

# The Runway Cliff Test

Four questions before you raise again — a pre-mortem for any AI startup.

## 1 Runway Denominator

Is runway cash ÷ burn — or cash ÷ free-cash-flow trajectory?

**Stability AI:** ~\$76M raised, ~\$4.4M left twelve months later.

**Your test:** recompute runway on your cash-flow trajectory, not last month's burn.

**Fix:** state runway as a range; add a 1.2x buffer on volatile costs.

## 2 Cost Elasticity

Can your single biggest cost shrink if demand misses plan?

**Stability AI:** ~\$300M of forward cloud commitments couldn't flex down.

**Your test:** can your largest cost fall within 30 days if demand drops 30%?

**Fix:** avoid all-in take-or-pay; route models; own one inference path.

## 3 Margin Truth

Does each extra unit of demand earn money — or burn it?

**Stability AI:** ~\$99M of compute to earn ~\$11M; 150M downloads, little paid use.

**Your test:** compute gross margin on AI revenue alone, not blended.

**Fix:** separate the AI-layer margin; price for consumption, not seats.

## 4 Capital Discipline

Is the round buying a fix — or buying denial?

**Stability AI:** capital funded new modalities; ~\$239M cumulative deficit.

**Your test:** write down what the round specifically changes in unit economics.

**Fix:** raise from strength (12+ months); track burn multiple, not just runway.