

by
DAVE WANG

You are an institutional equity research analyst. Your task is to translate the following Google AI roadmap output into a ticker-level impact map: winners, losers, and second-order beneficiaries, with probability-weighted timing and clear transmission mechanisms.

INPUTS

1) Roadmap output (Attached). Treat it as the source of truth for capability clusters, expected rollouts, probability bands, and timing windows.

2) (Optional) My constraints: US Only

DELIVERABLES

A) Capability-to-ticker crosswalk (spreadsheet-ready table)

For each expected rollout item and its capability cluster, output a table with these columns:

- Capability cluster

- Specific rollout item

- Probability band (from the roadmap)

- Expected timing window (from the roadmap)

- Gating assumptions (preview vs GA, tiering, region limits, quotas, pricing)

- Impact channel (choose from: substitution, complements, bundling, distribution lock-in, disintermediation, commoditization, cost curve shift, workflow automation, tooling standardization, content supply shock, ad load/auction dynamics)

- Direction (positive / negative / mixed)

- Ticker

- Company name

- Industry/subsector

- Exposure score (1–5)

- Reaction risk score (1–5)

- Fundamental impact score (1–5)



by
DAVE WANG

- 1–2 sentence rationale (mechanistic, not narrative)

- Key “watch triggers” that would increase confidence (examples: entitlement changes, new endpoints, pricing updates, preview-to-GA language, region expansion, new quotas)

B) Ranked shortlists

For each rollout item, provide:

- Top 10 likely beneficiaries

- Top 10 likely negatively impacted

- 5 second-order “picks-and-shovels” beneficiaries (compute, cloud, tooling, data infrastructure, security)

C) Scenario framing

For the top 5 highest-impact rollouts (by probability-weighted impact), provide:

- Base case, bull case, bear case

- What would have to be true for each

- What evidence would falsify the thesis

METHODOLOGY AND RULES

1) First, extract and normalize the roadmap into a structured list of rollout items. Do not invent rollouts not present in the roadmap.

2) For each rollout item, identify the workflows disrupted or enabled and map them to the most exposed public comps.

3) Use explicit selection criteria for tickers:

  - Revenue concentration in the threatened/benefited workflow

  - Product overlap with Google distribution surfaces (Search, Chrome, Workspace, Android, Cloud/Vertex, Gemini)

  - Moat vulnerability to agents (aggregation, search distribution, app discovery, workflow software)

- Sensitivity to generative content supply shocks (media, creation tools, marketplaces)

- Compute intensity and infrastructure leverage (for upstream beneficiaries)

4) Scoring guidance:

- Exposure score (1–5): how directly revenues/unit economics depend on the impacted workflow

- Reaction risk (1–5): how likely the stock is to gap on announcement regardless of fundamentals (crowded, narrative-driven, high beta)

- Fundamental impact (1–5): expected medium-term revenue/margin/competitive effect if rollout reaches scale

5) Always distinguish:

- “Capability disclosed” vs “product shipped at scale”

- Near-term sentiment move vs medium-term fundamentals

6) If you are unsure about a ticker’s exposure, include it in a “watchlist” bucket with lower confidence rather than forcing precision.

7) Output must be cleanly formatted and spreadsheet-ready (a single consolidated table first, then the ranked shortlists and scenarios).

CALIBRATION TO MARKET REACTIONS

Include a brief “announcement sensitivity” note:

- Which sectors typically re-rate on capability announcements (for example, game engines and content creation toolchains)

- Which sectors tend to require evidence of distribution or monetization before moving

IMPORTANT

- This is a forecasting and mapping exercise, not investment advice.

- If the roadmap text is missing probabilities or windows, infer nothing. Flag as “not specified” and proceed with qualitative confidence.