DAYBLINK CONSULTING

AI-Powered SaaS Risk Assessments at Scale

How we used AI to launch a security assessment program for hundreds of SaaS applications without overburdening application teams

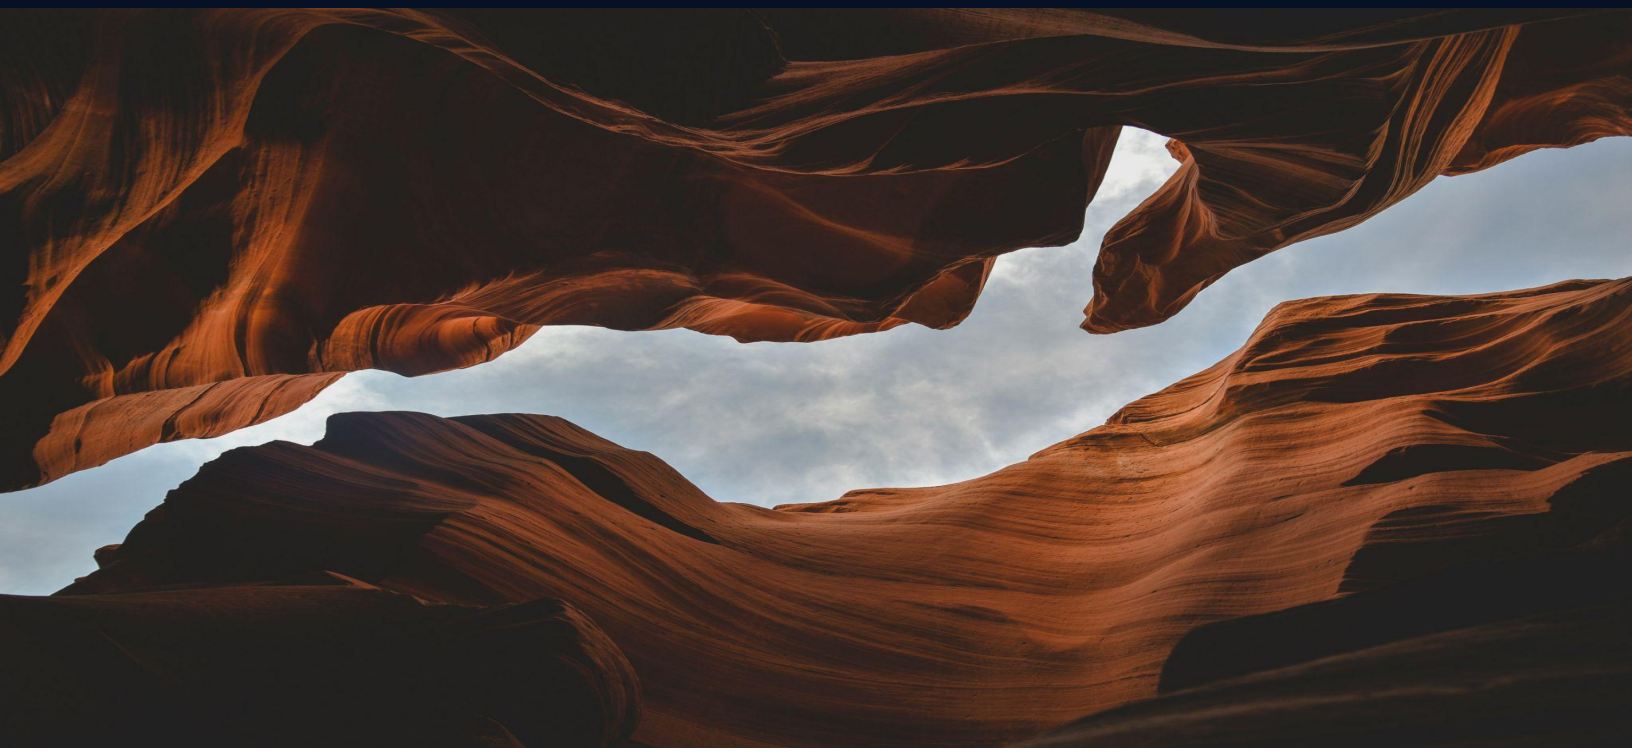

Introduction

Modern enterprises depend on hundreds of SaaS applications to run everything from core operations to customer-facing services. Ensuring each one meets security and compliance expectations requires a thorough, consistent assessment of its risk posture — a substantial undertaking when the estate numbers in the hundreds. One major enterprise set out to do exactly that.

The client was standing up a brand-new program to assess the security posture of several hundred SaaS applications. From the outset, leadership was concerned that a manual effort of this scale would place a heavy and repetitive burden on application teams already stretched thin across competing priorities.

Problem

A new program of this scale risked overburdening application teams before it could deliver value.

The program was designed to assess the security posture of more than 400 SaaS applications. The intended process required each application team to answer roughly 50 questions spanning 15 distinct cybersecurity domains, a meaningful time commitment that, multiplied across hundreds of applications, represented an enormous volume of effort.

That burden was the client's central concern. Application teams were routinely pulled in many directions by competing requests and were often asked to supply the same information more than once. Launching a manual assessment program of this size risked wasting time, breeding frustration, and undermining the goodwill the program needed to succeed.

Capacity was a constraint on the other side as well. With a finite assessment team and hundreds of applications to evaluate, a fully manual approach could not realistically assess this many applications concurrently or at the pace the organization required to stand the program up. Without acceleration, the effort risked stalling or far outgrowing its available resources.

Recognizing these challenges from the start, the team set out to build AI into the program from day one — accelerating the assessment workflow and easing the burden on application teams and assessors alike, all without compromising the depth and rigor the program demanded.

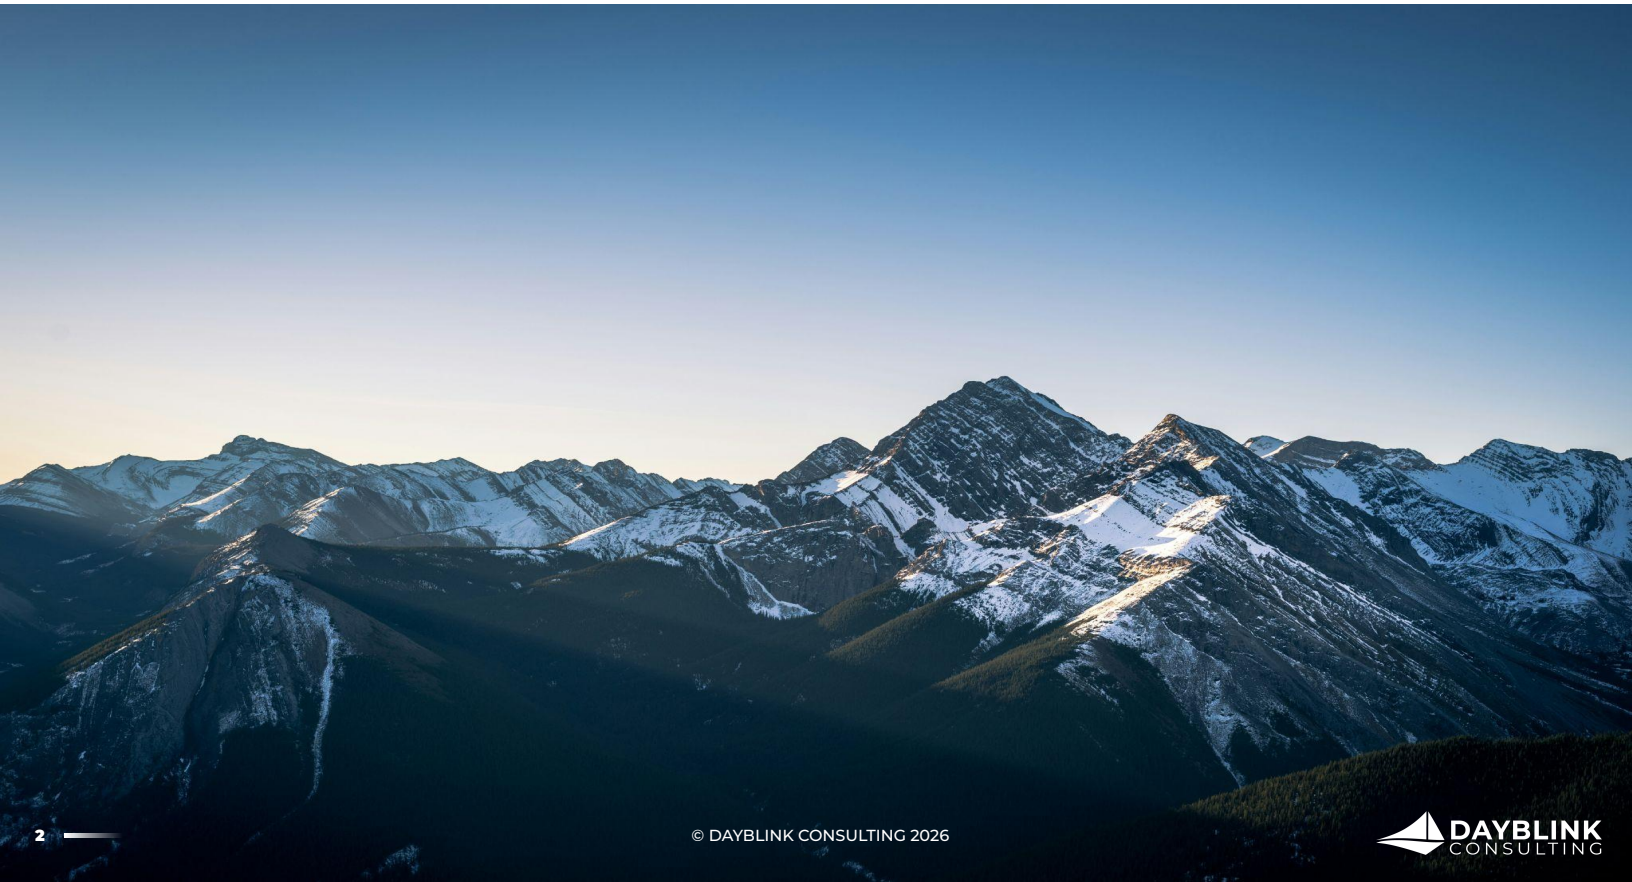

Solution

Our team built an AI agent that pre-populates assessments from existing evidence

DayBlink Consulting designed the program around its Pre-Assessment phase, where the heaviest preparatory work happens before an application team is ever engaged. By building intelligence into this phase from the start, we could remove effort from the process at its source rather than bolting on efficiency after the fact.

For each application, the team first assembles the relevant source documentation — SOC 2 reports, contracts, policies, existing assessments and questionnaires, and other supporting artifacts, typically amassing more than 50 documents per application. DayBlink Consulting then developed an AI agent that evaluates and synthesizes this entire body of evidence and pre-populates draft answers to the assessment questionnaire, reasoning across dozens of disparate documents to find the information each question requires.

The agent was engineered with rigor and traceability as first principles. It was given detailed background and instructions and required to cite its evidence down to the specific file, page number, and section, so every answer could be traced back to its source. It evaluates each of the roughly 50 questions, across 15 cybersecurity domains, against both the client's policies and standards and industry best practices, all powered by an advanced reasoning model.

Once its analysis is complete, the agent automatically generates an Excel file containing the full set of pre-populated, evidence-cited responses. Critically, it does not replace human judgment — it accelerates it. Application teams review the AI-generated answers and use them to inform their own, with some independent research still required, but to a far lesser degree than a blank questionnaire would demand.

This deliberately placed the heaviest lifting in an automated step at the very front of the process. Rather than confronting a blank, 50-question questionnaire, application teams begin from a researched, fully sourced draft, enabling faster assessment completion, with less duplicated effort, and allowing the assessment team to evaluate far more applications concurrently.

Reduced Burden on App Teams

Replaces blank questionnaires with researched, evidence-cited draft answers, saving hours on every application.

Faster, Concurrent Assessments

Shifts the heavy lifting to an automated front-end step, letting the team assess far more applications at once.

Traceable, Evidence-Based Answers

Every response cites its exact source file, page, and section, and is evaluated against client policy and best practice.

Outcome

The AI agent successfully enabled a high-velocity, scalable security program

The agent allowed the program to transform a traditional evaluation effort into a streamlined, automated operation. By shifting from a standard approach to an AI-supported, evidence-led workflow, the enterprise achieved the necessary scalability to assess its entire SaaS portfolio. This shift not only eliminated the anticipated bottlenecks and friction for application teams but also established a new standard for compliance and security oversight across the organization.

This transformation delivered clear, measurable business value across four key pillars:

1. **Operational Efficiency:** The agent is estimated to save 6 to 8 hours per application versus a fully manual approach across roughly 400 applications, returning thousands of hours of capacity to the program and its teams.
2. **Reduced Team Burden:** Application teams begin from a researched, sourced draft instead of a blank questionnaire, sparing them the time and repetitive data-gathering a manual program would have demanded.
3. **Consistency and Traceability:** Every answer is evaluated against the client's policies, standards, and industry best practices, and cites its exact source evidence down to the page, making each response auditable.
4. **Increased Throughput:** By automating the pre-assessment step, the team can run more assessments concurrently and assess the full portfolio of 400-plus applications at the pace the rollout required.

What could have been a slow, burdensome rollout instead launched as a scalable, AI-driven program for assessing the entire SaaS portfolio faster, more consistently, and with far less manual effort.

6-8 hrs
Saved per application versus a manual approach

400+
Applications assessed by the program

~3,000 hrs
Estimated capacity returned to the App Teams

50+
Source artifacts evaluated per application

WASHINGTON D.C. (Headquarters)

8609 Westwood Center Dr., Suite 110,
Tysons Corner, VA 22182

CONTACT US

Matthew Grocock, Manager
matthew.grocock@dayblinkconsulting.com

Dustin Clark, Consultant
dustin.clark@dayblinkconsulting.com

DAYBLINK CONSULTING

www.dayblinkconsulting.com