



Four features, each answering a problem the data found

One live-ops feature for each of my mobile games. None of them starts from “what would be cool”: each starts from a weakness the game’s own balance data or playtests show, and ends with how I would know it worked.

How I write a feature proposal

1. **The problem, measured:** a number from the simulator, a playtest note, a gap in the loop.
2. **The feature in one paragraph,** as a player would describe it.
3. **Rules and numbers** concrete enough to build and to simulate.
4. **Economy impact:** sources and sinks, and what it does to the existing curve.
5. **Player flow:** where it appears, how it is discovered, how it ends.
6. **Risks,** each with a mitigation built into the design.
7. **An A/B test** with a primary metric, a sample size and guardrails.
8. **Build cost,** naming the existing systems it reuses.

Priority (ICE: impact, confidence, ease, each out of 10)

GAME	FEATURE	IMPACT	CONFIDENCE	EASE	SCORE
Cloudhaven Skyport	The Porter's Hour	7	8	9	5.0
Brushfall	Commissions	8	7	8	4.5
Arcana Match	Seasonal Spreads	7	7	6	2.9
Arcana Siege	Arcana Companions	8	6	5	2.4

What I'd ship first: the Porter's Hour. It targets the earliest drop-off point (minutes 5–15 of the first session), it reuses an existing staff type and the existing rewarded-ad flow, and a two-week test answers it. Commissions comes next: it is the biggest long-term lever, but it needs art and a second economy pass.

All impact figures are hypotheses to test, not results. Baselines are the KPI targets in each game's Ad and KPI plan (D1 40%, D7 15%); sample sizes are computed for the stated lift.

Commissions: the gallery becomes a home

Turn the paint players can't spend into something they can see: patrons commission frames, lights and rooms for the gallery, paid in exactly the colours the paintings left over.

The problem, from the data. The economy works too well late on. After chapter 3 an average player is sitting on **2,146 paint** (a strong player 2,488), mostly off-colours the paintings never needed. The only sink is eight skins (300–2,500 paint), and the meta has no reason to come back between chapters.

The feature

A patron arrives at the gallery with a **commission**: “a gilt frame for the Wave, in blue and yellow”. Pay the paint and the gallery visibly changes: a frame, a bench, lighting, a new room. Each finished room raises the gallery's **prestige**, which unlocks the next patron and a cosmetic for the hero. It is Project Makeover's decorating meta, fed by the paint the climbs already produce, so it needs no new currency.



The gallery today: one frame per chapter, and nothing to do between paintings.

Rules and numbers

Commission size	150–600 paint, two or three colours, chosen from the player's three largest piles
Cadence	One open commission; a new one 4 h after the last is done (max 3 a day)
Rooms	4 rooms × 6 items; a room complete = a prestige rank
Rewards	Each rank: a hero skin variant or a frame style; no power, no paint
First appears	After the first painting is named (level 13)

Economy impact

A pure **sink**, aimed at the surplus: an average player's leftover 2,146 paint funds about six commissions by the end of chapter 3, a weak player's two or three. Because commissions draw from the biggest piles, they never take the colours a restoration needs. The workbook gets a Commissions sheet; the existing chapter targets don't move.

Build: M. Reuses the gallery view, the wallet, the skin unlocks and the save profile; new: 24 items of art, a patron card, a commission rule in Core with tests.

Player flow



Risks

RISK	MITIGATION
Players hoard paint for restorations and ignore it	Commissions only ever ask for the top three piles; a warning shows if a colour is needed by the next quarter
Decoration feels off-genre	Rooms are museum-like, not a home makeover; a test of the room art in the ad creative first
Paint surplus shrinks in later chapters	Commission size scales with the wallet: 15–25% of the largest pile

A/B test			
Control	Gallery as today	Variant	Gallery with commissions from level 13
Primary	D7 retention of players who reach level 13: 30% → 33% (+3 points)	Sample	≈ 3,759 new players per arm (two-sided 5%, 80% power), two weeks
Guardrails	paint short-falls in restorations (extra climbs) must not rise; session length; skin purchases		

The Porter's Hour

A rewarded helper for the stretch where a new island feels slow: a porter who carries coins for two minutes, so the player can stay at the busy counter.

The problem, from the data. The simulator shows a new player's first island has its longest flat stretch in **minutes 5–15**: the Terminal is open but staff aren't affordable yet, so the harbourmaster walks back and forth between counters and coin piles. The roadmap already names this plateau as the first thing to playtest.

The feature

Once per 20 minutes a **porter** waves from the dock: watch an ad (or skip it) and for two minutes a porter follows the coin piles and carries them to the next build pad by himself. The harbourmaster stays where the queue is longest, which is the game's best play, and sees two systems at once: serving and building. After the island's first real porter is hired (pad 11), the offer turns into a bonus “porter's tip” and fades out.



A busy port: the walking between counters and coin piles is what the porter takes over.

Rules and numbers


Offer	From minute 4 (Terminal open) until the Porter pad is bought
Duration	120 s; the porter carries 30 coins a trip at the harbourmaster's walking speed
Cooldown	20 min; at most 3 a day; the first one is free (no ad) to teach it
Pays into	The cheapest visible pad, so it never spends coins the player was saving
Island 2 and 3	Offered again only in each island's first 15 minutes

Economy impact

No new currency and no new income: the porter moves coins that already exist, so the timetable cap on income is untouched. My estimate, to check in the simulator before any test: it saves 15–25% of the walking in the plateau and pulls pads 5–9 forward by 60–90 seconds each. The island still takes 45–55 minutes; the middle just stops sagging.

Build: S. Reuses the staff walker and the coin-stack code (the Porter pad's NPC), the rewarded-ad interface and the goal panel; new: an offer bubble, a timer, simulator support to re-run the pacing.

Player flow



Risks

RISK	MITIGATION
Feels like the game is playing itself	Two minutes, and only the walking; serving still needs the player
Ad fatigue on day 1	First use free; cap 3 a day; never interrupts, the player taps the porter
Removes the satisfying coin stack	The harbourmaster still carries his own; the porter's stack is smaller and visibly separate

A/B test			
Control	No porter	Variant	Porter's Hour from minute 4
Primary	D1 retention: 40% → 43% (+3 points)	Sample	≈ 4,229 new players per arm (two-sided 5%, 80% power), one week
Guardrails	time to first staff hire; rewarded-ad views per DAU must not fall elsewhere (rush hour, offline x2); session length		


Arcana Companions

Each of the 22 major arcana becomes a companion with its own ultimate: earned by playing, levelled by repeat sieges, and chosen before a siege.

The problem, from the data. After the campaign there is little to chase: tower levels max out at about **25,600 gold** for all four towers (a typical siege pays 300–550), and the Moon's ultimate is the same in every run. The career bot finishes the map in about 28 runs; the daily siege is the only reason to return after that.

The feature

Before a siege the player picks one **companion**, one of the Arcana, who replaces the Moon's ultimate with their own: the Star heals the gate, the Tower drops lightning on the strongest enemy, Death kills the front of the road outright, Temperance slows everything. Companions are earned as card shards from wins (3 stars = more shards) and the daily siege, and level up to 5 with shards. It adds a collection and a pre-siege decision without touching the board's rules.



Judgement, the campaign's last wall; after it, companions give the game somewhere to go.

Rules and numbers

Roster	8 at launch, 22 over three updates
Earning	A win: 1–3 shards of that siege's companion
Levels	5 levels; each +15% ultimate effect; 10 / 20 / 40 / 80 shards
Balance rule	Every companion within ± 5 win-rate points of the Moon for the first-time bot
Monetisation	Cosmetic companion art only; shards are never sold

Economy impact

A second, long sink beside tower gold: a full roster at level 5 needs 3,300 shards (150 each), about 1,100 three-star wins, so companions carry the game well past the campaign. Because shards are not sold and each companion is bot-checked against the Moon, the campaign's tuned sawtooth (Judgement at 7 tries) does not move.

Build: L. Reuses the ultimate system, the profile save, card art from Arcana Match and the balance runner; new: 8 ultimates with VFX, a roster screen, shard rewards.

Player flow



Risks

RISK	MITIGATION
A must-pick companion	Each one bot-tested alone over 500 sieges, like the draft cards
More to juggle (the playtest's concern)	One choice before the siege, with a default; nothing new during play
Pay-to-win perception	Shards are earned only; only art is sold

A/B test

Control	Moon ultimate only	Variant	Companions (8), unlocked after siege 3 (the Devil)
Primary	D30 retention: 6% → 7.5% (+1.5 points)	Sample	≈ 4,386 new players per arm (two-sided 5%, 80% power), five weeks
Guardrails	first-time win rate on sieges 1–12 (must stay within ± 3 points); daily siege participation		


Seasonal Spreads

A four-week event with its own small deck: hand-made levels on a themed map, each awarding an event card, and a three-card spread from the event deck at the end.

The problem, from the data. The collection is finite: the 85-level journey completes all **78 cards**, and the GDD lists “the collection ends at 78” as a risk. The daily draw brings players back once a day but offers nothing new to collect after the deck is done.

The feature

Each season (for example *The Lantern Festival*) opens a short side road on the map: 12 hand-made levels over four weeks, three a week, each awarding one of 12 **event cards** with new art and insights. Completing the set unlocks a three-card **spread** (past, present, future) drawn from the player's own cards, and a card back to keep. Seasons repeat yearly, so a missed card can be earned the next year.



The collection: 78 cards, finished by the end of the journey.

Rules and numbers

Levels	12 per season, hand-made, at the difficulty of the player's journey position ±5 levels
Cadence	3 levels a week for 4 weeks; levels stay open until the season ends
Rewards	12 event cards, a foil for 5+ moves left, a card back for the full set
Spread	Unlocked by the set; a three-card reading from owned cards, once a day during the season
Monetisation	None new; the existing extra moves and boosters apply

Economy impact

No currency changes: it is content, not economy. It does use the existing level formula as a floor and the bot to check each hand-made level (60–95% for the goal-focused bot), so a season can be authored and validated in about a week.

Build: M. Reuses the map, the level loader (LevelCatalog's hand-made levels), the card renderer and the daily draw; new: 12 levels, 12 card arts, an event panel.

Player flow



Risks

RISK	MITIGATION
Content treadmill for a small team	12 levels a season and yearly repeats; the level editor and the bot make each level cheap to check
Tarot readings read as fortune-telling	Same tone rule as the daily draw: reflection and art, never prediction
Splits attention from the main journey	Event levels match the player's journey difficulty and are only three a week

A/B test

Control	No season	Variant	Season 1 for players past level 20
Primary	D30 retention: 6% → 7.5% (+1.5 points)	Sample	≈ 4,386 new players per arm (two-sided 5%, 80% power), the four-week season plus one week
Guardrails	journey progress per week must not fall; daily-draw use; extra-moves uptake		