Non-Human Account Management (v4)

By Graham Williamson, André Koot, Gloria Lee

To comment on this article, please visit our GitHub repository and submit an issue .

Table of Contents

Abstract

Non-human accounts are often the "Achilles' heel" of a robust IAM environment. While IAM professionals concern themselves with managing identities, authentication, RBAC, ABAC, governance, and auditing of user accounts, other IT staff are deploying devices and services that are given access to protected resources via hard-wired accounts, exposed services, and APIs.

The management of non-human account control should be consistent with user-based account management, and controls placed on user account access to high-assurance applications should also be applied to non-human accounts.

There is no single solution for dealing with non-human accounts. Some IAM professionals suggest all accounts should be managed via the same processes and same infrastructure to ensure consistent policy deployment. This consistency, they argue, should ensure that non-human accounts are not 'left-out' when IAM deployments occur. Others consider this impractical and recommend that purpose-specific processes be deployed for non-human accounts. But regardless of the mechanism(s) used to manage non-human accounts, ensuring that they are managed is paramount. Otherwise, non-human accounts will continue to be a cybersecurity attack vector favored by hackers for gaining access to corporate facilities.

Introduction

A non-human account is usually associated with a service or device rather than a human user. An example is a machine-to-machine service, such as a backup routine that runs during non-business hours to create an offline copy of production data. In this instance, the account permissions should be restricted, i.e., they should not have standard user access nor general Administrator privileges.

Devices such as sensors that provide data to be monitored are sometimes deployed with access to an account so that they can write to a database. Again, such an account should have limited privileges.

Fortunately, the use of such accounts is diminishing as the use of APIs becomes more sophisticated, providing better security and eliminating the practice of hardcoding usernames and passwords in connection routines.

While IAM professionals typically focus on user accounts, these non-human accounts represent a potential attack vector for organizations. These accounts should be considered when formulating policies for access to computer systems.

A comparison between the characteristics of these accounts is shown below:

	Person Identity	*Non-human Identity*
Usage	Multi-faceted, must accommodate multiple access requirements to many applications or protected resources	Purpose-specific, with a single requirement for each deployment
Lifecycle	Created during the 'joiner' process, modified when 'moves' occur, continually monitored for compliance, disabled, and then deleted according to the 'leaver' process.[i]	Created on deployment of the device/service, deleted on termination.
Access control	Dynamic – continual risk-assessed authentication matched to the assurance level requirements of the requested application or protected resource. MFA is used for authentication elevation.	Static – determined at the time of account creation. No MFA requirement.
Access endpoints	Users typically access computer services from smartphones, PCs, and laptops on an interactive basis.	Endpoints are typically devices or device controllers. They can also be computer applications, service routines, or Internet bots.

Table 1 - Account type characteristics

There are two broad categories of non-human accounts that IAM practitioners should differentiate:

- Machine-to-machine accounts used by devices or services to perform a specific function; these 'server' accounts should be monitored and alarm on any incident that is an anomaly to the expected operation.
- Accounts that have access to system functions but are not assigned to a specific individual; these 'system' accounts include administrator accounts with elevated privileges.

Terminology

- Bot – sometimes called an Internet bot, short for 'robot' but referring to a software routine that performs automated tasks over the Internet, a web robot referring to an autonomous network application, or simply a 'bot' referring to an automated, typically repetitive, task used for a specific purpose.
- Identity – defining attributes for a human user that may vary across domains, e.g., a user's digital identity will have a different definition in a work environment as opposed to the user's bank. A device identifier is sometimes referred to as its identity.
- CIA Triad - the fundamental Information security concepts of risk classification of resources from the perspectives of Confidentiality, Integrity, and Availability.
- Non-human/person account – any account not used by a person, including accounts used for devices, services, and servers.

- Server account – an account established with access rights to a specific server operation; this includes service accounts used by a computer application to access another application or service or an account used for a device connection. Note: these accounts are username accounts typically secure via a password.
- System account – a generic term for a privileged account that has extensive permissions that enable system configuration changes.

Non-human Access Control

A significant concern for the IAM practitioner is how to manage access control to and from devices, particularly with services not used interactively by humans. This includes bots that are increasingly being used for automated processes.

IoT Devices

IoT devices can be either a sensor or an actuator. In some cases, sensors provide a continuous stream of data that is displayed in real-time or discrete readings that are written to a database for periodic analysis. Actuators are devices typically used to control a process, turning something on or off. They may be used to open or close a valve by pulsing a servo motor a sufficient number of times until the desired aperture is reached. In many cases, devices are remotely located and connected via a controller to the supervisory system located in a central location. It is noted that IoT devices are becoming increasingly sophisticated with control capabilities and communication facilities built-in. This eliminates the need for a username/password account as IoT devices typically communicate to an API with encryption and digital signing functionality.

In a typical IoT configuration, there are three zones:

1. IoT devices (sensors & actuators). Managing access to and from devices should be governed by a policy that imposes requirements for encryption of the communications channel, such as DNP3, MQTT, and/or digital signature technology (e.g., PKI), to suit the required security level. In low-security environments, static passwords might be used that remain in service until the equipment is decommissioned. In higher-sensitive applications, the security credentials (passwords, certificates, etc.) will be periodically rotated. The selected security requirement must match the capability of the devices, but technical limitations often constrain IoT devices. “Terminology for Constrained-Node Networks" (RFC 7228) nominates three classes of devices:[ii]
    a. Class 0 – no capacity to support configurable authentication.
    b. Class 1 – limited capacity for key management, token support, etc.
    c. Class 2 – fully configurable and able to support dynamic authentication mechanisms.
2. The Controller (to which the devices are connected). If sensor device data is aggregated by a device controller that maps each sensor or actuator to its control logic, providing access control to actuators and protection on writing collected data to a database is required (see Service Accounts, below).

3. Human-Machine interface application (HMI) such as a controller app or a SCADA app monitoring or controlling the IoT devices. In some cases, sensors will write data directly to a database that is read by another application, such as a SCADA app or similar human-machine interface (HMI). Access to these applications will be by humans and should be managed via the IDM environment.

Historically IoT environments have been managed by a team responsible for operational technology (OT) and have had little to do with the information technology (IT) environment within an organization. The specialist nature of IoT technology has justified this organizational structure, and it is often corporate policy to isolate OT from potential compromise via the IT environment. But the requirement for isolation is diminishing as security technology improves. Integrating IoT systems with the IAM environment will improve access control capabilities and provide better corporate governance over operational technology deployments for most industrial applications.

If allowed by regulatory controls, best practice is to integrate the OT environment with the IT IAM environment. This enables the OT to set system entitlements via the IAM system and for OT staff to use their corporate credentials for authorization, potentially via a Privileged Access Management system.

There is increasing concern regarding the provenance of IoT devices and tracking devices throughout the supply chain to ensure no modifications have been made that could potentially deploy 'back-door' access.[iii] The IAM practitioner may wish to ensure corporate policy defines the certification processes to be employed for IoT devices and ensure that compliance with software supply chain policy is in place. This is increasingly important in regulated industries.

Just as important as securing the device itself is protecting the IoT device data. In many cases, databases with IoT devices are not adequately secured. A risk management approach should be employed to determine the adequacy of protection; building environment device data might be low risk but plant production data that is not adequately protected from industrial espionage might be considered critical. The IAM professional should ensure appropriate access controls are placed on industrial data stores. It is good practice to assign a data controller role to an industrial database.

Vulnerability Mitigation

There is no 'correct answer' when it comes to deciding the involvement of IAM practitioners in the management of IoT devices. At one end of the spectrum is the use case whereby all IoT deployments and management are the domain of OT personnel. In this case, the IAM involvement will be restricted to the human accounts that access the OT systems. Group management of entitlements to accounts that can configure IoT systems will heighten the level of security.

At the midpoint of the spectrum, components of the IoT configuration and operation will fall under IAM services. The IAM provisioning workflow will route configuration requests and potentially password rotation requests, to the responsible person. The IoT devices will participate in both attestation reporting to the responsible manager and

compliance management with integration to the Security Operations Center (SOC) and possibly the Security Information and Event Management (SIEM) system.

At the other end of the spectrum, the provisioning of devices is included in the identity management infrastructure. IoT devices are treated the same way as individuals, applying a 'digital identity' to devices. Their entitlements can be set via the normal account provisioning workflows, and their access control can use the same protocols. Most modern API systems, including gateways, use OAuth 2.0 for machine-to-machine communications, while Open ID Connect can be appropriate for IoT device controller authentication.[iv]

Service Accounts

There is a wide variety of service accounts. They are typically used in processes that are periodically run on an automated basis, e.g., via a UNIX cron job or Windows Task Scheduler. Auditors often overlook the accounts used by these processes because they are not accessed by users interactively. Since users do not log into them, they are typically basic, single-purpose accounts with restricted privileges.

Examples include:

- An account used to perform a nightly backup of data
- An account providing access to the HVAC system for monitoring purposes
- An account used for replication of data between directory instances.

The term 'batch account' is sometimes used for a service account. These often refer to one or more utility operations that run periodically during non-production hours to perform a system function. Multiple batch processes may use a single batch account.

Vulnerability Mitigation

Service accounts are a significant source of concern for many organizations because they are often established with a static password that, if not encrypted, can be read by any system administrator. If their access rights are not tightly scoped, these accounts can then be used interactively by a malicious actor and possibly used for lateral movement to other servers in the organization's network. If corporate data loss protection extends to service accounts, tools such as authentication monitoring for anomalies can guard against such vulnerabilities. User behavior analysis tools baseline the normal activity on an account; any deviation from this will generate an alert to the event monitoring system. Alternatively, static service accounts can be migrated to APIs that typically impose a strict security and monitoring regime.

Note: the term 'service account' is sometimes used to describe an account accessed periodically by a service person, e.g., an HVAC technician. Such accounts are user accounts and should be addressed in a company's IAM strategy. They are not addressed in this document.

Bots

The term 'bot' has come from the Robotic Process Automation (RPA) sector that had its genesis in plant automation, where software routines are deployed for repetitive

processes.[v] Bots are now used for everything from website crawlers to retrieve usage information to denial-of-service malware. Increasingly they are being used by organizations to automate repetitive tasks such as retrieval of building information management data or consolidating customer transaction data. In these cases, access by bots will be restricted to a specific purpose.

Bots typically use the Internet to access remote services or resources. A publicly available website should apply mechanisms to limit bot activity and avoid malicious access. These mechanisms might include applying screen-scraper controls, human verification checks, and DDOS protection. A common form of malicious activity is 'credential stuffing,' whereby a hacker alters login credentials to take control of a session.

Organizations need to prepare for the external use of bots. Bots will exhibit different characteristics compared to 'normal' non-human access to a process or service. For the IAM practitioner, user behavior analysis can be used to identify access anomalies.
A process for reviewing the use of bots should be established, testing their functionality prior to deployment and analyzing their usage patterns. Monitoring is a continuous task since malicious corruption of bots is a constant concern.

Vulnerability Mitigation
For the corporate application of bot technology, the IAM practitioner's task is to ensure that appropriate controls on credentials are observed and that PKI signatures and encryption are used as appropriate. Only sanctioned activities should be allowed.

For instance, a bot accessing website data will typically authenticate via HTTPS using an assigned session token. It is a good practice to expire session tokens periodically. The length of time a token should be valid should depend on the sensitivity of the service or resources being accessed.

Client Devices

Traditionally identities are people; they have identifiers stored in an identity datastore and then used to authenticate users to protected resources. It is increasingly necessary to also track the endpoint devices that users employ to access corporate resources, such as laptops, tablets, or smartphones. To track those devices, an object is created in the organization's directory or other data stores that record the detail for each device. This data allows us to grant access to a resource based on the device being used to access it.

There are several benefits to registering client devices:

- It can provide a second factor during a human authentication event, thus reducing the risk score associated with the authentication.
- It can be used to customize the presentation and improve the user experience by passing the details of a user's device to an application.
- It can enable unattended device authentication to support scheduled events such as device updates or data retrieval.
- It can remove a vulnerability and improve governance options when client device objects from the data store are disabled or removed when the time period from the LastLogonTimestamp has been exceeded.

Whether your environment is on-premise, hybrid-cloud, or multi-cloud, managing the client device identity lifecycle is key to reducing the organization's attack surface and maintaining compliance per corporate policy.

Vulnerability Mitigation
With the ubiquity of client devices these days, managing client devices can improve an organization's cybersecurity profile. For instance, a smartphone can be a valuable device for multi-factor authentication (MFA). It can provide a 'possession' factor, e.g., the user is using their registered mobile phone. It can also be used to provide a biometric check for an 'inherence' factor.

Some organizations use a Mobile Device Management (MDM) tool to manage client devices. MDM facilitates the tracking and management of devices and will typically include a self-service module to allow users to register and deregister their devices as new devices are acquired or old devices are lost or retired.

Selecting and deploying the appropriate solution for managing client device 'identities' is a core capability in enabling non-human access control.

System Account Access Control

System accounts give humans access to physical or virtual systems or servers and grant entitlements to privileged system functionality. While not strictly non-human accounts, system accounts are included here because they have no single individual to which they are assigned. System accounts typically refer to administration accounts that are established when a system/server is commissioned. Since this type of account is not directly associated with a single person, they are generally not managed via an organization's joiner–mover-leaver HR processes. IAM practitioners must concern themselves with the management of these accounts.

Admin or Root Account

The admin or root account of Windows and Linux or Unix servers is a highly privileged account with access to system-level operations on the respective platform:

- It is authorized at the highest level.
- It has access to every file and process running on a platform.
- It has permissions to configure the system operation and thereby influence the behavior of the platform.
- Logs from a system will typically display commands that have been run and responses that have been viewed
- Operational use of the account should be continuously monitored.

Note: virtualization and hypervisor platforms (VMware, Citrix, Xen) and container platforms (Docker, Openshift, DCOS, Kubernetes) have administrative accounts that provide an attack vector if not properly managed.

Superuser Account

The term Superuser applies to a business information system or application account that has elevated privileges over standard user accounts. It is generated as part of the system commissioning process when the system is deployed. The Superuser account has permission to modify a configuration, making it a mission-critical account in an information system.

Server Account

Accounts for middleware processes like DBMSs, ESBs, or other ICT components that run in the Windows or Linux operating system environments, are sometimes called server accounts. These are privileged accounts in an application such as a DBMS to give administrative access to a resource owner.

Consumer Devices

There is increasing concern regarding the vulnerability of consumer devices that have connections to the Internet. Recent incidents include:

- Privacy violations by devices that have audio or video capture capabilities and that are sending sensitive data back to a monitoring agency.
- Security incidents such as DDoS attacks as common, published administrative passwords are used, giving hackers access to consumer devices that are then used to conduct Distributed Denial of Service (DDoS) attacks.

Most jurisdictions are now requiring products to adhere to an appropriate set of standards that typically include:[vi]

- Ending the use of default passwords. All devices are shipped with a unique password that is not resettable to a common default setting.
- Enabling support for software updates. Devices are shipped with firmware that can be readily updated in the event that a vulnerability is detected.
- Supporting the secure storage of credentials. All credentials should be stored securely with encryption protection and/or a trusted storage mechanism.
- Shipping with a more secure default configuration. Attack surfaces are minimized by closing unused ports, restricting exposed services to only the functionally necessary, and running software with the lowest level of privileges necessary for the system operation,
- Restricting the storage of Personal Identifiable Information (PII). PII is never stored on the device, as per requirements of privacy regulation in the target geographies.

System Account Characteristics

Since system accounts are not assigned to a single identity, they cannot be wholly managed by an IAM solution, e.g., when the person with administrative privileges leaves an organization, it is not appropriate for such an account to be deleted. A common practice is to provide access to privileged accounts via a managed group so that all users in the group are granted access to the account. But management outside the IAM environment is still required. Good practices include:

- Using a configuration management database in which the server/service is registered as an attribute of the identity it belongs to.
- Assign an account owner to be accountable for the use of the account, typically the owner of the system/service that it belongs to. If no system owner has been defined, a responsible person in the IT department should be the accountable party.
- Interactive accounts should only be used for infrastructural changes or calamities. Admin privileges should be granted via a user's account, e.g., via membership in the appropriate Admin group.
- Passwords for Admin/root accounts must be closely managed. They can be secured via a manual procedure, a password vault, or a Privileged Account Management system.

Vulnerability Mitigation
IAM practitioners should assist in the protection of access to all system accounts. In a UNIX environment, this might be via the removal of the 'etc/passwd' file and the use of SUDO for privilege escalation. In a Microsoft Windows environment, a privileged access management (PAM) system is a common solution. In this case, system passwords are made specifically complex and rotated as appropriate. Access to such an account is via a PAM system, which restricts access to specific individuals with the appropriate entitlements and logs all access events.

If a PAM is not used, Windows supports the time-limited elevation of account privileges, with notification to management. Manual intervention that ensures appropriate use and management of system and server accounts is also good practice, as is including server accounts in corporate audits. This level of management will require corporate policy to be established for server accounts which will heighten the visibility of account management practices.

Increasingly, applications are being deployed on cloud services requiring an access control environment that suits each deployment. This type of deployment might mean configuring a resource manager to protect master account privileges or setting policies that ensure applications do not use the master account for database access.

The Future

The ubiquity of IoT devices will become more prevalent. Devices will span both the corporate and the consumer world, and integrating IoT devices and dataflows will be a new corporate risk. Automation will increasingly be deployed with Machine Learning and Artificial Intelligence, adding to the complexity of the access control environment. Integration with the IAM environment via the use of API gateways, database gateways, service meshes, and Policy-Based Access Control solutions should be considered.

Increasingly APIs are being used for machine-to-machine (M2M) communication. APIs provide the ability to apply consistent security controls on a communication channel and also to monitor it for management purposes. Companies adopting a gateway approach have the ability to provide consistency across M2M communications which is virtually impossible if each service instance is deployed individually.

As the adoption of cloud services continues to accelerate, the use of microservices and containerization will become prevalent. The IAM practitioner should ensure that the appropriate information security solutions are put in place to protect communications between services that communicate identity data.

The use of bots will also continue to accelerate; deployment of behavioral analytics and gateway technology should be considered. The US Department of Homeland Security[vii] advises the following:

- Nefarious bot developers will target new IoT devices for vulnerabilities as they are released to the market and will compete with each other to deploy malware.
- Bot code size will get smaller and more sophisticated to avoid detection and frustrate defenses.
- Botnets will be extended and better monetized, likely through interfaces to social media platforms.
- Botnet operators will operate increasingly globally, taking advantage of regional vulnerabilities. Attacks from foreign nation-state operators will increase.

Access control for non-human entities is a critical competence for risk-averse organizations. It is increasingly important to make sure devices and bots adequately identify themselves, move to APIs with consistent security and monitoring controls, and deploy data-loss prevention technologies such as behavioral analysis tools.

Conclusion

All too often, IAM practitioners are sequestered from non-human account management and only focus on the provisioning and access control associated with user accounts. This is unfortunate because it fragments the host organization's risk management approach to cybersecurity and frustrates the governance task. At the very least, the IAM practitioner should ask the appropriate questions as to how IoT devices are being secured, how server accounts are being managed, and what defenses are in place to thwart malicious bots. It is preferable that the IAM and InfoSec teams within an organization work together to ensure the consistent application of cybersecurity controls that are aligned with corporate policy.

Author Bios

 Graham Williamson

Graham Williamson is an IAM consultant working with commercial and government organizations for over 20 years with expertise in identity management and access control, enterprise architecture and service-oriented architecture, electronic commerce, and public key infrastructure, as well as ICT strategy development and project management. Graham has undertaken major projects for commercial organizations such as Cathay Pacific in Hong Kong and Sensis in Melbourne, academic institutions in Australia such as Monash University and Griffith University, and government agencies such as Queensland Government CIO's office and the Northern Territory Government in Australia and the Ministry of Home Affairs in Singapore.

Graham holds an electrical engineering degree from the University of Toronto and a Master of Business Administration from Bond University. As a member of the IDPro Body of Knowledge Committee, he looks forward to helping create the definitive body of knowledge for the IAM sector.

 André Koot

André Koot is IAM Strategist and Chief Customer Success Officer at Sonic Bee. His IAM experience comes from a financial accounting and auditing background. This background in anti-fraud detection and prevention business processes led to research in the area of authorization principles.

 Gloria Lee

Gloria Lee is a Senior Program Manager in the Azure AD Engineering team at Microsoft. As part of the customer experience team for Identity and Network Access, her role is driving customer success in the Azure Identity division. Gloria is focused on helping

customers increase security posture with the deployment of Azure Active Directory, Azure hybrid cloud-based solutions to provide identity management.

Prior to joining Microsoft, Gloria was a seasoned engineer/architect with 18+ years of experience in the areas of Identity, security, deployment of Microsoft O365 services as well as messaging and collaboration. She had previously spoken at various events such as Microsoft Identity Driven Airlift Conference for partners, GrayHat 2020, and Texas Security Summit. Outside of technology, she enjoys spending time with her kids/family and travel bargain hunting.

Change Log

Date	Change
2020-10-30	V1 published
2021-04-19	Author affiliation change
2022-02-28	Added a section on client devices; added Gloria Lee as an author
2023-04-10	Various changes to improve the clarity of the article such as the addition of device vs. service information

---

[i] Cameron, Andrew and Olaf Grewe, "An Overview of the Digital Identity Lifecycle," IDPro Body of Knowledge, 30 October 2020, https://bok.idpro.org/article/id/31/.

[ii] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <https://www.rfc-editor.org/info/rfc7228>.

[iii] Hashemi, Soheil, and Mani Zarei. "Internet of Things Backdoors: Resource Management Issues, Security Challenges, and Detection Methods." Transactions on Emerging Telecommunications Technologies. Wiley, October 12, 2020. https://doi.org/10.1002/ett.4142.

[iv] See section 'Mobile & API Innovation Gave Us OAuth & Delegated Authorization Frameworks' in Dingle, Pamela, "Introduction to Identity - Part 2: Access Management," IDPro Body of Knowledge, 17 June 2020, https://bok.idpro.org/article/id/45/.

[v] "What is a bot in RPA?," n.d., https://www.nice.com/guide/rpa/what-is-a-bot-in-rpa/.

[vi] For example, see Fernandez, Angel, "New IoT security regulations: what you need to know," Allot blog, 30 January 2020, https://www.allot.com/blog/new-iot-security-regulations-what-you-need-to-know/#.

[vii] Botnet Roadmap Status Update, Department of Commerce and Homeland Security, July 2020, https://www.commerce.gov/sites/default/files/2020-07/Botnet%20Road%20Map%20Status%20Update.pdf.