

Guide for the DAM Team: Recognising SQL Change Guard Connections

This document is for the team that runs database activity monitoring (DAM) and SIEM. The aim is that SQL Change Guard connections to the target databases do not raise "unknown source" alarms, and that access which does not come through the product can be told apart.

1. What will be done, in short

SQL Change Guard runs approved changes and query results on the target database. These connections come from two machines: the Worker machine and the API machine. On the DAM side, the addresses of these two machines are added to the allowlist together with the product's database account.

When the steps in this document are done, this question can be answered: "Is this session the work of SQL Change Guard, or did it come from outside the product?"

2. What to get from the organisation officer before starting

- A SQL Change Guard user with access (Auditor, DbAdmin, ChangeManager or SystemManager role).
- The list of monitored target database servers.
- The machines on which the Worker and the API run.

3. Step by step

1. Sign in to the SQL Change Guard web interface. In the Reports menu, open the "Expected Access Principals" report in the Access & Authorization group.
2. Select a date range that covers the last 30 days and run the report. The report gives one row for each registered target server. A server with no work is also listed (window count 0).
3. Take these values from each server row: server name (ServerName), the product account (ProductLogin), the kinds of work seen (SeenRoles), machine names (SeenHosts), and the client addresses seen by the target (SeenClientAddresses). Several values are separated by a semicolon. The report contains no passwords.
4. For each server, add these two things together to the DAM rule: the product account and the product addresses. Do not add only the address or only the account. If both match, the session is the work of the product.
5. Get the address of the API machine separately (section 4 below). This address does not appear in the report.
6. Test the rule (section 6 below) and report the result to the organisation officer.

7. The organisation officer sets "DamAllowlistMaintained" to "Yes" in SQL Change Guard under Settings > Parameters > External Controls. If the organisation does not use DAM, "NotUsed" is selected.

4. How to find the address of the API machine

The report comes only from work run by the Worker. The API machine reads object definitions, takes column previews and does "test connection". These connections do not run change commands and do not enter the report. Find the address like this:

8. On the API machine, run "test connection" for a server in SQL Change Guard under Settings > Servers.
9. On the target server, find the session whose name is "SQL Change Guard/Test". The queries are in section 5. The client address column is the address of the API machine as the target sees it.

Use the address the target server sees, not the address in the machine's own network settings. The two can differ because of network address translation (NAT), a VPN or the front layer of a cloud service.

5. Queries to see the session on the target

Connections are opened with this name: SQL Change Guard/<Role>/<Number>. Example: SQL Change Guard/Execute/123.

Role	When	Pool	Opens a window
Execute	While a change request runs	Closed	Yes
Query	Query result request, per server	Closed	Yes
Release	Release package, single connection mode	Closed	Yes
Sandbox, SandboxRelease	Isolated trial	Closed	Yes
Health	Connection health check (Worker)	Open	No
Preview, Definition, Test	Preview, definition read, connection test	Open	No

The roles with an open pool (Health, Preview, Definition, Test) keep an idle session open for a while. This is normal and these connections do not run scripts.

SQL Server:

```
SELECT es.session_id, es.program_name, es.host_name, es.login_name,
c.client_net_address
FROM sys.dm_exec_sessions es
LEFT JOIN sys.dm_exec_connections c ON c.session_id = es.session_id
WHERE es.program_name LIKE N'SQL Change Guard%';
```

PostgreSQL:

```
SELECT pid, application_name, client_addr, username, state
FROM pg_stat_activity
WHERE application_name LIKE 'SQL Change Guard%';
```

Oracle:

```
SELECT sid, serial#, audsid, username, machine, module, action, client_identifier
FROM v$session
WHERE module LIKE 'SQL Change Guard%' OR client_identifier LIKE 'SQL Change Guard%';
```

On an installation where everything runs on one machine, SQL Server may write "<local machine>" instead of an address. This is expected.

6. Testing the rule

On every execution, SQL Change Guard writes a "started" and an "ended" event to the audit trail. These events also go to the SIEM. An event contains these fields: windowKind (kind of work), serverName, executionLogin (account), executorHost (machine), clientAddress (address seen by the target), targetSessionId (session number), sessionApplicationName (session name), executionStartedAt and executionFinishedAt (UTC). The start event also carries the ITSM record of the request: changeTicket (for example CHG0030001) and, for an emergency request, emergencyTicket (for example INC0012345). This way a session in the DAM record is tied through the window directly to the ITSM record. In the end event, windowId is the number of the start event and the result is in windowOutcome: executed, failed, cancelled, rolledBack or abandoned.

To test:

10. Ask the organisation officer to run a small change request on the target server (for example, adding one row to a test table).
11. In the DAM record, find the session from that moment with the product account and from the product address.
12. Check that the session number and the time match the start and end events of SQL Change Guard.
13. The rule must count this session as "the work of the product" and must not raise an alarm.
14. Also test that an alarm is raised when a session with the same account is opened from somewhere other than the product address.

7. Seven example SIEM rules

The rules below are written independent of any vendor. Translate them into the language of your own SIEM product. The logic is this: a session on the target server that appears inside a window of the product, with the product account, from the expected address and with the same session number, has matched. A session that does not match is an alarm candidate.

Rules K1 to K7 look at change commands. There are three separate rules for reads (R1 to R3, below) and they cover only the reads (SELECT) made with the account of the product. The monitoring of reads by the other accounts of the organisation is the organisation's own policy. The Health, Preview, Definition and Test connections do not run change commands and do not open a window.

Prerequisites:

- Allowlist: the account and addresses of the product from the "Expected Access Principals" report, the address of the API machine by hand, the other accounts of the organisation from the organisation.
- The SIEM send threshold (SiemMinSeverity) must be Notice or lower and the Audit stream must be open. Window events have Notice severity.
- DAM, the database server, the machines of the product and the SIEM must use the same time source.
- Window table: the SIEM writes start events to a table keyed by the event number (windowId). When the end event arrives, it closes the same row. The rules look at this table.
- Rules that say "no event" (K2 and K5) must look at the time of the event itself (occurredAt) and be evaluated with a delay of a few minutes. During a collector outage events pile up, and looking at the arrival time produces a flood of alarms.

No	Rule	Logic	What it catches
K1	Change with an unlisted account	Change command on a production server and the account is not on the allowlist	Direct access, shared account
K2	Product account without a window	Change command with the product account and there is no open window on that server at that moment	The product password used elsewhere
K3	Product from a foreign address	Session with the product account or the product name and the client address is not in the address list of the product	Faking the name, machine name and account
K4	Request number mismatch	The number in the session name differs from the number of the window (request or package) open on that server at that moment	Working under the name of another request
K5	Window that	There is a start and no end with the	A crash where the

	never closes	same windowId within a set time. The time must be larger than the crash recovery threshold (default 120 minutes) plus the longest query timeout (MaxQueryTimeout, default 3600 seconds). Recommended: 4 hours	Worker does not return, a stuck job
K6	Emergency change not reviewed	isEmergency is true in the ChangeRequest.Executed event and there is no ChangeRequest.EmergencyReviewed within 24 hours	An emergency job that is never closed
K7	Another session inside a window	Inside an open window, a change command with the product account from a session different from the targetSessionId of the window	Someone using the account from a second session while a window is open

Read rules (query result requests): a query result request opens a separate Query window on each target server. The read is done inside a transaction on the target and rolled back without a commit; the window events are written for these reads as well.

No	Rule	Logic	What it catches
R1	Read with the product account without a window	SELECT with the product account and there is no open window on that server at that moment (unless the name is a product role that does not open a window)	Reading data with the product password
R2	Read from another session inside a window	While a Query window is open, a SELECT with the product account from a different session	Someone using the account from a second session while a window is open
R3	Query window exists, no DAM read record	The Query window closed with executed and DAM saw no SELECT in that interval	A gap in the read monitoring of DAM

The Health, Preview, Definition and Test connections also read but do not open a window. To avoid an R1 alarm for them, write an allow rule of "the product account, the product address and one of these role names". Because the name can be faked, this allowance must be given only to reads that come from the addresses of the product (API and Worker). The evidence of reads for the auditor is also in the product itself: the Query Result Activity, Masking Activity and Data Access and Export Log reports and the query result section of the evidence dossier.

Window that cannot be verified: if targetSessionId or clientAddress is empty in the start event (the session information could not be read), K3 and K7 cannot be evaluated for that window. Show this as a separate information row.

Limit: the people who can use the account of the product are the administrator of the API machine, the administrator of the Worker machine, and anyone who obtains the Worker access key. If they change something outside a window, K2 catches it; if they connect from another address, K3 catches it; if they act inside a window from another session, K7 catches it. No rule catches a person who enters the process of the product itself and sends the command from the session of the product. For this reason, access to the API and Worker machines must also be monitored.

8. Rules to pay attention to

- The application name is a label. The client says it itself, and someone can give the same name to their own connection. Do not treat the name alone as proof. Look for the account, the address seen by the target, the session number and the time range of the event together.
- The SQL account of the product must belong to this product only: it is not given for personal use, its password is in a vault and it is not used by another application. If the account is shared, the work of the product cannot be told apart from other work. The organisation officer records this with the "TargetAccountDedicated" declaration.
- The account name must not contain the words "token", "secret" or "password". Values like these are masked when they go to the SIEM.
- For a database behind a cloud service or a proxy or load balancer, the address seen by the target may be the address of an intermediate layer instead of the real address of the product machine. In that case do not treat the address alone as proof. The account, the session number and the time range can still be used.
- Window events that go to the SIEM have "Notice" severity. If the send threshold on the SIEM side is above Notice, or the Audit stream is closed, these events do not arrive. The organisation officer sees this on the Governance Readiness screen.
- If the end event of a window cannot be written or the process crashes, that window is closed later as "abandoned". The real result is in the request record.
- For organisations that use Oracle, the targetAuditSessionId field is not carried in CEF format; JSON, RFC 5424 or LEEF format is recommended. In CEF the ticket is in cs3 (changeTicket),

the incident record in flexString1 (emergencyTicket) and the request list of a package in flexString2 (requestIds).

9. If the addresses change

If the Worker or the API moves to another machine, the network changes or the database is moved to a different network, the allowlist must be updated. Take the new addresses from the report above each time, do not guess.