



Free assessment tool

AI Workflow Readiness Scorecard

Score one workflow across seven dimensions before you commit budget to automating it.

Most AI pilots stall for reasons that were visible before the pilot started: the workflow was a poor candidate, the data was not usable, nobody owned the result, or there was no defined point at which a person approves the output. This scorecard scores a single named workflow across seven dimensions, 35 questions in total, using published scoring anchors so two people assessing the same workflow arrive at close to the same number.

Version 1.0. Published 24 August 2026.

Prepared for Operations, IT and business leaders who have been asked to apply AI to a specific process and need to establish whether that process is a sound candidate before funding it.

simplicityitinc.com

Contents

	What this scorecard is for
	What this scorecard does not do
	How to score it
1	Dimension 1. Workflow suitability
2	Dimension 2. Data readiness
3	Dimension 3. Security and access
4	Dimension 4. Governance and accountability
5	Dimension 5. Human approval design
6	Dimension 6. Operational ownership
7	Dimension 7. Implementation readiness
	Scoring sheet
	Result bands
	A worked example
	Claims and their basis
	Sources
	About Simplicity IT

What this scorecard is for

This scorecard assesses one workflow at a time. Not a department, not a strategy, not a technology. Pick a single process that somebody has proposed applying AI to, name it at the top of the scoring sheet, and answer the 35 questions about that process only.

The reason for the narrow scope is that readiness is not a property of an organization. The same company can be entirely ready to automate a document classification step and entirely unready to automate a customer commitment, because the two differ in how visible the data is, who has to approve the output, and what happens when it is wrong.

Each question scores 0, 1 or 2 against a published anchor. The maximum is 70. The score puts the workflow into one of four bands, and each band has a specific next action, which is usually not "start building".

What this scorecard does not do

Being explicit about the limits is what makes the rest of it usable.

- It does not predict that a project will succeed. It measures conditions observable before work begins. A high score means the common preventable failures have been addressed. It is not a forecast.
- It does not assess a vendor, a model or a product. Two implementations of the same workflow on the same platform can differ enormously in quality, and nothing here distinguishes them.
- It does not establish regulatory compliance. Where a workflow touches regulated data or regulated decisions, the questions here will surface that fact, and the answer is a specialist review, not a higher score.
- It does not replace a security review. Dimension 3 asks whether the basics are in place, at the level a business owner can answer. It is not a threat model.

How to score it

Twenty to forty minutes with the right people in the room. It goes wrong in predictable ways, so the method matters more than it looks.

- Name the workflow in one sentence that includes what starts it and what ends it. If that sentence needs an "and also", you have two workflows. Score them separately.
- Get three people: someone who does the work, someone who owns the system the data lives in, and someone who is accountable for the outcome. Most low scores come from the second and third of those never having been asked.
- Score against the anchor text, not against your impression. If the anchor for 2 says a document exists, and no document exists, the score is not 2 however well understood the process is.

- When the room disagrees, score the lower value and write down the disagreement. A contested 1 is more useful information than an agreed 2.
- Where nobody in the room knows the answer, score 0. An unknown is not a neutral. It is a thing that will be discovered later at a worse moment.
- Total each dimension out of 10, then total the seven dimensions out of 70.

SECTION 1

Dimension 1. Workflow suitability

Whether this particular piece of work is a sound candidate at all. This dimension fails more pilots than the technology does, and it fails them before anyone writes code.

W1 Question 1

The workflow has a defined trigger and a defined end state.

0	The workflow is described as an area of activity rather than a process. Nobody can say what starts it.
1	The trigger is clear but the end state is contested, or the process branches into several endings that have not been enumerated.
2	A single sentence names what starts the work and what condition means it is finished, and the room agrees with it.

W2 Question 2

Volume and frequency justify the build.

0	The workflow runs occasionally or unpredictably, and no one has counted it.
1	Volume is known and moderate. The case rests on quality or turnaround rather than on quantity, and that case has not been written down.
2	Volume and frequency are measured, and the measurement is written down. Whether the case rests on volume, on turnaround, or on consistency, it is stated in numbers somebody stands behind.

W3 Question 3

The work is language, document or data shaped rather than dependent on actions the solution cannot take.

0	Completing the work requires physical action, or changes in systems that offer no interface, or judgement that depends on information nobody has recorded anywhere.
1	The core of the work is language or document based, but one or more steps depend on a system with no usable interface, and no plan exists for those steps.
2	Every step either operates on text, documents or structured data, or has a documented interface. Steps that fall outside are identified and deliberately left with a person.

W4 Question 4

The current process is documented well enough that a new joiner could follow it.

0	The process lives in the heads of the people who do it. Two practitioners would describe it differently.
1	Documentation exists but is out of date, partial, or describes the intended process rather than the one actually performed.
2	A current written description exists, and someone who does the work has confirmed it matches what they actually do, including the exceptions.

W5 Question 5

A good output can be described in advance and recognised when produced.

0	Quality is a matter of taste, or is judged only by whether a downstream person complains.
1	The people who do the work can recognise a good output reliably, but the criteria have not been written down or agreed between them.
2	Written criteria exist, and a set of past examples has been marked good or not good against them, so the criteria have been tested on real cases.

If this dimension scored low

- If W1 or W3 scored 0, stop here. The workflow as described is not the unit of work to automate. Split it until each piece has one trigger, one ending and a usable interface, then score the pieces.
- If W4 scored 0 or 1, run a process walkthrough with two practitioners before anything else. It is the cheapest work in this entire list and it changes the design of everything that follows.
- If W5 scored 0, build a set of twenty past cases labelled good or not good. Without it there is no way to tell whether the built solution works, and no way to tell when it stops working.

SECTION 2

Dimension 2. Data readiness

Whether the information the workflow depends on can actually be reached, trusted and used. An assistant grounded in organizational content is only as good as the state of that content, and the state of that content is usually worse than the people responsible for it believe.

D1 Question 6

The information the workflow needs sits in systems the organization controls.

- | | |
|---|--|
| 0 | Key inputs live outside the organization, in personal accounts, in a supplier system with no export, or on individual machines. |
| 1 | Most inputs are in controlled systems, with known exceptions that have no plan attached. |
| 2 | Every input has been traced to a named system the organization administers, and the exceptions have been either eliminated or accepted in writing. |

D2 Question 7

The content is findable and current.

- | | |
|---|---|
| 0 | Several versions of key documents circulate and nobody can say which is authoritative. |
| 1 | An authoritative location exists but superseded copies remain reachable alongside it. |
| 2 | There is one authoritative location per input, superseded material has been removed or archived out of reach, and someone owns keeping it that way. |

Source [14](#)

D3 Question 8

The quality and coverage of the data have been looked at, not assumed.

- | | |
|---|---|
| 0 | Nobody has inspected the data. Quality is asserted from memory. |
| 1 | A sample has been reviewed informally and problems are known anecdotally. |
| 2 | A sample has been reviewed against the criteria from W5, and the gaps, error rates and missing fields are written down. |

Source [9](#)

D4 Question 9

Permissions on that content are correct today, before anything is switched on.

0	Sharing is broad or unreviewed. Nobody has looked at who can currently reach the content in scope.
1	Permissions have been reviewed for the primary locations, but sites with company wide or anonymous access have not been enumerated across the estate.
2	An oversharing review has been run across the locations in scope, over permissioned content has been remediated or restricted, and the review is scheduled to repeat.

Sources [11](#), [13](#)

D5 Question 10

Sensitive information inside the workflow is identified and labelled.

0	No classification exists. Whether the content includes regulated or confidential material is unknown.
1	A labelling scheme exists but is applied inconsistently, or only manually, across the content in scope.
2	Sensitive content in scope is labelled, the labelling is at least partly automatic, and the labels carry protection that applies wherever the content goes.

Sources [7](#), [8](#)

If this dimension scored low

- If D4 scored 0 or 1, treat it as the blocking item for the whole initiative. An assistant that can search everything a user can reach will find over permissioned content quickly and at scale, and will present it as a helpful answer.
- If D2 scored 0, the fix is content lifecycle work rather than AI work, and it delivers value whether or not the AI project proceeds.
- If D5 scored 0 and the workflow touches personal, financial, health or contractual data, get a classification exercise done before design starts. Retrofitting labels onto a live solution means reprocessing everything it has already touched.

SECTION 3

Dimension 3. Security and access

Whether the identity and access foundations underneath the workflow are sound. These questions are pitched so a business owner can answer them with help from IT. They are not a substitute for a security review.

S1 Question 11

The people who will use this are protected by modern authentication and access policy.

- | | |
|---|---|
| 0 | Multifactor authentication is optional or partially deployed, or legacy authentication protocols remain enabled. |
| 1 | Multifactor authentication is enforced for most users, but device and location conditions are not applied, or exclusions have accumulated without review. |
| 2 | Multifactor authentication is enforced, legacy authentication is blocked, access policies apply conditions beyond the password, and the exclusion list is short and reviewed. |

Sources [1](#), [2](#), [3](#)

S2 Question 12

The identity the solution itself runs as is scoped to the minimum it needs.

- | | |
|---|--|
| 0 | The plan is to run it with broad administrative rights, or with a shared account, or nobody has considered what identity it uses. |
| 1 | A dedicated identity is planned, but its permissions have not been enumerated and will be granted broadly to begin with. |
| 2 | A dedicated identity exists, the specific permissions it needs are listed, and there is an owner and a review date for that grant. |

Sources [4](#), [5](#)

S3 Question 13

Where the data goes, and where it is processed, is known and acceptable.

- | | |
|---|---|
| 0 | Nobody has established where the data is processed or which service handles it. |
| 1 | The processing location is known, but it has not been checked against the contractual or regulatory constraints that apply to this data. |
| 2 | Processing location and data handling are documented, and have been checked against the specific contractual and regulatory constraints on this data by someone qualified to do so. |

Sources [10](#), [12](#)

S4 Question 14

Access from outside the organization is understood.

0	External sharing settings and guest access have not been reviewed for the content in scope.
1	External access is known but broad, or guest accounts exist that nobody has reviewed recently.
2	External sharing is set deliberately for the locations in scope, guest access is reviewed, and cross tenant access settings are configured rather than left at defaults.

Sources [6](#), [15](#)

S5 Question 15

The solution will produce an audit record of who asked for what and what it returned.

0	No logging is planned, or logging is assumed to exist without anyone confirming it.
1	Platform level logging exists but nobody has confirmed it captures this workflow, and nobody has said how long it is kept.
2	The specific events this workflow produces have been identified, the retention period is set deliberately, and somebody knows how to retrieve a record when asked.

Source [12](#)

If this dimension scored low

- If S1 scored 0, that is a tenant wide issue and it is larger than this workflow. It should be resolved on its own merits regardless of what happens to the AI initiative.
- If S2 scored 0, decide the identity model before build starts. Broad rights granted for convenience during a pilot are rarely narrowed afterwards, because narrowing them risks breaking a thing that works.
- If S5 scored 0 and the workflow touches a regulated process, resolve it before launch. Reconstructing who saw what, months later, without a log, is not possible.

SECTION 4

Dimension 4. Governance and accountability

Whether a named person is answerable for what this produces. Governance here means a small number of decisions that have actually been made, not a policy document that exists.

G1 Question 16

One named person is accountable for the outputs of this workflow.

- | | |
|---|--|
| 0 | Accountability sits with a committee, a department, or nobody. |
| 1 | A name has been proposed but the person has not accepted it, or does not know. |
| 2 | A named individual has accepted accountability in writing and can state what they are accountable for. |

G2 Question 17

There is a written statement of acceptable use for this workflow.

- | | |
|---|---|
| 0 | Nothing written. Users will decide for themselves what to put in and what to do with what comes out. |
| 1 | A general organizational AI policy exists but says nothing specific about this workflow or this data. |
| 2 | A short written statement covers what may be entered, what the output may be used for, and what to do when the output looks wrong, and the users of this workflow have seen it. |

G3 Question 18

The consequence of a wrong output has been assessed, and the assessment is proportionate to the consequence.

- | | |
|---|--|
| 0 | Nobody has asked what happens if the output is wrong and acted on. |
| 1 | The question has been discussed informally. The worst case is understood in the room but not recorded. |
| 2 | The worst realistic outcome is written down, and the controls in place are proportionate to it, with someone senior enough to accept that risk having done so. |

G4 Question 19

Regulatory, contractual and client obligations touching this workflow are identified.

- | | |
|---|--|
| 0 | Nobody has checked whether client contracts, sector regulation or data protection obligations constrain this. |
| 1 | Obvious obligations are known, but the client contracts covering the data in scope have not been read for restrictions on processing or on automated handling. |
| 2 | Someone qualified has confirmed which obligations apply, has read the relevant contract terms, and has recorded the constraints they impose on this design. |

G5 Question 20

Changes to the solution will go through a defined change process.

0	Whoever built it will change it when asked. There is no record of what changed or when.
1	Changes are informally reviewed, but prompts, instructions or configuration can be altered without a record.
2	Every change including prompt and instruction changes is recorded, reviewed by someone other than the author, and reversible.

If this dimension scored low

- If G1 scored 0, resolve it before any build. An unowned solution in production is one where nobody is answerable for a bad output, which means the answer gets improvised under pressure.
- If G5 scored 0 or 1, note that prompt and instruction text is production configuration. Treating it as content that anyone can edit is how a working system silently becomes a different system.
- If G4 scored 0 and the organization handles client data under contract, get the contracts read before design. Restrictions on automated processing appear in ordinary commercial agreements more often than people expect.

SECTION 5

Dimension 5. Human approval design

Where a person reviews the output, and whether that review is real. A human approval point that a person cannot practically exercise is a rubber stamp, and a rubber stamp is worse than no control because it creates a record suggesting review occurred.

H1 Question 21

The point at which a person approves the output is defined.

0	Not designed. Either the output goes straight out, or review is assumed to happen somewhere.
1	Review is expected but its position in the flow is informal, and it can be skipped when the queue is long.
2	The approval point is explicit in the design, the flow cannot proceed past it without an action, and it is documented.

H2 Question 22

The approver can see enough to make a real judgement.

- | | |
|---|--|
| 0 | The approver sees only the final output, with no indication of what it was based on. |
| 1 | Some context is shown, but checking a questionable output means leaving the tool and searching manually. |
| 2 | The output is presented with its sources or inputs attached, so the approver can verify a specific claim without leaving the task. |

H3 Question 23

The approver has the authority, the expertise and the time.

- | | |
|---|---|
| 0 | The approver is junior to the decision, lacks the domain knowledge to spot a wrong answer, or would need more time per item than the process allows. |
| 1 | Two of the three hold. Time pressure or expertise is a known concern that has not been addressed. |
| 2 | The named approver has the standing to reject, the expertise to detect a wrong output, and enough time per item, and the time has been estimated rather than assumed. |

H4 Question 24

What the approver did is recorded.

- | | |
|---|---|
| 0 | Approvals leave no trace. There is no way to tell later whether an output was reviewed. |
| 1 | Approval is recorded, but edits and rejections are not, so the record shows only that something was accepted. |
| 2 | Approvals, rejections and edits are all recorded with who and when, which also produces the data needed to tell whether quality is improving. |

H5 Question 25

There is a defined path for when the output is wrong or the system is unavailable.

- | | |
|---|---|
| 0 | No fallback. If it fails or produces nonsense, people will improvise. |
| 1 | The manual process still exists and could be resumed, but nobody has confirmed the staff and access to run it are still in place. |
| 2 | The fallback is written down, someone has confirmed it still works, and users know how to trigger it and who to tell. |

If this dimension scored low

- If H3 scored 0, redesign before building. Volume that exceeds what an approver can review is the mechanism by which a human approval point becomes a formality, and it usually appears in the numbers before launch if anyone does the arithmetic.
- If H2 scored 0, require source attribution in the design. An approver who cannot check a claim without leaving the task will stop checking within weeks.
- If H4 scored 0 or 1, add rejection and edit capture. It is the only cheap source of quality data you will get, and it is not recoverable retrospectively.

SECTION 6

Dimension 6. Operational ownership

Who runs this after the people who built it have moved on. Pilots are staffed by their enthusiasts. Production is not, and the handover is where most of them quietly stop being maintained.

O1 Question 26	
A named team or person will run this after launch.	
0	Not discussed. The implicit assumption is that whoever builds it keeps it.
1	A team has been identified but has not agreed, has not been resourced, or has not seen what they are being asked to take on.
2	A named owner has accepted the running of it, knows what the work involves, and has capacity allocated for it.

O2 Question 27	
Users have somewhere to report a problem.	
0	No route. Users will message whoever they remember was involved.
1	An informal route exists that depends on specific individuals being available.
2	Problems go through the normal support channel, the people staffing it know this workflow exists, and there is a documented path to whoever can act on it.

O3 Question 28

Output quality will be measured after launch, not just before.

0	No measurement planned. Quality will be assessed by whether anyone complains.
1	Usage will be measured. Quality will not, or only through occasional informal review.
2	A defined measure of output quality, a sampling schedule and a named reviewer exist, and a threshold has been set at which the workflow is paused.

O4 Question 29

Running costs have an owner and a budget.

0	Costs have not been estimated, or sit on a card belonging to a project that ends.
1	Costs are estimated and funded for the first year, with no plan beyond it and no alerting on overrun.
2	Consumption is understood, a budget line owns it beyond the project, and someone is alerted when usage moves outside the expected range.

O5 Question 30

There is a way to turn it off.

0	No plan. Disabling it would be improvised and its effects on downstream processes are unknown.
1	It could be disabled, but the manual process it replaced has degraded and the impact of stopping is not documented.
2	A documented rollback exists, what happens to in flight work is defined, and someone has the authority to invoke it without convening a meeting.

If this dimension scored low

- If O1 scored 0, the pilot has no path to production regardless of how well it performs. Identify the run owner before the build, and have them present during it.
- If O3 scored 0, add quality sampling now. The inputs and the content this depends on will change, and a solution that quietly degrades looks exactly like one that is working until someone acts on a bad output.
- If O5 scored 0, write the rollback. It is an hour of work that determines whether a bad week is a bad day.

SECTION 7

Dimension 7. Implementation readiness

Whether the organization can actually execute this now. A workflow can be an excellent candidate and still be the wrong thing to start this quarter.

I1 Question 31

Success is defined as a number, and the current value of that number is known.

- | | |
|---|--|
| 0 | Success is described in adjectives. Faster, better, more efficient. |
| 1 | A measure has been chosen but the baseline has not been taken, so improvement will be unprovable. |
| 2 | A measure, its current value, and the target are written down, and the person who will be held to the target has agreed to it. |

I2 Question 32

There is a sponsor with the authority to change how the work is done.

- | | |
|---|---|
| 0 | The initiative is being driven from a level that cannot change the process it is automating. |
| 1 | A sponsor exists but is not engaged, or can fund the work but cannot direct the affected team. |
| 2 | The sponsor owns the process being changed, is funding it, and is available to resolve decisions during delivery. |

I3 Question 33

The people who do the work today are involved and are not being surprised by it.

- | | |
|---|---|
| 0 | They have not been told. The initiative is being designed around them. |
| 1 | They have been informed but not consulted, and their concerns about accuracy or role have not been addressed. |
| 2 | Practitioners are in the design, their objections have been recorded and answered, and what changes about their role has been said plainly. |

I4 Question 34

A place to build and test exists that is not production.

- | | |
|---|--|
| 0 | No non production environment. Testing would happen against live data and live processes. |
| 1 | An environment exists but does not have representative data, so testing will not tell you much. |
| 2 | A non production environment exists with representative data, obtained in a way that does not itself create a data protection problem. |

I5 Question 35

The subject matter experts have committed specific time.

0	Expert availability has not been discussed. It is assumed people will make time.
1	Time has been promised in principle but not scheduled, and the experts have competing commitments.
2	Named experts have specific time in their calendars, and their managers have agreed to it.

If this dimension scored low

- If I1 scored 0 or 1, take the baseline before starting. It takes days and it cannot be reconstructed afterwards, which means without it the project can never be shown to have worked.
- If I3 scored 0, address it before anything visible happens. The practitioners are the people who know the exceptions, and the exceptions are what determine whether the solution survives contact with real work.
- If I5 scored 0 or 1, expect the schedule to slip by roughly the amount of expert time that was assumed rather than booked.

Scoring sheet

Record the workflow name, the date, and who was in the room. A score without those three is not comparable to any other score, including a later score of the same workflow.

Dimension	Q1	Q2	Q3	Q4	Q5	Total
Dimension 1. Workflow suitability						
Dimension 2. Data readiness						
Dimension 3. Security and access						
Dimension 4. Governance and accountability						
Dimension 5. Human approval design						
Dimension 6. Operational ownership						
Dimension 7. Implementation readiness						
Overall score, out of 70						

Print this page, or copy the grid. Score every question before totalling, and record the reason for any 0 next to it. The reasons are the working document; the number is a summary of them.

Result bands

The band is a starting point for a conversation, and the two rules underneath it matter more than the total. Any dimension scoring 4 or below is a blocking issue whatever the total says, and any 0 on D4, G1, H3 or O1 should be resolved before build regardless of band.

0 to 27 Not ready	The preconditions for this work are not present. Building now produces a demonstration that cannot become a production system. The useful next step is the two or three lowest scoring dimensions, treated as their own piece of work with their own value, most of which is worth doing whether or not the AI project ever proceeds.
28 to 41 Foundational gaps	The workflow may be a reasonable candidate but the ground underneath it is not solid. Identify every dimension scoring 6 or below and close those specific items first. A limited proof of concept on non sensitive data can run in parallel, provided nobody treats it as a pilot and it has an agreed end date.
42 to 55 Conditionally ready	Ready for a scoped implementation with named conditions attached. Write down every item that scored 0 or 1, decide for each whether it is closed before launch or accepted as a known risk with an owner, and get that list agreed by the sponsor. The conditions are the deliverable here, not the score.
56 to 70 Ready to implement	The known preventable failures have been addressed. Proceed to design, keeping the remaining 1s on the risk register with owners. This band means the workflow is a sound candidate and the organization can execute. It does not mean the implementation will be good, which is a separate matter of engineering and of the people doing it.

A worked example

An illustrative example, constructed to show the method. It is not a Simplicity IT client and describes no real organization.

The workflow: when a supplier invoice arrives in a shared mailbox, extract the supplier, amount, purchase order reference and dates, check them against the purchase order, and either queue the invoice for payment or route the exception to a buyer. Roughly 900 invoices a month, three staff, two days of lag at month end.

Workflow suitability scored 9. Clear trigger, clear ending, document shaped, well documented, and the team could produce a hundred past invoices already marked handled correctly or not. Only W2 lost a point, because the two day lag had never been measured, only complained about.

Data readiness scored 6. Invoices and purchase orders were in controlled systems, but the supplier reference data existed in two places that disagreed, and the finance shared drive had company wide access nobody had reviewed in four years. D4 scored 0.

Security and access scored 8. Multifactor authentication enforced, legacy authentication blocked, logging in place. S2 scored 1 because the intended service identity was going to be granted broad access to the finance site to save time during the build.

Governance scored 7. A named finance systems manager accepted accountability without hesitation. G4 scored 1, because nobody had checked whether the outsourced payment provider contract said anything about automated processing, and it turned out to.

Human approval scored 5. The design routed only exceptions to a person, which was sound, but H3 scored 1: the estimate was 40 exceptions a day against one buyer who had 90 minutes, and nobody had done that arithmetic until this question was asked.

Operational ownership scored 4. No run owner, no quality measurement, no rollback. This is the ordinary state of a project that has been scoped as a build rather than as a service.

Implementation readiness scored 8. Sponsor engaged, practitioners involved, experts booked. I1 scored 0 because success was described as "faster month end".

Total 47, conditionally ready. The three things that changed as a result: the oversharing review on the finance drive was done first and took a week; the exception volume was recut so that only two categories routed to the buyer, bringing it to twelve a day; and a run owner was named before the build started rather than after. The measured two day lag became the baseline for I1. None of that required new technology, and all of it would have surfaced six months later at considerably greater cost.

Claims and their basis

Every claim this document makes about itself is listed here with what it rests on. Nothing here is a claim about your organization, a guarantee of compliance, a predicted saving, or a promised outcome.

The scorecard produces a reproducible score.

Every question publishes the observable condition that earns 0, 1 and 2. Reproducibility is a property of the published anchors, which anyone can inspect. It is not a measured inter rater result and is not presented as one.

The seven dimensions cover the common causes of stalled AI work.

The dimensions are Simplicity IT delivery judgement, informed by Microsoft guidance on Copilot and agent readiness. They are presented as a practitioner framework, not as a Microsoft framework and not as research.

A high score does not predict a successful project.

Stated in the document. The scorecard measures readiness conditions that are observable before work starts. No outcome, saving, or success rate is claimed.

Sources

Every source cited in this document, numbered in the order it is first referenced. The numbers on the items above point here. Each source was checked on the verification date shown. Microsoft documentation changes, so re-check anything you are about to act on.

1. Common Conditional Access policies

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/entra/identity/conditional-access/concept-conditional-access-policy-common>

2. Common Zero Trust identity and device access policies

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/security/zero-trust/zero-trust-identity-device-access-policies-common>

3. Block legacy authentication with Conditional Access

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/entra/identity/conditional-access/block-legacy-authentication>

4. Start using Privileged Identity Management

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/entra/id-governance/privileged-identity-management/pim-getting-started>

5. Best practices for Microsoft Entra roles

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/entra/identity/role-based-access-control/best-practices>

6. Configure cross-tenant access settings for B2B collaboration

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/entra/external-id/cross-tenant-access-settings-b2b-collaboration>

7. Learn about sensitivity labels

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/purview/sensitivity-labels>

8. Learn about sensitive information types

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/purview/sit-sensitive-information-type-learn-about>

9. Get started with content explorer

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/purview/data-classification-content-explorer>

10. Microsoft 365 Copilot requirements

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/microsoft-365/copilot/microsoft-365-copilot-minimum-requirements-data-compliance>

11. Configure a secure and governed data foundation for Microsoft 365 Copilot

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/microsoft-365/copilot/configure-secure-governed-data-foundation-microsoft-365-copilot>

12. Microsoft 365 Copilot architecture, data protection and auditing

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/microsoft-365/copilot/microsoft-365-copilot-architecture-data-protection-auditing>

13. Data access governance reports for SharePoint sites

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/sharepoint/data-access-governance-reports>

14. Site lifecycle management

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/sharepoint/site-lifecycle-management>

15. Turn external sharing on or off for SharePoint and OneDrive

Microsoft Learn. Verified 2026-08-24.

<https://learn.microsoft.com/sharepoint/turn-external-sharing-on-or-off>

About Simplicity IT

Simplicity IT Inc. is a Microsoft-focused managed services, security and implementation company. We began in 2011. We assess, implement, migrate, integrate, secure, govern and operate Microsoft cloud environments, and we design the human approval points that keep automated work accountable.

If this document was useful

Book an engineering discovery session at simplicityitinc.com/contact/

Related reading

- AI implementation and governance, simplicityitinc.com/solutions/ai-implementation/
- Forward Deployed AI Engineering, simplicityitinc.com/services/forward-deployed-ai-engineering/

Privacy

We use your details to send you this resource and occasional related guidance. We do not sell your information and you can ask us to delete it at any time by emailing hello@simplicityitinc.com.

What requesting this document does not do

Requesting a resource does not create an account, a trial, a subscription, partner status, or an engagement. It sends you a document.

Version 1.0. Published 24 August 2026. Copyright 2026 Simplicity IT Inc. This document may be shared unmodified.