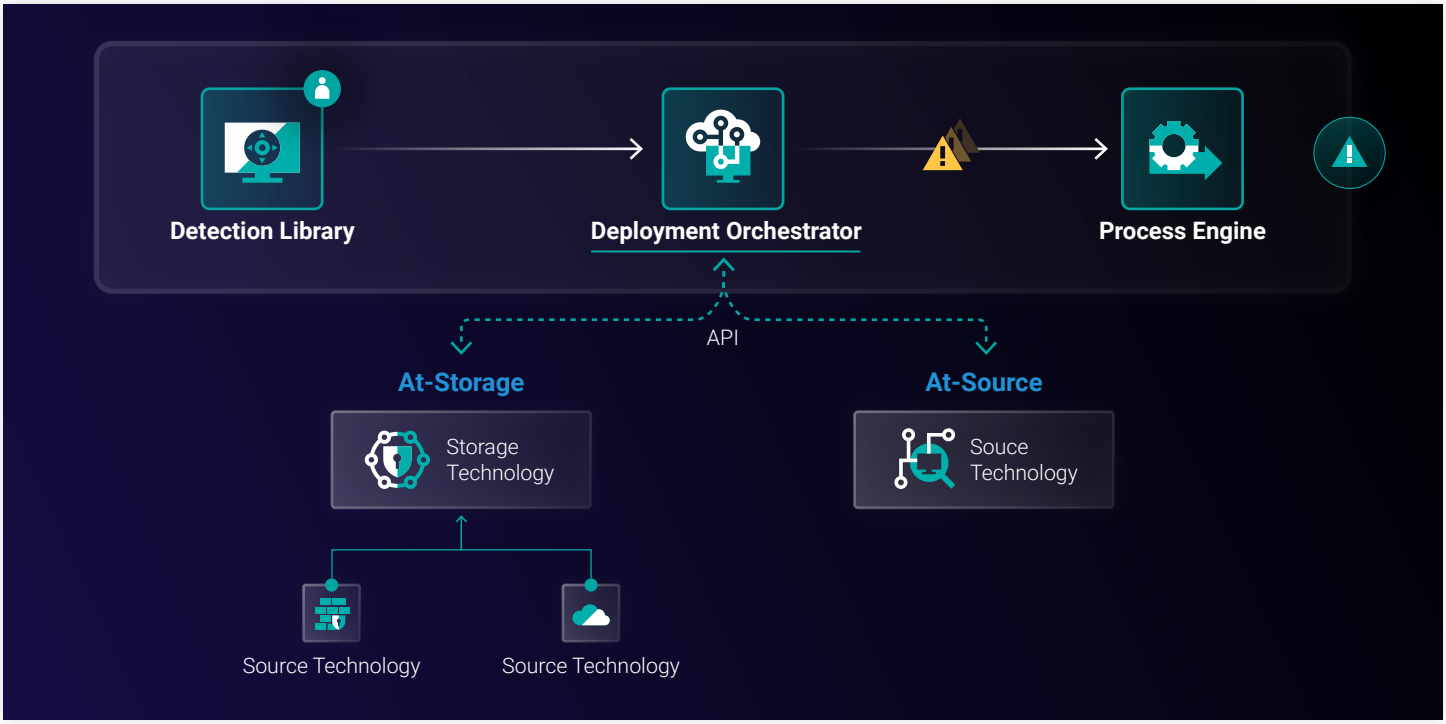

Key Components of Detection Orchestration

Detection Library

The detection library is where detections are added, modified, and deprecated. This allows for centralized management and transparent version control for detections.

A detection library consists of curated detections that are both technology-specific and threat-specific (tech-neutral).

Developed by detection researchers and tailored to the organization's needs, this library encompasses all detection authors and serves as the backbone of the orchestration process.

Technology-specific detections detect threats that only a particular tool or system can detect, such as using an iOS specific Mobile Device Management (MDM) technology to identify jailbreaking or rooting of devices that are bypassing security controls.

Threat-specific detections identify threats based on behavior and tactics, techniques, and procedures (TTPs) using various technologies. This could be detections for insider threats, compromised accounts, malware, or other harmful activities that could be detected using different technologies.

Orchestrator

An orchestrator is a pivotal technology that connects various security tools, threat intelligence sources, and IT service management (ITSM) systems. It deploys detection logic from its centralized detection library and ingests alerts for deduplication and analysis. This ensures seamless interaction and efficient operation across the security ecosystem.

Source Technologies

These technologies track and monitor activities for security purposes. Often referred to as log sources when their data is sent to storage technologies, they provide the necessary data for detecting potential threats. They are typically the source of truth as they are the closest security technology to the threat.

Examples: Endpoint detection and response (EDR), firewalls, cloud security solutions

Storage Technologies

Storage technologies are designed for the long-term storage of data reported by monitoring source technologies. They ensure that the data is readily available for analysis and threat detection.

Examples: Security information and event management (SIEM) systems, data lakes

Utilizing these components together provides a robust framework for centralized management and orchestration of effective threat detection within for an organization.

Process Engine

The process engine serves as the initial filter for alerts generated by the orchestrator before they are reviewed as incidents. This engine suppresses, de-duplicates, correlates, and consolidates related alerts into a single incident. By reducing noise and preventing multiple alerts for the same issue, it ensures a more organized and manageable alert response.

Detection-as-Code Foundations

Detection orchestration leverages “Detection-as-Code” to ensure an effective detection strategy. Detection-as-Code (DaC) brings the principles of software development to cybersecurity to create, manage, and maintain detection rules for identifying security threats and anomalies. The purpose of this approach is to bring the same rigor, collaboration, and automation found in software engineering to the process of writing and managing detection rules.

Key Pillars:

Collaboration

At its core, DaC encourages collaboration between security analysts, detection developers, and other stakeholders.

This involves peer reviews of detection rules to catch issues and improve quality and using a central repository where team members can contribute to new detection rules and collaborate on best practices.

Version Control

Using version control within detection management ensures that every change is tracked and can be audited. This allows for stakeholders to keep a detailed history of what changes that were made, by whom, and why. It also fosters better collaboration between multiple team members for merging changings simultaneously and seamlessly. With better version control, if a new detection rule causes issues, the organization can easily roll back to a previous version.

Documentation

Comprehensive and up-to-date documentation is essential for understanding and maintaining detection rules. Within DaC, it is required to keep detailed descriptions of what each detection rule does as well as how to deploy, test, and modify the detection rules. It’s also crucial to document the changes and updates that are made to the detection rules over time.

CI/CD Pipelines

Probably the most important pillar of Detection-as-Code is the ability to deploy detections from a centralized hub across the environment. Using continuous integration and continuous deployment (CI/CD) practices, detections can be automatically tested inline and updated based on changes. This deployment model is the key for modern SOCs to scale and maintain their detection rules efficiently.

Expanded Testing

Testing is crucial to ensure that detection rules function as intended without false positives or negatives. Applying the testing concepts from software engineering into cyber creates a more expansive view for testing detections. It requires the detection developer to not only unit test the validity of the logic, but also test how the detection will interact with other systems and ensure that there are no adverse effects.

Metrics and Monitoring

Monitoring the effectiveness of detection rules and gathering metrics to guide improvements.

In DaC, the number and types of alerts generated by detection rules are tracked, along with the accuracy of the detection rules to measure the performance and overall impact.

The Detection Engineering Lifecycle

The detection engineering lifecycle is a structured approach for detection engineering that aims to help an organization maintain a proactive stance against cyber threats while minimizing risks and enhancing their overall security posture. This lifecycle is fundamental to the detection design framework laid out in this guide.

There are four phases of the detection engineering lifecycle, each playing a critical role in the development, deployment, and optimization of detection capabilities.

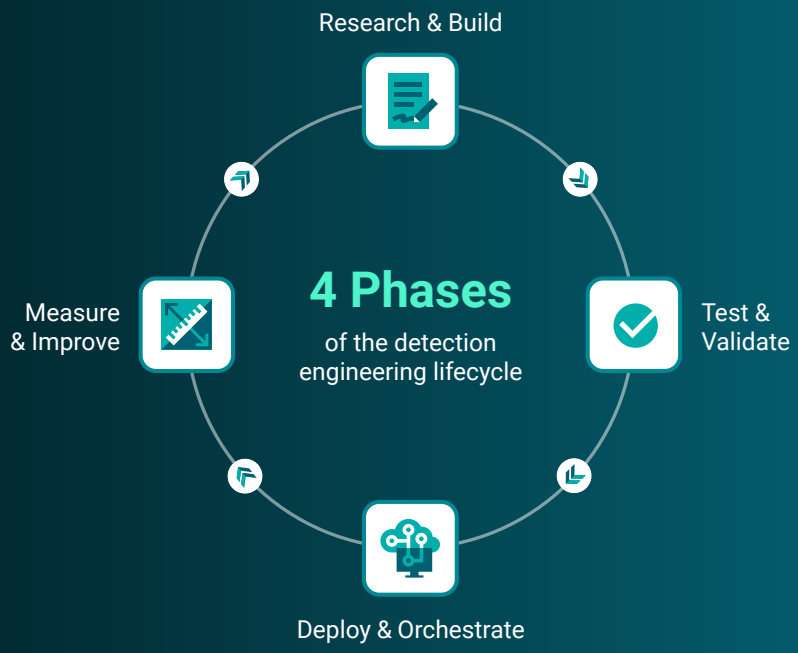

Research and Build

The first step involves identifying and understanding potential threats that could impact the organization by leveraging threat intelligence, cyber frameworks like MITRE ATT&CK, and historical data. Based on this research, robust detection logic and rules can be developed into a library tailored to the organization's environment. This step is crucial because it establishes a strong foundation, ensuring that all subsequent steps are informed by accurate, relevant, and well-researched detection logic.

Test and Validate

Once the detections are built, they should be rigorously tested and validated to ensure they accurately identify threats in real-world scenarios without generating excessive false positives. This stage involves several levels of testing, from basic syntax validation to advanced threat simulations. Testing and validation are essential because they confirm the reliability and effectiveness of the detections, reducing operational risks and ensuring that the detections are ready for deployment.

Deploy and Orchestrate

After validation, the detection rules can be deployed across the organization's security infrastructure, either directly at the source (e.g., EDR, cloud) or through a storage solution (e.g., SIEM, data lake). This step puts the validated detections into action, ensuring they work cohesively across various technologies to provide comprehensive threat coverage.

Measure and Improve

The final step involves continuously monitoring the performance of the deployed detections and analyzing metrics such as detection accuracy, false positives or negatives, and response times. Based on this ongoing evaluation, the detection logic can be refined and updated to improve its effectiveness and adapt to emerging threats. Continuous improvement is vital for maintaining high security levels over time, ensuring the organization stays ahead of evolving threats.

The detection engineering lifecycle lays the foundation for a resilient and adaptive defense strategy, essential for modern SOCs. Each phase of the lifecycle reinforces the others, creating a robust process that ensures the organization is always prepared to detect, respond to, and mitigate security incidents effectively.

The following sections of this guide dive deeper into each phase, exploring how detection orchestration plays a critical role in optimizing every aspect of the lifecycle. You'll find detailed insights, key focus areas, practical recommendations, specific guidelines, examples, and more—all designed to equip your SOC with the knowledge and tools needed to build or refine a comprehensive detection strategy.

Phase 1: Researching and Building Your Detection Library

With all the necessary components of detection orchestration in place, the first critical step is developing a comprehensive detection library tailored to your organization's needs. An effective detection library forms the backbone of a robust detection strategy, enabling your organization to respond swiftly to new threats without the need to start from scratch. By centralizing and continuously refining detections, you capture the expertise of various detection authors, ensuring a comprehensive, well-rounded defense.

This section will guide you through building out your detection library step-by-step, enhancing your overall security posture and resilience against an ever-evolving threat landscape.

Step 1: Prioritize Detections

With the vast array of threats and the type and number of detection authors involved, determining which detections to create first can be daunting. Organizations should always prioritize their detections based on the needs of the business to help them effectively allocate resources while addressing the most critical threats first.

Performing this analysis in advance maximizes the impact of the security operations team's efforts and optimizes security investments to yield the maximum return. A phased approach to detection also helps narrow focus, ensuring accurate tuning and reducing the noise from false positives, which allows the security operations team to focus on genuine threats and respond more efficiently while significantly decreasing the risk of successful attacks.

Key Areas to Prioritize:

Organizational risks: Understand the specific risks your organization faces to help you prioritize detections that protect the most-critical assets and processes that affect business operations.

Past incidents: Analyze previous security incidents to help identify vulnerabilities and threat patterns that have already affected your organization. Learning from past incidents allows you to implement targeted detections to prevent recurrence and improve resilience.

Penetration testing results: Penetration tests reveal weaknesses in your defenses by simulating real-world attacks, providing valuable insights into potential vulnerabilities. Address the findings from pen tests through detections to help close security gaps and strengthen your organization's defenses.

Compliance audits: Compliance with industry regulations and standards is essential for avoiding penalties and maintaining trust with stakeholders. Implement detections based on compliance requirements to ensure adherence to legal and regulatory obligations, reducing the risk of noncompliance. Addressing compliance requirements through prioritized detections also enhances trust with customers and partners.

Threat research:

Threat hunting results: Threat hunting proactively identifies unknown threats within your environment, uncovering sophisticated or stealthy attacks. Use the results from threat hunts to enhance detection capabilities by addressing threats that might otherwise go unnoticed.

Emerging threats: Stay informed about emerging threats to help the organization prepare for new attack vectors and techniques. Use both closed- and open-source threat intelligence to ensure your detection strategy evolves with the threat landscape and maintains its effectiveness against the latest threats.

Industry-specific threats: Different industries face unique threats. Instead of building detections that focuses on all industries, first focus on your specific industry and unique business model for more tailored and effective detections. This provides targeted protection against the most relevant threats to your industry, improving overall security.

Security technology detections: Leverage the prebuilt detections from your technology vendors to provide a quick and efficient way to cover common attacks specific to the technologies in use. This offers a baseline level of protection with minimal configuration, quickly and cost-effectively enhancing your detection capabilities.

Starting with these key areas and maintaining a focus on the organization's specific needs can help you develop a robust and effective detection strategy that enhances overall security posture and resilience against threats. Once the foundational areas are covered, you can move on to more advanced detection areas.

Step 2: Identify Necessary Data Sources

After prioritizing detections based on your business's risk profile and needs, next, you should ensure you have the data you need to create those detections. Effective threat detection relies on a variety of data sources to provide comprehensive visibility into an organization's environment. Utilizing diverse data inputs ensures that potential threats are identified quickly and accurately, minimizing risks, and enhancing security posture. Below are the key data sources required for effective threat detection.

Foundational Security Technology

The first data set to focus on are the foundational security technologies that will provide immediate visibility across the critical parts of your business. Implementing detections using these data sources first offers the most upfront coverage to reduce risk. Those technologies are:

- Operating systems (Windows, Linux, etc.)
- Network (firewall, proxy, DNS, etc.)
- Identity and access management (IAM)
- Endpoint detection and response (EDR / antivirus)
- Email security
- Business-critical applications
- Cloud security
- Industry-specific (operational technology, electronic health records [EHR], point-of-sale [POS] systems, etc.)

Historical Data

Historical data from past incidents and investigations can provide contextual insights that enhance accuracy and reduce noise. This helps identify patterns from previous incidents, improving the accuracy of current detections. It also provides a baseline for normal behavior, making it easier to detect anomalies. Use the data in your ITSM systems, or wherever historical data is housed, within detections to detect repeat offenders and similar attack methods and respond to new threats faster.

Threat Intelligence

Threat intelligence provides timely, relevant, and actionable information about threats and adversaries that enhances threat detection effectiveness. It is crucial to use this data in your detections to uncover emerging threats and enrich fidelity and accuracy. Threat intelligence is everywhere, so it's important to start with broad coverage and then incrementally bring in more granular threat intelligence. We recommend implementing threat intelligence in this:

- Indicators of compromise (IOC) threat feeds
- Industry-specific threat intelligence (e.g., Information Sharing and Analysis Centers [ISACs] data)
- Threat profiles (threat actors, campaigns, ransomware-as-a-service [RaaS] groups, etc.)
- Deep and dark web intelligence

Business-Specific Data

To ensure heightened monitoring and protection for high-value targets within the organization, it's important to encompass business-specific data, like critical assets or executive accounts, within your detection strategy. Below is a list of common business-specific data that should be included to prevent high-impact breaches:

- Critical assets
- Executive usernames
- Brand assets
- Registered domains
- IP ranges
- Business operational knowledge

Step 3: Determine Detection Authors and Complexity

Once you've prioritized detections and have all the required data to start building, you need to determine who will create the detections using your data. Detection authors are responsible for creating the logic that identifies specific threat activities. For a comprehensive detection strategy, it is essential to include contributions from three main types of detection authors.

Because each type brings unique benefits and perspectives, ensuring robust and diverse threat detection coverage, it is crucial to include detections from a combination of authors. That said, always ensure they are utilizing the required data to build out a detection library that aligns to your priorities. Multiple authors will help you achieve a comprehensive, multi-layered detection orchestration strategy that is adaptable and capable of addressing diverse and evolving threat landscapes.

Detection Complexity

Different types of threats require different detection approaches. For instance, advanced threats like fileless malware and Living-off-the-Land (LotL) attacks require sophisticated detection capabilities such as behavioral analysis. In contrast, detecting a brute-force attack against a Windows machine is more straightforward and can be handled with simpler detection methods. Therefore, the detection library of modern SOCs should encompass not only a variety of detections from multiple authors, but detections with varying levels of complexity to ensure comprehensive coverage against evolving and diverse threats. Detection complexity can be range significantly, however every detection can fall under one of three categories of complexity: basic, intermediate, and advanced.

Basic Detections

Basic-level detections rely on straightforward, well-defined methods to identify known threats. These methods typically reference only one technology source to capture a specific threat.

Benefits:

- Easier and faster to build
- More efficient at identifying common attack tactics

Detection Types:

- Signature-based detections: Using known signatures of malware and threats
- Rule-based detections: Using widely known patterns to flag common suspicious activities

Example: (Successful brute force)

- Data sources: Windows OS
- Logic: Three failed logins followed by successful login within three minutes

Considerations:

- Higher chance of false positives
- Easily evaded by sophisticated attacks

Intermediate Detections

Intermediate-level detections incorporate additional data sources, use threat intelligence to uncover sophisticated threats, and allow for cross-correlation between tools. Intermediate-level detections may detect multi-step, multi-tool activities.

Benefits:

- Identifies more evasive threats
- Higher accuracy and fidelity than basic level
- Verifies threat based on gathered intel

Example: (Phishing link clicked followed by successful MFA)

- Data sources: Email security, threat intel, multifactor authentication (MFA) tool
- Logic: One clicked event where the URL was classified as Phishing by threat intelligence, followed by a successful MFA event from the same user within five minutes.

Detection Types:

- Threat-enabled detections: Detect when IoCs like an IP address, hash value, or domain collected from trusted threat intelligence are present in the network.
- Cross-correlation: Correlates events from multiple source technologies to identify potential threats. Detections can consist of multiple basic level detections or building blocks to generate an even higher-fidelity detection.

Considerations:

- Requires more time and skill, additional data sources, and collected threat intelligence

Advanced Detections

Advanced-level detections leverage cutting-edge technologies like machine learning and artificial intelligence (AI) to identify complex and evolving threats. They also often integrate cross-domain data for enhanced accuracy.

Benefits:

- Identifies more evasive threats earlier in the attack lifecycle
- Allows for faster adaptation to emerging threats
- Uses larger subsets of data for comprehensive visibility

Example: (Abnormal data transfer volume)

- Logic: Significant deviations in the amount of data being transferred to or from a system compared to the usual volume to the system or user.

Detection Types:

- Anomaly detection: Identifies deviations from normal behavior using statistical models
- Risk base: Leveraging data about users and assets to calculate risk profiles that help detect unwanted gaps

Considerations:

- Depends on an accurate baseline
- Requires a heavy amount of tuning for continuous learning
- Can generate high false positives if not tuned or using the right baseline
- Requires advanced expertise and time

Detection Authors

Technology Vendor Authors

Vendors of security technologies provide prebuilt, default detections specific to their technology. These out-of-the-box detections typically require minimal configuration and cover a broad range of common attacks detectable by the vendor's technology.

Benefits:

- Ready-made and quick to deploy
- Provides a baseline level of protection
- Cost-effective

Best for:

Basic-level detections

Considerations:

- Detection customization capabilities vary by technology.
- Not initially tuned to specific environments or industries, leading to high false-positive rates.
- Limited to threats only that technology can detect (i.e., email technologies are best to detect email attacks, but not lateral movement on the endpoint).

Best Practice Recommendations

- Focus on high-severity detections to ensure critical threats are addressed promptly. Filter out low-value detections such as blocked activities, PUPs (Potentially Unwanted Programs), PUAs (Potentially Unwanted Applications), and non-security events (e.g., policy updates).
- Customize and tune technology vendor–authored detections where possible to better fit your specific environment and threat landscape. Adjust detection thresholds and parameters to reduce false positives and improve accuracy.
- Regularly update detection rules based on the latest vendor releases and threat intelligence. This keeps your detection capabilities current and effective against emerging threats.
- Analyze detection effectiveness and adjust as necessary to maintain optimal performance and relevance.
- Leverage the documentation provided by vendors to understand detection logic, expected outcomes, and configuration options. This helps in better managing and customizing detections.

Internal Security Operations Team Authors

Your security operations team can create custom detections based on the organization's specific needs, such as compliance audits, previous incidents, penetration testing, and industry-specific attacks.

Best Practice Recommendations

- Document and standardize processes, including the rationale, expected outcomes, and any tuning parameters of detection logic. Standardize detection development processes to ensure consistency and repeatability.
- Collaborate cross-functionally with other departments, such as IT, compliance, and operations, to ensure detections are aligned with overall organizational objectives and are comprehensive.
- Continually improve detection logic based on new threat intelligence, changing business requirements, and feedback from incident response activities. This helps in maintaining relevance and accuracy.
- Provide ongoing training for detection researchers and engineers to help them stay current with the latest threat landscapes and detection methodologies.
- Invest in advanced tools that support custom detection development and testing.
- Rigorously test new detections in a controlled environment to validate their effectiveness and minimize false positives before deploying them. Use a variety of test scenarios to ensure comprehensive coverage.

Benefits:

- Tailored to the organization's environment and technologies
- Full control over detection logic
- Highly specific to organizational threats and requirements

Best for:

Basic- and intermediate-level detections

Considerations:

- Resource-intensive, requiring dedicated detection researchers and engineers with a wide range of expertise.
- Testing and ongoing maintenance demand additional expertise, tools, and time.
- Simple detections have a higher chance of false positives, and sophisticated attacks can easily evade them.

Third-Party Security Provider Authors

Third-party security providers—such as co-management providers, managed security service providers (MSSPs), or managed detection and response (MDR) services—also create custom detections. These detections add diversity and redundancy, addressing evolving threat actor tactics that may bypass vendor-specific detections and go beyond internal security operations team expertise.

Benefits:

- Broad and granular detection coverage across various tool stacks
- Tailored to the organization and industry-specific threats, providing a more targeted and accurate detections
- Crowdsourced from other companies and industries, enhancing detection quality and relevance
- Reduces the need for extensive in-house expertise
- Extends security operations team while freeing up resources for other critical tasks
- Performs detection testing, validation, and measurement

Best for:

Intermediate- and advanced-level detections

Considerations:

- Requires investment and commitment to a third-party partner
- Detections may not be transferable if contract is terminated or expires
- Detection logic is not always shared
- Certain providers may work only within a certain toolset
- Longer time to coverage if provider doesn't use detection orchestration

Best Practice Recommendations

- Select a partner with detection orchestration capabilities and a strong track record in providing effective threat detections that will be tailored to the organization.

Phase 2: Testing and Validating Detections

Having a large library of detections means nothing if they don't work or if they produce high volumes of false positives. It will only give an organization a false sense of security while slowing down their response times. Therefore, it's important to test and validate that detection logic is effective, reliable, and capable of identifying the intended threats in real-world scenarios without excessive false positives or negatives.

There are four stages of testing detections, each more complex than the next, with a specific objective and methodology for validating:

Stage 1: Syntax Validation

Objective: Ensure that the detection logic is syntactically correct and error-free.

Process: Verify that the detection rule translates accurately into the native language of the security tool from the orchestrator (e.g., SIEM query language, EDR rules).

Validation criteria:

- The logic executes without syntax errors.
- The logic aligns with the intended threat identification or behavior.

Outcome: Confirmation that the detection rule is correctly formulated and deployable.

Stage 2: Data Visibility Verification

Objective: Confirm that the required event types and data sources are available and integrated for the detection logic to function. Even if the logic is flawless, a detection will not trigger if the required data is not available.

Process: Validate that the necessary data types (e.g., logs, events) are available for the detection platform to query and generate a detection.

Validation criteria:

- The data sources are correctly configured and ingested.
- The event types required for detection are present and correctly formatted.

Outcome: Confirmation that the detection rule has the necessary data inputs to function effectively.

Stage 3: Threat and Attack Simulation

Objective: Test the detection logic against real-world scenarios to validate its effectiveness in identifying genuine threats.

Process: Simulate the threat in a controlled environment or use breach and attack simulation tools to replicate the threat. This tests the detection's ability to identify and respond to the threat as expected.

Validation criteria:

- The detection triggers correctly upon threat simulation.
- The detection logic identifies the threat with minimal false positives or negatives.

Outcome: Validation that the detection logic is operationally effective and can accurately detect the intended threat in real-world scenarios.

Recommendation: Regularly run simulations to validate the ongoing effectiveness of your detection rules as the threat landscape evolves. Be sure to test across different environments—including cloud, on-premises, and hybrid infrastructures—to ensure coverage.

Stage 4: Operational Validation

Objective: Validate the detection's performance over time in the live environment.

Process: Monitor the detection after deployment to ensure it operates as intended and does not produce excessive false positives or negatives. This stage involves continuous assessment and feedback loops.

Validation criteria:

- The detection consistently identifies threats with high accuracy.
- False positives and false negatives are within acceptable limits.

Outcome: Verification that that the detection remains effective and relevant in the operational environment, enabling continuous refinement and adjustment.

Each of these stages builds on the previous one, ensuring that the detection logic is not only syntactically correct and supported by the necessary data but also capable of accurately identifying threats in both test and live environments.

This structured approach ensures the detections within the library are reliable, actionable, and effective in protecting the organization against real-world threats.

Phase 3: Deploying and Orchestrating Detections

After rigorous testing, the next step is deploying the detections from your library across the organization. Detection orchestration simplifies this process by allowing you to build a detection once and deploy it remotely across the environment. By centralizing and streamlining deployment, the process maintains consistency and scalability while freeing up teams to focus on deeper analysis and response. The entire deployment process is facilitated by the orchestrator, which pulls detection logic from the library to run remote queries at the designated location for a new detection or ingest any existing triggered alerts, ensuring comprehensive and efficient threat detection.

At-Storage vs. At-Source Detections

To maximize the effectiveness of your detection strategy, you'll need to determine whether to execute detection activities at the source technology or through a centralized storage tool. There are advantages and disadvantages to both approaches, which we'll go into below.

Option 1: Detecting at Storage

For at-storage detection, security operations teams can centralize the necessary data in a SIEM or data lake. The detection orchestrator will then execute the relevant detection logic from the detection library to the storage tool and generate an alert if the queries trigger a result.

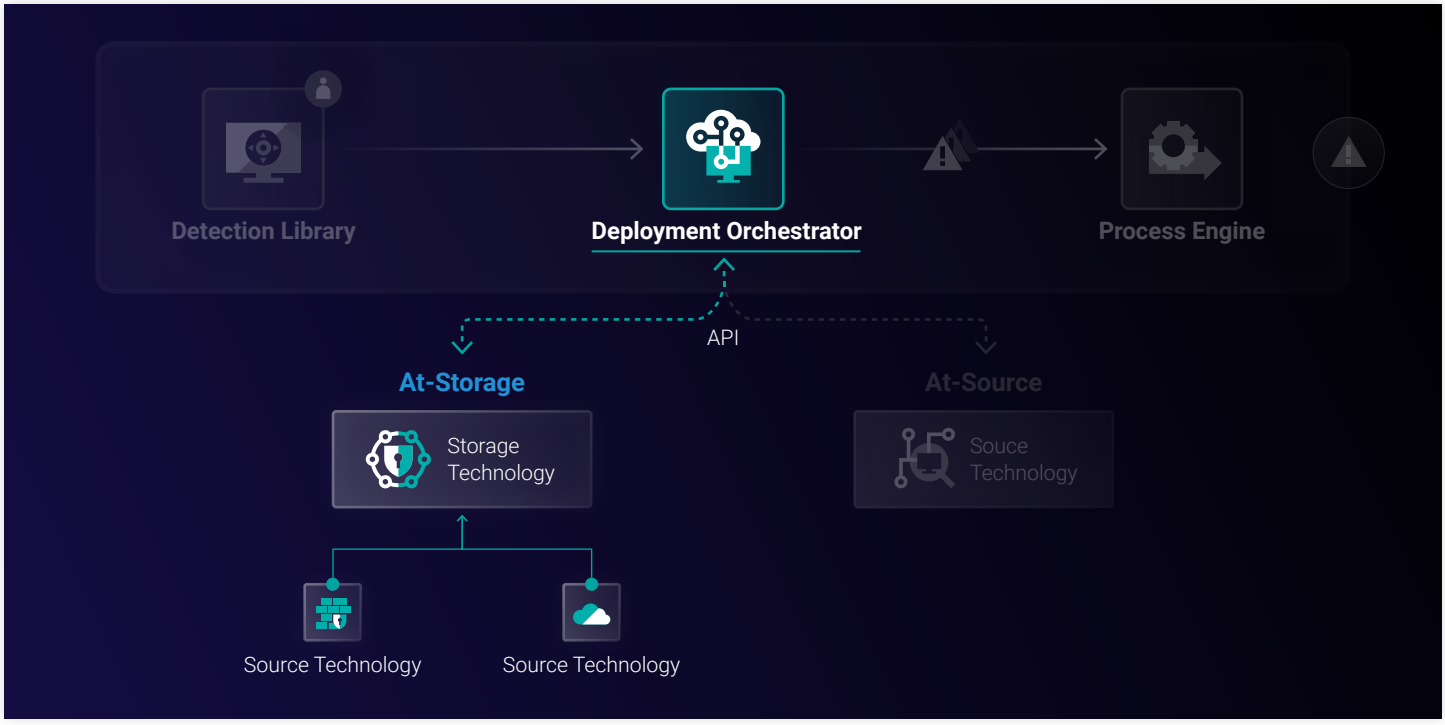

Benefits:

- Allows for more-complex detections
- Built-in correlation capabilities available in certain technologies
- Helps to meet compliance requirements

Best for:

- Intermediate-level detections (cross-correlation, intel-enabled).
- Custom applications or compliance-driven detections.
- Detections requiring high volume of data (operating system, raw network telemetry, etc.).
- When a source technology's APIs are not robust enough for at-source detection.

Considerations:

- Ingesting various data sources comes with additional cost.
- New detection use cases require additional ingestion, impacting costs, deployment time, and resources.
- Misconfigured or inaccessible log sources can slow down detection capabilities, drawing out mean time to detect.
- Incomplete or compromised logs due to multi-hop normalization or field parsing issues can create security coverage gaps.

Option 2: Detecting at Source

Security operations teams can also execute detections directly at the source technology reporting the activity, bypassing the need for a storage tool. In this case, the detection orchestrator will pull the detection logic from the library and execute remote queries at the source technology, like an EDR or cloud tool.

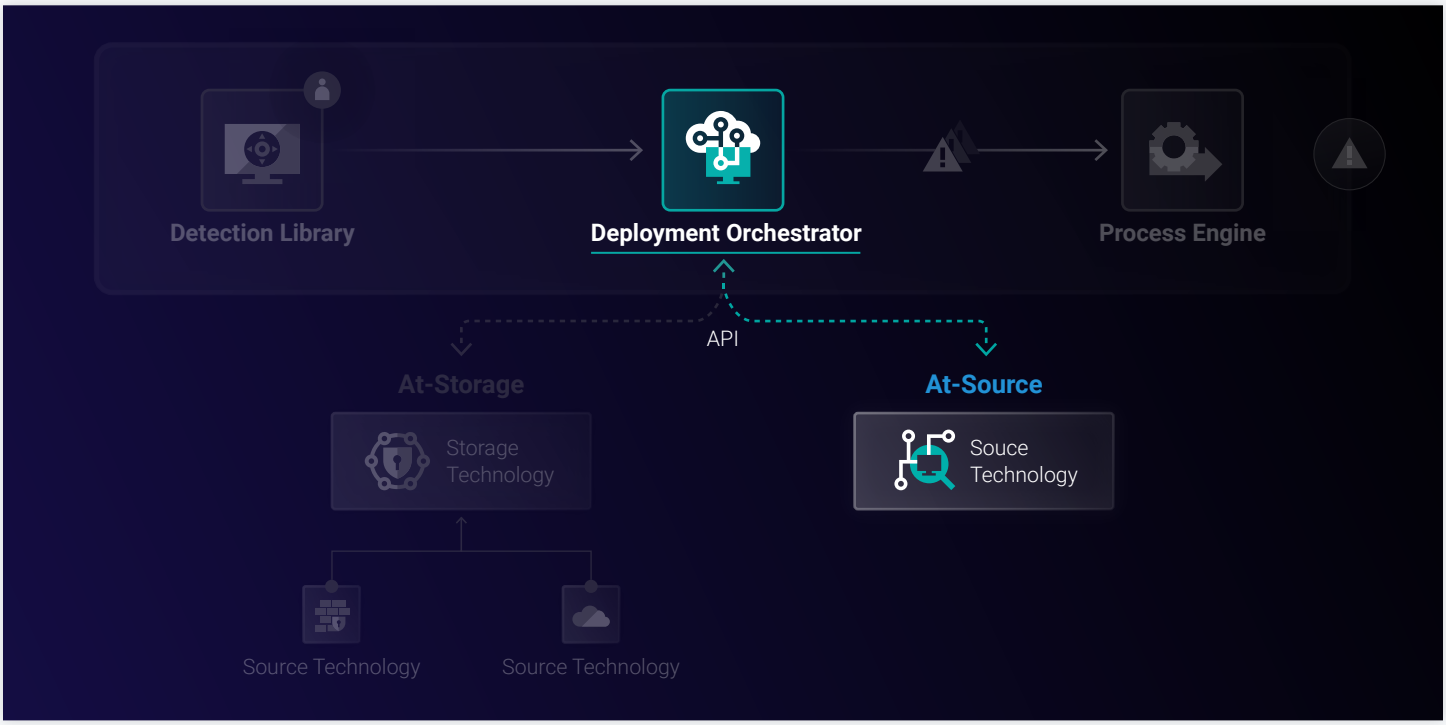

Benefits:

Low latency	Cost efficient	Avoids data quality issues through one-time data normalization

Provides optionality and scale: supports quick replacement or addition of source technologies while maintaining consistent detection coverage

Best for:

- Basic-level detections (technology vendor authored, signature-based, etc.).
- Non-cross-correlation, advanced-level detections.
- When a source technology has robust APIs.

Considerations:

- Some technologies may charge querying fees.
- Limited capability for more advanced detections.

Deciding on the Right Deployment Strategy

The detection orchestrator plays a crucial role in enhancing threat detection by executing remote queries either at the storage level or directly at the source. Detecting at storage offers comprehensive capabilities for more complex detections and compliance but comes with higher costs and scalability challenges. Detecting at the source, on the other hand, provides lower latency and cost efficiency, suitable for simpler and more straightforward detections. Given the two options, it's best to understand the approach that a modern SOC should adopt. The specifics can vary from company to company, but the key is to leverage the strengths of both options for the most effective strategy. Assess each technology and make the determination using the guiding questions below:

1. Does the technology's APIs support robust querying?

If the technology has robust APIs for remote querying, it's a prime component for at-source detection. If not, the data required for detection should be sent to a storage tool.

2. What risks or threats does this technology help to detect?

Not every technology detects the same threats. Understanding which threats and risk does the technology help to uncover will let you know which detection you want, and the complexity level needed.

3. What other use cases are available for this technology's data?

Other than threat detection, there are many other use cases for a technology's data—investigations, data enrichment, threat hunting, compliance, reporting, etc. It's important to understand these use cases to help determine whether you need the data stored long-term, if it's OK to be ingested on-demand directly from the source, or if storing only a partial amount of the data for those specific use cases is required.

4. What are the costs associated with detecting at source vs detecting at storage for this technology?

Regardless of tech, **the most effective threat deployment strategy** should be deployed using the following structure:

Basic-level detections: **At Source**

- Signature base
- Simple rule base

Ingestion of technology vendor–authored detections

- Basic level
- Advanced level

Intermediate-level detections **At Storage**

- Cross-correlation
- Threat intel–enabled

Phase 4: Measuring and Improving

The final phase of an effective detection strategy focuses on measuring the effectiveness of your detections to ensure they are both actionable and reliable. This involves a thorough assessment of coverage using actionable metrics that will lead to continuous improvement. Below are the three steps involved in measuring detection effectiveness.

Step 1: Align to a Framework

To gauge the effectiveness of your detections, utilize established frameworks that provide a structured approach to evaluating detection capabilities and identifying security gaps. Frameworks like MITRE ATT&CK, the NIST Cybersecurity Framework (CSF), and industry-specific models offer comprehensive taxonomies of tactics, techniques, and procedures (TTPs) used by adversaries.

Align your detections with these frameworks to assess current coverage and pinpoint areas needing improvement. Additionally, cybersecurity models such as the Cyber Kill Chain (CKC) by Lockheed Martin, the Pyramid of Pain, and the Diamond Model of Intrusion Analysis provide insights into attack methodologies relevant to your industry. These models help in identifying weaknesses in your detection strategy and guiding future enhancements.

Step 2: Develop Key Metrics

Once the detections are aligned to a framework that works best for your business, develop and monitor a set of key performance indicators (KPIs) to evaluate the effectiveness of your detections comprehensively. These metrics should provide clear, relevant, and timely information about the overall detection strategy. This helps in guiding future detection efforts and prioritizing actions. Essential metrics to track include:

- **Coverage and visibility:** Measure alignment with frameworks to understand detection gaps.
- **Accuracy rate:** Track false positives, true positives, and false negatives to assess detection precision.
- **Attack simulation pass rate:** Evaluate the effectiveness of detections through simulated attacks.
- **Mean time to detect (MTTD):** Measure the average time taken to detect threats to gauge responsiveness.

Step 3: Continuously Improve

Regularly review and refine detection rules based on the metrics and feedback gathered. Using insights from attack simulations and real-world incidents, adjust detection logic and close identified gaps. Finally, implement a feedback loop where insights gained from metrics and real threats directly inform updates to the detection library. This ongoing process ensures that your detection strategy evolves in response to emerging threats and changes in the threat landscape.

The only way to ensure that your detection strategy is both robust and adaptable is to consistently measure the effectiveness of your detections. Establishing a repeatable measurement process helps continually enhance your organization's ability to identify and respond to threats effectively.

Conclusion

Detection orchestration represents the pinnacle of detection strategies for the modern SOC, offering a unified and efficient approach. By centralizing the management and deployment of detection rules, this methodology enables seamless implementation and updates across diverse environments—whether on-premises, in the cloud, or hybrid.

This means security operations teams can scale their organization by building detection logic once and deploying it universally, ensuring no coverage gaps when integrating new technologies or transitioning between different platforms. Additionally, detection orchestration enhances threat detection accuracy and operational efficiency which provides faster response times, reduced operational overhead, and a robust defense mechanism that adapts quickly to emerging threats.

Detection orchestration is the best approach for maintaining continuous protection and strengthening your organization's overall security posture.

The ReliaQuest GreyMatter Security Operations Platform

ReliaQuest exists to Make Security Possible, allowing security teams to detect, contain and respond to threats within minutes—anytime, anywhere.

Built on over a decade of experience, ReliaQuest is the only cybersecurity technology company at scale, at the enterprise, that provides optionality and modularity—making the customer the platform.

ReliaQuest's security operations platform, GreyMatter, empowers enterprise security teams to leverage their current or future technology stack to drive greater visibility and automation without the need to centralize data or standardize tools.

This allows security leaders to achieve outcomes specific to their business by:

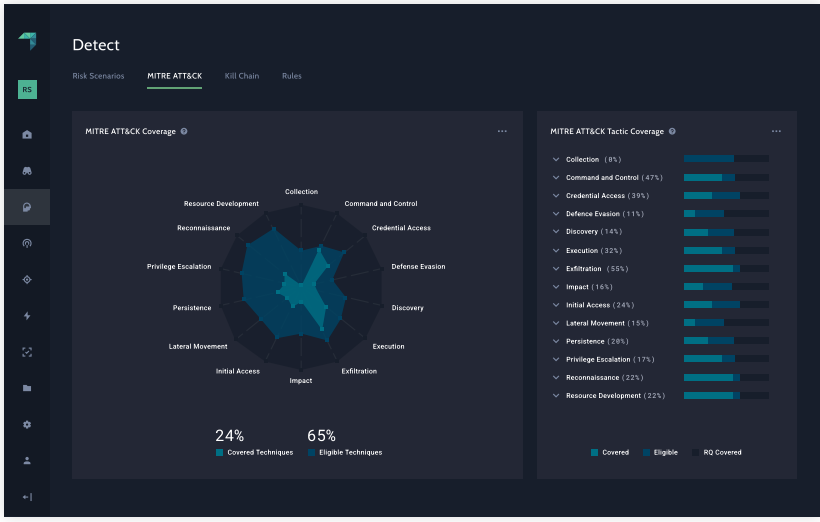

- Detecting threats no matter where their data lives, whether at-source or at-storage
- Containing threats faster with AI and automation
- Responding to threats within minutes
- Eliminating the need for Tier 1 and Tier 2 security operations activities

Case Study:

Leading Information Services Company Uses GreyMatter to Migrate Over 250 Detections in Under 48 Hours

The Problem:

A leading information services company was faced with the challenge of replacing their existing technology stack with the Microsoft Suite.

This transition necessitated the migration of over 250 finely tuned threat detections from Splunk and CrowdStrike EDR to Azure Sentinel and Microsoft Defender EDR, aiming to maintain robust threat detection without any disruption.

Migration Process Challenges:

- Manually recreating and testing hundreds of detections would be time-consuming and complex.
- The security team was still undergoing training on the new Microsoft technologies.
- The close expiration of existing technology contracts increased the urgency for expedited coverage.

Migration Process Solution:

To address these challenges, the ReliaQuest Detection team leveraged GreyMatter to review and repoint existing detections to the new Microsoft technologies.

They conducted extensive testing and validation to ensure existing tuning was preserved and alert volumes remained manageable. Following validation and approval, the detections were incrementally enabled within 48 hours until all were fully operational on the Microsoft technologies.

This seamless migration ensured continuous threat coverage, avoiding gaps and exposure to threats. It also streamlined the process, reducing the complexity and time required for manual recreation and testing. The security team was able to focus on their training without the pressure of managing the transition, ensuring their proficiency in the new technologies. The migration, validation, and deployment were completed within 48 hours, meeting the urgent need for expedited coverage and maintaining a scalable and adaptable security architecture ready to address future technology changes.

> "Our transition to the Microsoft suite was seamless thanks to GreyMatter. We were under pressure with expiring contracts and GreyMatter enabled us to migrate all our detections **in under 48 hours**, which allowed my security team to focus on learning our new technologies."
>
> **— The Company's CISO**

With over 1,000 customers worldwide and 1,200 teammates across six global operating centers, ReliaQuest delivers security outcomes for the most trusted enterprise brands in the world.

Learn more at www.reliaquest.com.

Appendix:

Examples and Use Cases

Below are several examples and use cases illustrating the value of detection orchestration:

Mergers and Acquisitions
When companies merge, it's critical to ensure interoperability and consistent threat detection across the different vendor technologies. Detection orchestration can integrate and coordinate threat detection across various vendor tools from both entities, providing rapid coverage with unified visibility and control.

Benefits:

- Quickly onboard and harmonize new security technologies without disrupting existing detection capabilities.
- Apply uniform detection logic across all integrated technologies, ensuring consistent threat detection standards.
- Centralized management provides comprehensive visibility into the entire security landscape, simplifying monitoring and incident response.

Hybrid Environments (On-Prem and Cloud)
Many organizations invest in both cloud technologies and on-premises technologies based on their functions and needs of the business. Maintaining consistent threat detection and response capabilities across diverse environments can be challenging. Detection orchestration can bridge the gap between on-premises and cloud security tools by connecting to both and deploying detection from a centralized location, providing unified threat detection.

Benefits:

- Ensure consistent detection rules and policies across on-premises and cloud environments.
- Simplify management and monitoring by centralizing detection orchestration.
- Easily scale detection capabilities as the organization expands its cloud footprint or adds new on-premises technologies.

Rapid Protection Against Emerging Threats
Keeping up with the fast-paced evolution of threat tactics and techniques can be challenging for any business. With detection orchestration, you can immediately deploy updates of detection logic based on new threats across all your technologies.

Benefits:

- Quickly implement and propagate new detection rules based on intel findings to address emerging threats.
- Leverage threat intelligence feeds and real-time updates to stay ahead of adversaries.
- Streamline the process of updating and managing detection rules from a single platform, reducing the burden on security operations teams.