Recorded Future®

Intelligence to Risk Framework

By Levi Gundert and Dylan Davis

Introduction

Recorded Future maintains a unique vantage point from interactions with a wide variety of security programs, which vary by industry, region, and maturity, and between the public and private sectors. Our goal with this white paper is to address a commonly observed pain point, evolving from security practitioner to executive advisor, with our "Intelligence to Risk" (I2R) framework.

In the public sector (military, law enforcement, and so on), the value of intelligence-driven decision-making is well understood, having been proven through centuries of war. In the private sector, however, the value of threat intelligence is less appreciated, despite being no less critical for enterprises seeking competitive advantages.

In the private sector, the process to transform threat intelligence into valuable risk recommendations is often opaque, resulting in intelligence capabilities and teams struggling to define workflows and outcomes that measurably reduce risk. Executives are typically tasked with making strategic business decisions, and none is more relevant than managing cyber risk. However, executive uncertainty around security control investment and the point of diminishing returns is pervasive. In our experience, executives reading "intelligence reports" are often left wondering "so what, now what?"

The I2R framework is a six-tiered progression of refining information from data to actionable intelligence. To avoid information overload, we begin with an event, pattern, or anomaly that should trigger a fact-finding process. We then explore the implications of the threat posed by this event, validate existing controls, make recommendations based on upside and downside risk, and take action based on those recommendations. We use the I2R framework internally at Recorded Future to better articulate risk to our clients. We hope you will use I2R to help your organization better understand intelligence production that meaningfully informs private sector risk and, by extension, executive decisions.

Background

Before launching into the framework, a tour of definitions and scope is helpful for context. For the purpose of discussing the I2R framework, "operational" means any workflow that contains an element of automation. "Strategic" denotes a workflow that requires a human brain.

Operational threat intelligence is itself a critical security control. It is a compliment and force multiplier to other security controls. However, it is outside the scope of the I2R framework, which is focused on strategic intelligence. Other case studies on creating and measuring operational intelligence can be found in the Recorded Future public library.

The word "actionable" is a loaded term frequently attached to intelligence objectives. "Action" is a subjective goal in the private sector. For our purposes, "actionable" equates to measurable outcomes that reduce risk and can be simply articulated. The concept is relatively intuitive to military organizations, in which intelligence is actionable if it supports decisions and kinetic results. The private sector still has plenty of work to realize actionable intelligence that informs risk.

To succeed in strategic intelligence, a team should produce high-quality communications through a variety of channels, such as presentations, reports, or conversations.

We define "high-quality communication" as intelligence communicated in a concise, simple, and relevant way:

- **Concise**: When communicating through the written word, keep your sentences short. They're never short enough. When communicating verbally, get to the point. Use top-down communication, starting with the main point, allowing your leadership to ask for clarification when necessary.
- **Simple**: When communicating through the written word, never use excessively large words. If it's not a word you would use in conversation, exclude it. If it's security jargon, explain it.
- **Relevant**: Anytime you write an email, report, or create a presentation, ask yourself, "So what? Now what?" The most common first questions from leadership are "So what?" and "Why do I care about this?" Ensure you're answering this question of yourself in all of your communications.

Equally if not more important is the skill of second-order thinking. Second-order thinking is the ability to think beyond the obvious. The insight provided to leadership should have second-order implications incorporated throughout. The most meaningful insights we've observed tend to be non-obvious or contrarian, with foundational logic and data.

"So what?" is a simple but powerful question that not only helps with relevant communication but can force teams to think past the obvious. When communicating an observation or recommendation, analysts should ask "so what?" to determine the main insights.

The I2R framework assists teams in building this skill by providing a structured process to iterate through when advising leaders.

The Intelligence to Risk (I2R) Framework

I2R is a pyramid because each successive layer, beginning from the bottom, refines information into actionable insight.

I2R Framework Input

Before stepping into the framework, let's first decide what information should be processed. With overwhelming amounts of information to process from multiple directions, it's important to prioritize data that matters, so our suggestion is to focus on time or scope (or both) to reduce the feeling of information overload.

- **Time**: We've observed teams run their threat intel research based on monthly sprints broken down by topic. For example, month 1 is Linux ransomware trends, month 2 is the initial access broker market, month 3 is the semiconductor threat landscape, and so on. These topics can change throughout the year as major events occur (such as the Russia-Ukraine war).
- **Scope**: Another approach is narrowing the scope of information each team focuses on, which helps both small and large security programs. The scope can be decided at the team or individual level. A team example would be that they decide to focus only on information pertaining directly to their industry until they've built internal capacity to expand into adjacent topics. For an individual example, each individual focuses on their functional area, as well as their industry, such as vulnerabilities targeting companies in my industry.
- **Both**: Depending on the security team's resource capacity and their willingness to adapt, we might recommend using both approaches. For example, individuals rotate out their functional focus on a monthly or quarterly basis, ensuring they're covering a broader surface area of information.

I2R Framework Explained

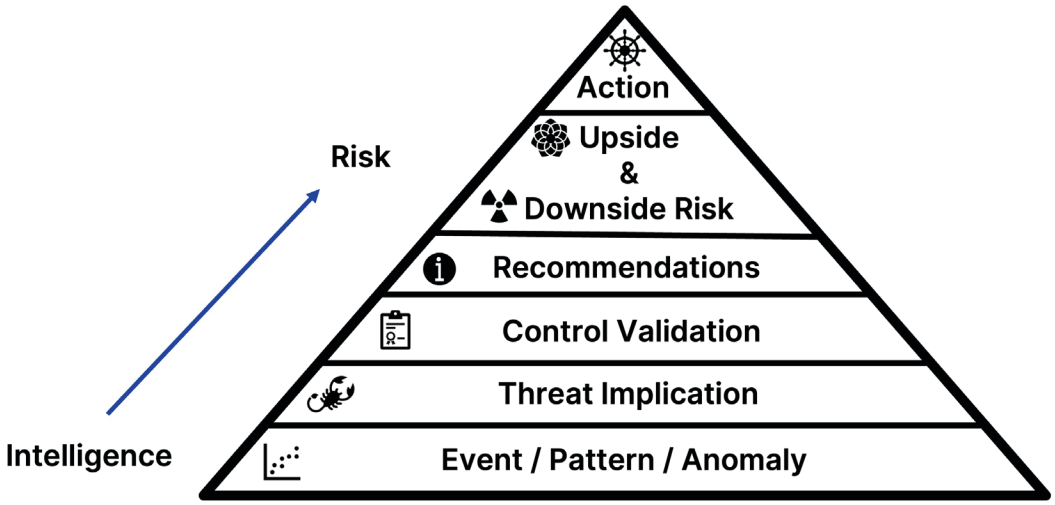

Tier 1 — Event/Pattern/Anomaly

Intelligence is produced from data analysis. The data is generated from a single event, a group of events with similarities (pattern), or an anomaly detected in the context of a larger pattern.

Events and anomalies are time-based observations. Events tend to be real-time observations. An example of an event in this context would be when the Log4j vulnerability was publicized, and leaders all over the world relied on security teams to advise on their posture against potential exploits.

Anomalies, by contrast, are backward-looking and found usually through retrospective research. We can use the same Log4j example. Imagine we have an analyst researching major vulnerabilities affecting open-source technologies in the previous two years, and through their research, Log4j stands out — this would be an anomaly.

Finally, in our experience, patterns tend to be the most frequent insight derived from our use of the I2R framework. To uncover a meaningful pattern, a team will need broad knowledge or willingness to construct hypotheses that can be disproven through follow-up research. Patterns are formed from disparate data sources and can potentially highlight non-obvious or contrarian insight. An example of this insight is listed below in the example section.

Tier 2 — Threat Implication

Following the identification of a relevant event, pattern, or anomaly, human analysis is needed to determine the presence of a threat implication. This step requires thinking about derivative (second-order, third-order, and so on) implications. In the private sector, valuable second-order cognitive skills build on broad business and geopolitical awareness.

We provide prompts to help with critical thinking and implications identification.

- What makes this threat unique?
- What are the plausible scenarios and associated implications of this threat?
- Is there an element of this threat that increases its likelihood or impact?
- Is this threat particularly relevant to my organization? If so, why?

Tier 3 — Control Validation

Completion of Tier 2 results in a hypothesis. Confirming a threat is a risk to your business requires control validation and determining granular existing control efficacy. To properly validate security controls, ideally, multiple stakeholders from across an organization are involved in discussions.

The depth of security control validation will determine the quality of recommendations generated. This tier requires technically capable resources, where the threat implication involves technical controls or multiple controls are affected by the threat implication.

Tier 4 — Recommendations

Now that we have a specific threat with security controls overlayed, we should think through the different recommendations. Recommendations can be broken down into two types: long-term and short-term. Long-term recommendations tend to be associated with higher cost, impact, and sustainability, while short-term recommendations may produce quicker value. When communicating recommendations to leadership, provide both long-term and short-term recommendations (if applicable).

Tier 5 — Upside and Downside Risk

Tier 5 concerns qualifying and quantifying the upside and downside risks associated with recommendations. This stage is tightly coupled with the recommendations stage.

Historically, we have observed analysts sharing the risks associated with a threat, but rarely do they highlight the risks associated with a recommendation. Highlighting the downside risk for a recommendation is insufficient; articulating an upside opportunity creates a more complete picture. When communicating insights to executive leadership, we recommend focusing on a single downside risk and upside risk per recommendation. This focused approach highlights the most significant risks and gains. However, think critically through all the gains and risks, then place them into either an appendix or notes for future reference.

Upside risk comes down to money saved or revenue generated. This monetary benefit is made of multiple categories, such as speed to market, market share, reduced churn, or new market entrance.

A downside risk is an event that causes a business to lose revenue (directly or indirectly). Downside risk from cyber events generally occurs via five broad categories: operational downtime, brand and reputation degradation, compliance failures, financial fraud, and intellectual property or trade secret theft.

These categories are helpful for cybersecurity professionals when thinking through the articulation of upside and downside risk, as part of a recommendation, derived from second-order threat implications.

Each one of the categories is a book unto itself with exhaustive examples. Risk, in all its forms, and financial loss are the language of businesses, particularly enterprises. Helping executives understand how adversarial intent and opportunity translate to the likelihood of occurrence and organizational impact is a critical process in this tier. While we advocate for risk and loss quantification where possible, in reality, most organizations are happy to qualify risk into categories (for example, risk registers that use “high/medium/low” or “red/yellow/green”), however imprecise.

Tier 6 — Action

An executive leadership team is likely to act on recommendations when they are articulated in a complete risk story that incorporates resource constraints. Remember, not acting is always an option, and probably the most common outcome.

Output

The most frequent I2R outputs are leadership presentations, executive summary reports, longer-form briefing emails, or preparation before conversations with leadership. The type of the deliverable shapes what risk insight is shared with executive leadership and when it is shared.

When preparing for a conversation with leadership, the emphasis of the work may focus on the threat implication’s second-order effects and the associated risks for each recommendation.

Conversely, when creating a larger presentation with more time articulating each layer you may decide to spend more time discussing the process behind the quantified upside and downside risks.

The quality of each output will vary by person based on adjacent skill sets (speaking, writing, presentation design, and so on), which is out of scope for this paper.

Applying The I2R Framework

In this section, we walk through two detailed applications of the I2R framework with real-world examples.

Example 1

Bots for Stealing One-Time Passwords Simplify Fraud Schemes (source)

In this scenario, an analyst comes across a recent research report from Recorded Future titled "Bots for Stealing One-Time Passwords Simplify Fraud Schemes". This report details how one-time password (OTP) bypass bots work, their increase in use, and the threats they pose to individuals and financial institutions. The I2R pyramid is applied in the following exercise.

- **Pattern**: Our analyst in this scenario is well-read and notices a pattern of actors using phone-native social engineering attacks to bypass MFA. The phone-native methods include MFA fatigue, SIM swaps, OTP bypass bots, malware injection into Android applications, and session cookies stolen via infostealers. Two recent attacks they've come across are the Roasting 0ktapus Campaign and the Uber attack.
- **Threat Implication**: After additional research and second-order thinking, our analyst creates two hypotheses:
    - These increased efforts have created a thriving ecosystem for the creation of easier and automated methods to bypass MFA controls. With lower barriers to entry for bypassing MFA, 1H 2022 had a higher baseline of MFA attacks than any previous year increasing the likelihood defenders will encounter this attack type. The associated impact of a significant breach could equate to an annual risk exposure of $1,537,546.
    - Threat actors are increasing their social engineering efforts toward phone-native channels due to their effectiveness. This is likely due to targets lowering their personal security posture while communicating in phone native channels, such as SMS, WhatsApp, Twitter, Discord, and YouTube comments.
- **Control Validation**: During this phase, the analyst will need to work collectively with other teams to figure out what existing security controls can defend against the threat implications above. In this hypothetical organization, 50% of critical applications have MFA enabled through SMS-based one-time passwords (OTP), with no additional controls.

- **Recommendations**:
  - **Short-term**: Expand existing MFA to all critical applications, as well as simultaneously transition from basic SMS-based MFA to push authentication-based MFA applications. The implementation of this push authentication MFA should include conditions disallowing more than (X) number of push authentications within (Y) time frame. This condition will help mitigate an MFA fatigue attack like the kind that affected Uber.
  - **Long-term**: Complete all the short-term recommendations, while initiating a larger rollout of FIDO2-compliant security keys such as YubiKey. This rollout should initially be tailored toward employees with access to critical applications, with the intent to drastically reduce the effects of phishing attacks. While presenting this recommendation, the analyst highlights that security-centric organizations such as Google and Cloudflare have publicly reaped the rewards of incorporating these security best practices.

For brevity, we'll choose a single recommendation to detail the associated downside and upside risks, but when completing this process, consider all recommendations. Remember, we're focusing our leader's attention on one downside and upside risk per recommendation while noting the others for future reference (if needed). In this scenario, we proceed with the long-term recommendation.

- **Downside and upside risks**:
  - **Downside**: Rolling out the FIDO2-compliant hard keys such as YubiKey will come at a higher cost (money, time, and human resources). The amount of time required to roll out this hard-key program could be quantified after speaking with different vendors. Once you've understood the average time to roll out the hard-key program, we can quantify both the person-hours to completion and the amount of time we're exposed to these MFA bypass methods. The numbers shared with your leaders should be prefaced with different caveats on accuracy and limited scope, but this shows your team's willingness to see the leader's perspective. We're not only making a recommendation but considering the tradeoffs our leadership is guaranteed to face.
  - **Upside**: There are many benefits to rolling this hard-key program out, such as improved employee experience and stronger security, but we're focusing on one. In this scenario, our analysts are aware that the company recently lost business due to not meeting CMMC Level 3 MFA security requirements, and this hard-key program will open gates to that lost revenue. We've done our research and know once this program is complete, we're fully compliant, which correlates to (X) clients and (Y) revenue lost in the previous 12 months. This information helps us highlight the revenue potential we're gaining by executing this security key rollout recommendation.
  - **Note**: Depending on the recommendation, there may be existing public materials that summarize the return on security investment, such as rolling out Yubikeys (ROI 203%).
- **Action**: After sharing this insight with your leadership, they're convinced to pursue the long-term recommendation, so the team goes off and executes.

Example 2

Insikt Research Note: Analysis of Emerging RAT "Asbit"

In this scenario, our analyst comes across a piece of research covering a new remote access trojan (RAT) called "Asbit". This RAT uses DNS over HTTPS (DoH) to communicate with its command-and-control (C2) server and infect its victims with an attachment via Discord.

Note: DoH encrypts the DNS requests between the client and server. This newer protocol is a partial win for consumer privacy, which protects against ISPs from monitoring their traffic, but not completely. On the other hand, this is a loss of visibility for defenders, hindering their ability to see what domains are being queried, allowing attackers to bypass network detections.

- **Event**: The analyst found this research note unique due to the RAT's ability to communicate to the C2 over DoH and the initial-access vector of Discord attachments.
- **Threat Implication**: Our analyst is proactive and decides to research additional trends around the two unique attributes mentioned above. This leads to two threat implications.
    - **DoH**: Our analyst realizes that DoH is one of many moves the overall private sector is taking to protect consumer privacy. These initiatives are reducing defender visibility over time while empowering attackers. A few examples would be DNS over TLS (DoT) and the QUIC protocol, which is the bedrock for HTTP3 and contains DNS over QUIC (DoQ). The analyst's hypothesis is that attackers will continue to embed their malicious activity into the protocols supporting consumer privacy. This newly encrypted malicious communication will only increase attacker activity, possibly allowing them to exfiltrate even more data from their victims.
    - **Discord**: This is the first time our analyst observed an actor using Discord as an initial-access vector. The analyst's hypothesis is that actors will continue to move where their targets spend most of their time. As more employees spend their time on TikTok, Discord, WhatsApp, and other similar platforms, there's an increased likelihood of social engineering attempts through these channels. The effectiveness of social engineering through these channels is likely higher than through traditional channels such as email due to the lack of employee awareness.
- **Control Validation**: In this scenario, our analyst concludes that there are no direct security controls in place to defend against this DoH communication bypassing network defenses.

- Recommendations:
    - **Short-term**: There are two simple steps the defender can take to make the use of DoH more difficult for the attacker. Additionally, these steps will expose the successful use of DoH, simplifying the identification of likely malicious traffic.
        - Step 1: For all the managed endpoints that can be centrally controlled, disable the use of DoH.
        - Step 2: Block known DoH providers for managed and unmanaged endpoints.
    - **Long-term**: Complete all the short-term recommendations, plus build out infrastructure that allows you to intercept, decrypt, and analyze TLS traffic to and from your network uncovering malicious behavior. It's important to ensure this infrastructure can sustain high throughput without harming the employee's daily work. With this recommendation there's a chance for pushback from leadership due to privacy concerns, so you'll want a strong business case as to why this will benefit the overall company.

As in the first example, we'll choose a single recommendation, the short-term one, to focus on for brevity.

- Downside and upside risks
    - Downside: The recommendation we're pursuing is a one-to-one approach, where one solution partially mitigates one problem, which has limited effects and relevance. The solution, which focuses on mitigating the use of DoH today and in the long term, becomes less relevant due to alternative protocols (that is, QUICs DoQ). Over time we can observe the malicious use of alternative privacy protocols (DoH, DoT, DoQ, and QUIC) showing a potential increase in risk over time. After a few months, you can report these observations back to your leadership, emphasizing the importance of pursuing longer-term security controls.
    - **Upside**: The most significant gain when pursuing this short-term recommendation will be saving money. We can contact vendors or conduct external research to understand the price ranges for infrastructure that intercepts decrypts and analyzes this traffic. Once the average cost is established, we can incorporate that quantitative gain into this recommendation. When communicating this to your leadership include a preface stating that we're saving money in the short-term, but opening ourselves up to continuous alternative protocol use, as well as playing whack-a-mole with this one-to-one approach.
- **Action**: Our leadership decides to execute the short-term recommendation, leaving the long-term recommendation for future consideration. This is still a win. It may not be ideal based on your research, but you're making progress and building trust with your leadership.

Conclusion

Security exists to reduce risk. Security practitioners can reduce risk at different levels of abstraction: operational and strategic. Few security teams are able to reduce risk at a strategic level due to an inability to acquire or maintain an executive table presence.

The I2R framework has a single goal: to help the security community move from practitioners to trusted leadership advisors.

To directly access this framework, use the link below for the reference guide.
https://l.ead.me/executive-insights

Levi Gundert is the Senior Vice President of Global Intelligence at Recorded Future, where he leads the continuous effort to measurably decrease operational risk for clients. Levi has spent the past 20 years in both the public and private sector, defending networks, arresting international criminals, and uncovering nation-state adversaries. He's held senior information security leadership positions across technology and financial startups and enterprises. He is a trusted risk advisor to Fortune-500 companies, and a prolific speaker, blogger, and columnist.

Dylan Davis is the manager of Strategic Intelligence at Recorded Future. Dylan leads many executive conversations with strategic customers, providing threat intelligence insights into their areas of interest. Before moving into Recorded Future's research team, Dylan acted as a Senior consultant working to improve the threat intelligence programs of some of Recorded Future's most complex customers. Prior to Recorded Future, Dylan worked as a strategic technical advisor to the executive team within Mastercard, helping shape the company's digital roadmap.

About Recorded Future®

Recorded Future is the world's largest intelligence company. Recorded Future's cloud-based Intelligence Platform provides the most complete coverage across adversaries, infrastructure, and targets. By combining persistent and pervasive automated data collection and analytics with human analysis, Recorded Future provides real-time visibility into the vast digital landscape and empowers clients to take proactive action to disrupt adversaries and keep their people, systems, and infrastructure safe. Headquartered in Boston with offices and employees around the world, Recorded Future works with more than 1,400 businesses and government organizations across more than 60 countries.

Learn more at recordedfuture.com and follow us on Twitter at @RecordedFuture.