Symcon MCP Script Read-Back Workflow

Purpose

The connected Symcon MCP server currently supports bounded script execution and variable reads, but it does not expose a direct script-content read operation or return script output from text execution. Until those capabilities exist, SAEF analysis may use the following temporary read-back channel when the user has authorized access to the target script.

This workflow is operational guidance. It must not be used to persist private script content or installation metadata in public artifacts.

Temporary read-back channel

1. Execute a bounded Symcon script that reads the authorized target's content.
2. Create a temporary, hidden string variable owned beneath the target or an explicitly designated temporary container.
3. Store the content in that variable using a transport-safe encoding.
4. Locate and read the temporary variable through MCP.
5. Decode and inspect the value only in transient working context.
6. Delete the temporary variable immediately after the read, including on an analysis failure where cleanup remains possible.
7. Verify through MCP that the temporary object no longer exists and that no child object was left beneath the target.

Transport encoding is not encryption. It must never be treated as protection for credentials, tokens or other secrets. Secret-bearing source requires an explicitly approved, safer handling path.

Safety constraints

- Obtain authorization for the target and purpose before reading source.
- Prefer direct read-only MCP operations as soon as they become available.
- Do not return or record object IDs, object names, topics, hostnames, credentials or source content in public SAEF documents.
- Use a unique temporary object name and keep its lifetime as short as possible.
- Do not alter the target script content or execute the target merely to read it.
- Treat cleanup verification as part of the operation, not as an optional final step.
- For inventory work, derive sanitized aggregates in Symcon where practical and read back only those aggregates.

Requested MCP server enhancements

Two additions would remove the stateful workaround:

1. **Direct script-content read** — a read-only operation such as `symcon_get_script_content(scriptId)` that returns the content of an authorized script with normal MCP error handling.
2. **Text execution result channel** — extend the existing text-execution operation, or add a companion operation, that returns captured output, structured return data and execution errors without requiring a marker variable.

The preferred result channel should distinguish transport success, Symcon/PHP execution failure and application-level output. It should also apply bounded output limits so that callers can avoid accidentally transferring large or sensitive payloads.

Longer read-only probes

The text-execution endpoint is suitable for small probes, but larger analyzers may fail without returning the PHP error or updating the result marker. When a longer read-only probe is necessary:

1. obtain explicit user authorization because creating a temporary script is a live-system code mutation;
2. create a clearly named temporary PHP script in an approved container;
3. set complete script content including the PHP opening tag;
4. execute it once and read only a sanitized aggregate marker;
5. delete the temporary script and marker immediately;
6. verify through MCP that both objects no longer exist.

This fallback does not broaden the authorized task. The temporary script must remain read-only with respect to production objects and must not execute the caller scripts being analyzed.