THE METRICS **THAT ACTUALLY MATTER**

- % tier-0 workloads with sovereign-ready plan
- Recovery outcomes for a “cloud region/vendor disruption” scenario
  (time-to-recover, ability to operate degraded, success rate of failover tests)
- # of findings related to data residency/sovereignty

ACTION BOX:
ADDRESSING “NEOCLOUD”, GEOPATRIATION RISK A 90-DAY “RESILIENCE+COMPLIANCE WITHOUT KILLING AGILITY” PLAN

1

Days 0–30:
Score and place workloads

Score workloads by sensitivity + geopolitical exposure, then decide placement:

public / isolated region / sovereign region / local provider / on-prem.

Start with tier-0 and the “quietly critical” dependencies (identity, DNS, logging, CI/CD, core SaaS). Those are often the real single points of failure.

2

Days 31–60:
Build a sovereign-ready playbook for tier-0

- For each tier-0 workload:
- Define a reference architecture for sovereign-ready deployment
- Document data flows and residency boundaries
- Identify the minimum viable secondary placement option
- Run a tabletop for a “region/vendor disruption” scenario

3

Days 61–90:
Add governance for AI hosting (GPU workloads included)

- Create a lightweight but enforceable standard:
- Where sensitive AI workloads can run
- Which providers are approved and under what controls
- How you validate residency, retention, and access
- What “exit” looks like before you sign

THE DAY 90 FINISH LINE:

Make tier-0 sovereign-ready by design—so resilience and compliance don’t show up as last-minute fire drills.