NHI
Mgmt
Group

The Holy Grail of

Non-Human Identities - Databases

Non-Human Identity Management Group

https://nhimg.org

info@nhimg.org

The Holy Grail of Non-Human Identities - Databases

This white paper will focus on challenges around Database Non-Human Identities.

Databases are valuable assets that need protection from external and internal threats. They are prime targets for acquiring sensitive data, causing operational impact, and performing fraudulent activities. From an insider threat perspective, many incidents, including accidents typically involve databases and non-human identities.

Databases are one of the most challenging areas to securely manage. Here’s why:

- **Local DB Accounts** – Modern databases and cloud-based infrastructure promote centralized identity/directory services to manage database accounts. However, many organizations still use local accounts created directly on the DB. This makes it difficult to discern what DB accounts exist and whether they are non-human or human accounts, especially without a consistent naming convention.
- **Password authentication** – Many local database non-human identities use password-based authentication instead of password-less authentication. Discovering these passwords is relatively easy, especially if hard-coded passwords exist in source code repositories. Enforcing password controls like complexity, length, and cycling can be challenging with local passwords.
- **Inventory & Ownership** – Developing account inventory is difficult because most organizations use multiple database platforms, each requiring custom connectors to extract inventory data. Identifying account ownership is challenging, as applications may have other applications connecting to their DB. Additionally, many organizations have technical debt around DB non-human identities, from inactive/orphaned accounts to secrets sprawl, credential sharing, and the creation of generic/shared accounts that allow humans to bypass privileged controls.
- **Password cycling** – Password cycling is an industry challenge, especially on DBs where the same logical account may exist on many DB instances and require simultaneous cycling. Because DB non-human identities get shared across applications, understanding and coordinating dependencies without causing operational impact requires significant planning and investment.
- **Account Permissions –** Identifying what accounts exist on a database is achievable, but understanding their privileges can be challenging e.g. does the account provide read (including sensitive read) or write privileges to data or database objects (DML vs DDL).
- **Interactive Logins -** disabling interactive logins is an industry best-practice to stop misuse of non-human identities. Unfortunately, databases by design need to support interactive logins, so it is very hard to stop someone using a non-human on a database interactively, if the credentials are known.
- **Monitoring Database Activity** – Monitoring operating system activity is generally straightforward with standard logs. However, the database activity monitoring space is challenging because this level of logging is not usually enabled by default. Understanding DB account usage and SQL activities is one of the most challenging areas to address, with risks of humans using database non-human identities to bypass controls.

So how do we tackle some of the challenges?

Understanding the key risks around database non-human identities can be challenging, but there are effective techniques to connect the dots and gain insights without a complete overview:

- **Active Directory** – When database non-human identities are managed in AD, you can extract key attributes from there. This metadata needs to be extracted regularly, using PowerShell scripts.
    - **Password Age** – the *pwdLastSet* attribute from AD will highlight old passwords that need to be cycled.
    - **Account Usage** – the *lastLogon* and *lastLogonTimestamp* attributes from AD indicate when the account last authenticated successfully to a domain controller or across multiple domain controllers.

  Note, whilst this information in AD can highlight key risks (old passwords, inactive accounts, usage), determining if the account is directly associated with a database is not that straightforward. Inspecting AD usually requires more contextual information to understand whether the AD account is connected to or permissioned on a database, as the account could be used for multiple purposes e.g. permissioned to a Windows Server.

- **Infrastructure Event Logs** – Inspecting the event logs of the operating systems hosting the DB servers can provide substantial information for database activity, such as login events (including failed logins), AD authentication events (including user details and IP address).
- **DB Logging & Monitoring** – allows you to track account usage, connection origin, targets, frequency, and queries being run, etc. The options:
    - **DB Platform Logging** – Each database instance logs key database activity/events, incurring a significant performance penalty. Ongoing maintenance by DBAs is also required. The available metadata may be limited, making it difficult to track incoming connections, to join the dots.
    - **DB Agents** – run with privileged permissions on the database server, monitoring all SQL activity, including direct database activity. They can intercept and block/drop suspicious activity based on rules. Challenges exist around deploying these agents on multiple database platforms/instances and ongoing maintenance. Native database logging may have limited metadata, hindering full details of incoming connections.
    - **Network Monitoring** – is the preferred option because network traffic data is generally available in SIEM logs/events. It can identify SQL traffic, understand incoming database connections (who, where, what, when), and use this information to spot suspicious access patterns.
- **UEBA (User Entity and Behaviour Analysis)** – is the glue that pulls everything together and identifies suspicious patterns of activity through anomaly detection and risk scoring. In the case of database non-human identities, there will be clear/regular patterns of usage for these accounts, to help identify the risks e.g. humans using these accounts.

It is critical that the capability utilized to solve these challenges/risks offers continuous monitoring, a hybrid solution for all platform and database types, and uses SIEM logs/events. The capability should cover real-time risks and provide a risk-based approach to drive broader risk reduction/hygiene activities.

If you are an organisation that would like to understand more about the risks around NHIs and how to go about establishing a risk program to manage and remediate the risks, contact us at the **NHI Mgmt Group -** info@nhimg.org.