ILTA'S QUARTERLY MAGAZINE

# PEER TO PEER

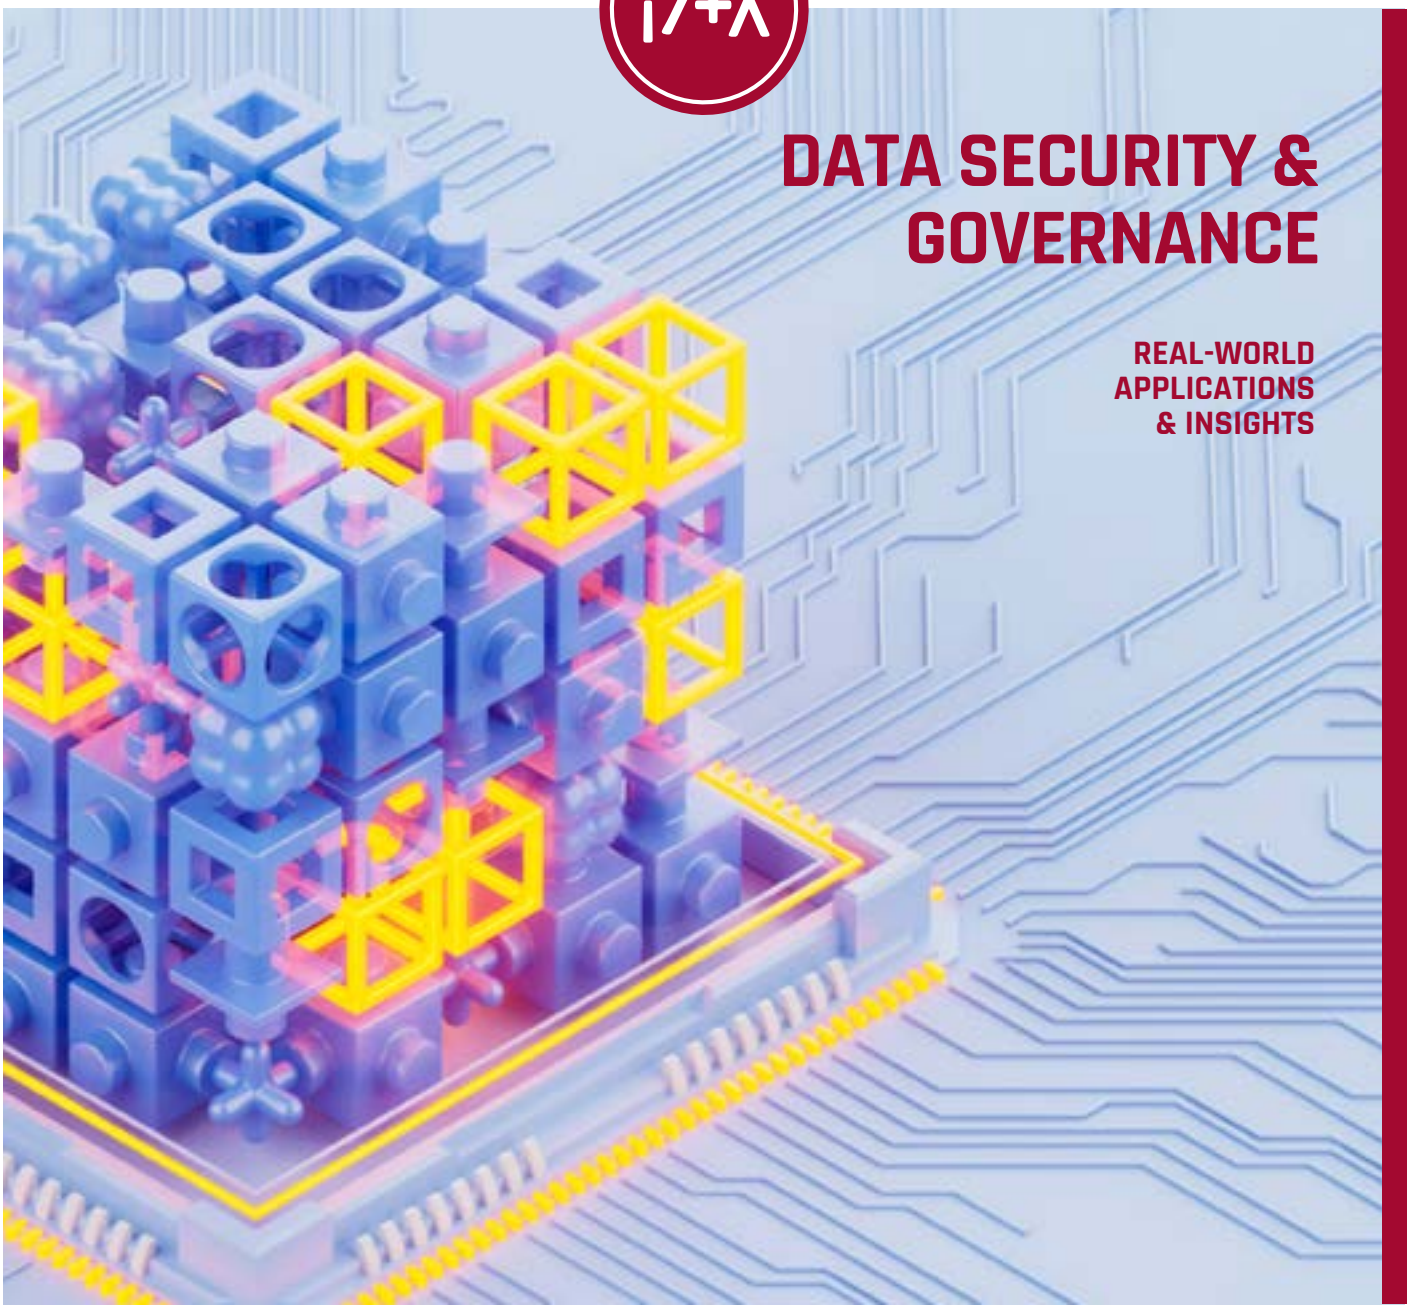

## DATA SECURITY & GOVERNANCE

REAL-WORLD APPLICATIONS & INSIGHTS

PEER TO PEER MAGAZINE / SUMMER 2026 / ISSUE 2 / VOLUME 42

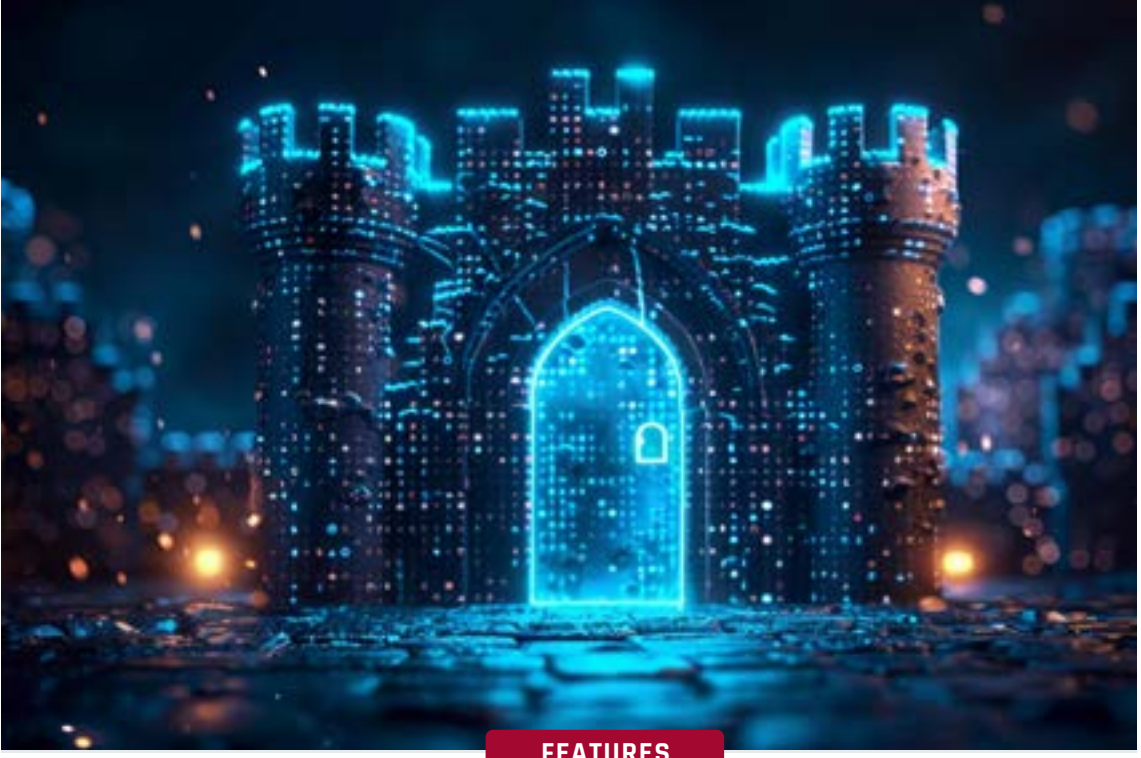

FEATURES

# WANT TO EFFECTIVELY EVALUATE MANAGED SOC PROVIDERS?

## LEARN HOW TO RUN YOUR OWN BAKE-OFF

by Ken Fishkin

Law firms are now major targets of nation-state and financially motivated threats because of the sensitive nature of the data they handle. According to IBM, the global average cost of a data breach has risen to $4.88 million. For professional services organizations, including legal, accounting, and consulting firms, the cost of a data breach is even higher, with an average cost of $5.08 million. As a result, security managers need to ensure that they have the appropriate resources to identify threats and respond to them as quickly as possible.

When it came time for Lowenstein Sandler to choose a new Security Operations Center (SOC) provider, we did not want to go with the traditional method of conducting Request for Proposals (RFPs), because this solution needed to be based on evidence, not marketing materials. It seemed that a bake-off model was the most appropriate way to determine which solution would work best in our environment.

This article provides a walkthrough of our experience evaluating five managed SOC providers, so other security professionals can use this method in their evaluations.

### Developing Success Criteria

Choosing a managed SOC provider can be a very daunting task if you are unfamiliar with the landscape. SecureWorld reports that operating a fully staffed 24/7 SOC in the United States costs approximately $2.86 million annually, excluding technology and training costs. As a result, outsourcing this service is rapidly growing and has produced a large number of vendors, making it difficult to differentiate genuine operational capabilities among them.

To determine which operational capabilities were most important to us, we identified three specific deficiencies with our incumbent managed SOC provider that we wanted to avoid with our new one:

- Incident detection and response rates were slow and frequently inconclusive, with alert triage that provided limited investigative value.

- An absence of meaningful AI-assisted capabilities, particularly around behavioral analytics and automated triage, that are now table-stakes features among leading providers.

- A Service Level Agreement (SLA) was vague regarding its responsiveness in identifying a potential incident.

As a result, we needed to ensure that vendors met our requirements in their SLAs for Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR).

**Why this matters:**
One potential competitor had marketing materials that claimed they would detect incidents within 5 minutes, but their actual SLAs were much higher than their competitors'. One also needs to distinguish contractual SLAs from marketed Service Level Objectives (SLOs) in a scoring rubric. Score only what vendors have committed to; disregard aspirational targets that carry no contractual consequence.

## Choosing the Participants

Previously, we relied on the vendor to supply security tools for identifying and responding to incidents. At the time of the bake-off, we had complete ownership of the tools to be used, such as Endpoint Detection and Response (EDR) and Security Information and Event Management (SIEM). Unfortunately, some vendors only had solutions based on utilizing their own security tools, so they were not considered for this bake-off. Through a combination of peer referrals at the Legal Services Information Sharing and Analysis Organization (LS-ISAO) and my own research, we were able to develop a short list of five vendors that met all of our requirements.

**Why this matters:** I wanted a mix of competitors along with some niche ones that showed promise. Word-of-mouth recommendations are great, but we wanted to give some lesser-known vendors with unique services a chance to compete as well.

## Designing the Evaluation: Simulating Real Threat Actor Behavior

The centerpiece of the bake-off was the red-teaming exercises conducted to audit the providers. We needed the tests to be effective, realistic, fair, and timely, so we hired a penetration tester to red team each vendor. Due to the high cost of retaining a pen tester, each vendor was given only one day to demonstrate their services. To do these tests safely, we added two test laptops and two test accounts to our existing network.

In that timeframe, they were all given the same tests:

- **Active Directory enumeration**: simulating an attacker mapping the domain environment following initial access, a technique observed in the vast majority of enterprise intrusions, according to MITRE ATT&CK data.

- **Risky sign-in attempt**: testing detection of anomalous authentication behavior consistent with credential compromise or adversary-in-the-middle attack patterns.

- **Deployment of multiple command-and-control (C2) payloads**: testing endpoint and network detection of post-exploitation tooling across multiple C2 frameworks.

- **External SSH connection establishment**: testing detection of unusual outbound connectivity from an internal host, a common indicator of tunneling or data exfiltration staging.

- **Simulated ransomware attack**: the highest-stakes scenario, testing detection speed, alert quality, and escalation response under conditions that approximate the most consequential threat facing law firms today.

Vendors were not given the specific scenarios in advance, though they were informed that red-team testing would occur. Each vendor's detection timeline, alert content, and escalation behavior were monitored throughout scenario execution and scored against the defined rubric.

**Why this matters:** We wanted to test the skill level of our competitors. Some tests can be easily detected by our tools, but others require a review of our SIEM logs to identify unusual behavior.

## Building a Defensible Process

The governance framework was designed explicitly to be audit-defensible and capable of withstanding scrutiny from firm leadership, outside counsel, or any other internal or external reviewer.

Key governance elements included:

- Defined Rules of Engagement establishing the precise scope and boundaries of red team activities, including what was and was not permitted during testing.
- Mutual Non-Disclosure Agreements (mNDAs) executed with all five vendors prior to any substantive engagement.
- Identical test conditions applied across all participants, including sequencing, timing parameters, and the information provided to each vendor about the environment.
- Contemporaneous documentation of all scoring decisions, including the rationale for judgment calls made during the evaluation.

**Why this matters:** In order to prevent any type of bias, having hard metrics to back up your final decision is essential.

## Lessons Learned

Since this was a new evaluation process for our team, we had some surprises.

**KEN FISHKIN** AAISM, CISSP, CCSP, CIPP/US, serves as Associate Director of Information Security at Lowenstein Sandler, where he leads the firm's cybersecurity strategy and operations. Since joining the firm in 2019, he has overseen its cybersecurity program to not only protect the organization but also advance its broader business mission. Ken is widely recognized for his leadership within the cybersecurity community. As President of the ISC2 New Jersey Chapter, he guides the chapter's strategic growth, has expanded regional engagement, and champions high-quality professional education for practitioners at all career stages. He strengthened partnerships across industry, academia, and the public sector and advanced mentorship initiatives. His contributions helped reinforce the chapter's role as a trusted hub for professional development and community advancement. He serves on the Executive, Threat Intelligence, and PHISH Committees of the Legal Services ISAO and sits on several advisory boards, including InfraGard NJ, HMG Strategy, and the New Jersey Gartner CISO community. In addition, he mentors students at Saint Peter's University and regularly delivers guest lectures at conferences and university events. Through his leadership, mentorship, and service, Ken demonstrates a sustained commitment to advancing both organizational resilience and the broader cybersecurity profession.

There were some threat-identification issues with our tools, as they did not report attacks that were closed the day before, since the tool assumed it was reporting about a known issue. We had to reset the testing for a vendor.

Another unexpected development occurred when an outside vendor heard about the bake-off and wanted to participate. To maintain integrity, we did not accept any more vendors to participate.

One vendor wanted to withdraw because they were not comfortable with our test scenario, but we convinced them to stay. As a result, they did not report the incidents to us appropriately. In hindsight, we should have accepted their withdrawal.

## Bake-Off Takeaways

The traditional RFP model has served organizations well in lower-stakes procurement projects. It is poorly suited to a managed SOC selection, where the difference between a strong and mediocre provider may not become apparent until a real incident occurs. The bake-off model is more demanding, as it requires internal investment in design, governance, and execution than a traditional RFP. It also delivers something the RFP process cannot: direct, observable evidence of how a provider performs under conditions that simulate real attacks.

As a result of this bake-off, we were able to find a vendor that satisfied our current and near-future needs. Our bake-off produced a selection decision based on observed performance, a governance component capable of withstanding scrutiny, and a set of institutional lessons that will improve every subsequent vendor evaluation. Organizations that invest in a rigorous procurement process get a rigorous result. That is precisely the standard the security professional community should hold itself to.