Introducing Ero, an AI Co-Pilot for Security Compliance Challenges

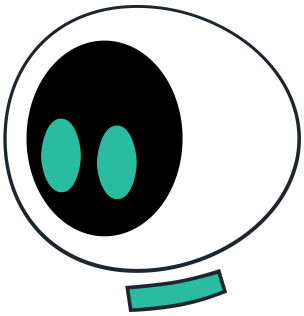

Come See Ero in Action

trustero.com/demo

Download the white paper:
go.trustero.com/erowhitepaper

Table of Contents

AI is Transforming Compliance

The transformative potential of Artificial Intelligence (AI) is far-reaching, and one area that stands to benefit significantly is Governance, Risk, and Compliance[1] (GRC), particularly in the realm of information security compliance.

Trustero recognizes that while robust information security policies and adherence to regulations are essential for business well-being, their implementation and validation have become increasingly complex. This growing challenge hampers many organizations from keeping up with the pace.

In their report, Why Digital Trust Truly Matters[2], McKinsey & Company said that 65% of B2B respondents said that they will only make online purchases or use digital services after making sure a company has a reputation for protecting its customers' data.

As more companies are required to meet these standards to foster trust, complexity is also increasing as compliance frameworks adapt to new security challenges. This situation proves daunting, particularly for smaller technology firms. It's not just technology companies feeling the heat, though. Sectors like healthcare, finance, and government are also grappling with the growing array of sector-specific regulations.

Large Language Models[3] (LLM) – with the capacity to manage and interpret the ever-increasing volumes of data – can enhance trust across industries by making compliance more accessible, regardless of size or resources. More accessible compliance fosters a Trust Graph of reputable entities. Trustero is using AI to build that Trust Graph, driving towards a secure and interconnected digital world that benefits everyone.

Meet Ero: A GRC AI Co-pilot

Powered by LLMs tuned for reliable expertise in information security, Ero is a trust and compliance robot. It analyzes, interprets, and acts upon vast amounts of GRC data and then provides organizations with a comprehensive view of their security posture and regulatory environment before they're audited. It provides an on-demand, in-depth scan before an audit which enables organizations to close compliance gaps while also saving significant time and money to achieve a pristine audit report.

---

[1] Governance, Risk Management, and Compliance Defined.
https://en.wikipedia.org/wiki/Governance,_risk_management,_and_compliance
[2] Why Digital Trust Truly Matters.
https://www.mckinsey.com/capabilities/quantumblack/our-insights/why-digital-trust-truly-matters
[3] AI Large Language Model. https://en.wikipedia.org/wiki/Large_language_model

Ero's Core Functionality

1. Conduct automated audits and notify departments of their compliance failures
2. Automatically map internal processes to the compliance requirements they satisfy
3. Recommend policies and processes based on industry best practices and a business's existing tools
4. Detect compliance requirement failures on-demand so departments can quickly remediate before losing customer confidence
5. Automatically collect relevant evidence to prove compliance

This paper will focus on the auditing capabilities of Ero. In the following sections, the paper provides a primer on GRC, discusses how human auditors approach information security audits, then dives into how Ero leverages LLM to mimic human auditors in conducting security audits.

GRC Primer

Modern SaaS businesses need to establish effective governance, manage risks, and ensure compliance as these tasks are foundational pillars that organizations rely upon. These components collectively form the basis for the organizational structure.

Governance: Blueprint for Direction

Governance provides the framework for organizational decisions. Effective governance sets roles, responsibilities, and decision-making processes. This system aligns the organization's efforts toward common goals. Through sound governance, organizations establish a solid foundation upon which every decision and action is built.

Risk Management: Safety Measures for Stability

Risk management within the GRC context identifies vulnerabilities and potential pitfalls. By addressing these risks, organizations enhance stability and resilience. A robust risk management strategy not only safeguards the organization but also ensures preparedness for any audit, demonstrating a proactive commitment to risk mitigation and, therefore, building trust.

Compliance: Adhering to Standards

Compliance within GRC ensures conformity with regulations that govern the industry and market. This adherence not only prevents penalties but also enhances the organization's reputation. A solid compliance framework aligns the organization with legal, ethical, and security standards, reinforcing its commitment to integrity and accountability. The process of compliance to laws and standards typically requires a third-party audit of adherence.

The following sections discuss audits and audit approaches in detail.

Tri-Layered Approach to Audits

A universal approach auditors use to conduct SOC 2[4], ISO 27001[5], HITRUST[6], PCI DSS[7], CMMC[8], and similar audits individually or at the same time consists of three foundational layers: Governance, Operational, and Technical. This audit approach is rooted in compliance standards such as AICPA's AU-C Section 230[9] (Audit Documentation), AU-C Section 530[10] (Audit Sampling for SOC 2 Audits), and ISO 27001 Auditing requirements (ISO 1702-1[11], ISO 9001[12], ISO 27006[13]). Each layer groups together an aspect of examination that needs to be conducted for a related set of operational processes. The following sections narrate audit procedures of each layer.

[4] American Institute of Certified Public Accountants. (n.d.). AICPA SOC 2 Report. https://us.aicpa.org/interestareas/frc/assuranceadvisoryservices/aicpasoc2report
[5] International Organization for Standardization. (n.d.). ISO/IEC 27001:2013. https://www.iso.org/standard/27001.html
[6] Health Information Trust Alliance. (n.d.). HITRUST CSF® Assurance Program. https://hitrustalliance.net/product-tool/hitrust-csf/
[7] PCI DSS v4.0 Quick Reference Guide: Understanding the Payment Card Industry Data Security Standard version 4.0. https://docs-prv.pcisecuritystandards.org/PCI%20DSS/Supporting%20Document/PCI_DSS-QRG-v4_0.pdf
[8] Department of Defense Chief Information Officer. (n.d.). Cybersecurity Maturity Model Certification (CMMC). https://dodcio.defense.gov/CMMC/
[9] American Institute of Certified Public Accountants. (2017). AU-C Section 230: Audit Documentation. https://us.aicpa.org/content/dam/aicpa/research/standards/auditattest/downloadabledocuments/au-c-00230.pdf
[10] American Institute of Certified Public Accountants. (2019). AU-C Section 530: Testing the Operating Effectiveness of Controls. https://us.aicpa.org/content/dam/aicpa/research/standards/auditattest/downloadabledocuments/au-c-00530.pdf
[11] ISO/IEC 17021-1:2015 Conformity assessment — Requirements for bodies providing audit and certification of management systems — Part 1: Requirements. https://www.iso.org/standard/61651.html
[12] ISO 9001:2015 Quality management systems — Requirements. https://www.iso.org/standard/62085.html
[13] ISO/IEC 27006:2015 Information technology — Security techniques — Requirements for bodies providing audit and certification of information security management systems. https://www.iso.org/standard/62313.html

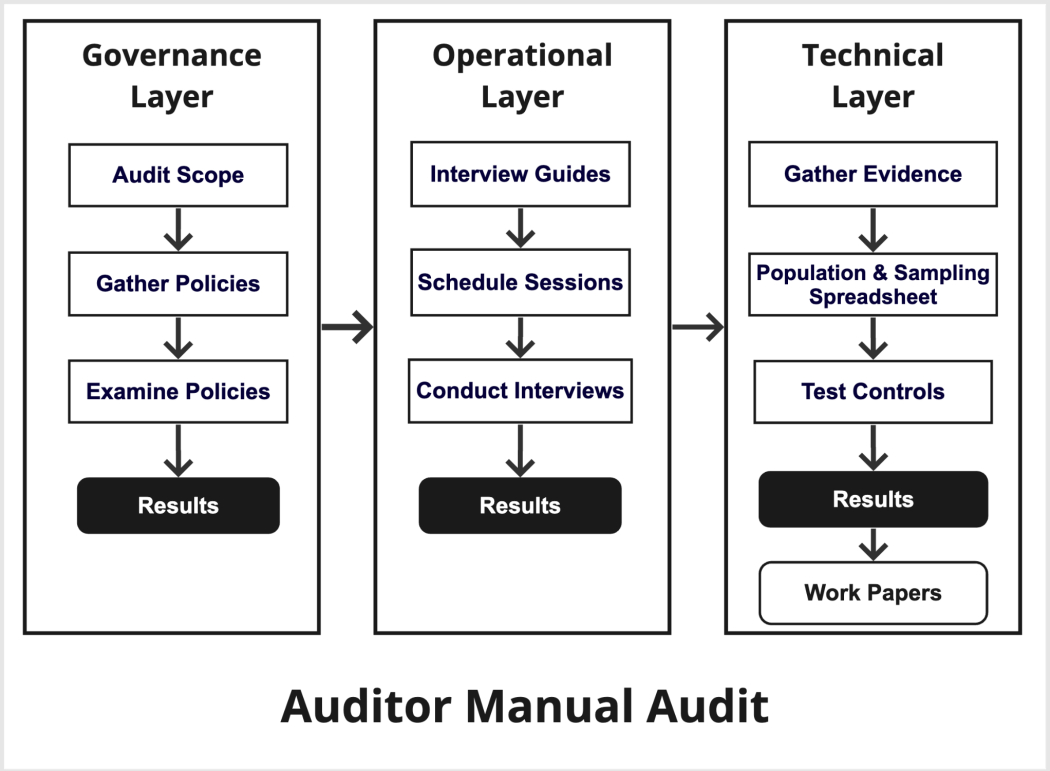

Examination Begins within the Governance Layer

In the Governance Layer, an auditor will first validate the audit scope and request to see additional formal documentation per applicable standards, laws, and regulations in scope. Formal documentation includes, but is not limited to, policies, business continuity, and incident response plans. SOC 2 specifically requires a written System Description with an architectural diagram and ISO 27001 requires the Statement of Applicability (SoA) within an Information Security Management System (ISMS) to be fully documented.

After the audit team has gathered all formal documentation and the Risk Register, a thorough examination of the policies will commence to ensure all controls with objectives are clearly identified and outlined within the corresponding policy. The goal is to establish that the organization has a good understanding of its risk profile and solid philosophy about managing those risks. During this process, the audit team documents and creates guidelines for conducting interviews within the Operational Layer. They also create a comprehensive request list for all supporting documentation and evidence required within the Technical Layer for testing of controls. The request list is derived from and based on the audit scope, policies, and additional formal documentation gathered in this layer.

The Risk Register, if done correctly, shows the auditor exactly how each individual control is mitigating inherent risk. The inherent risk is quantified by identifying all applicable risk drivers (threats) to the organization per predisposing condition or vulnerability with the likelihood of occurrence and potential adverse impact.

Unearth Evidence in the Operational Layer

Based on the identification of individual job roles and responsibilities within the organization, the audit team will schedule interview sessions with individuals to validate that an operational process or procedure identified within a policy has been implemented and is being followed. During these live interviews or working sessions, the audit team records the meetings to further gather screenshots to meet the corresponding audit requirements. The goal is to understand and verify that the practices and processes (i.e., controls) are properly designed and aligned with the risk management philosophy and decisions/policies of the organization.

ISO 27001 certification bodies are required to spend an *x* number of face time hours, based on the size of the organization, with each organization being audited. The auditor may also provide a complete request list showing what additional supporting evidence and documentation is required for them to complete additional testing of controls within the technical layer.

Test Controls in the Technical Layer

After completing the two initial audit layers, the audit team generally has strategically gathered enough evidence through interviews and evidence request lists. The auditors have sufficient evidence populations (user lists, change management tickets, audit logs, etc.) and begin random evidence sampling (users, tickets, etc.). The evidence samples are used to test each control validating the control's operating effectiveness. The results of the control test ensure the associated risk is being mitigated and controls operate as designed. Only when the chain of understanding risks, making decisions, designing processes, and operating processes is functioning can an organization claim that it's properly managing its risks.

Concurrently to control tests, the audit team of testers are creating "work papers" to demonstrate the accuracy and completeness of the audit. The work papers are used to track and identify all the formal documentation, recorded interview sessions, and a comprehensive request list with supporting documentation and evidence. The audit team who conducted the testing then submits the detailed work papers to their Audit Manager who performs a quality assurance (QA) audit check. Finally, the Lead Auditor, typically an Audit Partner or Principal, validates that the Audit Manager has completed the QA audit check and determines if a Certification can be issued, or if a Report is qualified or unqualified.

Ero Audit Deep Dive

One of Ero's core functionalities is to conduct audits using the three-layered approach to audits described earlier. Ero performs the audit on a control-by-control basis. While this approach comes with some current limitations, it allows for better focus on the problem of each control and a better understanding of what context Ero needs to provide in order to verify the effectiveness of the control. Understanding the context of each control is the key to Ero's effectiveness. On a high level, Ero examines the objective of each control and attempts to

classify it into test categories (i.e., which tests are applicable to each control). Subsequently, each test category attempts to identify the relevant population sample needed, and the relevant context, and test the control objective against the population sample as well as other contexts to cross-reference.

Ero's approach to conducting the three-layered audit involves four major parts: data collection, control classification, context identification, and test definition. The following sections examine each part in detail.

George Totev:
Enterprise Risk and Compliance Executive and former Head of Risk and Compliance at Atlassian

"Ero is the intelligent companion that every GRC person needs. It performs the tedious, heavy work and comes up with some very interesting and valuable recommendations. It will definitely shorten your time to audit readiness and will get you closer to continuous compliance – the key to building and maintaining trust."

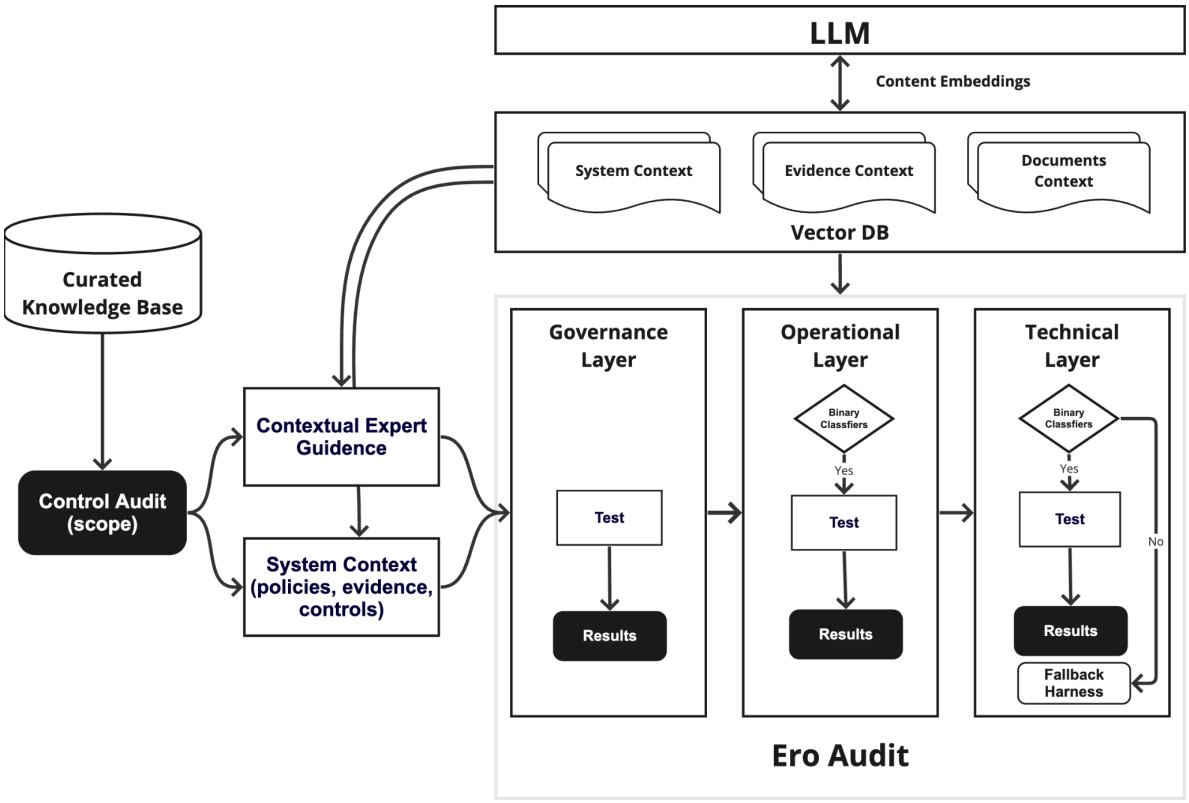

Data Collection and Organization

In order to audit a business, all necessary data about the business needs to be available to Ero. Similar to human-conducted audits, the more data an auditor has access to, the more accurate the audit. Trustero automatically gathers information about a business's operations. The gathered information includes relevant business policies such as business continuity and access management policies; information about the business's purpose and how it's organized;

information about the software and services used; information about what controls are in place; and evidence proving the business is following its policies and controls.

Some information is difficult to gather automatically. In this case, Ero provides tailored guidance based on the known context to help Trustero users identify and enter information to Trustero. An example of this is a piece of control evidence such as a screenshot of Jira[14] tickets showing certain task assignments being completed. This tailored guidance effectively mimics auditor interaction in the tri-layered audit process.

The collection of gathered information about the business forms the system context for Ero. Trustero automatically organizes the system context linking related information. The business's policies are linked to the controls that implement the policy. Policies and controls are linked to the compliance framework requirements it helps satisfy. Evidence documents are linked to controls proving controls are being implemented. These linkages form a graph of related information instructing Ero on how to navigate the audit environment.

The linkages between policies and controls describe how controls implement a policy. These linkages are formed in two ways: expert curation of templated data; and AI similarity matching using learned linkage patterns. A core functionality of Ero is to form linkages based on similar patterns. Ero learns from expert curated policies, controls and their linkages, then applies them to similar patterns of related policies and controls. This similarity technique is also used to form linkages between controls and the framework requirement it satisfies, between controls and the risks it mitigates, and between controls and supporting evidence.

Control Classification

A classic approach to classification would categorize the controls into different categories and attempt to train a classifier based on the control objective to accurately predict these classes, creating a simple multi-class model. However, the problem within the compliance space is that a given control might cover multiple categories depending on the wording. These imply that different types of tests must be performed in order to validate the control's effectiveness, thus the problem becomes a binary classification of control to test applicability, training T classifiers where T is the number of possible applicable tests. More concretely, how can Ero tell if a specific control requires the cross-referencing of the list of newly hired employees against onboarding procedure step, or does it require the verification that access to systems is granted following a different process, or it could be both depending on the wording of the control. This is compounded by the fact that the number of positive examples is relatively small per test type. Various techniques such as data augmentation can help with this problem but they are not enough. Recent developments in LLM have shown that these problems can be solved by combining two techniques: few-shot learning[15] and chain-of-thought prompting[16]. Essentially

[14] Jira is a project and issue tracking software provided by Atlassian. https://www.atlassian.com/software/jira
[15] AIMultiple: What is Few Shot Learning https://research.aimultiple.com/few-shot-learning
[16] Google Research: Large Language Models Perform Reasoning via Chain of Thought https://blog.research.google/2022/05/language-models-perform-reasoning-via.html

giving examples and asking the LLM model to reason about each type of test with examples allows the models to increase the accuracy and identify exactly what is essential for the test. Ero employs the combination of binary classification to identify relevant tests to apply to a control and the LLM few-shot learning and chain-of-thought techniques to conduct control tests. Thus by combining human expertise on a small curated dataset of objective to test pairings as well as reasons for the classification, Ero can achieve 100% accuracy on multiple variance sets.

Context Identification

Providing appropriate, quality context is key to improving the accuracy of a LLM system. In order to achieve this, Ero combines an expert-curated knowledge base and extracts specific contextual information about the user's systems in order to reduce the hallucination problem[17]. Depending on the particular stage in the process, various contextual information is required, thus the context information is organized around these purposes. The document context stores policy and supporting documents, these documents usually come in the form of text, PDF, etc. They are then chunked appropriately based on simple heuristics and the embeddings of these chunks are stored in a vector database. Similarly, the description of the user's system needs to be embedded and stored, particularly the free-form texts describing the system alongside the services used and the description of these services to provide better context for the LLM.

The most complex of the context systems is the evidence context. Evidence collected in order to support controls is heterogeneous, (i.e., some of it is structured data that comes from our evidence collectors - receptors), some of it is unstructured data provided in the form of images such as screenshots, some are MSWord documents, etc. All of this data needs to go through a text chunking and embedding process in order to be able to identify which aspects are relevant at the particular step of the process. The structured data also needs to be embedded, as Ero needs to be able to identify the relevant rows, or pieces of information at any given time.

Lastly, the curated knowledge provides its own context. While it does not need to be embedded and stored in the vector db, as the system already knows which expert guidance is relevant where, it does need to be provided through key steps in the process.

Test Definition

Ero follows a layered approach to testing. Every control undergoes 3 layers: governance, operational, and technical.

The governance layer examines the policy linked to each control, and the contextual information is extracted from these documents based on similarity to the control objective. This allows the LLM system to reason about particular aspects of these documents and if the policy and control are in agreement.

[17] Hallucination (artificial intelligence) https://en.wikipedia.org/wiki/Hallucination_(artificial_intelligence)

The second layer is the operational layer. Similar to the governance layer, its primary purpose is to validate a set of specific documents linked to controls. The particular context needed for this set of tests is the system description and the supporting documents, extracting relevant information from both in order to verify specific requirements are met within the documents.

The last layer is the technical layer, tasked with verifying the particulars of each control. Each test attempts to identify the sample population that needs to be tested (e.g., a set of employees), and then to verify that the control is followed by cross-referencing this sample against other provided information. Each piece of evidence can provide some information. All of this needs to be aggregated in the end and summarized in order to make a reasoned decision about the control. At each step of this process, Ero uses the context database to augment the prompts with necessary contextual information to increase the reliability of the LLM. Furthermore, the chain of thought prompting throughout the entire process increases.

Putting it Together

Ero conducts an audit by examining a collection of relevant controls required to satisfy a compliance framework requirement such as SOC 2 Type 2. Let's dig into how Ero assesses one particular control's operational effectiveness. From Ero's perspective, a control is defined by its objective statement. The control objective examined in this example is:

> During coding and prior to code being operational, secure programming techniques are used: peer review and approval, security iterations and test-driven development.

Ero begins the assessment by gathering data in support of the control and creating a context for examinations. Based on this control objective sentence Ero automatically links this control to its related policy, the Secure Development policy. It further links the control to the framework requirement it satisfies, the SOC 2 Trust Service Criteria CC8.1. Next, Ero gathers evidence required to prove the control is implemented. Ero then gathers the system description provided by the user, itemizes SaaS and software used by the system, and Trustero expert knowledge on procedures various organizations use to satisfy the control. Based on the objective and tools used by the user, Ero understands that a list of Pull Requests (PR) should be gathered from the source code version control system used to be examined. Further, Ero understands the user uses Jira to track security iteration events and a list of Jira tickets should be gathered for examination. PRs and Jira tickets are considered evidence and evidence can be gathered automatically by Trustero or can be manually provided by the user. Ero is multi-modal in that it will understand evidence in structured format collected by Trustero or unstructured format such as a screenshot of Jira uploaded by a user.

After all the evidence has been collected Ero is ready to perform the examination. First, the governance layer examination is performed. Ero begins by examining the Secure Development policy in order to understand that related control objectives are in agreement. Ero indexes the gathered context and extracts relevant portions for examination. Concretely, the extracted information is any portion of the policy that relates to peer review, security iterations, test-driven development, services the system uses to perform these tasks, and general guidance around

this type of control from the Trustero knowledge base. All this information is fed into the LLM to assess: if the policy and the control are in agreement, if the policy contains all the necessary details, and if the policy is appropriately scoped to the system.

Next comes the operational layer examinations. This layer is concerned with examining supporting documents. In this case, the control is specific to the development process, it does not include specific supporting documents such as a disaster recovery plan, business continuity plan, etc. Ero is able to classify this control as not related to these documents and thus this layer does not need to perform a detailed examination.

Last is the technical layer examinations. Ero runs a series of classifiers to determine which tests are applicable. For example: is this a control related to asset management? Ero classifies this control as a change management control and thus it performs a series of tests designed to examine change management control.

Ero begins the change management control examination by assessing if the evidence set is sufficient to perform a more detailed assessment. Once again Ero uses the indexed system description, Trustero expert knowledge base, and the policy to assess if the evidence supplied is sufficient. This allows Ero to understand:: if a user is using GitHub, Ero should expect GitHub Pull Requests (PR) as evidence, and if they use Jira, Ero should expect associated tickets. Once Ero establishes that the evidence will most likely be enough, it moves on to the technical verification of the control. Each technical control is structured into 3 parts: the applicability classifier, the population sample, and the cross-reference verification. In this case, Ero already classified the control as a positive so it moves on to the population sample. Ero attempts to identify which is the set of PRs in question by estimating which is the most likely PR evidence, and sampling a few PRs from this PR evidence list.

Next Ero uses this sample to cross-reference against other pieces of evidence. In this example, Ero would examine the list of PRs and cross-reference it against the Optical Character Recognized (OCR) Jira tickets, validating that the ticketing system corresponds to the PRs and checking that the PRs have an appropriate number of reviewers and code tests applied. The cross-reference between GitHub PR and Jira ticket satisfies the detailed procedure described in the Secure Development policy where a change requirement is tracked in a Jira ticket and corresponding remediation is tracked in a GitHub PR.

After completing the three-layer examinations, Ero provides a detailed analysis of these documents citing the specific pieces of information used to support the assertion that the control is satisfied or not. Below is the Ero assessment for the examined control:

> The evidence from 'Test-driven Sprints & Security Iterations.png', 'Github Security checks for incoming PRs', and 'GitHub Pull Requests and Reviewers - GitHub' all indicate that there is a process of peer review and approval in place, satisfying one aspect of the control. The evidence from 'Test-driven Sprints & Security Iterations.png' also supports the use of test-driven development and security iterations. The evidence from 'List of merged pull requests.png' does not provide

explicit details about the programming techniques used or the process followed to ensure the code's security, but it does not contradict the other evidence. Therefore, the control is considered satisfied.

In a complete audit, Ero will examine all relevant controls using the steps described above. The sum of the control assessments determines the audit result.

Ero Data Security and Privacy

Ero is built on the Trustero multi-tenant distributed system, where each tenant's data is logically partitioned from other tenants. All data is encrypted at rest and in transit. Trustero maintains four categories of data: a tenant graph describing a tenant's operations configuration, a content store of gathered evidence, an index store of evidence embeddings, and a learned model store. Each data category uses a distinct type of persistent data store services provided by AWS. As Trustero learns about a tenant's environment, that learned knowledge is isolated to the tenant and not used as global knowledge.

Trustero utilizes a layered approach to network security. All of Trustero's core data services reside in a virtual private network with limited outbound access. All inbound traffic goes through a virtual load balancer and firewall.

Trustero also makes use of 3rd party hosted LLMs. In this case, Trustero only uses hosted LLMs that guarantee prompt data, and fine-tuned models are not shared.

How Ero Helps Your Business

Streamlining audits
Examination of the environment to be audited and collecting and organizing the data that are needed can take considerable time. Ero does both automatically in real-time, mapping evidence to controls and examining the audit environment before the human audit. Not only does this save time and money leading up to the audit, but it also ensures higher levels of confidence.

Recommending actions that lead to certification
Prescribing the right set of policies, controls, and procedures is critical to satisfying the expectations of an auditor and achieving certification. The sort of deep understanding needed to do so is typically the domain of expensive consultants or in-house compliance experts. Through the normal course of data entry or running digital audits, Ero can prescribe an effective set of policies, controls, and procedures to not only maximize the potential for a clean report but also minimize effort on the way.

Avoiding the complexity problems that come with growth
As organizations grow, they are likely to deal with additional regulations. For example, a company with a SOC 2 Type 2 that is expanding internationally will need an ISO 27001. If they

accept payments, they'll need PCI and if they work in healthcare, they'll need to comply with HIPAA. Adding multiple frameworks adds significant complexity. Ero can automatically map single pieces of evidence to multiple controls and policies that satisfy multiple framework requirements. Ero can also recommend the ideal single set of policies and controls to satisfy multiple frameworks with minimized effort.

Conclusion

Digital trust is table stakes for doing business. Governance, Risk, and Compliance (GRC) is the blueprint for building it, and compliance with security frameworks is the accepted indicator of that trust. But, meeting the robust requirements of security frameworks is a complex and difficult undertaking. It's only getting more challenging as common security frameworks like SOC 2 Type 2 evolve to keep pace with bad actors and new regulations, and as more sector-specific frameworks emerge and proliferate.

The rise of LLMs is impacting every business function. In GRC, it is driving the potential to significantly improve the security coverage of thousands of companies regardless of their size and resources. But, LLMs by themselves are unable to provide sound and effective guidance that is free of misinformation.

Trustero's GRC co-pilot, Ero, uses AI trained in information security and compliance to mimic the actions of human auditors. It automatically provides a comprehensive view into an organization's security posture long before they are finally audited by the human auditors who issue their report(s). For information security professionals who deal with security compliance and audits, Ero provides the baseline data they need before embarking on a compliance journey and gives them the ability to check their progress along the way. For auditors, it serves as a powerful efficiency and accuracy tool that can quickly verify pre-audit work, freeing them up to focus on the strategic aspects of the audit process.

Ero's mission is ultimately to grow the Trust Graph by making stronger security posture and active compliance much more accessible to more organizations. The Trust Graph extends beyond an organization to an interconnected network of trustworthy organizations plus their vendors, partners, and customers. When organizations have greater fidelity in their security posture early on, they focus on closing their most critical security gaps. When they know what they will be tested on and how they're currently tracking against those criteria, they can do a better job developing the controls that lead to pristine and exceptionless audits. As more organizations do those things and achieve compliance, the Trust Graph grows and the digital world becomes a safer and more secure place.

Glossary

- Evidence. Evidence is information, which may be gathered automatically or manually, that supports an organization's assertion that a Control is in place and effective. Evidence could include, but is not limited to, configuration information about systems, screenshots from systems that can provide relevant information, documents, links to documents, emails, or in information gathered through person-to-person discussions.

- Controls. Controls are ways that an organization ensures it is living up to its promises. A control can take various forms such as: limiting access, documenting a process, or ensuring contingency plans that are up to date. Ultimately, controls exist to bring an organization into alignment with auditing frameworks, such as SOC 2, which industry trusts to make companies better business partners.

- Policies. A policy is a deliberate system of guidelines. In the context of compliance, policies are like promises, which must be backed with routine actions tracked by Controls. Policies establish credibility and will be the basis on which an auditor assesses controls.

- Frameworks. Framework is a shorthand for "compliance framework," such as ISO27001, SOC 2, or HIPAA.

- Statement of Applicability (SoA). A SoA is a document required for ISO 27001 certification. It's a document that states the Annex A controls that an organization determined to be necessary for mitigating information security risk and the Annex A controls that were excluded.

- Information Security Management System (ISMS). An ISMS represents the collation of all the interrelated/interacting information security elements of an organization so as to ensure policies, procedures, and objectives can be created, implemented, communicated, and evaluated to better guarantee the organization's overall information security. This system is typically influenced by an organization's needs, objectives, security requirements, size, and processes.

- Risk Register. A risk register is a document used as a risk management tool and to fulfill regulatory compliance acting as a repository for all risks identified and includes additional information about each risk, e.g., nature of the risk, reference and owner, mitigation measures. It can be displayed as a scatterplot or as a table.