ADS Cheat Cards

All five reference cards

Card 1 — The Views at a Glance

Tier 1: Architectural Views (Section 3)

#	View	What it describes	Primary audience
3.1	**Logical View**	Components, services, patterns — *what the solution is*	Architects, Developers
3.2	**Integration & Data Flow**	Data flows, APIs, integrations — *how components talk*	Integrators, Architects
3.3	**Physical View**	Hosting, compute, networking, environments — *where it runs*	Infrastructure, DevOps
3.4	**Data View**	Data stores, classification, retention — *what data it holds*	Data Architects, Compliance
3.5	**Security View**	IAM, encryption, monitoring, threats — *how it's protected*	Security, CISO, Compliance
3.6	**Scenarios**	Use cases, ADRs — *how it behaves and why*	All stakeholders

Remember: No single view tells the full story. Together, they answer every reasonable question from any reasonable reviewer.

Card 2 — Quality Attributes at a Glance

Tier 2: Quality Attributes (Section 4) — cross-cutting

#	Attribute	Assess across every view	Key questions
4.1	**Operational Excellence**	Can we operate it in production?	Observability, runbooks, alerts, incident response
4.2	**Reliability & Resilience**	What happens when something fails?	RTO, RPO, DR strategy, backup, fault tolerance
4.3	**Performance Efficiency**	Does it meet its targets at load?	P95/P99 latency, throughput, concurrency, growth
4.4	**Cost Optimisation**	Is it priced to the value delivered?	Capex/opex, unit economics, commitment discounts, exit cost
4.5	**Sustainability**	Is it environmentally responsible?	Carbon baseline, efficient resources, SCI metric

Document tradeoffs explicitly — when one quality attribute constrains another, name the trade-off and the rationale.

Security is a View (3.5) — it's cross-cutting enough to warrant its own structure, so don't look for it here.

Card 3 — RAID Template

Section 6.1–6.5 Decision Making & Governance

Constraints (6.1) — fixed limitations the design must work within

ID	Constraint	Category	Impact	Last Assessed
C-001		regulatory / technical / commercial / organisational / time		

Assumptions (6.2) — things treated as true but not verified

ID	Assumption	Impact if false	Certainty	Status	Owner
A-001			high / medium / low	open / closed	

Risks (6.3) — potential events with negative impact

ID	Risk	Severity	Likelihood	Owner	Mitigation	Residual
R-001		H/M/L	H/M/L			H/M/L

Dependencies (6.4) — external factors this design relies on

ID	Dependency	Direction	Status	Owner
D-001		inbound / outbound	committed / not-committed / resolved	

Issues (6.5) — problems that have already materialised

ID	Issue	Impact	Owner	Resolution Plan	Status
I-001		H/M/L			open / in-progress / resolved

Golden rule: Every row needs an owner and a date. Anonymous = unaccountable.

Card 4 — Compliance Scoring Scale

0–5 Compliance Scoring

Score	Level	What it means
0	Not Addressed	No evidence for this area
1	Acknowledged	Concern recognised, no design or evidence
2	Partial	Some requirements met, significant gaps
3	Mostly Addressed	Most requirements met, minor gaps (passing grade)
4	Fully Addressed	All requirements met with evidence
5	Exemplary	Reference-quality — best practice

Overall score

The overall score is the **lowest individual section score**, not an average. A solution with 5 in Performance but 1 in Security is a "1", not a "3" — the Security gap is the concern.

What to score

14 sections: 1, 3.1–3.6, 4.1–4.5, 5, 6 (Executive Summary, the six Views, the five Quality Attributes, Lifecycle, Decision Making).

Organisation thresholds (suggested defaults)

- All sections >= 3 -> production approval
- All sections >= 4 -> Tier 1/2 critical systems
- Section = 0 or 1 -> remediation plan required

Card 5 — Documentation Depth Chooser

Pick your depth

Depth	When to use	Typical effort	Governance gate
Minimum	PoC, dev/test systems, Tier 5 internal tools	1–3 hours	Development review
Recommended	Production-bound systems, Tier 3–4	1–2 days	Production approval
Comprehensive	Regulated, Tier 1–2 critical	1–2 weeks	Enterprise review

Matching Business Criticality to Depth

- **Tier 1 Critical** -> Comprehensive
- **Tier 2 High** -> Comprehensive
- **Tier 3 Medium** -> Recommended
- **Tier 4 Low** -> Recommended
- **Tier 5 Minimal** -> Minimum

Rule of thumb

Over-documenting a simple tool burns goodwill. Under-documenting a critical system hides risk. Calibrate to risk, not to tradition.

Card 6 — Sustainability Quick Hits

Where carbon footprint is actually decided

Most of a SAD's environmental impact is locked in by a small number of decisions. Get these right and Section 4.5 becomes evidence rather than aspiration.

Decision	Where in the SAD	Quick win
Cloud region	3.3 Physical View	Pick a region with high renewable energy share. Carbon intensity varies 5–10x across regions of the same provider.
Non-prod runtime	3.3 + 5.5 Operations	Auto-shutdown dev/test out of hours. Typical saving: 60–70% of non-prod compute cost and carbon.
Compute family	3.3 Physical View	ARM (Graviton, Ampere, Cobalt) and latest-generation x86 deliver 20–40% better performance-per-watt than older instance types.
DR posture	3.3 + 4.2 Reliability	A warm standby running 24x7 doubles compute footprint. Match DR mode (cold / pilot light / warm / hot) to the actual RTO.
Data retention	3.4 Data View	Set a retention policy and automate expiry. Indefinite "keep just in case" is the most common waste.
Cold storage tiering	3.4 Data View	Move logs/snapshots/backups to archive tiers (S3 Glacier, Azure Archive). Cost and carbon both drop sharply.
Caching & async	3.1 Logical View	Avoid recomputation. Cache responses, debounce polls, replace busy-wait with events.
Right-sizing cadence	5.5 Operations	Quarterly right-sizing using cloud advisor tools catches drift. Without a cadence, over-provisioning grows linearly.

The three Green Software principles

Principle	Meaning	Typical action
Carbon efficiency	Emit the least carbon for a given output	Region choice, ARM compute, retention
Energy efficiency	Use the least energy for a given task	Caching, async, code-level efficiency
Carbon awareness	Do more when the grid is clean, less when it isn't	Time-shift batch jobs, defer non-urgent work

When to skip sustainability prompts

A Tier 5 internal tool running on shared SaaS doesn't need a Software Carbon Intensity baseline. Match the depth of treatment to the system's actual footprint. A 50-instance e-commerce platform deserves the full Section 4.5 treatment; a fortnightly batch job on shared compute does not.
