



CLOUD & AZURE ARCHITECTURE

Auditing your Azure architecture from review to remediation

A complete method to run a Well-Architected Review on Azure, turn findings into a prioritised remediation plan and manage residual risks over time.

11 pages · Risk scoring grid, prioritisation matrix and 30/90/180-day remediation plan included

5

Azure Well-Architected pillars assessed during the review

30/90/180

days: the three horizons of the remediation plan

100%

of findings tied to a tracked risk and a named owner

CONTENTS

What you will find in this document

This document starts from the context and triggers of a Well-Architected audit, details the method for collecting, assessing and scoring risks, then leads to a prioritised remediation plan and an operational follow-up framework.

01	Executive summary — the essentials on one page	p.3
02	Why audit your architecture now	p.4
03	Scoping and collecting evidence	p.5
04	Assessment tools: what to choose	p.6
05	Scoring and prioritising risks	p.7
06	Remediation plan and follow-up	p.8
07	Self-assessment grid	p.9
08	A 5-step roadmap	p.10
09	Conclusion & about Move2Cloud	p.11

Who is this document for? CTOs, cloud architects, platform leads and CISOs in organisations already running production workloads on Azure, whether a mature landing zone or an estate migrated without a formal review. Also useful to SRE and FinOps teams tasked with carrying out the corrective actions.

EXECUTIVE SUMMARY

The essentials on one page

Turning an architecture review into dated, costed and tracked engineering decisions — rather than a report filed away.

The Microsoft Azure Well-Architected Framework structures the assessment of a workload around **5 pillars**: reliability, security, cost optimisation, operational excellence and performance efficiency. Yet most organisations run architectures designed during accelerated migrations, where these pillars were never explicitly examined: year after year, Flexera's State of the Cloud surveys put estimated waste at around **30% of cloud spend**, a direct symptom of architectures that have never been reviewed.

The gap is widening for two reasons. First, the platform evolves faster than the architectures running on it: service deprecations, new SKUs, new availability guarantees — a single-instance VM stays at **99.9%** SLA while a multi-zone deployment reaches **99.99%**. Second, the shared responsibility model places most security failures on the customer configuration side, not the provider's: Gartner has long estimated that nearly all cloud security incidents stem from customer errors.

The good news: a Well-Architected Review is toolled and repeatable. Between Microsoft's WAF assessment, Azure Advisor, Defender for Cloud and platform metrics, most findings can be measured without additional instrumentation. This document provides the method to run the review in a few weeks, rate each gap according to its impact and effort, build a remediation plan prioritised in waves, and maintain a register of residual risks accepted at executive level.

5

pillars of the Azure Well-Architected Framework

~30%

of cloud spend estimated as wasted (Flexera, State of the Cloud)

99.99%

SLA for a multi-zone VM, against 99.9% for a single instance

The question this document answers is not "is our Azure architecture good?", but "which gaps do we fix first, at what cost, and which risks are we consciously accepting until the next review?". The sections that follow set out the full method, from scoping the review to the residual risk register, with grids and checklists that can be reused directly.

SECTION 02

Why audit your architecture now

An Azure architecture is not fixed: it drifts, under the effect of fast migrations, platform changes and trade-offs made under pressure. The question is therefore not whether to audit it, but at what point the cost of delay exceeds the cost of the review.

1. 1. Architecture debt inherited from migrations

Most Azure estates were built through successive migrations, often lift & shift, without any formal architecture review. The choices made then — single instance, local redundancy, flat network — remain in production long after the workloads have become critical.

2. 2. Design dictates real availability

Availability gaps come down to design details: a single-instance VM on Premium disk is covered at 99.9%, against 99.99% as soon as it is deployed across several availability zones. An audit makes these gaps explicit before the incident, not after.

3. 3. A platform moving faster than the architectures

Microsoft evolves and deprecates its services continuously (retirement of TLS 1.0/1.1, classic resources, legacy APIs). Without periodic review, an architecture that was compliant yesterday becomes unsupported, non-compliant with NIS2/DORA, or simply more expensive than its modern equivalent.

“ You don't audit an architecture to produce a report, but to decide what you fix on Monday morning. ”

The following sections detail the method: scoping the review, examining the five pillars, rating the findings, building the remediation plan and steering residual risks.

SECTION 03

Scoping and collecting evidence

A Well-Architected Review never applies to "all of Azure": it applies to a bounded workload, with its subscriptions, its dependencies and its service commitments.

SCOPING ITEM	WHAT MUST BE AGREED BEFORE STARTING
Technical scope	The named list of subscriptions, resource groups and external dependencies (SaaS, on-premises systems, DNS, identities) included in the review — and what is explicitly excluded from it.
Service objectives	The RTOs, RPOs, availability and latency SLOs expected by the business, formalised per workload: without a target, no reliability finding can be demonstrated.
Expected evidence	For each pillar, the authoritative exports: Resource Graph inventory, Azure Policy compliance, Defender for Cloud secure score, Advisor recommendations, Cost Management data and existing architecture documentation.

Key takeaway: Set the scope on a named workload, with its owner and its service objectives, and require exportable evidence for every finding: without a dated artefact, a finding is only an opinion.

Automate collection before opening the workshops

Collection should be largely automated and reproducible: an Azure Resource Graph export of the inventory (SKUs, zonal redundancy, runtime versions), the Azure Policy compliance state, the Defender for Cloud secure score, the Azure Advisor recommendations and a twelve-month Cost Management export form the factual baseline. These extractions can be replayed identically at the end of remediation, which lets you measure the delta rather than declare it. Workshops with the teams come afterwards: they serve to explain the observed gaps (accepted debt, regulatory constraint, budget trade-off) rather than to discover the existing estate, which typically reduces the interview phase to two or three two-hour sessions per pillar.

SECTION 04

Assessment tools: what to choose

No single tool covers all five pillars on its own: the structured assessment provides the framework, Azure telemetry provides the evidence, and custom queries fill the remaining gaps.

TOOL	COVERAGE	HOW IT WORKS	BEST SUITED TO
Azure Well-Architected Review (Microsoft Assessments)	The 5 pillars, via questionnaire	Guided questionnaire on a workload, returning a score per pillar and recommendations exportable to CSV	Structuring the review and obtaining a baseline of recommendations aligned with the Microsoft framework
Azure Advisor + Defender for Cloud	Reliability, security, cost, performance	Continuous analysis of the actual configuration of resources, with secure score and prioritised recommendations	Objectifying configuration gaps and measuring progress after remediation
Azure Resource Graph + Azure Policy	Custom compliance and inventory	KQL queries at tenant scale and evaluation of rule sets defined by the organisation	Verifying the internal and regulatory requirements that standard tools do not cover

Common pitfall: Mistaking a secure score or an Advisor score for a level of architectural maturity: these scores rate the configuration of existing resources, not design choices. A perfectly configured single-region architecture can display an excellent score while being incapable of meeting its RTO.

Consolidate the three sources into a single repository of findings

The right setup consists in making the three levels converge into a single repository of findings. In practice: you import the CSV export from the Well-Architected assessment, attach to it every Advisor recommendation and every non-compliant Defender for Cloud control, then complete it with the Resource Graph queries that surface the blind spots (resources outside availability zones, databases without geo-redundant backup, storage accounts still open to the public Internet). Each line then carries an identifier, a pillar, an owner, a piece of evidence and a rating — it is this table, not the PDF report, that becomes the working object of the following sections.

SECTION 05

Scoring and prioritising risks

An unrated finding is a finding that will never be arbitrated: rating turns a list of gaps into an order of play that can be defended before a steering committee.

Rate impact and likelihood, not technical annoyance

Rate each finding on two axes: the business impact if it materialises (unavailability, data leak, budget overrun) and the likelihood of occurrence given the actual exposure. The product of the two gives the raw severity.

Weight by effort to isolate the quick wins

Two findings of equal severity are not handled in the same order: relate severity to the estimated remediation effort in person-days. The highest ratios form the quick-win wave, to be started within 30 days.

Reserve a priority lane for security and restoration

A finding related to authentication, encryption or backup restoration goes first, whatever its rating: these defects have no acceptable degraded mode and determine the ability to absorb a major incident.

CRITICALITY LEVEL	ASSOCIATED COMMITMENT RULE
Critical	Data loss, public exposure or failure to meet the RTO: fix started within 30 days or waiver signed by the sponsor.
Major	Significant degradation of service, cost or operability: scheduled within the quarter, with an owner and a target date.

SECTION 06

Remediation plan and follow-up

The value of a review is measured by the number of findings actually closed six months later, not by the number of pages in the report.

Turn findings into deliverable work packages with a closure criterion

Group findings into coherent work packages by team and technical domain — "harden network exposure", "make data restorable", "make the recovery plan reliable" — rather than into a list of isolated tickets. Each package is given a named owner, an estimated effort, a target date and a verifiable closure criterion: the Resource Graph query or the Policy control that must turn green. This formulation makes it possible to close a package without debate and to feed the items directly into existing sprints, instead of creating a parallel workstream that competes with the product roadmap.

Maintain a living residual risk register

Not all findings will be fixed, and that is a legitimate trade-off. What is not legitimate is leaving the risk untracked: every untreated gap joins a register stating the nature of the risk, the estimated impact, the reason for acceptance, the person accountable for it and the review date. This register is read in committee, quarterly, together with the progress indicators — Azure Policy compliance, Defender for Cloud secure score, monthly cost drift and, above all, the result of the last restoration test actually performed.

Key takeaway: Formalise every untreated risk in a dated, named register with a review date: a risk accepted explicitly is a risk under control, a forgotten risk is an on-call surprise.

The remediation plan is not an end-of-engagement deliverable but an item of the product backlog, reviewed at the same rhythm as other engineering commitments. A three-hour quarterly review, based on the same exports as at scoping, is enough to hold the trajectory: it measures the compliance delta, closes the delivered work packages, re-examines the accepted risks and incorporates the findings arising from changes since the previous review. It is this repetition, more than the quality of the initial report, that distinguishes a useful audit from a shelved one.

SECTION 07

Self-assessment grid

Review these six questions before launching your next Well-Architected Review, or in the weeks following the debrief of an audit already carried out.

- ☐ **Does each critical workload have a named owner and service objectives (SLO, RTO, RPO) that are written down and validated by the business?** Without an owner or a quantified objective, no architecture finding can be arbitrated or prioritised.
- ☐ **Do you have, for every finding from your last review, a dated and exportable artefact (ARM/Bicep export, Advisor screenshot, diagram, policy extract)?** A finding with no dated evidence is debated endlessly and loses all value at the next review.
- ☐ **Can you distinguish, in your reporting, what relates to resource configuration (secure score, Advisor) and what relates to design choices?** An excellent configuration score can mask an architecture structurally incapable of meeting its RTO.
- ☐ **Are your findings scored against an explicit scale combining impact and likelihood, rather than ranked by expert gut feeling?** Only consistent scoring makes it possible to defend a running order before an investment committee.
- ☐ **Are remediation batches embedded in the product backlog, with an estimated effort and a target sprint, rather than kept in a separate document?** A plan that lives outside the backlog is never weighed against features and ends up shelved.
- ☐ **Are the risks you have decided not to address formally accepted, dated, signed and reviewed at least once a year?** An accepted risk with no review date becomes a blind spot as soon as the technical or regulatory context changes.

How to read your answers: Two "no" answers signal a point of vigilance to address within the quarter; three "no" answers or more indicate that your audit approach will produce findings with no lasting effect.

SECTION 08

A 5-step roadmap

This trajectory unfolds over one to two quarters and builds on existing forums — product committees, sprint reviews, security committee — without creating parallel governance or freezing development in progress.

- 1 Scope a workload**
Choose a named workload, identify its business owner and formalise its service objectives. Delimit the subscriptions, resource groups and external dependencies included, and what is explicitly out of scope.
- 2 Collect the evidence**
Export configurations (Bicep/ARM, Azure Policy policies, network topologies), Advisor and Defender for Cloud recommendations, availability metrics and costs. Date and archive every artefact: it will serve as the reference for the next review.
- 3 Run the review**
Test the evidence against the five pillars of the Well-Architected Framework in workshops with the build and run teams. Systematically distinguish configuration gaps from design gaps, which automated scores do not detect.
- 4 Score and arbitrate**
Apply an impact × likelihood scale to each finding, tie it to a measurable service objective and rank it. Validate the running order in committee, formally recording the risks you choose to accept.
- 5 Industrialise the follow-up**
Turn priority findings into costed engineering batches, injected into the product backlog with a target sprint. Set up a three-hour quarterly review based on the same exports to measure the delta, close out delivered batches and incorporate new findings.

CONCLUSION

An audit is only worth the trajectory it triggers

Organisations that derive real benefit from a Well-Architected Review are not those that produce the most exhaustive report, but those that can convert a handful of well-scored findings into delivered engineering batches, then verify that the gaps closed stay closed. They gain on three fronts: visibility shared between IT and the business on what holds up and what does not, a documented budget trade-off between architecture debt and new features, and the ability to answer an auditor or a demanding client within hours, evidence in hand.

Beyond Azure, the discipline described here — named perimeter, dated evidence, explicit scoring, batches in the backlog, quarterly review — is transposable to any cloud foundation and to adjacent exercises: FinOps review, NIS2 or DORA preparation, qualification of a disaster recovery plan. The framework changes, the mechanics stay the same. The point is not to add one more ritual, but to reconnect architecture to the steering forums that already exist.

Move2Cloud	is a consulting firm specialised in infrastructure, cloud and IT transformation. We support organisations in the audit, remediation and ongoing governance of their cloud environment.
Contact	Move2Cloud — Paris · move2cloud.fr
Next step	If several answers in the self-assessment grid are negative, a targeted two- to three-week diagnostic on a single critical workload is enough to produce a scoped perimeter, a list of scored findings and a first costed remediation batch.

Sources

Microsoft — Azure Well-Architected Framework, Microsoft Learn documentation — Microsoft — Azure Well-Architected Review and Azure Advisor (Microsoft Learn) — Microsoft — Cloud Adoption Framework for Azure, Govern and Manage — CIS — Microsoft Azure Foundations Benchmark — ANSSI — Recommendations for cloud computing architectures and usage

Document written by Move2Cloud as an example and illustration, for educational purposes. Azure tooling and Well-Architected Framework recommendations change frequently: check the current Microsoft documentation before any architecture decision.