

The Digital Mujtahid

Product Thinking from First Principles

10 Principles for Building Products That Matter

A Product Lead & Mujtahid Playbook

Jauhari Che Wan

PRODUCT LEAD × MUJTAHID × MAQASID

PRODUCT LEAD × MUJTAHID PLAYBOOK

The Digital Mujtahid
Jauhari Che Wan

Based on Classical Usul al-Fiqh
Marty Cagan's Product Framework
Maqasid al-Shariah

Copyright © 2025
All rights reserved

Typeset in Times / Helvetica
Printed in Malaysia

*For every Product Lead who's ever asked: "What would the
Mujtahid do?"*

Contents

Principle #1: Principle #1: The Problem Statement

What Problem Are We Solving?

Principle #2: Principle #2: Evidence Gathering

Continuous Discovery as Istiqsa'

Principle #3: Principle #3: Stakeholder Consultation

Shura as Discovery

Principle #4: Principle #4: The Product Decision

The Principle as Definition of Done

Principle #5: Principle #5: Minimum Viable Principle

MVP as Minimum Viable Principle

Principle #6: Principle #6: Technical Debt

Israf vs Investment

Principle #7: Principle #7: Pricing

Adl (Justice) in Pricing

Principle #8: Principle #8: Growth Metrics

Barakah vs Vanity Metrics

Principle #9: Principle #9: Team Organization

Shura at Scale

Principle #10: Principle #10: Legacy Product

Sadaqah Jariyah as Compound Impact

Principle #1

Principle #1: The Problem Statement

What Problem Are We Solving?

EVIDENCE (ISTIDLAL)

Qati al-Thubut (certain transmission): data that is indisputable—e.g., server logs showing drop-off rates.

Zanni al-Thubut (speculative transmission): user quotes—they might be lying or mistaken.

Qati al-Dalala (certain meaning): the interpretation is obvious—e.g., “I stopped because I felt punished” clearly indicates emotional pain.

Zanni al-Dalala (speculative meaning): a high drop-off rate could mean multiple things (bad UI, low intent, technical failure).

I weigh the evidence: The quantitative drop-off is qati al-thubut (the data is fact) but zanni al-dalala (the cause is ambiguous). The qualitative quote is zanni al-thubut (one user’s opinion) but qati al-dalala (the emotional job is clear). To strengthen, I apply Istishab (presumption of continuity): we presume no problem exists until evidence contradicts. Here, the drop-off contradicts the assumption that users want auto-save. So the default state is: the current solution is not solving the real problem.

I also apply Sad al-Dhara’i (blocking the means): if we build auto-round-up without solving visibility, we may actually increase anxiety (users feel they lost control). That harms Hifz al-Aql. The evidence does not support moving forward.

Key lesson from evidence: Distinguish anecdote from data. One user’s “I felt punished” is an anecdote—but if it resonates with patterns across 8 of 10 users, it becomes mustanid (supported evidence). The Product Lead triangulates; the Mujtahid demands tawatur (multiple independent chains) before accepting.

PRODUCT_LEAD:

I call a cross-functional sync: Engineering, Design, Data, and the CEO. I share the Opportunity Solution Tree and the evidence. I start with the user quotes—let them feel the pain. Then I ask: What do we think the root cause is?

The CEO says: “But auto-round-up works for everyone else.” That’s an appeal to popularity (istihsan bi al-‘adah). I counter: “We need to solve our users’ job, not copy features.”

The Data Lead points out: “Our retention curve flattens after 30 days—users who save manually stay longer than those who use auto-features.” That’s qati al-thubut evidence.

Design says: “What if we build a dashboard first? Show spending breakdown before any saving feature.”

I record all input. I don’t dismiss the CEO—I note his concern as a maslaha (speed to market). But I weight it lower because it contradicts user evidence. I then ask: Who else should we consult? We need to talk to users who churned. I schedule 5 more interviews.

MUJTAHID:

SHURA (CONSULTATION)

I establish adab al-shura (etiquette of consultation):

No rank-based weighting. The CEO’s opinion gets the same initial weight as the intern’s. True evidence trumps authority.

Record dissent explicitly. If Engineering says “this is technically hard to roll back,” that becomes a shart (condition) for the decision.

Seek those with khibra (expertise). The Data Lead’s analysis is stronger than a generalist’s intuition. I assign a weight of 3x to domain experts.

The key mas’ala (issue) is: Is the problem we’re solving the right one? The shura reveals consensus that auto-round-up addresses a symptom, not the root. But we need one more round: consult the ahl al-khibra## THE PRINCIPLE (HUKM)

HUKM: We will not build any solution — not a single line of code, not a prototype, not an A/B test — until we have a validated problem statement that passes the Maqasid Problem Filter and is framed as a Jobs-to-be-Done.

DALEEL:

Quantitative data from our analytics shows a 40% drop in daily active users within the first week of new feature launches — not because the features were broken, but because they solved problems nobody had. Qualitative interviews (n=12) revealed users couldn’t articulate what job they were hiring our app to do; they described tasks, not outcomes. The Shura session with stakeholders confirmed that we have been building from “cool idea” rather than “real struggle.” The evidence is zanni al-dalala (probable in meaning) but qati’ al-thubut (certain in occurrence): we are wasting engineering capacity on assumptions.

MAQSAD: Primary: Hifz al-Aql (Preservation of Mind) — by forcing ourselves to think clearly about the true problem, we protect the team’s cognitive energy from being scattered across low-impact work. Secondary: Hifz al-Mal (Preservation of Wealth) — we stop burning budget on features that don’t move the North Star metric.

SHURUT (Conditions that must be met before the principle is considered fulfilled):

The problem statement must include a JTBD functional job (e.g., “When I read Quran, I want to capture the verse that struck me so I can reflect on it later”) — not a feature request.

The problem must pass the Maqasid Filter: does solving it preserve Aql, Din, Nafs, Mal, or Nasl? If none, deprioritize.

The problem must be validated with at least 5 user interviews where the user’s struggle is observed, not just stated.

The problem must have a measurable outcome (e.g., increase in “verse capture rate” by 20%) before any solution sketch is drawn.

MUNKATHIRAT (Nullifiers — conditions that invalidate this principle):

If a stakeholder (including the CEO) bypasses the problem statement and says “just build this feature,” the principle is nullified — we must immediately call a Shura to re-establish discipline.

If the team starts wireframing before the problem statement is written in the JTBD format, the principle is broken — halt and rewrite.

If the problem statement changes mid-sprint without a new validation cycle, the principle is violated — rollback to discovery.

THE PROTOCOL — Do This This Sprint

STEP 1: Write the JTBD Problem Card (by Tuesday)

STEP 2: Take the feature request your team is most excited about. Write it on a physical index card. Flip it over. On the back, answer: What job is the user trying to get done? What is the struggle? What progress are they trying to make? If you cannot write a clear functional job in one sentence, do not proceed. This is your “istidlal al-mas’ala” (inference of the problem).

STEP 3: STEP 2: Run the Maqasid Filter (Wednesday)

STEP 4: Bring the problem card to a 30-minute standup with the product triad (PM, designer, lead engineer). For each Maqсад (Din, Nafs, Aql, Mal, Nasl), ask: If we solve this problem, which Maqсад do we serve? How? Give one concrete example. If the problem serves zero Maqсад, kill it. If it serves one but with negative externalities (e.g., increases user anxiety), apply Sad al-Dhara'i (blocking the means to harm) and deprioritize.

STEP 5: STEP 3: Validate with 5 Struggles, Not 5 Opinions (by end of Sprint)

STEP 6: Conduct 5 user interviews. Do not ask “Would you use X?” Ask: Tell me about the last time you tried to do [the job]. What made it hard? What did you do instead? Record each struggle as a quote. After the interviews, rewrite the problem statement based on what you observed. If the statement changes, you have validated the real problem. If it stays the same, you may have confirmation bias — interview 3 more users.

What would our product look like if we deleted every feature that was built without a validated problem statement? Take that list. Burn it in your mind. Now ask: What is the one feature we would keep — the one that actually solves a real struggle for a real person? That is your North Star. Everything else is decoration. Start there next sprint.

Principle #2

Principle #2: Evidence Gathering

Continuous Discovery as Istiqsa'

DISCOVERY (ISTIQA)

PRODUCT_LEAD:

MUJTAHID:

EVIDENCE (ISTIDLAL)

Quantitative (what: metrics, dashboards, funnels)

Qualitative (why: interviews, diary studies, usability tests)

Behavioral (what users actually do: logs, clickstream, feature adoption)

Attitudinal (what users say: surveys, NPS, feedback forms)

You need all four to triangulate. A high drop-off in onboarding (quantitative) tells you where people leave. Interviews (qualitative) tell you why — maybe they're confused by the savings rules. Behavioral data (logs) shows that 70% of users who attempt to set a rule abandon after step 3. Attitudinal data (survey) reveals 60% find the rule builder "too complex." Now you have evidence. But the habit is: you don't wait for a perfect dataset. You gather minimum viable evidence — just enough to reduce uncertainty to a level where you can make a decision. If the risk is low, a single interview with a clear pattern may suffice. If the risk is high (e.g., launching a new payment rail), you need a full investigation: user research, competitive analysis, prototype testing, and a staged rollout.

MUJTAHID:

In Istidlal (■■■■■■■■), we classify evidence by strength. The hierarchy:

Qati' al-Thubut wa Qati' al-Dalala — Definitive in transmission and meaning (like a mutawatir hadith). In product: a clear, replicable A/B test with statistical significance, or a consistent pattern across 20+ user interviews.

Qati' al-Thubut wa Zanni al-Dalala — Definitive transmission but ambiguous meaning. Example: a metric that is reliable but could be interpreted multiple ways (e.g., “time on page” could mean engagement or confusion).

Zanni al-Thubut wa Qati' al-Dalala — Probable transmission but clear meaning. Example: a single user interview where the user's words are unambiguous, but you need more data to confirm it's not an outlier.

Zanni al-Thubut wa Zanni al-Dalala — Both weak. Example: a gut feeling from the CEO, or a survey with a biased sample.

Your job as product Mujtahid is to weigh the evidence and assign a confidence level. The maqsad of Hifz al-Mal means you protect resources from being wasted on low-confidence decisions. You do not build a feature based on zanni evidence alone — unless the cost of delay is higher than the cost of a mistake (Maslaha principle). Signal vs. noise: a single spike in support tickets about a bug is signal. A low NPS score with no qualitative context is noise. You must filter noise by triangulating across quadrants. The minimum viable evidence for a principle (a product decision with high impact) is at least two independent sources from different quadrants, with no contradictory evidence. For an MVP experiment, one strong qualitative signal plus a behavioral data point is sufficient.

PRODUCT_LEAD:

SHURA (CONSULTATION)

Identify the right stakeholders — those with relevant expertise (not just hierarchy). The engineer who built the legacy code, the customer support agent who hears complaints daily, the finance analyst who knows the unit economics.

Present the evidence neutrally — do not lead the consultation with your preferred conclusion. The Mujtahid states the facts, the opportunities, the risks.

Weigh the input — not all voices are equal. The person closest to the user's pain may have more weight than the VP who hasn't spoken to a customer in a year. Weight by relevance, not rank.

Record dissenting opinions — if a respected stakeholder disagrees, note it in your product decision log. That dissent becomes a munkathir (nullifier) if later evidence proves it correct. It also protects you from groupthink.

A Muslim PM should see Shura as a sunnah that protects Hifz al-Mal — the collective wisdom of the team prevents costly mistakes. The output of Shura is not consensus; it's a refined set of shurut (conditions) and munkathirat (nullifiers) that you will carry into the principle.

End of Part 1. Continue to Part 2: The Principle (Hukm), The Protocol, and The Retrospective.## THE PRINCIPLE (HUKM)

HUKM: We will institutionalise a weekly Continuous Discovery habit that systematically collects evidence across all four quadrants of the Evidence Quadrant (Quantitative, Qualitative, Behavioral, Attitudinal) before committing any engineering resources to a new feature or experiment.

DALEEL: Analysis of our last three sprints shows 67% of shipped features failed to move the North Star metric because decisions were based on attitudinal data alone (customer “I want X” interviews) without behavioral validation. The Evidence Quadrant framework, when applied

rigorously, reduces waste (Hifz al-Mal) by grounding decisions in triangulated signals. Our own A/B test history confirms that features validated with ≥ 3 quadrants have a $4\times$ higher success rate.

MAQSAD: Primary: Hifz al-Mal (Preservation of Wealth/Data) – avoiding wasteful engineering spend and protecting the integrity of our data pipeline. Secondary: Hifz al-Aql (Preservation of Intellect) – ensuring decisions are made with sound reasoning, not guesswork.

SHURUT:

Evidence must be refreshed within the last 14 days before any sprint commitment.

At least two quadrants must agree before a decision is taken to build.

Behavioral data (actual user interaction logs) takes precedence over attitudinal data (stated preferences) in case of conflict.

The Evidence Quadrant board must be visible to the entire team and updated weekly.

MUNKATHIRAT:

Skipping the weekly evidence review without a documented emergency.

Making a build decision based on a single quadrant (e.g., only quantitative metrics or only a CEO's opinion).

Ignoring disconfirming evidence that emerges after a decision is made – the principle is invalidated if the team refuses to revisit the evidence.

THE PROTOCOL — Do This This Sprint

STEP 1: Every Monday morning, block 90 minutes for the Continuous Discovery huddle. The PM brings three pieces of evidence: one quantitative (e.g., funnel conversion), one qualitative (e.g., user interview quote), and one behavioral (e.g.,

session recording clip). No slide decks – just raw evidence on the Evidence Quadrant board.

STEP 2: For any feature being considered for the next sprint, the team must map it to at least two quadrants within that same week. If only one quadrant has evidence, the feature is deferred to the next cycle. Use the Opportunity Solution Tree to trace the evidence back to a specific opportunity.

STEP 3: Every Friday, run a 30-minute “Evidence Audit” where the team asks: “What did we learn this week that disconfirms our current assumptions?” If any Munkathirat condition is triggered, the build is paused and the Hukm is re-evaluated before Monday.

Where in our last sprint did we ship something based on a single quadrant, and what would it have cost us if we had been wrong – both in wasted engineering hours (Hifz al-Mal) and in misleading data that now pollutes our model (Hifz al-Aql)?

Principle #3

Principle #3: Stakeholder Consultation

Shura as Discovery

EVIDENCE (ISTIDLAL)

Build your evidence stack:

Quantitative: Funnel drop-off rates, feature usage heatmaps, survey NPS segmented by stakeholder group. What does the data say about the problem's prevalence?

Qualitative: 5 user interviews this week, 2 stakeholder shadow sessions (e.g., sit with compliance during a manual review). What story does the data not tell?

When stakeholders conflict (e.g., users want speed, compliance wants more checks), weight evidence by frequency, severity, and reversibility. A 2% conversion loss from adding a step is zanni until you validate it with a cheap experiment. A regulatory fine of \$10M is qati' (definitive) — it must be addressed first.

MUJTAHID:

Daleel in Usul is ranked: Qati' al-Thubut wa Qati' al-Dalala (definitive source and meaning) vs Zanni in either. In Shura, stakeholder input is rarely qati'. It is zanni evidence that must be weighed.

Qati' evidence: Explicit regulatory requirements, binding contracts, clear user harm (e.g., bug causes financial loss). These override everything.

Zanni evidence: CEO's strategic hunch, engineer's "I think this is easier," user's stated preference (which may differ from behavior).

The Mujtahid uses tarjih (preponderance) rules:

Who holds the burden of proof? The stakeholder advocating a change must bring daleel. If they say “competitor does it,” that’s zanni — insufficient.

What is the maqsad? If the feature protects Hifz al-Din (e.g., ensuring riba-free transactions), it takes precedence over Hifz al-Mal (revenue).

Sad al-Dhara’i* (blocking the means to harm): If a feature opens a path to mafsada (e.g., easier fraud), it is avoided even if users request it.

Istishab (presumption of continuity): Assume the current state is valid until evidence proves otherwise. Don’t build a feature just because someone asked — the default is no change.

Use a simple Tarjih Matrix: For each stakeholder input, assign a score (1-3) for Qati’/Zanni, Maqsad weight (Daruri=3, Hajji=2, Tahsini=1), and Mafsada risk (high=3, medium=2, low=1). The highest total wins — but only if the maqsad is preserved.

PRODUCT_LEAD:

SHURA (CONSULTATION)

Weekly 30-min sync with each quadrant (Users, Team, Business, Regulators). Use a consistent format: “What problem are we solving? What evidence do you have? What outcome do you need?”

Monthly Shura Council — one hour, all four quadrants present. Each presents their top opportunity (backed by evidence). The goal is not to vote but to surface trade-offs.

Record dissent explicitly. When a stakeholder disagrees, write their dissent and the daleel they offered. This is tadwin al-shura — documenting consultation. It protects Hifz al-Karama: their voice was heard, weighted, and preserved.

MUJTAHID:

The Mujtahid consults by tamyiz (discernment), not democracy. The Prophet ■ consulted Badr strategy with the Companions, but the final decision was his. Shura is advisory, not binding — but the advisor’s dignity is sacred.

How to consult as a Mujtahid PM:

Start with the maqsad. Open the Shura by stating: “We are here to protect \$MAQSAD for \$STAKEHOLDER. What evidence do you have that \$FEATURE serves or harms this?”

Weight by expertise, not hierarchy. The engineer’s technical feasibility is zanni but high weight in that domain. The CEO’s strategic vision is zanni but high weight in market context. The user’s pain is zanni but weight increases with frequency.

Use the Shura Quadrant to identify gaps. If only Business speaks, ask: “Where is the user’s voice?” Silence is also a signal — it may indicate hifz al-Karama is being violated (their input wasn’t solicited).

End with a ruling, not a poll. After gathering all input, you (the Product Lead / Mujtahid) issue a hukm: “We will proceed with \$DECISION, based on \$DALEEL, with \$SHURUT (conditions), and \$MUNKATHIRAT (rollback triggers).” Record the reasoning. This is tadwin al-ijtihad.

End of Part 1. Principle, Protocol, and Retrospective continue in Part 2.##
THE PRINCIPLE (HUKM)

HUKM: We will build a recurring, structured Shura Council for every major product decision, with rotating representation from Users, Team, Business, and Regulators, and we will publish a summary of input and outcomes.

DALEEL: Discovery interviews across 12 stakeholder interviews revealed that 8 felt their voice was either ignored or collected too late.

Quantitative survey (n=50) showed 74% of stakeholders would increase trust with a formal consultation process. Maqasid analysis confirms that Hifz al-Karama (dignity) is violated when stakeholders are informed after decisions are locked. Qiyas: the Prophet ﷺ consulted even in minor matters (Khandaq, changing strategy based on Salman al-Farsi's input). The daleel is strong: recurring Shura preserves dignity AND improves decision quality.

MAQSAD: Hifz al-Karama – the preservation of stakeholder dignity through genuine voice, which is a necessary condition for trust, alignment, and sustained collaboration. Secondary Maqasid: Hifz al-Mal (business outcomes improve when blind spots are surfaced early).

SHURUT:

The Shura Council must include at least one representative from each quadrant (Users, Team, Business, Regulators) per decision.

A written “Shura Brief” (problem statement, constraints, options) must be distributed 48 hours before the session.

The Product Manager must explain how each piece of input was used or why it was not used, and publish this in a public log.

Rollback trigger: If any quadrant misses two consecutive sessions, the decision is paused until they are represented.

MUNKATHIRAT:

When stakeholder input is never referenced in the final decision document (nullifies dignity).

When sessions become monologues from the product team (no genuine listening).

When the same individuals dominate every session without rotation (entrenchment, not shura).

THE PROTOCOL — Do This This Sprint

STEP 1: This sprint, identify the next major product decision (e.g., feature prioritization, pricing change, or sunseting a module). Draft a one-page Shura Brief with the problem, three options, and one clear question: “What would you do if you were me?” Share it with at least one person from each quadrant by Wednesday.

STEP 2: Hold a 45-minute Shura session on Thursday. Start with the brief, then each quadrant speaks uninterrupted for 5 minutes. Take verbatim notes. End with: “I will share my decision and how your input shaped it by Monday.”

STEP 3: By Monday, publish the decision log in a visible place (Slack, wiki, or product board). Include: what we decided, which inputs changed our thinking, which inputs we deprioritized and why. Then schedule the next Shura session two weeks out.

When did you last change a product decision because a stakeholder disagreed with you? If you can't remember, your Shura is performative. Dignity is not about being heard—it's about being able to change the outcome. The real test of Hifz al-Karama is whether your product roadmap has scars from stakeholder voice.

Principle #4

Principle #4: The Product Decision

The Principle as Definition of Done

THE PROTOCOL — Do This This Sprint

STEP 1: THIS SPRINT – SECURE THE FATWA

STEP 2: Write a one-page executive summary of the crypto trading proposal, including the exact token mechanics, smart contract details, and user flow. Submit it to the Shariah board with a deadline of 10 business days. Prepare a sign-off document that explicitly states the permissible conditions (e.g., no leverage, no interest, instant settlement). If the board demands changes, iterate within the sprint.

STEP 3: STEP 2: BUILD THE REVERSIBLE INFRASTRUCTURE

STEP 4: Engineers implement a feature flag toggling the entire crypto module. Write a rollback script that can reverse any pending transactions and restore the previous app state within 60 minutes. Document the rollback procedure and assign an on-call engineer for the beta period. Test the flag and script in staging with simulated transactions.

STEP 5: STEP 3: LAUNCH THE BETA AND MONITOR DAILY

STEP 6: Invite the first 100 users from a waitlist (pre-qualified for high trust). Send them a one-time consent screen with a summary of the Shariah board's ruling and a "I understand the

risks” checkbox. Monitor daily: log every transaction, flag any that deviate from the approved mechanics, and run a daily compliance check against the Shurut. At the end of two weeks, hold a retrospective before deciding on broader rollout.

What single assumption about our users’ religious commitment did we make that, if wrong, would invalidate our entire decision? If the assumption is that users will accurately self-report when they feel a transaction is impermissible, but in reality they stay silent out of fear of missing out, then our “zero complaints” metric becomes meaningless. We would ship a feature that violates Hifz al-Din while believing it is safe. The uncomfortable question: Do we truly know our users’ threshold for religious compromise, or are we relying on a proxy that can easily break? Next sprint, add a post-transaction satisfaction survey that explicitly asks: “Did this transaction feel in line with your Islamic values?” Use the answer to calibrate our Shurut, not just as a checkmark.

Principle #5

Principle #5: Minimum Viable Principle

MVP as Minimum Viable Principle

Principle #6

Principle #6: Technical Debt

Israf vs Investment

THE PROTOCOL — Do This This Sprint

STEP 1: Audit Your Debt Quadrant (This Week)

STEP 2: Take 2 hours as a team. Open your bug tracker and code review history. For each recent shortcut, answer: “Did we plan it? Do we have a repayment date?” Assign it to one of the four quadrants. Create a single “Debt Ledger” document (Google Sheet or Notion) with columns: Type, Description, Date Incurred, Payoff Date, Interest Cost (estimated hours lost per month). **PRODUCT_LEAD:** Do this during your Tuesday tech review. **MUJTAHID:** This is *istiqsa’* (exhaustive examination) of your Mal—do not skip the “why” behind each entry.

STEP 3: STEP 2: Create a Debt Repayment Board (This Sprint)

STEP 4: In your project management tool, add a “Debt Repayment” swimlane. Move all Reckless and Inadvertent items into it. Assign each item a story point value (based on effort to fix) and a DRI (directly responsible individual). Set a sprint goal: “Reduce Reckless debt by 30% (or at least fix the top 3 items).” **MUJTAHID:** This is *tazkiyah* (purification) of your codebase—treat it as a non-negotiable worship of trust (*amanah*).

STEP 5: STEP 3: Define a “Debt Covenant” for Future Sprints (This Sprint Retro)

STEP 6: Write a 1-page team contract:

STEP 7: “We will never merge a PR without logging its debt impact in the ledger.”

STEP 8: “We will never approve a deliberate debt without a written payoff sprint ID.”

STEP 9: “We will review the ledger every retro and escalate any overdue item.”

STEP 10: PRODUCT_LEAD: Make this part of your Definition of Done. **MUJTAHID:** This is your shart (condition) that prevents israf from recurring—treat it like a legal stipulation in a contract.

One piercing question for your next retro: “Where are we knowingly calling Israf an ‘investment’ because admitting it’s waste would force us to change our roadmap?” **PRODUCT_LEAD:** Look at the last three features you shipped. How many shortcuts did you tell yourself were “necessary to hit the deadline”? Now ask: Did that deadline actually deliver better outcomes, or did you just shift the cost to next quarter? **MUJTAHID:** The mujtahid knows that niyyah (intention) alone does not justify israf. If the shortcut wasted more time than it saved, it was haram even if you meant well. Measure the actual interest paid (hours lost to bugs, refactoring, onboarding friction). Compare it to the time you “saved.” If the debt’s interest exceeds its principal, you have committed israf—and the only repentance is repayment.

Principle #7

Principle #7: Pricing

Adl (Justice) in Pricing

What would it take for us to cut our price by 50% and still sustain the business? If we can't answer that question honestly, we haven't truly discovered the value we deliver — we're just guessing. A just price isn't the highest the market will bear; it's the price that reflects genuine worth while preserving the relationship. Did our pricing increase trust or erode it? Did we protect the user's wealth or exploit their need? The answer will surface in the churn rate and the unsaid resentment in every cancelled subscription.

Principle #8

Principle #8: Growth Metrics

Barakah vs Vanity Metrics

DISCOVERY (ISTIQA'S)

Vanity metrics—total signups, raw MAU, paid acquisition numbers—are gharar (uncertainty) because they don't reflect genuine benefit. They are like counting seeds that never sprout. Barakah growth is increase that carries blessings—it multiplies benefit without corrupting the intention.

We must map the problem space using the five Maqasid as filters:

Hifz al-Din – Does this metric encourage ethical behavior?

Hifz al-Nafs – Does it protect users from financial stress?

Hifz al-Aql – Does it promote sound decision-making?

Hifz al-Mal – Does it preserve wealth (not waste it on acquisition)?

Hifz al-Nasl – Does it sustain family continuity?

Acquisition-only metrics fail Hifz al-Mal (waste) and Hifz al-Nasl (no continuity). Retention and referral metrics that come from genuine value—those have Barakah. Our istiqsa' reveals that the real opportunity is not “more users” but “more users who stay and bring their family.” That is growth with Barakah.

PRODUCT_LEAD:

Let's gather daleel—evidence. Quantitative: We pull cohort analysis. Users acquired through paid ads have a 30% 90-day retention. Users acquired through organic referrals have 70%. The median lifetime value (LTV) of a referred user is 3x higher. The paid users churn faster and cost more to acquire. That's quantitative daleel that paid acquisition is less

Barakah for this product.

Qualitative: User interviews reveal a pattern. “I downloaded because of a discount code, but I never used it after the first month.” “My sister told me this app helped her save for hajj—so I trust it.” The qualitative daleel confirms: Barakah growth is rooted in trust and real outcomes, not incentives.

We also look at negative evidence: churn spikes after promotional campaigns. Users who joined via paid ads are 50% more likely to delete the app within 7 days. This is sad al-dhara’i—blocking the means to harm. The harm here is wasted resources (israf) and misleading metrics that cause product teams to optimize for the wrong thing.

MUJTAHID:

Now we apply the daleel hierarchy from Usul al-Fiqh.

Qati al-Thubut (certain in transmission) – We don’t have revelation, but we have reliable analytics data from our own backend. That is strong evidence, albeit zanni (speculative) in interpretation.

Qati al-Dalala (certain in meaning) – The meaning is clear: referred users stay longer. That’s a qiyas (analogy): if organic referral leads to higher retention, then investing in referral features is analogous to preserving wealth (Hifz al-Mal) and lineage (Hifz al-Nasl).

Istishab (presumption of continuity) – We presume that the current pattern (paid users churn) will continue unless proven otherwise. So we should not scale paid acquisition without stronger evidence.

Maslaha (public benefit) – Sustainable growth benefits the community: users save money, families are less stressed, the company survives long-term. Vanity growth harms the company (burnout, wasted budget) and users (poor experience).

What is “churn” in spiritual terms? Churn is inqita’—a severing of the relationship. The user leaves the blessing. A metric that measures “users

who leave” is good, but more important is “users who stay and refer.” That is iltizam—commitment. The Barakah filter asks: Does this metric measure genuine commitment or fleeting attention?

Evidence conclusion: The daleel stack points away from paid acquisition as a primary growth lever. The strongest evidence (quant + qual + qiyas) supports investing in referral programs, activation improvements, and retention loops. Those metrics survive the Barakah filter.

PRODUCT_LEAD:

Time for Shura—structured consultation. I schedule a 90-minute cross-functional sync. Attendees: CEO, VP Marketing, Head of Engineering, Customer Support Lead, and two user researchers. I frame the decision: We need to choose between doubling down on paid acquisition or building a referral + retention engine.

First, I present the evidence from Discovery and Istidlal: the cohort data, user interview clips, and the Barakah filter. Then I open the floor.

CEO argues: “We need scale to raise Series B. Paid ads are proven. Referrals take time.”

VP Marketing pushes back: “Our paid CAC has risen 40% this quarter. The quality is dropping.”

Customer Support Lead shares: “Most complaints come from paid users who don’t understand the product.”

User researcher shows a clip: “I only used it because of the \$5 bonus—now I forgot I even have the app.”

I weight the inputs based on proximity to evidence and expertise. The CEO’s opinion is important but based on a fear of missing growth targets. The Support Lead and researcher have direct daleel from users—their voices carry more weight. I record all opinions, including dissent.

MUJTAHID:

SHURA (CONSULTATION)

Identify the stakeholders – Everyone present has a right to speak, but we weight their wilaya (authority) based on knowledge and proximity to the truth. The user researcher and support lead are the shuhud (witnesses) to reality.

Record each opinion with its daleel – CEO’s daleel: “We need speed to survive.” That is a maqsad of Hifz al-Mal (company survival), but it conflicts with Hifz al-Nasl (sustainable growth). We note the tension.

Seek consensus (ijma’) – We don’t have full consensus. The CEO dissents. We document that dissent as mukhalafa (opposition) but proceed based on the preponderance of evidence.

Issue the ruling – The Shura concludes## THE PRINCIPLE (HUKM)

HUKM: We will only ship growth initiatives that pass the Barakah Filter — meaning the metric must be sustainable, user-centric, and directly tied to long-term retention (Hifz al-Nasl), not just short-term acquisition or vanity spikes.

DALEEL:

Quantitative: Cohort analysis shows that users acquired via “growth hacks” (discounts, viral loops, referral spam) have a 60% lower 90-day retention than organic users.

Qualitative: User interviews reveal that aggressive referral prompts erode trust — “I felt like a product, not a person.”

Qiyas (analogy): Just as riba (usury) creates apparent wealth but destroys real economic stability, vanity metrics create apparent growth but destroy product lineage (Nasl). The Barakah filter is analogous to sad al-dhara’i — blocking the means to harm, even if the means look beneficial in the short term.

MAQSAD: Hifz al-Nasl (Preservation of Lineage / Continuity) — ensuring the product grows in a way that sustains its user community and value over generations, not just this quarter.

SHURUT:

Each growth experiment must define a Barakah ratio: (Retained users after 30 days) / (Total users acquired). Target: $\geq 40\%$ for any experiment to continue.

All growth tactics must be opt-in, transparent, and avoid manipulation (no dark patterns, no pre-checked boxes, no shame-based prompts).

Monthly Barakah audit: Review AARRR funnel through the lens of maslaha (public benefit). If any stage has a negative maslaha score (e.g., activation causes user anxiety), the initiative is paused.

Rollback trigger: If support tickets from growth features exceed 2% of new users within 7 days, automatically halt the experiment.

MUNKATHIRAT:

Nullifier one: Any growth tactic that violates user data privacy or consent (even if legal) — this nullifies the principle immediately. Hifz al-‘aql (preservation of intellect) requires informed choice.

Nullifier two: If the growth metric improves but the core product experience (NPS, task success rate) declines more than 5% relative to baseline — the principle is broken.

Nullifier three: If the team cannot articulate how the growth initiative serves a real user job (JTBD) beyond the company’s revenue goal — the principle is suspended until a clear maqsad is defined.

THE PROTOCOL — Do This This Sprint

STEP 1: Audit Your Current Metrics

STEP 2: Monday–Tuesday.

STEP 3: Pull your last 3 months of AARRR data. For each stage, write two numbers: the raw metric (vanilla) and the Barakah-adjusted metric (retention at 30 days ÷ acquisition cost or effort). Highlight any stage where the Barakah ratio is below 40%. Flag those as “vanity growth.” Discuss with your team: “Which of these are we secretly proud of but actually harming our Nasl?”

STEP 4: STEP 2: Build and Calibrate Your Barakah Filter

STEP 5: Wednesday–Thursday.

STEP 6: Create a simple 3-question filter for every growth experiment before it ships:

STEP 7: Does this serve a real user job (JTBD) or just our revenue goal?

STEP 8: Can a user opt out without penalty?

STEP 9: Will this increase thiqah (trust) or just traffic?

STEP 10: If any answer is “no,” the experiment is blocked. Write the filter as a checklist in your PRD template.

STEP 11: STEP 3: Apply the Filter to One Upcoming Growth Experiment

STEP 12: Friday.

STEP 13: Take your next planned growth initiative (e.g., a referral program, a discount campaign). Run it through the Barakah filter. If it passes, define the specific shurut (conditions) and munkathirat (rollback triggers) for that experiment. If it fails, redesign it — or kill it. Ship only if it passes. Measure the Barakah ratio after 30 days.

Where are we currently optimizing for a vanity metric that we know — deep down — is misleading our team and eroding our product’s long-term trust? Sit with that question for 10 minutes. No defensiveness. Write down the metric name. Then ask: What would it cost us to drop this metric entirely and replace it with a Barakah-adjusted one? The answer is usually far less than the cost of chasing ghosts.

Principle #9

Principle #9: Team Organization

Shura at Scale

SHURA (CONSULTATION)

Record the dissent from the Engineering Lead: “We will lose velocity building fraud alert first.” That becomes a shart (condition) for the decision: if velocity drops below X, revisit.

Now document the shura process in a log:

Participants: CEO, Eng Lead, PM, User Researcher.

Inputs: Outcome (reduce complaints), User JTBD (self-audit), Technical constraints (2 weeks vs 3 months).

Weighted evidence: User research (2x weight), Engineering feasibility (1.5x), CEO outcome (1x).

Decision: Squad maqsad = “Reduce fraud complaints by enabling self-audit via transaction history.”

Dissent recorded: Eng Lead wants fraud alert; rationale noted. Condition: If after 4 weeks fraud complaints don’t drop, re-open shura.

This is shura at scale — not a committee of equals, but a structured consultation where the wali al-amr (Product Lead) holds final authority, bounded by the maqsad and shurut.

[End of Part 1 — Scenario, Discovery, Evidence, Shura. Continue to Part 2: Principle (Hukm), Protocol, Muhasaba.]## THE PRINCIPLE (HUKM)

HUKM: We will organize our product teams into autonomous squads with embedded Shura circles, where each squad owns a full opportunity space and makes binding decisions through structured consultation, not hierarchy.

DALEEL: Three evidence streams converge. First, continuous discovery interviews across 12 teams showed that cross-functional squads with decision authority ship 3x faster and produce higher user satisfaction than functionally siloed groups. Second, the prophetic model of Shura at Medina—where the community was organized into small, autonomous units (e.g., the nuqaba’) that consulted regularly and then acted—demonstrates that preserving community (Hifz al-Nasl) requires distributed decision-making, not central command. Third, our own A/B test of squad vs. functional organization in the payments domain revealed a 40% reduction in time-to-learning and a 22% increase in team Net Promoter Score.

MAQSAD: Hifz al-Nasl (Preservation of Community) is served directly. Squad Shura preserves the community of makers, users, and stakeholders by ensuring that every decision emerges from local, authentic consultation rather than remote authority. It also preserves the product’s community of users: squads stay close to their user segment, preventing the alienation that comes from one-size-fits-all product decisions. Secondary maqasid: Hifz al-Aql (intellectual property and sound reasoning) through diverse input; Hifz al-Mal (efficient resource use) through faster iteration.

SHURUT:

Each squad must include at least one product manager, one designer, and two to three engineers (cross-functional minimum).

Every squad holds a weekly Shura huddle (45 min max) where all members have equal voice; decisions are recorded and shared with a Guild-level Shura for coherence.

Squads must publish their opportunity solution trees and weekly Shura notes to a shared repository — transparency is a shart (condition) for autonomy.

No squad may ship a change that violates the product's North Star metric or the company's ethical framework (Maslaha check) without escalating to the Guild Shura.

MUNKATHIRAT:

If a squad's opportunity solution tree goes unvisited by any user interview for two consecutive sprints, the Shura circle is invalidated and must be reconvened with fresh evidence.

If the Guild Shura observes that two or more squads are solving overlapping problems without coordination, the squads must merge or re-scope; continued overlap nullifies the autonomy principle.

Any squad that fails to meet its OKR threshold for two quarters will be dissolved and its members redistributed — autonomy is conditional on outcome accountability.

THE PROTOCOL — Do This This Sprint

STEP 1: Map your current org to the Org Quadrant (this sprint).

STEP 2: Draw four boxes: Functional (teams by skill), Cross-Functional (teams by feature), Squad (teams by opportunity space), Guild/Chapter (cross-squad communities). Plot every team you have. Identify which teams are pure functional and which are already squad-like. For each functional team, list the top three opportunity spaces they touch. This map is your baseline. Deliverable: a single A3 page with the quadrant and annotations.

STEP 3: STEP 2: Select one opportunity space and form a pilot squad (next sprint).

STEP 4: Pick the opportunity space with the highest user pain and the clearest JTBD (Jobs-to-be-Done). Pull one PM, one designer, two engineers from existing functional teams. Give them a single OKR: “Improve [metric] by X% in 8 weeks.” Declare them autonomous: they do not need VP approval for any decision that stays within their opportunity tree. They must, however, hold a Shura huddle every Wednesday at 10 AM sharp. No exceptions.

STEP 5: STEP 3: Run the pilot for 8 weeks, then run a Muhasaba retrospective.

STEP 6: The pilot squad ships at least one experiment per week. The Guild Shura (composed of all squad leads) reviews the pilot’s opportunity solution tree biweekly. At week 8, answer: “Did the squad preserve community? Did it ship faster? Did users feel heard?” If the pilot clears all three Shurut (cross-functional composition, weekly Shura, transparent notes, no North Star violation), scale to a second squad. If any Munkathirat fire (e.g., no user interviews for two sprints), dissolve and try a different opportunity space.

What did we actually fear when we gave a squad full autonomy, and how much of that fear was rooted in evidence versus ego? This question cuts to the heart of Hifz al-Nasl. Most leaders resist Squad Shura not because it fails empirically, but because it requires surrendering control. The community cannot be preserved by a single guardian; it must be preserved by a thousand small, responsible shuras. If your retrospective reveals that your hesitation was about authority, not outcomes, then the principle isn’t broken — your heart is. Fix that first.

Principle #10

Principle #10: Legacy Product

Sadaqah Jariyah as Compound Impact

Where is your product currently acting more like a sadaqah ‘adiyah (one-time charity) than a sadaqah jariyah (ongoing charity)? Be honest: list every feature that dies when you stop paying attention. Then ask yourself: If I died tomorrow, would this product continue to benefit people without me? If the answer is no, you have not built a legacy — you have built a job. The challenge is not to add more features, but to design the exit. A Sadaqah Jariyah product is one that makes its creators unnecessary. That is the ultimate measure of compound impact.