

MCP write-access scorecard

What to require before an AI agent can write to your systems

Ask each vendor to show every behaviour live. Mark each cell **D** demonstrated, **S** shown on a slide only, **M** missing. Treat whatever they cannot demonstrate as missing.

Requirement	What to ask	A good answer	Demo test	Vendor A	Vendor B	Vendor C
1. Exact approval	Who approves, seeing what?	A person approves the final parameters, outside the agent's reach. Any change voids it.	Approve a write, change one parameter, resubmit. It must be refused.			
2. Write classification	Who decides a tool writes?	Your versioned manifest in the gateway. Unknown tools are denied.	Relabel a write tool as read-only. It must still need approval.			
3. Single execution	What does a retry do?	One approval, one idempotency key, at most one execution.	Send one approved call twice, then with a parameter changed, then have the agent re-propose it. Expect one execution and two refusals.			
4. Provable log	What can the log prove?	Append-only records in storage nobody can rewrite.	Replay one approved write from the log, then ask who can delete that record.			
5. The user's rights	Whose permissions does the call carry?	The user's, via token exchange. No broad agent account.	Have someone without refund rights ask the agent for a refund. The target system must refuse.			
6. Failure behaviour	What if approval is down?	Writes stop; pending writes have an owner.	Stop the approval service and ask the agent to write. Nothing should execute.			

Source: Rami Osman, AI and cloud architect, Paris. "What to require from an MCP server before an agent can write to your systems", 5 October 2026. <https://osmanrami.fr/writing/mcp-server-security-write-access/>
Licence: CC BY 4.0. You may share and adapt this sheet with attribution and a link to the source page.