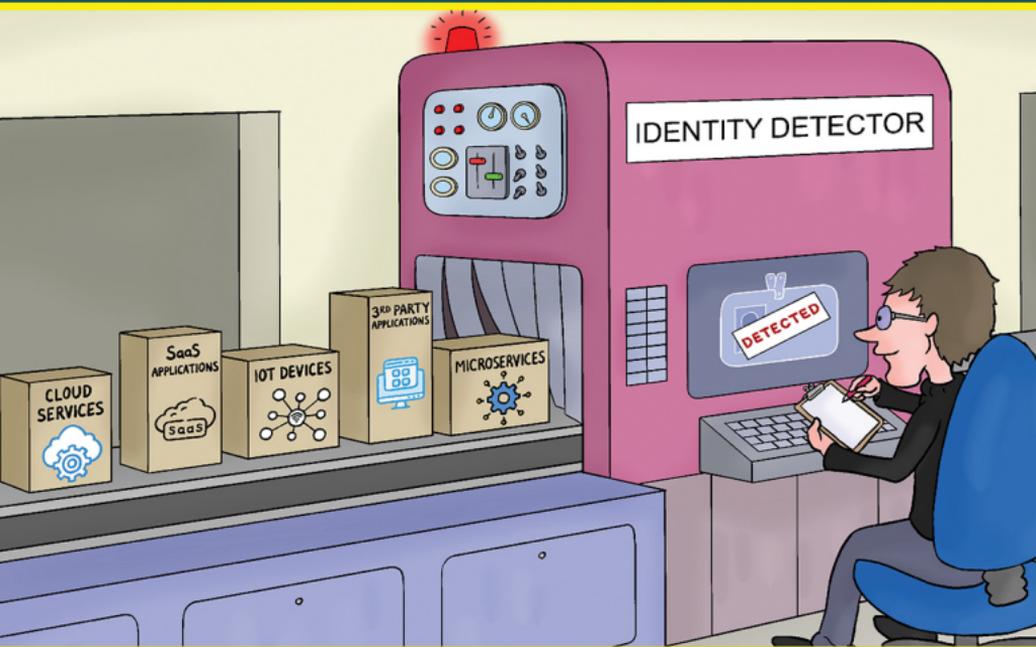


Conversational Securing Machine Identities

Tony Goulding (Cybersecurity Expert and Evangelist)



Learn about:

- Machine identities and the growing security risk they present
- How to protect your organisation by mitigating this threat

Sponsored by

Delinea

Sponsored by Delinea

Delinea is a pioneer in securing human and machine identities through intelligent, centralized authorization, empowering organizations to seamlessly govern their interactions across the modern enterprise. Leveraging AI-powered intelligence, Delinea's leading cloud-native Identity Security Platform applies context throughout the entire identity lifecycle – across cloud and traditional infrastructure, data, SaaS applications, and AI. It is the only platform that enables you to discover all identities – including workforce, IT administrator, developers, and machines – assign appropriate access levels, detect irregularities, and respond to threats in real-time. With deployment in weeks, not months, 90% fewer resources to manage than the nearest competitor, and a guaranteed 99.99% uptime, Delinea delivers robust security and operational efficiency without compromise.

Delinea

Learn more about Delinea on
delinea.com, LinkedIn, X, and YouTube.

Conversational Securing Machine Identities (Mini Edition)

by Tony Goulding

© 2025 Conversational Geek



ConversationalGeek®

Conversational Securing Machine Identities (Mini Edition)

Published by Conversational Geek® Inc.

www.ConversationalGeek.com

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Trademarks

Conversational Geek, the Conversational Geek logo and J. the Geek are trademarks of Conversational Geek®. All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. We cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or programs accompanying it.

Additional Information

For general information on our other products and services, or how to create a custom Conversational Geek book for your business or organization, please visit our website at www.ConversationalGeek.com.

Publisher Acknowledgments

All of the folks responsible for the creation of this book:

Author:

Tony Goulding

Project/Copy Editor:

Ian Whiteling

Content Reviewers:

Robert Sawyer

Alex FitzGerald

The “Conversational” Method

We have two objectives when we create a “Conversational” book. First, to make sure it’s written in a conversational tone so that it’s fun and easy to read. Second, to make sure you, the reader, can immediately take what you read and include it into your own conversations (personal or business-focused) with confidence.

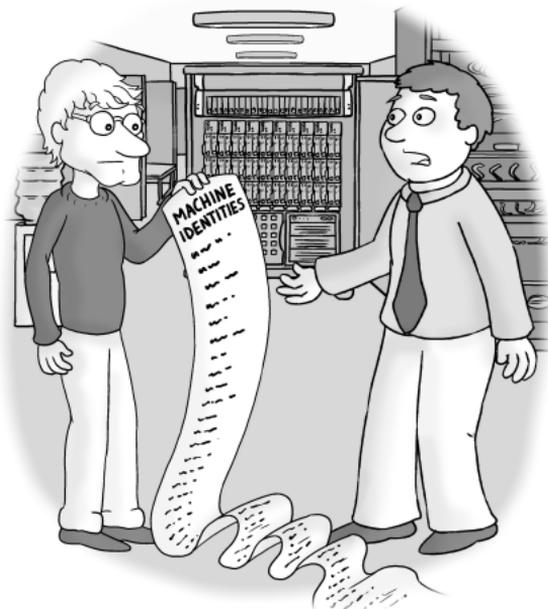
“Geek in the Mirror” Boxes

We infuse humor and insight into our books through both cartoons and light banter from the author. When you see one of these boxes it’s the author stepping outside the dialog to speak directly to you. It might be an anecdote; it might be a personal experience.



Within these boxes I can share just about anything on the subject at hand. Read 'em!

Managing Machine Identities



*“We have ***how many*** machine identities???”*

Machine identities (also known as Non-Human Identities or NHIs) enable IT systems and workloads to securely authenticate, operate, and perform tasks automatically without direct human interaction. More and more businesses are relying on them to run their operations, but they come with a risk...

Why Should You Be Concerned About Machine Identities?

Because, without strict oversight, machine identities can be compromised to bypass security controls, escalate privileges, or exfiltrate data.

An astounding 80% of data breaches are tied to compromised identities that fall into this category¹. When you consider that machine identities have high levels of privileged access to critical data, it's no wonder these types of breaches can have a tremendous impact on an organization.

It's essential to protect machine identities and their associated credentials with the same level of rigor as human users, with policy-based access controls and continuous monitoring.

Identities and Their Credentials

An identity is a unique representation of an entity, such as a person, application, system, or device. By

¹ Non-Human Identity Group, *40 NHI Breaches* (2025)

associating attributes like usernames, credentials (passwords, biometrics, tokens), permissions, and roles, that identity can be authenticated and authorized to access specific resources and perform designated actions. Thus, a human identity might include attributes like a first and last name, address, job title, etc., and something tied to that identity that can be used to log into an enterprise IT resource. These are credentials such as an email address (an ID) plus a password (i.e., an account).

Similarly, machine identities usually take forms like service principals, workload identities, or IAM roles, depending on the environment (e.g., AWS IAM roles, Azure Managed Identities, Kubernetes service accounts). Managing the lifecycle of credentials associated with machine identities is crucial to reducing the risk of a compromise. Best practices include credential rotation, limiting permissions (principle of least privilege), and using vaults or dedicated identity security solutions.

Where Can You Find Machine Identities?

Machine identities significantly outnumber human identities in an enterprise environment.



For every human identity in the workplace there are 20 machine identities – and they're growing by 24% each year!²

Machine identities exist throughout your IT environment, especially if you've been heavily adopting SaaS platforms, cloud services and DevOps workflows. You'll find them operating in containers, microservices, databases, third-party application integrations, CI/CD pipelines, AI LLMs, Operational Technology (OT), Internet of Things (IoT) devices, robotic process automation (RPA), and sensors.

All this has implications for security and management. The rapid proliferation of machine identities expands your attack surface, making effective management and security of these

² ESG, *Managing Non-human Identities for an Effective Cybersecurity Program* (2024)

identities crucial to prevent potential breaches. The sheer volume and diversity of machine identities present challenges in discovering, tracking, and managing them effectively, requiring a robust identity management strategy.

Agentic AI is Increasing Risk

The rise of Agentic AI – AI systems capable of making independent decisions – is increasing the use of machine identities and expanding the attack surface at many organizations. By 2028, agentic AI will be incorporated into 33% of enterprise software applications, a significant increase from less than 1% in 2024.³

With the right access, these powerful tools can be weaponized. Attackers can leverage AI agents to mix harmful or deceptive data into the dataset used to train a machine learning model.

Unlike malware injections that can be identified by checking code, AI poisoning is much more difficult to

³ Gartner, *Intelligent Agents in AI Really Can Walk Alone*

detect and remove because it operates within the model's training data.

A few recent examples of attackers leveraging AI-related machine identities and credentials include:

- Significant security flaws in the DeepSeek AI system allowed attackers to access and leak sensitive data like chat histories, API keys, and backend information, as well as manipulate the outputs by injecting malicious prompts.
- LLMjacking targets AI applications hosted on major cloud platforms. They leveraged AWS keys that provided unauthorized access to GenAI services. With their access they produced explicit celebrity images and other harmful content, including running rogue chatbot services for Dark Roleplaying.

Compliance Auditors and Cyber Insurance Providers are Asking Questions

Compliance auditors and cyber insurance providers are asking questions! Machine identities, including

AI agents, are hot button areas in cyber threat assessments. As organizations increasingly rely on automated processes and AI-driven operations, the number of machine identities has surged, making them a focal point in security evaluations.



Some auditors perform red team attack exercises to simulate real-world cyberattacks and expose vulnerabilities in an organization's security defenses, highlighting unmanaged, vulnerable machine identities. You don't want to end up in that report!

Why Are Machine Identities So Difficult to Manage and Secure?

There are several reasons why machine identities are particularly ripe for exploitation.

Fragmented Management Workflows

Traditionally, IT operations and security teams have focused on managing and protecting human identities. Structured joiner/mover/leaver workflows involve different processes, tools, and people than machine identities do.

For example, human identities are generally created when a new employee joins the organization and is set up in a system like Active Directory, via an Identity and Access Management (IAM) process. As an employee's role changes, the associated identity may receive greater permissions and entitlements. Ultimately, when they leave the organization, the identity is deprovisioned through centralized processes.

As awareness of identity-based attacks has increased, organizations have gotten very good at

implementing security controls such as vaulting credentials, establishing role-based access, requiring user access reviews, and enforcing multi-factor authentication (MFA). To increase their privileges, human users often ask for permission and get approved by another human. Even when they log in with the correct username and password combo, they're often met with MFA challenges to validate they are who they say they are, adding a layer of identity assurance against stolen credentials.

Operating in the Shadows

Machine identities are provisioned and deprovisioned rapidly, often off the radar of centralized IT operations and the IAM team. Many are created by DevOps teams, cloud services, and even applications during cloud adoption, software development, and integration of third-party services.

Traditional identity and access tools like password vault weren't built to manage machine identity secrets at scale, especially given the speed and automation of elastic cloud environments. Many systems, especially in the cloud, are dynamic and

can self-update, adapt, and replicate, requiring continual discovery and management.

This makes it hard to maintain an accurate and up-to-date inventory or centrally manage machine identities in traditional vaults. As a result, machine identities are often created with very weak security controls and orphaned.

No Human Ownership

Not only are some machine identities created automatically, but they're often shared among several applications and services. Without a designated human owner responsible for ongoing governance, machine identities can easily accumulate excessive or outdated permissions.

Plus, the lack of human ownership means that no one knows exactly why the machine identity was created or what business processes, applications, and workflows depend on it. Systems may break and business may be disrupted if the wrong machine identity is removed, or credentials are rotated. So, people simply decide to keep machine identities running, rather than take the risk.



In many cases, the same machine identity is shared among systems, such as production and non-production environments, increasing exposure. Or a single machine identity uses the same secret (a password, token, certificate or key) across multiple systems and environments. This is a well-documented risk that increases the blast radius of a compromise.

Code and Repository Storage

An application developer looking for a shortcut may hardcode credentials tied to machine identities in source code and neglect to remove them before release. It's also common for developers to accidentally commit credentials to public repositories such as GitHub.

According to GitGuardian's *State of Secrets Sprawl 2024* Report, there are nearly 13 million exposed secrets on public GitHub. They also shared that private code repositories are four times more likely to include stored secrets. Under the gun to ship code quickly, developers may also share secrets across internal tools like Slack, Jira, or Confluence, increasing the likelihood of accidental leaks or

insider threats. When threat actors or malicious insiders discover those credentials, they can unlock machine identities, and the systems tied to them.

Dynamic Activity Patterns

Today, numerous detection models exist to identify potential identity threat patterns. For example, a threat detection system may recognize that a human identity that normally accesses the network from a certain IP address and visits certain resources is behaving in unusual ways that could indicate an insider threat or credential theft.

However, these models aren't built for machine identities, which have vastly different activity patterns and are more challenging to baseline. Traditional behavioral analytics often faces challenges in monitoring machine identities. However, advancements in AI-driven security tools and machine learning models are enhancing the detection of anomalies associated with these non-human entities.

With all this in mind, next, let's talk about how you can reduce your risk of machine identity compromise.

Lifecycle Governance for Machine Identities

The good news is the first step to solving any problem is recognizing that you have a problem. So, now that you understand the scope and risk of machine identities, you get to move forward.

Treat machine identity management as a lifecycle, just as you would for human identities. These best practices can help you reduce your attack surface and your risk.

- 1) Discover Machine Identities Operating in Your Environment** – Start by getting a lay of the land. Understand your current machine identity attack surface and evaluate your risk. Continuous discovery of machine identities is essential, as they are frequently created across your identity fabric (SaaS platforms, cloud services, and on-premises environments). Implement code scanning to discover credentials in code. Utilize tools to detect exposed secrets in code commits. Because machine identities are constantly

being created, your discovery needs to be regular, ideally, continuous, so your inventory is always up to date.

- 2) Create secrets using standard protocols** – Effectively managing and securing machine identities, especially for AI agents, bots, and services, requires authentication mechanisms tailored to non-human interactions. This may include:

Delegation Tokens: Users can securely delegate specific permissions to AI agents using delegation tokens such as OAuth 2.0, ensuring agents operate within defined boundaries while maintaining accountability.

Ephemeral Credentials: Using short-lived, non-persistent credentials for AI agents such as OpenID Connect or SAML minimizes the risk of credential compromise and unauthorized access.

- 3) **Vault all credentials, including for machine identities** – For human identities, you can store usernames and passwords as secrets in an enterprise vault. For machine identities, you can similarly store an identifier and a static token, key, or certificate in a vault.
- 4) **Modern identity security solutions** – that incorporate elements of Privileged Access Management (PAM), secrets vaulting, and Cloud Infrastructure Entitlement Management (CIEM). These are designed to manage both human and machine identities effectively. They facilitate the continuous discovery of machine identities across your identity fabric, allowing you to bring them under central control. This centralization allows you to implement automated processes for credential rotation and expiration, enhancing security by reducing the risk of credential misuse.
- 5) **Right size entitlements** – Just as you limit privileges for human identities, you can do the same for machines. They should have

only the rights to interact that they require for specific workflows and business requirements (aka least privilege). If entitlements are rightsized to least privilege, even if a credential or client is compromised, the attacker is contained, and the impact is lower.

- 6) Document ownership and dependencies** – Assigning a human owner for each machine identity increases accountability, so you can confidently know when they're provisioned, updated, or deprovisioned, and by whom. Machine identity credentials should be rotated, reassigned, or deprovisioned as necessary.
- 7) Implement real-time monitoring and anomaly detection** – Identity Threat Detection and Response (ITDR) and Cloud Infrastructure Entitlement Management (CIEM) solutions that track and analyze the usage of cloud services can alert administrators to suspicious activities, such as unusual API requests or excessive invocation patterns. Advanced anomaly

detection mechanisms should be deployed to flag potential hijacking attempts.

- 8) Audit activity of machine identities, including AI** – Ensure that all actions performed by machine identities, including AI models and their interactions with cloud services are logged. You can use services like AWS CloudTrail or Google Cloud Logging to monitor API calls, track access, and investigate suspicious events.
- 9) Transition to dynamic, ephemeral secrets** – Ultimately, your goal is move away from static secrets to adopt a just-in-time (JIT) model, with policy-based, automated secrets that use non-static tokens like OAuth tokens.

Convergence of Solutions for Human and Machine Identity Management

As of this writing, no single vendor provides a fully holistic solution that covers every aspect of human and machine identity security across IAM, PAM, ITDR, CIEM, IGA, GRC, and secrets management.

That said, a unified identity security platform with integrated solutions can provide visibility and control over both human and machine identities. A platform approach consolidates identity and access data across your environment, enabling consistent security policies, risk insights, and automated controls.

It will help you understand the full scope of your risk exposure so you can report accurately and prioritize your security controls and risk mitigation activities. Teams responsible for human and machine identities can operate from a unified, risk-aware source of truth.

The Big Takeaways

You can't afford to sweep machine identity risk under the rug. Chances are, you have hidden or orphaned machine identities operating in your environment without human oversight or compliant governance. The problem is only going to get worse as cloud and AI-driven adoption increases.

Start now to set up a systematic approach to discovering and managing machine identities. Embed security best practices across teams and automate controls where possible – developers who create code, cloud administrators, cloud architects, IT operations teams, compliance and risk management, and incident response. Everyone has an important role to play in getting machine identity sprawl and excessive permissions under control.

The more your processes and systems are aligned, the easier it will be to demonstrate compliance, improve cyber resilience, and have a positive impact on your business goals.

Managing Machine Identities with Delinea

Now that you have a solid understanding of machine identities, you know that securing them is essential. Delinea offers a powerful solution.

- The Delinea Platform secures machine identities in a centralized vault and helps you transition from static to dynamic credentials.
- The policy-based approach ensures authentication and authorization of machine-to-machine communications, all while maintaining least privilege and just-in-time access for all machine identities, including AI.
- By continually discovering and inventorying secrets, accounts, and credentials across on premise, cloud environments, and applications, you'll have a comprehensive and up-to-date understanding of your attack surface and risk exposure.

- With Delinea, you can audit and monitor machine identity access and activity to detect and remediate threats.

As a result, you enhance security for applications and infrastructure to lower your risk, automate to reduce operational costs, accelerate application development, and improve compliance.

Learn more at

<https://delinea.com/solutions/machine-ai-solutions>



Delinea

Unlock AI's potential, not your defenses.

AI is transforming the enterprise, unleashing new possibilities for greater efficiency, rapid innovation, and sustained growth. It's also greatly expanding the attack surface.

Machine identities now outnumber humans, making them prime targets for attackers seeking to exploit privileged credentials.

Secure AI with Delinea so you can:

- Build an AI strategy with confidence
- Secure your AI stack against sophisticated threats
- Gain complete visibility and control of both sanctioned and unsanctioned AI use

delinea.com

More and more businesses are using machine identities to enable IT systems and workloads to securely authenticate, operate and perform key tasks. However, they come with a growing security risk. In this eBook we look at the threats and how to mitigate them.



About Tony Goulding

Tony Goulding is a cybersecurity expert with over 25 years in IT security. His expertise includes identity management and threat detection. Tony is ITIL and CISSP certified and actively shares insights at global events.



ConversationalGeek®

For more content on topics geeks love visit

conversationalgeek.com