TAG
2025
Security Annual
SPECIAL REPRINT EDITION

HUMAN-CENTERED SAAS SECURITY

AN INTERVIEW WITH RUSSELL SPITLER, CO-FOUNDER & CEO, NUDGE SECURITY

ROBOTS AND AI: PERFECT TOGETHER

TECHNOLOGY HAS CHANGED HEALTHCARE. NOW HEALTHCARE NEEDS TO CHANGE ITS SECURITY

SECURING SMART CONTRACTS IN A DECENTRALIZED WORLD

TAG DISTINGUISHED VENDOR
nudge

The need to reduce cyber risk has never been greater, and Nudge has demonstrated excellence in this regard. The TAG analysts have selected Nudge Security as a 2025 Distinguished Vendor, and such an award is based on merit. Enterprise teams using Nudge's platform will experience world-class risk reduction—and nothing is more important in enterprise security today.

The Editors,
TAG Security Annual
www.tag-infosphere.com

REPRINTED FROM THE TAG SECURITY ANNUAL

©TAG INFOSPHERE, INC. 2025

AN INTERVIEW WITH RUSSELL SPITLER,
CO-FOUNDER & CEO, NUDGE SECURITY

HUMAN-CENTERED SAAS SECURITY

With the explosion of SaaS usage and increasingly decentralized workforces, security teams face mounting challenges in regaining visibility and control, without disrupting productivity. We recently spoke with Nudge Security, whose patented SaaS discovery and behavior-based interventions are redefining how organizations detect shadow IT, guide secure user behavior, and scale governance. By combining just-in-time nudges with agentless discovery and responsible AI, Nudge is helping enterprises secure their SaaS environments from the ground up, starting with people.

TAG: What motivated Nudge Security's "people-first" philosophy in SaaS security, and how does this shape your long-term vision for securing decentralized workforces?

NUDGE SECURITY: When reducing risk, you must either remove the threat or change the environment. As most of us cannot do much to remove adversary groups, we must focus on changing our environments to reduce risk. The challenge with risk as it relates to SaaS is that it is created through the daily decisions of every employee in the organization. Seemingly mundane choices, such as creating OAuth grants to sync data, trialing new apps, sharing data with coworkers, and inviting colleagues to the platform, all create unique risks that are not centrally controlled. Our vision for the people-first approach is rooted in this reality: we need to engage those who generate this risk in order to remove it. Our approach enables organizations to gain centralized visibility and create a scalable solution that provides the business context necessary to resolve and mitigate risk.

TAG: Nudge's patented SaaS discovery uncovers the "long tail" of shadow IT that others miss. Can you explain how it works—particularly without requiring network or endpoint agents—and why this matters?

NUDGE SECURITY: Traditional approaches to detecting Shadow IT / SaaS rely primarily on network analysis. The limitations of this approach not only come head-to-head with the modern distributed workforce, but also the modern architecture of apps. We no longer have the convenience of a few dozen apps all routing back to known domains, such as salesforce.com or github.com. Rather, we now see dozens of network connections going to various cloud providers, PaaS providers, and other services. To address these challenges, part of our discovery method utilizes a design pattern commonly used in SaaS apps—the use of email to drive user engagement. By analyzing incoming, machine-generated emails from SaaS apps, it not only enables the detection of SaaS use that is not dependent on device or network, but also provides historical insight into SaaS use across the entire enterprise. When combined with browser activity and direct API connectivity, this approach offers broad and deep insight into the use and activity within every SaaS app, without requiring lengthy rollouts or prior knowledge of the app to detect it.

The result of using a Nudge as an intervention means that you end up with more than 200% compliance rates and substantially improved outcomes.

TAG: Your platform uses just-in-time "nudges" via browser, email, or Slack to guide users toward secure behavior. How effective are these compared to traditional blocks, and can you share success metrics?

NUDGE SECURITY: We conducted a study involving 1000 participants, applying modern research principles to evaluate the effectiveness of nudges compared to traditional blocking approaches. The results were impressive. In a conventional blocking approach, we found that only 33% of participants complied with the intervention; an astounding 67% of people reported that they would work around it to complete their tasks. With an intervention in the form of a Nudge, we found that 77% of people would comply. Ultimately, the only thing that happens when you block is that you end up blocking productivity. The result of using a Nudge as an intervention means that you end up with more than 200% compliance rates and substantially improved outcomes.

TAG: Generative AI is now part of your platform: building SaaS vendor profiles and monitoring employee tool usage. How do you ensure AI adds value while keeping customer data private?

NUDGE SECURITY: We leverage many AI techniques in the delivery of our platform, utilizing them across various aspects of functionality to support secure and efficient operations. The most important design principle in our application of AI is to ensure that the models we use do not train on our customer data under any circumstances. Rather, we train the models independently—outside of any customer environment—and use them afterward to derive outcomes from the customer dataset. This approach allows us to deliver value from AI-driven insights while maintaining strict boundaries around customer data privacy and control.

TAG: Looking ahead, what are the next frontiers for Nudge—deeper integrations, new automations, expanded compliance support, or global scaling?

NUDGE SECURITY: As we work with our customers, we are realizing that the problem we are solving—the behavior of the workforce—is the root cause of many modern security scalability challenges. We are currently working with customers to develop further automated TPRM processes, implement deeper AI governance controls, enhance identity governance automation, and implement automated SaaS cost controls. As we look to the future, we are excited about the value we can help our customer realize and how we can remove many of the roadblocks organizations face as they try to scale their security programs.

ROBOTS AND AI: PERFECT TOGETHER

DR. EDWARD AMOROSO

ROBOTS AND MECHANICS

A few years ago, I read Walter Isaacson's biography of Leonardo da Vinci and was struck by how smoothly the great master glided between engineering, art, motion, and anatomy. His **mechanical knight**, for example, was one of the first semi-modern designs of a robot.[1] It was sketched out to sit, wave its arms, and move its head and jaw. This was an amazing thought experiment, done at a time when people still believed that the heavens circled the Earth.

Most people view the Industrial Revolution as ushering in the era of the robot. We've all seen those black-and-white images of early machines clunking along to perform human tasks like weaving textiles and assembling parts. And every computer science student will remember hearing how that early **Jacquard loom** introduced machine programmability using punch cards. These advances were truly amazing, but they were mechanical.

Figure 1. Early robots were mechanical.

[1] *According to my GenAI assistant, the word robot entered the lexicon in a 1920 play called R.U.R. (Rossum's Universal Robots). The term derived from the Czech word "robota," meaning forced labor. I must say that this image of work drudgery is kind of funny now, since we usually do not think of robots being forced to do anything. But it is an interesting notion.*

Progress in robotics seemed to accelerate in the 1940s and 1950s with feedback systems, a principle essential to robotic motion and control. Older readers might still remember Isaac Asimov formulating his famous **"Three Laws of Robotics"** in a series of science fiction stories. These laws, which were intended to prevent robots from harming humans, reflected growing public anxiety and wonder about machine autonomy. (Sound familiar?)

ROBOTS AND COMPUTING

The next wave of industrial robots showed up in the early 1960s, coinciding with advances in programming and computer science. George Devol and Joseph Engelberger are now credited with having developed the first true robot for industrial use. Their robotic arm, installed at a General Motors plant in 1961 (the year I was born), performed repetitive welding tasks with greater speed and precision than any human.

In retrospect, this was a transformative moment. In fact, that welding arm might be worthy of being called the birth of the modern industrial robot. It operated in a structured environment with pre-programmed instructions, which set the foundation for the widespread robotic adoption in manufacturing that we see today. But again, this was a structured environment, and everything was pre-programmed.

Figure 2. Subsequent robots were programmed.

As one would expect, computer technology matured and robots became more capable. The 1970s and 1980s saw the rise of machines powered by microprocessors and integrated with computer vision and sensors. Companies began mass-producing them for automotive and electronics factories. In Japan, there was emphasis not only on industrial robots but also on humanoids designed for entertainment (maybe a little weird, but I'm just saying).

Outside the factory, robots found their way into exploration. NASA's series of Mars rovers demonstrated how semi-autonomous systems could navigate, observe, and experiment in extreme and unpredictable environments. As anyone with a YouTube account can watch today, these rovers incorporate mobility, autonomy, and simple decision-making logic to deal with delays in communication. But they are still just programmed computers.

ROBOTS AND ARTIFICIAL INTELLIGENCE

Today, the definition of robot is expanding to include mobile, networked, and intelligent machines that operate in a variety of domains. The key advance has been the introduction of autonomy to their operation—and such autonomy comes from artificial intelligence (AI) technology. It is hard to overestimate the power derived from the confluence of AI software and robotic operation.

Simple devices like the **Roomba**, a robotic vacuum cleaner released in 2002, became one of the first widely adopted consumer robots to use AI. Meanwhile, researchers and companies pushed the boundaries of humanoid robotics, with Boston Dynamics producing machines like Atlas and Spot that demonstrate agility, balance, and adaptability. Their **dancing robots** have been quite the hit online. A more recent video shows a robot helping a human with construction.

Robotic prosthetics, exoskeletons, and surgical robots using AI have also now emerged, transforming healthcare, giving surgeons more dexterity and precision. This autonomy has enabled machines to adapt to their environment, plan paths, and even learn from past interactions. In warehouses, for example, **Amazon** has revolutionized fulfillment using Kiva robots, coordinating hundreds of units to move goods quickly and efficiently.

Figure 3. Modern robots are autonomous.

The key advance for humanoid operation is that integration of AI with robotics has allowed these devices to interpret images, recognize faces, understand natural language, and make decisions. Voice assistants like Alexa and Siri now blur the line between robotics and ambient AI. Self-driving cars, a form of mobile robotics, depend on complex systems of sensors, GPS, LIDAR, and AI to navigate roadways.

WHAT IS NEXT FOR ROBOTS?

This article began with a salute to Leonardo's ability to combine art and science so smoothly. Any reasonable observer should now see that AI-enabled robots, with their ability to reason and make decisions, will shift like him toward softer human tasks, such as creating art, showing empathy, and even (gulp) adopting religious views. (Dogma might be the key to a human-controlled kill switch for out-of-control robots, but that is for another article.)

My view is that modern AI-enabled robotics will no longer exist as a single domain but will expand to involve neuroscience, psychology, philosophy, and cognitive science. Robotics labs now include work on social robots, collaborative robots (or cobots), swarm robotics, bio-inspired robots, and soft robotics. And the best machines can navigate unstructured environments, grasp delicate objects, and interact with humans in natural ways.

This idea of robots being poised to integrate into society does have some nice advantages. In agriculture, autonomous drones and harvesters can improve efficiency. In elder care and education, humanoid robots can offer companionship and tutoring. On the battlefield and in disaster zones, unmanned aerial and ground vehicles can perform reconnaissance and rescue. Increasingly, robots can become partners to humans, not just tools.

Yet this progress also raises some challenging questions. What ethical frameworks should govern robots? What jobs will be displaced by robotic systems? How do we ensure safety, transparency, and trust in autonomous machines? As robots become more capable and more human-like, these questions shift from theoretical to immediate, including for adjacent issues such as security. These are interesting times.

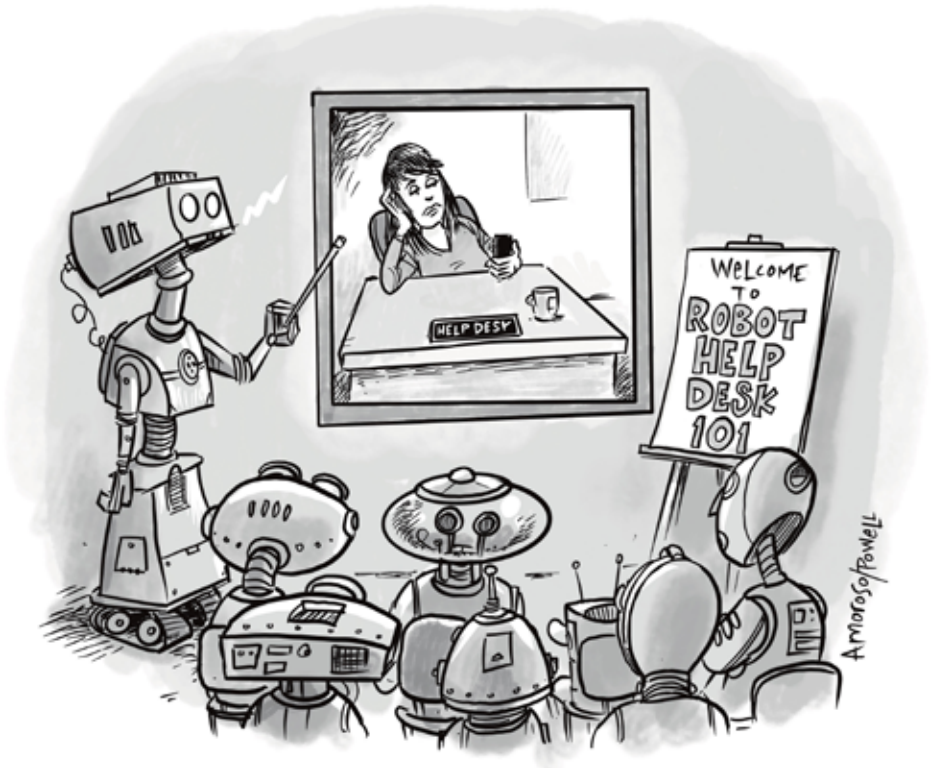

"You'll need to learn to mimic the lethargic disinterest of the human help desk operator."

TECHNOLOGY HAS CHANGED HEALTHCARE. NOW HEALTHCARE NEEDS TO CHANGE ITS SECURITY

JOHN RASMUSSEN

As a practicing chief information security officer (CISO) in the healthcare sector, I have given many presentations regarding the **"fourth wall"** of healthcare technology. The term refers to plays in which actors turn directly to the audience and address them, thus making them part of the performance. You can see this in recent Deadpool movies, where Deadpool brings you in on the joke, directing comments to you.

I've adopted the term to help describe health technology and its interaction with patients. Over many years the technology used in the provision of care was essentially flat. Computers were used to collect information and allow the physician to make clinical decisions based on that data. But that has changed. With the introduction of interactive technologies like surgical robots that appeared in the 1990s, technology now has the ability to have a physical impact on the patient. Included in this technology stack are implanted cardiac monitoring devices, insulin delivery systems, point-of-care devices, and the aforementioned surgical robots.

When they work, they're great. They save lives. But when these systems fail to meet their purpose, they can physically harm a patient. They can even cause death or irreparable harm. The integration of these physically impactful and computer-controlled medical devices to the hospital network has fundamentally changed the role of cybersecurity in healthcare in the last 20 years. Security is no longer all about data protection. It's now a critical component of patient safety.

THE CHALLENGES OF WORKING IN HEALTHCARE SECURITY

When I changed my career as a security contractor for USAID and the IRS to focus on healthcare two decades ago, I quickly realized that information technology had a much larger reach than the normal business support systems I had dealt with in government. In my old job, the IT staff could be disconnected from the business mission, providing support and systems to their co-workers without having to know what functions the systems supported. All I needed to know was whether it was a vital business system in order to prioritize our response. We were mainly concerned with the availability and confidentiality of government data.

In healthcare, understanding the mission is the cornerstone of supporting it. As a security professional, you must be acutely aware that you are there first to serve the patient. I had to shift my perspective to focus on patient safety and optimizing healthcare operations while maintaining the security of the organization. Very often we healthcare CISOs are accused of creating barriers to the workflow and provision of care by clinicians. This has never been the desired approach. We understand the value of the services provided by our organizations and do not want to cause injury or delay care. We work hard to make security controls frictionless.

Learning about the mission and shadowing physicians and rounding with nurses woke me up to the importance of the work that I needed to do, and made me realize how limited my view was. With a new perspective, I also realized that there was much more technology in the workplace than I was aware of. And new tech was being introduced on a daily basis.

A BOOM IN TECHNOLOGY

My experience seeing technology used in real time in a cardiac catheterization lab broke that fourth wall for me. Physicians were using an imaging device to guide a wire through the circulatory system. It was truly mind-blowing to see how reliant they were on the technology, how successfully they had mastered it, and how quickly they could maneuver it! This also made me realize that these systems were a priority over the classic IT systems I had supported earlier in my career.

Medical devices are used for diagnostics, monitoring, and treatment. This area of technology has boomed in the past decade and continues to expand its reach. The growth in use has improved patient outcomes and has greatly enhanced the quality of life for patients who would have suffered or even died without this intervention. **One recent example** of the advance was the use of a robot to complete a heart transplant. Its use was minimally invasive. The typical heart transplant requires the breastbone to be split and makes the recovery more susceptible to post-op complications like infection. The use of the robot greatly improved recovery time and the patient was discharged about a month after the surgery.

More recently, federally funded research performed by Johns Hopkins University utilized a robot to conduct an **autonomous operation** on a lifelike patient. This device used AI and was trained using videos of the procedure. The surgery was a success and was said to have been performed with the same competence as a skilled surgeon.

VULNERABLE PATIENTS, EXPANDED ATTACK SURFACES

These benefits do not come without a cost. They bring new risks. Infusion pumps that automatically deliver medication can harm patients if misconfigured. Temperature monitors on vaccine refrigerators can fail and spoil a batch of vaccines. A malfunctioning air handling system can impact infection control in a facility. These are dangers that can occur inadvertently, but they are also part of the growing attack surface and need to be included in the cybersecurity strategy of today's healthcare professionals.

AS WE INTRODUCE NEW TECHNOLOGIES AND METHODOLOGIES TO DELIVER CARE, HOW DO CISOS LIKE ME PRACTICALLY EVALUATE RISK AND APPLY SECURITY CONTROLS TO THESE DEVICES?

There has been a lot of coverage in the media reporting that the U.S. healthcare sector is a favorite target of threat actors. Most commonly, we are the target of ransomware attacks. More nefariously, we are subject to compromise by nation-states or other organizations that may want to utilize this technology to cause harm or sow mistrust in the medical system.

As we introduce new technologies and methodologies to deliver care, how do CISOs like me practically evaluate risk and apply security controls to these devices? Given these advances in devices and treatment methods, and the web of technology used to support them, healthcare organizations face a lot of challenges to ensure their use does not compromise patient safety. Many of these technologies require connectivity to networks or third parties for monitoring and adjustment, thus increasing the attack surface in hospitals. A lot of this technology lies beyond the direct control of healthcare technology specialists. It can be physician-prescribed, third-party managed tech that uses a patient's home network or cellular networks for communication. Medical devices can be contracted to the manufacturer to provide on-site support on using the hospital network and the internet.

As security professionals, we are tasked with protecting the confidentiality, integrity, and availability of data within the healthcare organization. As such, we have a broad responsibility to protect these medical devices along with other systems, like the electronic medical records, the financial records, and the healthcare facilities themselves. We have a number of regulatory requirements that apply to the business that don't necessarily align with requirements that apply to medical devices. For example, HIPAA requires that patient information be encrypted at rest. However, a legacy portable ultrasound device may not be capable of encryption at rest. How do we square this requirement with the use of technology that does not support it? We must apply technical, physical, and administrative controls (or compensating controls) across the entire technology stack to protect them from compromise, breach, or harm.

WHAT WE NEED IS A NEW SECURITY MODEL

The old security model no longer seems sufficient. Given the different technology support scenarios, the availability and application of security controls is going to vary widely across the provider space. I suggest to the healthcare security community that we review our current risk assessment methodologies and create a multidisciplinary risk evaluation approach across the provider space so

that we understand the real physical cybersecurity risks. And we then weigh them against the practical outcomes of the use of the technology to develop a balanced and cost-effective approach to security.

Historically, the acquisition and use of clinical technologies has been led through hospital departments where the technology is used, and implemented through the clinical engineering department or staff within departments like radiology or cardiology. IT involvement has usually been limited to providing an IP address, creating firewall rules, and stepping aside. This has left the cybersecurity team on the sidelines for many years and has resulted in work that requires the team to retrofit security controls only much later to protect these systems.

That's much too late. The CISO's team should be involved in all aspects of the acquisition process and on a strategic planning level to help inform where the greatest risks are going to come into play. There is no need to delay the capital approval process through post-proposal review. Strategically, the CISO should play an important role in helping to assess the risk of introducing new technology, especially technology that is categorized by the FDA as important to patient safety!

Hospitals and health systems should reconsider their reporting structures, or at least their cybersecurity reporting structures, to include additional departments like clinical engineering, facilities, and specialty departments. And they should pave a way for groups to work together to accomplish the mission. One approach may be the addition of business information security officers that support these groups and specialize in technology.

Healthcare security practitioners will be the first to admit that they are not clinicians, but they do have a role in protecting patient safety. If the cybersecurity team is seen as a cost center and not as a strategic business partner, the organization is going to place itself at greater risk for the compromise of one of these systems, which could harm the people we pledge to serve. Taking a fresh look at cybersecurity strategy should help open the eyes of technologists who are integral to the mission of delivering safe and effective patient care.

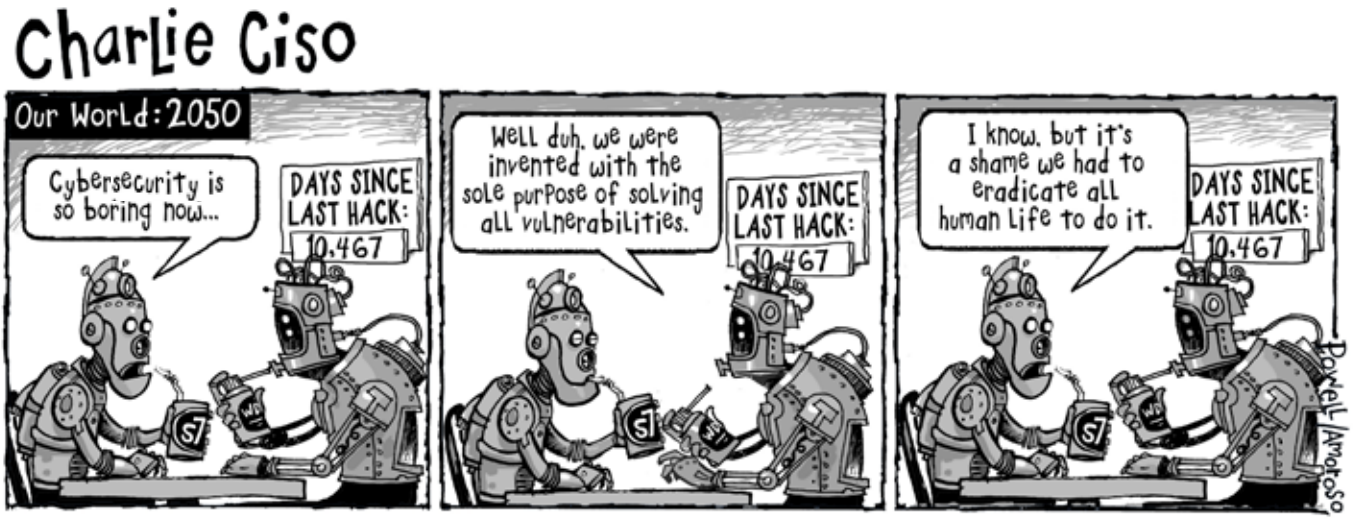

THE IMMUTABLE PROMISE: SECURING SMART CONTRACTS IN A DECENTRALIZED WORLD

DAVID NEUMAN

In 2016, a hacker found a tiny loophole in the code of **"The DAO,"** a decentralized venture capital fund built on the cryptocurrency Ethereum. This single vulnerability resulted in the theft of approximately $60 million worth of Ether. The attack wasn't executed through sophisticated malware or by breaching a centralized database; it simply took advantage of how the smart contract's code was written. This incident highlights the sobering reality that in the world of smart contracts, code isn't just law; it's absolute.

Smart contracts represent a fundamental shift in how we create and enforce agreements. Unlike traditional contracts that require interpretation and enforcement by human parties or legal systems, smart contracts operate deterministically—they will always produce the same output given the same input, with no room for interpretation or negotiation once deployed.

This concept was first envisioned by computer scientist and legal scholar Nick Szabo in the 1990s. In 1996, he defined smart contracts as "a set of promises, specified in digital form, including protocols within which the parties perform on these promises" without the use of artificial intelligence. However, it wasn't until the launch of Ethereum in 2015 that smart contracts found a practical implementation platform. Ethereum expanded on Bitcoin's

blockchain capabilities by introducing a programming language that allowed developers to write complex, conditional logic into blockchain transactions.

What makes these digital agreements revolutionary isn't just their automation but their foundation in blockchain technology. This architecture provides them with three game-changing qualities: code immutability, transparency of the code, and decentralization of the underpinning network.

Once deployed, a smart contract's code cannot be altered. This immutability ensures that agreement rules remain constant, providing certainty to all parties involved that the terms won't unexpectedly change. All network participants can view smart contract code, allowing for public verification of the contract's logic and execution. This transparency builds trust among parties who might be reluctant to enter agreements without familiar intermediaries. Smart contracts operate across decentralized networks, eliminating the need for centralized authorities to oversee contract execution. This distribution of trust across the network fundamentally changes how we think about agreements.

THE VALUE PROPOSITION: WHY SMART CONTRACTS MATTER

Smart contracts have evolved from theoretical concepts to practical technology with transformative real-world applications. As Nick Szabo predicted, "Today, as smart contracts develop and replace some traditional contracts, they are reducing costs and speed up execution."

This real-world application demonstrates just one of the transformative benefits smart contracts offer. Their rapid adoption across industries stems from several advantages they provide. Traditional contracts often require lawyers, notaries, and other intermediaries who add time and expense to transactions. Smart contracts drastically reduce these costs through automation. For example, Insurwave transformed marine insurance by connecting clients, brokers, insurers, and third parties through smart contracts on the Corda Blockchain. Their system manages asset data, links it to policy contracts, enables real-time pricing, and validates loss data—all without manual intervention.

Smart contracts also eliminate human error in contract execution. Each transaction follows precisely coded instructions, reducing mistakes and misinterpretations that plague manual processes. This precision becomes particularly valuable in complex financial operations, where even minor errors can have significant consequences.

Most importantly, smart contracts remove the need for trust between parties. In decentralized finance (DeFi), protocols like Compound demonstrate this by allowing users to lend and borrow cryptocurrency without traditional financial intermediaries. As Vitalik Buterin, co-founder of Ethereum, explains: "A smart contract is a mechanism involving digital assets and two or more parties, where some or all of the parties put assets in, and assets are automatically redistributed among those parties according to a formula based on certain data that is not known at the time the contract is initiated."

This trustless nature enables entirely new business and governance structures. Decentralized Autonomous Organizations (DAOs) operate through code rather than traditional hierarchical management, allowing for novel forms of collaboration and decision-making. MakerDAO, for instance, governs a decentralized stablecoin system through smart contracts and community voting mechanisms.

SECURITY VULNERABILITIES: THE DARK SIDE OF SMART CONTRACTS

On the morning of June 17, 2016, the Ethereum community woke to a nightmare. Someone was draining The DAO of its funds at an alarming rate. Developers scrambled to understand what was happening, but the attack couldn't be stopped—the contract was executing exactly as it was written, just not as it was intended.

Philip Daian, a prominent smart contract security researcher, has studied these vulnerabilities extensively. In analyzing The DAO attack, he noted that "this type of problem wasn't well understood, and people weren't looking for it actively." He further observed that developers and reviewers need to stop examining code in isolation and assuming other components are secure, as vulnerabilities often emerge from the interaction between multiple parts of a system. This fundamental challenge has led to several spectacular failures beyond The DAO hack.

When the Parity multi-signature wallet suffered a bug in 2017, it wasn't due to a sophisticated attack but a curious user who accidentally triggered self-destruct functionality in a shared library. This simple mistake froze approximately $300 million worth of Ether permanently.

In April 2018, multiple tokens built on Ethereum, including BeautyChain (BEC), suffered from an integer overflow vulnerability. This allowed attackers to generate an astronomical number of tokens by exploiting how the contract handled mathematical operations, temporarily crashing their value.

"MANY BLOCKCHAIN PROJECTS NOW OFFER SUBSTANTIAL REWARDS FOR IDENTIFYING VULNERABILITIES. ETHEREUM'S BUG BOUNTY PROGRAM OFFERS REWARDS UP TO $250,000 FOR CRITICAL VULNERABILITIES..."

Vulnerabilities can sometimes arise from external dependencies. In November 2020, Harvest Finance lost approximately $34 million when an attacker manipulated price data from liquidity pools upon which the protocol's smart contracts relied, highlighting the risks associated with depending on external data sources, even when the contract code itself is secure.

These failures illustrate the unique security challenges smart contracts face. Vulnerabilities generally fall into several categories:

Reentrancy Attacks: These attacks "exploit coding vulnerabilities that enable external contracts to reenter functions before updating contract states." When contracts make external calls before updating their states, malicious contracts can exploit this to repeat actions like withdrawals or introduce harmful code.

Access Control Flaws: Insufficient permission settings can lead to unauthorized access to critical functions. This occurs "when contract code fails to apply user permission levels to restrict access," allowing unauthorized users to access or modify contract data or functions and potentially steal funds or assets.

Logic Errors: Sometimes, the business logic contains flaws that aren't technical exploits but functional misalignments with intended behavior.

Oracle Manipulation: Smart contracts often rely on external data sources called "oracles" to trigger execution, creating vulnerability at these integration points.

BEST PRACTICES FOR SMART CONTRACT SECURITY

The immutable nature of smart contracts means that security must be a foundational consideration rather than an afterthought. Philip Daian's research has highlighted that "to mitigate security issues that were quickly evident in the deployment of smart contracts, developers have tried a wide variety of

security techniques." He notes that standard practice now includes manual review by external security firms, often supplemented with verification tools.

This security-first mindset has evolved into established best practices across the industry. Rather than rushing to deploy, professional security audits by specialized firms like OpenZeppelin, Trail of Bits, or Consensys Diligence have become standard practice for any contract handling significant value. Aave, a leading DeFi lending platform with billions in total value locked, undergoes multiple independent audits before implementing substantial upgrades.

Beyond traditional code reviews, formal verification employs mathematical methods to prove contract behavior matches specifications. MakerDAO used this approach for its core contracts, using mathematical proofs to ensure critical functions operate correctly under all conditions.

The principle of least privilege, providing only necessary access to sensitive operations, has become a foundational security concept. To reduce attack surfaces, contracts should perform only required functions and limit interaction with other contracts.

While immutability is a feature of blockchain, developers have created upgradeability patterns that preserve data integrity while allowing future security improvements. OpenZeppelin Upgrades Plugins provide standardized patterns for creating upgradeable contracts, enabling developers to fix vulnerabilities without sacrificing security.

BUILDING SECURITY INTO THE SMART CONTRACT FOUNDATION

As blockchain adoption grows beyond early enthusiasts to mainstream industries, the stakes for smart contract security continue to rise. Financial institutions, supply chain managers, and government agencies are all exploring blockchain solutions requiring unprecedented security assurance.

Looking to the future, Vitalik Buterin acknowledges that "the main advantage of blockchain technology is supposed to be that it's more secure, but new technologies are generally hard for people to trust, and this paradox can't be avoided." This tension between technological security and human trust will continue to shape the evolution of smart contracts.

Many blockchain projects now offer substantial rewards for identifying vulnerabilities. Ethereum's bug bounty program offers rewards up to $250,000 for critical vulnerabilities, incentivizing white-hat hackers to find and report issues rather than exploit them. Automated tools like Mythril, Slither, and Echidna can detect common vulnerabilities through static and dynamic analysis, complementing manual audits with scalable security checking. As the technology matures, we'll likely see continued innovation in security tooling, insurance models, and governance mechanisms designed to protect the growing value secured by smart contracts.

The technology offers unprecedented efficiency, transparency, and autonomy. Yet its immutable nature means that security vulnerabilities can have severe and irreversible consequences.

The future of smart contracts lies in their technical capabilities and our collective ability to develop them responsibly. Nick Szabo's vision from decades ago is finally becoming reality. But as he noted, "Trustworthy third parties are security holes." Smart contracts aim to eliminate those third parties, but in doing so, we must ensure that the contracts themselves are trustworthy.

As we build this new foundation for digital agreements, finding the right balance between innovation and security will remain a critical challenge. The most successful projects will be those that recognize security not as an obstacle to innovation but as its essential partner, because in a world where code is law, we must ensure it's law we can trust.

Nudge Security helps modern, cloud-first organizations discover, secure, and govern shadow SaaS and GenAI tools. Its patented SaaS discovery provides full visibility from day one, while AI-driven risk insights prioritize threats. Built for human-centric security, Nudge automates user nudges and governance workflows—enabling scalable SaaS security with minimal friction.

TAG DISTINGUISHED VENDOR

REPRINTED FROM THE TAG SECURITY ANNUAL

©TAG INFOSPHERE, INC. 2025