

Palantir: Where Should the Next Deployment Team Go?

Internal ID: S&F-STR-440

Topic: Growth Strategy, Customer Strategy, Operations Strategy

Industry: Application Software

Source: Strategy & Finance

Publication date: 5 October 2026

License: Not specified

Decision date and information cutoff: October 2, 2026, 4:00 p.m. Eastern Time

Decision scope: Allocation of the next increment of customer-facing deployment and product capacity between U.S. commercial AIP opportunities and government programs

Case convention: This is an author-created classroom decision exercise based on public information. Palantir has not disclosed this specific decision meeting or the alternatives presented here. The assumed capacity constraint is for analysis; public evidence does not establish that capacity is binding. Later outcomes are excluded.

The decision

Palantir Technologies has two large, fast-growing customer markets. U.S. commercial revenue is accelerating amid customer interest in its Artificial Intelligence Platform (AIP), while government programs continue to account for slightly more than half of company revenue. Both markets may benefit from Palantir's software and implementation model, but they compete for some of the same scarce people: forward-deployed engineers, product specialists, security expertise, and experienced deployment leaders.

Assume Palantir's leadership can direct its next increment of customer-facing deployment and product capacity. Should it:

1. **Prioritize U.S. commercial AIP scaling:** direct the marginal teams toward commercial pilots, production deployments, and reusable workflows;
2. **Prioritize government programs:** direct more of the marginal capacity toward government deployments, including complex mission workflows and program expansion; or
3. **Allocate through customer-level gates:** retain teams in both markets and release additional capacity as deployments meet defined evidence gates for production, reuse, customer value, and economics.

The alternatives are constructed for this case; they are not disclosed management plans. Recommend an allocation rule and explain what evidence would cause leadership to shift capacity. The choice is not whether to leave either market. It is where the next deployment team is most likely to produce durable customer value and a return to Palantir.

What the recent results show

For the quarter ended June 30, 2026, Palantir reported \$1.935 billion of revenue, up 93% year over year. Government revenue was \$990.0 million, up 79%; commercial revenue was \$945.4 million, up 110%. Those are global customer-type segments. In the United States, government revenue was \$809 million (+90%) and commercial revenue was \$764 million (+149%). Government and commercial classifications should not be confused with U.S. and non-U.S. geography. [1]

Expansion of existing accounts contributed materially to both segments' growth. Palantir attributed \$428 million of the year-over-year Q2 government revenue increase and \$407 million of the commercial increase to customers that existed at December 31, 2025. The filing does not break that growth into price, seats, usage, modules, or new workflows, and does not identify how much was caused by AIP. [1]

Commercial showed a higher reported contribution margin in Q2. Commercial contribution was \$734.9 million, or 78% of commercial revenue; government contribution was \$701.1 million, or 71% of government revenue. Palantir defines these as non-GAAP segment measures. The calculation excludes stock-based compensation and allocates cost of revenue and sales-and-marketing expenses using account headcount and time, with direct cloud costs assigned to accounts. The company says these costs include customer deployment, operation, acquisition, and expansion. Contribution is not GAAP operating income and does not measure the return on moving one engineer or product team from one segment to another. [1]

There are costs beneath the growth. Q2 cost of revenue rose 54% year over year to \$296.9 million; \$89 million of that increase came from third-party cloud hosting. Sales and marketing expense rose 39% to \$339.5 million, and research and development rose 43% to \$192.5 million. These are consolidated expense trends. They do not show the marginal cost of an additional commercial or government deployment, or how much hosting cost is recoverable from customers. [1]

The company describes a customer model that can require upfront effort. It says pilots and bootcamps are generally undertaken at Palantir's expense without guaranteed future return. It also says it manages investment at the account level rather than by industry or sector. Public filings do not disclose bootcamp cost, conversion rates, time to production, deployment hours, or team utilization. [1]

Demand signals are not the same as realized business

Palantir reported \$2.132 billion of U.S. commercial total contract value (TCV) for Q2, up 153% year over year, and \$6.238 billion of U.S. commercial remaining deal value (RDV) at June 30, up 124% year over year and 27% quarter over quarter. These figures indicate commercial demand under Palantir's definitions, but they are not recognized revenue, cash collected, or a guaranteed backlog. The company's methodology presumes available options will be exercised and contracts will not be terminated; most contracts permit termination for convenience. Timing, scope, exercise, and termination affect how much value is ultimately realized. [2]

At year-end 2025, company-wide RDV was \$11.2 billion: \$6.8 billion commercial and \$4.4 billion government. The 2025 filing separately excluded \$12.3 billion in government indefinite-delivery, indefinite-quantity contract ceilings where funding was undetermined or unguaranteed. RDV and government contract ceilings therefore require different interpretations and should not be added together as committed revenue. [3]

In July 2025, the U.S. Army announced an enterprise agreement consolidating 75 contracts into a framework with a performance period of up to ten years and a purchase-option cap of up to \$10 billion. The Army explicitly described that figure as maximum potential value, not an obligation or commitment. The agreement may simplify procurement and create potential reach, but the public announcement does not show funded task orders, take-up, realized revenue, or deployment economics. [4]

From proof of concept to repeatable platform

A useful question is whether a promising pilot becomes a production workflow and then a repeatable platform. Haefner and colleagues' review describes AI implementation as connected stages of proving concepts, productionizing, and platformizing. It emphasizes technical foundations such as data pipelines and infrastructure as well as organizational capabilities, vision, and structure. Their framework came from thematic analysis of 22 practitioner podcast interviews; it is a way to organize questions, not a representative or causal estimate of Palantir's performance. [5]

Two experiments offer additional, limited insight into AI-enabled work. In a randomized study of 758 BCG consultants using GPT-4 on designed knowledge-work tasks, AI assistance improved performance on tasks within the tested capability frontier, but participants were less likely to produce a correct answer on a task outside that frontier. The experiment involved one consulting firm, short tasks, and a model available in June 2023. It does not measure AIP, Palantir's deployment productivity, or customer returns. [6]

A separate study's abstract reports a 14% average increase in issues resolved per hour after a generative-AI assistant was introduced to 5,179 customer-support agents at one company; gains were concentrated among less experienced agents. The researchers' full text was not available for this case preparation, so this is provisional context only. Its customer-support setting and intervention are not estimates of AIP's effect or Palantir's implementation capacity. [7]

The studies suggest questions, not a forecast: Which workflows are suitable for AI assistance? Where is human review necessary? What data, security, and organizational changes are needed to move from a demo to production? How much of a solution can be reused, and which customer requirements remain bespoke?

Customer evidence: a promising but selected example

Airbus and Palantir launched Skywise as a data platform joining aircraft, fleet, maintenance, and operational information with Airbus's aerospace expertise and Palantir's platform deployment. Airbus later described a supplier rollout that began with ten selected early adopters, partly chosen for strategic representation and digital readiness. It reported 25% time savings across five functions at Premium Aerotec and a "typically" 30% reduction in part removals due to non-conformities at Thales on a per-component basis. [8]

These examples illustrate how a platform could connect data, domain knowledge, and workflows across an ecosystem. The figures are Airbus/customer-reported and scoped to selected participants. Public accounts do not disclose denominators, independent measurement, counterfactuals, or durability. They are not estimates of typical Palantir customer impact, AIP performance, or contribution margin.

Three ways to allocate the next team

Prioritize U.S. commercial AIP scaling

A commercial-first allocation follows the fastest recent growth and the higher reported contribution margin. A team could work on promising U.S. commercial pilots, accelerate suitable workflows into production, and identify reusable components across customers. U.S. commercial TCV and RDV add demand signals, and existing-customer expansion indicates that growth is not confined to new logos.

The case for this option depends on the growth becoming repeatable. Public disclosures do not identify AIP's share of revenue or contract value, commercial customer cohorts, pilot conversion, renewal, deployment effort, or reusable-workflow economics. Faster growth could reflect several products, customer expansions, concentrated demand, or a temporary adoption window. Higher reported contribution margin does not prove that the marginal commercial deployment produces the higher return after its full incremental costs.

Prioritize government programs

A government-first allocation builds on a market that generated \$990 million of Q2 revenue and grew 79%. Government programs can involve mission-critical workflows, complex data, and a need for deep deployment expertise. Existing-customer contribution to growth was similar in dollar scale to commercial, and the Army's enterprise agreement demonstrates a broader procurement framework.

This path also brings risks. Agencies and programs can face procurement delays, budget changes, option decisions, and complex security or deployment requirements. Contract vehicles and caps are not funded purchases. The 71% reported contribution margin is lower than commercial's, but company disclosures do not show whether deployment type, customer mix, contract timing, or allocation methods explain the difference. Government capacity may support durable customer relationships and product learning, but the public data do not quantify those returns.

Allocate through customer-level gates

A gated portfolio could keep both markets staffed while directing the next available team toward the opportunity with the strongest evidence at that moment. The gates might ask whether the use case has a measurable operational objective; whether data, security, and organizational owners are ready; whether the workflow has reached production; whether users repeatedly rely on it; whether adjacent workflows can reuse product components; and whether customer expansion, renewal, and contribution justify further effort.

This design can make investment more disciplined without assuming each segment needs identical skills or has the same time horizon. It also has costs: gates require reliable customer-level data, can add process, and may favor nearer-term commercial revenue unless government cycles are handled fairly. The public record does not provide the data to set thresholds, prove a capacity bottleneck, or establish one optimal allocation rule.

Evidence the leadership team would need

Before moving more capacity, the team would want comparable, customer-level evidence, while preserving differences in procurement and mission cycles. Useful information could include:

- pilot-to-production timing and conversion, by customer type and use case;

- customer outcomes defined before deployment, with baselines and credible comparison methods;
- deployment hours, specialist requirements, compute usage, and costs through production and expansion;
- repeated use, renewal, expansion, and contraction by customer cohort;
- product components reused across deployments and effort required for customer-specific work;
- funded government task orders and option exercise, separated from contract ceilings and optional value;
- cloud and support costs, including whether and how they are recovered from customers;
- customer concentration, procurement dependencies, and the skills or security clearances required for each opportunity.

The case does not establish whether customer-facing capacity is the binding constraint. It is possible that demand, customer readiness, procurement, cloud availability, product roadmap, or another input is more limiting. Any thresholds, review intervals, and staffing percentages proposed by students are hypothetical management assumptions.

Assignment

Assume you advise Palantir's hypothetical leadership team at the October 2, 2026 cutoff. Recommend where to direct the next increment of customer-facing deployment and product capacity.

1. Choose a commercial-first, government-first, or gated portfolio policy. What would the policy allocate now, and what would cause it to change?
2. How should leadership compare growth, contribution margin, customer value, deployment and cloud costs, and contract visibility?
3. What evidence would distinguish repeatable AIP adoption from demand driven by other products, existing-account expansion, or other factors?
4. How should the team account for government procurement timing and mission value without treating options or ceilings as realized revenue?
5. Which assumption most affects your recommendation, and what evidence would change your view?

Do not attribute reported growth to AIP without evidence, treat contribution margin as GAAP profit or a marginal return, infer a binding capacity constraint, or count TCV, RDV, or the Army cap as guaranteed revenue. Clearly label student assumptions and proposed thresholds.

Sources and conventions

[1] Palantir Technologies Inc., [Q2 2026 Form 10-Q](#), filed August 4, 2026. Revenue and existing-customer comparison, printed pp. 29-30; contribution definition and allocation, pp. 25-26; Note 12 contribution table, pp. 18-19; cost trends, pp. 30-31; customer model and geographic mix, pp. 22-24.

[2] Palantir, [Q2 2026 earnings release](#), August 3, 2026. U.S. commercial TCV/RDV highlights and metric definitions, including options and termination language.

[3] Palantir Technologies Inc., [2025 Form 10-K](#), filed February 17, 2026. “Total Remaining Deal Value,” pp. 67-69, and related risk factors.

[4] U.S. Army Public Affairs, [“U.S. Army Awards Enterprise Service Agreement to Enhance Military Readiness and Drive Operational Efficiency,”](#) July 31, 2025. Contract consolidation, term, option cap, and explicit no-obligation qualification.

[5] Naomi Haefner, Vinit Parida, Oliver Gassmann, and Joakim Wincent, “Implementing and scaling artificial intelligence: A review, framework, and research agenda,” *Technological Forecasting and Social Change* 197 (2023), 122878. [Full text](#). Exploratory qualitative analysis of 22 practitioner podcast interviews; used for implementation stages and mechanisms, not a performance estimate.

[6] Fabrizio Dell’Acqua et al., “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality,” *Organization Science* 37, no. 2 (2026), 403-423. [Article](#). Preregistered randomized experiment involving 758 BCG consultants; single firm, selected tasks, GPT-4, not an AIP or deployment-productivity estimate.

[7] Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” *Quarterly Journal of Economics* 140, no. 2 (2025), 889-942. [NBER abstract and summary](#). The case preparation accessed the abstract/summary only; the 14% figure is provisional context, not a Palantir estimate.

[8] Airbus, [“Airbus launches Skywise - aviation’s open data platform,”](#) June 2017; and [“Airbus extends Skywise to suppliers,”](#) July 2018. Customer/counterparty descriptions and selected, customer-reported supplier examples; not independent causal evidence.