Non-Human Account Management

By Graham Williamson, André Koot

Table of Contents

Abstract

Non-person 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.

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 be applied to non-human accounts as well.

There is no single solution for dealing with non-person accounts. Some IAM professionals suggest all accounts should be managed via the same processes and the same infrastructure to ensure consistent policy deployment. This common management, they argue, should ensure that non-person accounts are not 'left-out' when IAM deployments occur. Others consider this to be impractical and recommend purpose-specific processes be deployed for non-human accounts. But regardless of the mechanism(s) used to manage non-person accounts, ensuring 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-person identity is associated with a service or device rather than a human user. Identity in this context is defined by the identifier(s) of the device. A device must be identifiable for the device to interact with corporate systems that record data from a sensor or send a signal to an actuator. For instance, a building management system might periodically log into a computer system to write environmental data into a corporate monitoring system database.

In this document, 'non-person accounts' include computer system accounts that are not associated with a person, such as a backup routine that runs during non-business hours to create an off-line copy of production data. Such accounts should be restricted to the specific purpose for which they are created and should be suspended if used interactively.
While IAM professionals typically focus on user accounts, these non-person accounts represent a potential attack vector for organizations and should be included when an organization formulates policy for access to computer systems.

	Person Identity	*Non-person Identity*
Usage	Multi-faceted, must accommodate multiple access requirements to many applications or protected resources	Purpose-specific, single requirement for each deployment
Lifecycle	Created on user-engagement, modified as requirements change, continually monitored for compliance, disabled when service is suspended, deleted on staff exit.	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 time of account creation. No MFA requirement.
Access end-points	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-person accounts that IAM practitioners should differentiate:

- accounts used by devices or services that will never be used interactively
- accounts with interactive access to system functions that are not assigned to any one individual.

Terminology

- Bot – sometimes called an Internet bot, short for 'robot' but referring to a software routine that performs automated tasks over the Internet or 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-Person account – any account not specifically assigned to a person, such as accounts used for devices, services, and servers.
- Server account – an account with privileged access rights to a server's operation typically used for configuration purposes.
- Service account – an account used by a computer application to access other applications or services for a specific purpose.
- System account – a generic term for a privileged account that has extensive permissions to system-level functions, typically used to install new applications, perform system updates, or make configuration changes.

Non-person Access Control

A significant concern for the IAM practitioner is how to manage access control to/from devices and services that are not used interactively by humans. This includes access control for the 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 displayed in real-time or discrete readings written to a database for periodic analysis. Actuators are devices that are 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 will be connected via a controller to the supervisory system located in a central location.

There are three zones in a typical IoT configuration:

- IoT devices (sensors and actuators). Managing access to and from devices should be governed by policy that imposes requirements for encryption of the communications channel (e.g., 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; IoT devices are often constrained by technical limitations. IETF RFC 7228 nominates three classes of devices:[1]
    - Class 1 – no capacity to support configurable authentication
    - Class 2 – limited capacity for key management, token support, etc.
    - Class 3 – fully-configurable and able to support dynamic authentication mechanisms.
- 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

[1] 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>.

logic, providing access control to actuators and protection on writing collected data to a database is required (see Service Accounts, below).
- 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 requirements for isolation is diminishing as security technology improves. For most industrial applications, the integration of IoT systems with the IAM environment will improve access control capabilities and provide better corporate governance over operational technology deployments.

If allowed by regulatory controls, integration of the OT environment with the IT IAM environment should occur. This allows the OT system entitlements to be set via the IAM system and for OT staff to use their corporate credentials for authorization, potentially via a Privileged Access Management system.

Increasingly there is 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. The IAM practitioner may wish to ensure corporate policy defines the certification processes to be employed for IoT devices.

Protecting the IoT device data is as important as protecting access to the device itself. In many cases, databases with IoT devices are not adequately secured. This lack of security is of little concern if it's building environment device data, but it could be critical if it is plant production data that is not adequately protected from industrial espionage. The IAM professional should ensure appropriate access controls are placed on industrial data stores. Most users will only need 'read' access; management of the database should be via an authorised data controller, authenticated to an appropriate assurance level.

IoT Solution Use Cases

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, a human identity is responsible for IoT devices. The IAM provisioning workflow will route configuration requests, and potentially password rotation requests, to the responsible person. The IoT devices will participate in attestation reporting

to the responsible manager and compliance management with integration to the SOC and possibly the SIEM.

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

Service Accounts

There are 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. The accounts used by these processes are often overlooked by auditors because they are not accessed by users interactively. Since users do not log into them, they are typically quite 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 service account.

Service Account Solution Use Cases

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. That service account, if compromised, can be used interactively by a malicious actor and possibly used for lateral movement to other servers in the organization (e.g., using these accounts for stepping-up authorizations to other environments). Including service accounts in the corporate data loss protection tools, such as authentication monitoring for anomalies, can guard against such vulnerabilities. A better practice is to migrate static service accounts to APIs that typically impose a strict security and monitoring regime.

Note: the term 'service account' (a non-person account) is sometimes used incorrectly to describe an account that is accessed periodically by a service person, e.g., an HVAC technician. Such accounts are user accounts and are not addressed in this document other than to note that because this personnel are often external to a company, and therefore not in the IAM data store, a 'generic account' is sometimes established for any person in the service company to use. This is an IAM issue. There is no place for generic accounts in the modern organization, and they should not be used. Options include:

- Federated authentication with the service company
- Self-service password management with approval workflow
- MFA device issuance to the service company
- Deployment of an API to manage and monitor service company access.

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. 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, using human verification checks and DDOS protection. A common form of malicious activity is 'credential stuffing' whereby login credentials are altered by a hacker to take control of a session.

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

Bot Solution Use Cases

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 validity time period of a token should depend on the sensitivity of the service or resources being accessed.

System Account Access Control

While not strictly non-person 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 not normally managed via the joiner–mover-leaver HR processes in an organization. System accounts give humans access to physical or virtual systems or servers and grant entitlements to privileged

system functionality. 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.
- The 'root' or 'admin' accounts have the 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
- The 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 DBMS's, ESB's, 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 either audio or video capture capabilities. In this instance, sensitive data is being fed back to a monitoring agency
- Common and published administrative passwords are used, giving access to consumer devices to a hacker. Malware is then installed on these devices, which can be used for DDOS attacks.

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

- No use of default passwords. All devices are shipped with a unique password that is not resettable to a common default setting.

- Provision for software updates. No device should be shipped with fixed firmware that cannot be readily updated in the event that a vulnerability is detected.
- All credentials should be stored securely with encryption protection and/or a trusted storage mechanism.
- Attack surfaces should be minimized. Unused ports should be closed, exposed services should be restricted to only those functionally necessary, and software should run with the lowest level of privileges necessary for the system operation.
- Personal Identifiable Information (PII) should not be stored on the device. Privacy regulation in the target geographies should be observed.

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, e.g., any user in the group is 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.

System Account Solution Use Cases

The IAM practitioner should assist in the protection of access to the 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 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, time-limited elevation of account privileges with notification to management is supported in the Windows environment. Manual intervention that ensures system and server accounts are appropriately used and managed is also good practice, as is including server accounts in corporate audits. This intervention will require corporate policy

to be established for server accounts, heightening the visibility of account management practices.

Increasingly, applications are being deployed on cloud services and requiring an access control environment that suits each deployment. This deployment-specific access control model might mean configuring a resource manager to protect master account privileges or setting policies to 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 integration of 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 can 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[2] 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 to avoid detection and more sophisticated to frustrate defenses.
- Botnets will be extended and better monetized, likely through interfaces to social media platforms.
- Botnet operators will increasingly operate globally, taking advantage of regional vulnerabilities. Attacks from foreign nation-state operators increase.

Access control for non-person entities is a critical competence for risk-averse organizations. Making sure devices and bots adequately identify themselves, moving to APIs with

[2] Botnet Roadmap Status Update, Department of Commerce and Homeland Security, July 2020.

consistent security and monitoring controls, and deploying data-loss prevention technologies such as behavioral analysis tools is increasingly important.

Conclusion

All too often, IAM practitioners are sequestered from non-person account management and only focus on the provisioning and access control associated with user accounts. This limited approach 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 about how IoT devices are being secured, how server accounts are being managed, and what defenses are in place to thwart malicious bots. The IAM and InfoSec teams within an organization should work together to ensure consistent application of cybersecurity controls that are aligned with corporate policy.

Author Bio

Graham Williamson,
André Koot, Security and IAM Consultant Nixu Oyj