PLTR gideon lens

Palantir Technologies Inc. (PLTR) - Gideon-Lens Equity Research

Report date: 2026-08-24 | **Analyst:** Codex, GPT-5.6 Sol High | **Price:** USD 179.94 as of the completed 2026-08-21 XNAS session | **Market capitalization:** approximately USD 462.21 billion on 2,568.694 million quote-basis diluted shares | **Fiscal year-end:** December 31

Verdict block

Verdict: PASS

Scorecard:

Gate	Result	Notes
Management	FAIL	Exceptional execution does not cure Class F founder control, chronic dilution, and weak minority-shareholder protection
Moat	YELLOW	Ontology/workflow switching costs and Government trust are real; Commercial durability and direct pricing power are not cycle-tested
2046 durability	YELLOW	A governed operational context layer is plausible, but hyperscaler bundling and modular architectures remain credible erosion paths
P/E <= 15 on confident EPS at TTM + 5 calendar years	FAIL	Base EPS USD 3.724686 implies USD 55.87 at 15x, a 68.95% shortfall to the current price
Compounder Score	Audited 55 / 100	Reject classification; hard vetoes and inadequate returns prevent dream-investment status
Business Quality Score	Audited 53 / 80	Exceptional runway and margins are offset by governance, dilution, disclosure, and durability limits
20% Return UW Score	Audited 2 / 20	Base and probability-weighted returns are negative; even Bull is below 20%

One-line thesis (or reason to pass): Palantir is executing extraordinarily well, but Management FAIL, a YELLOW 2046 moat, a 153x reported P/E, and negative Base and probability-weighted returns make PLTR a PASS at USD 179.94.

Security return underwriting (mandatory even when the verdict is PASS):

Measure	Result	Notes
Primary valuation method	P/E	One independently locked method is used across all scenarios, probability weighting, and net-SGD
Gross local-currency Bear / Base / Bull IRR	-26.29% / -6.24% / 11.90%	Five-year annualized shareholder return in USD
Probability-weighted gross local-currency IRR	-4.31%	CAGR of USD 144.386481 probability-weighted terminal wealth using 28.4353% / 50.0000% / 21.5647% weights
Return bridge	USD 179.94 -> USD 39.15 / 130.36 / 315.67	Exact horizon 2026-06-30 to 2031-06-30; terminal statutory diluted EPS times 25x / 35x / 45x P/E; zero distributions

Gross local-currency underwriting is in USD. The price anchor is the completed 2026-08-21 session and the operating anchor is TTM ended 2026-06-30. Bear, Base, and Bull terminal equity values are USD 102.178712 billion, USD 338.666063 billion, and USD 818.277069 billion on 2,610.088747 million, 2,597.849227 million, and 2,592.215058 million diluted shares. Current cash and marketable securities are USD 9.409099 billion, financial debt is zero, operating lease liabilities are USD 0.211400 billion, and current net cash after leases is USD 9.197699 billion. Under the locked P/E method, cash-driven interest income, lease cost, NCI attribution, and capital structure already affect statutory EPS and the equity multiple; no terminal net cash is added and no lease claim is subtracted again. Distributions are explicitly zero because Palantir pays no current dividend and no source-bound distribution path was underwritten.

Bear-case operating and return bridge

Bear case	Start (TTM 2026-06-30; price 2026-08-21)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.155941B	USD 12.753092B	15.68%	Government and Commercial decelerate sharply as procurement, customer concentration, and competition constrain adoption
Operating margin	42.80%	42.81%	+0.02 pp	Contribution pressure and a 24% revenue burden for R&D, G&A, and SBC eliminate incremental statutory leverage
Operating earnings	USD 2.634652B	USD 5.460195B	15.69%	Exact Contribution bridge flows once through shared corporate and statutory costs
Diluted share count	2,568.694M point-in-time	2,610.088747M	0.32%	Scenario dilution retains SBC and outstanding award effects; no unsupported buyback offset
EPS	USD 1.173334 TTM	USD 1.565904	5.94%	Statutory diluted EPS after tax, below-line items, NCI, and scenario shares
Valuation denominator	153.36x reported / 203.11x normalized	Terminal EPS USD 1.565904	Comparable normalization disclosed	Current extremes are disclosure, not terminal anchors
Valuation	153.36x reported / 203.11x normalized	25.00x P/E	-178.11x vs normalized	Mature-software peer evidence and governance/durability risks justify severe normalization
Share value	USD 179.94	USD 39.15	-26.29%	USD 1.565904 EPS times 25x; USD 102.178712B terminal equity divided by diluted shares
Annual dividend	USD 0.00	USD 0.00	0.00%	Explicit zero-distribution diagnostic
Five-year distributions	-	USD 0.00	cumulative	Counted once at zero
Total ending wealth	USD 179.94	USD 39.15	-26.29% gross USD IRR	Terminal share value plus zero distributions

Base-case operating and return bridge

Base case	Start (TTM 2026-06-30; price 2026-08-21)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.155941B	USD 20.896081B	27.69%	U.S. Commercial expansion and Government demand remain exceptional but decelerate from current rates
Operating margin	42.80%	57.55%	+14.75 pp	Contribution scale and lower R&D/G&A/SBC ratios produce extraordinary leverage while costs still rise in dollars
Operating earnings	USD 2.634652B	USD 12.026335B	35.48%	Government and Commercial Contribution reconcile through one statutory bridge
Diluted share count	2,568.694M point-in-time	2,597.849227M	0.23%	Persistent equity issuance remains; the path is not erased by an assumed repurchase program
EPS	USD 1.173334 TTM	USD 3.724686	25.99%	Statutory diluted EPS after 21% tax, below-line income, NCI, and dilution
Valuation denominator	153.36x reported / 203.11x normalized	Terminal EPS USD 3.724686	Comparable normalization disclosed	Normalized current earnings do not support retaining today's multiple
Valuation	153.36x reported / 203.11x normalized	35.00x P/E	-168.11x vs normalized	A premium to mature peers remains, but governance FAIL and YELLOW moat cap it
Share value	USD 179.94	USD 130.36	-6.24%	USD 3.724686 EPS times 35x; USD 338.666063B terminal equity divided by diluted shares

Base case	Start (TTM 2026-06-30; price 2026-08-21)	Horizon (2031-06-30)	CAGR / change	Justification
Annual dividend	USD 0.00	USD 0.00	0.00%	Explicit zero-distribution diagnostic
Five-year distributions	-	USD 0.00	cumulative	Counted once at zero
Total ending wealth	USD 179.94	USD 130.36	-6.24% gross USD IRR	Terminal share value plus zero distributions

Bull-case operating and return bridge

Bull case	Start (TTM 2026-06-30; price 2026-08-21)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.155941B	USD 32.798045B	39.74%	AIP-led expansion, durable NRR, Government demand, and international scaling sustain category leadership
Operating margin	42.80%	65.62%	+22.82 pp	Platform reuse and revenue density overwhelm inference, deployment, and corporate cost growth
Operating earnings	USD 2.634652B	USD 21.520588B	52.20%	Statutory bridge retains all modeled corporate costs and SBC
Diluted share count	2,568.694M point-in-time	2,592.215058M	0.18%	Bull still includes dilution rather than assuming capital allocation cures it
EPS	USD 1.173334 TTM	USD 7.014825	42.99%	Statutory diluted EPS after 18% tax, favorable below-line income, NCI, and shares
Valuation denominator	153.36x reported / 203.11x normalized	Terminal EPS USD 7.014825	Comparable normalization disclosed	Terminal denominator is the same audited statutory diluted EPS basis
Valuation	153.36x reported / 203.11x normalized	45.00x P/E	-158.11x vs normalized	Reserved for sustained low-20s terminal growth and exceptional margins, below ServiceNow and PLTR extremes
Share value	USD 179.94	USD 315.67	11.90%	USD 7.014825 EPS times 45x; USD 818.277069B terminal equity divided by diluted shares
Annual dividend	USD 0.00	USD 0.00	0.00%	Explicit zero-distribution diagnostic
Five-year distributions	-	USD 0.00	cumulative	Counted once at zero
Total ending wealth	USD 179.94	USD 315.67	11.90% gross USD IRR	Terminal share value plus zero distributions

CIO actionability (separate from ER2 underwriting):

Measure	Result	Notes
Status	PASS	Current-price actionability result, separate from the exhaustive-rule ER2 verdict
Bear / Base / Bull net-SGD IRR	-27.90% / -6.34% / 13.94%	Same P/E terminal wealth translated at 1.142784 / 1.269760 / 1.396736 USDSGD and reduced by a 0.5% terminal-wealth cost
Probability-weighted net-SGD IRR	-3.66%	CAGR of SGD 189.613489 weighted terminal wealth from SGD 228.480614 starting value
10% mandate hurdle	FAIL	Base and weighted returns are negative
15% stretch hurdle	FAIL	Even Bull remains below 15% net SGD
Downside and portfolio fit	FAIL / UNAVAILABLE	Bear loses most capital; holdings, factor exposures, liquidity needs, and risk budget were not supplied
Proposed sizing ceiling	none	Missing portfolio context constrains fit and sizing only, not finite returns

The net-SGD overlay uses live 2026-08-21 USDSGD of 1.269760, symmetric -10% / flat / +10% terminal FX diagnostics, zero distribution withholding because distributions are zero, and no personal-tax assumption. It uses the same primary P/E method and scenario probabilities as gross USD underwriting.

Trade execution authorized: NO

Kill / trigger criteria:

- PASS trigger persists while Management remains FAIL, Moat remains below PASS, or Base gross return remains below 10%.
- Downgrade evidence: diluted shares continue compounding above 3% annually without FCF-per-share or ROIC-based compensation and repurchase accountability.
- Downgrade evidence: NRR falls below 120%, U.S. Commercial customer additions stall, or Contribution margins fall more than 500 bps without a verified investment payback.
- Downgrade evidence: Government award concentration, procurement delays, or program loss causes two consecutive quarters of segment contraction.
- Upgrade condition: Class F control gains a credible sunset, minority-shareholder protections improve, and compensation binds to ROIC/FCF per share.
- Upgrade condition: audited Commercial retention, churn, payback, platform economics, and pricing-power evidence clear the moat/2046 gate.
- Price-only diagnostic: USD 52.39 is the Base maximum entry price for a 20% gross USD IRR under the locked 35x exit; it is not a BUY trigger while qualitative gates fail.

Research sizing if actionable: none. No ranking, sizing, portfolio mutation, order, message, or BUY-queue action is proposed or authorized.

Stage 0 - Time anchor

Today is 2026-08-24 SGT. Palantir's fiscal year ends December 31. The latest complete period is Q2 FY2026 ended 2026-06-30. TTM is Q3 FY2025 plus Q4 FY2025 plus Q1 FY2026 plus Q2 FY2026, ending 2026-06-30; the horizon is exactly five calendar years later, 2031-06-30.

Stage 1 - Disqualifier sweep

ROIC: YELLOW - conventional ROIC is depressed by excess cash and R&D expensing; a directional three-year ROIIC proxy is strong but not a full cycle. **FCF conversion: GREEN** - TTM FCF is USD 3.358 billion and exceeds statutory earnings, though SBC is a material add-back. **Leverage: GREEN** - USD 9.409 billion cash/marketable securities and no financial debt. **End markets: GREEN** - Government modernization and enterprise AI adoption have long runways, with procurement and concentration risks. **Governance: RED** - Class F founder control, minority-shareholder protection, compensation design, and dilution fail the management gate. **AI disruption: YELLOW** - Palantir can benefit as the governed action layer, but model and platform commoditization remain credible. **Pricing power: YELLOW** - NRR and margins imply expansion, but direct price/churn evidence and contract termination rights limit confidence.

Management and moat

Management is FAIL. Seven completed quarterly revenue guides were exceeded and operating execution is exceptional, yet Class F can preserve 49.999999% founder control as economic ownership falls. Diluted shares rose about 33% from FY2021 to FY2025; buybacks were immaterial versus SBC. Compensation lacks explicit ROIC or FCF-per-share hurdles, and no named lead independent director offsets a non-independent founder chairman.

Moat and 2046 are YELLOW. Government has the strongest protection through classified deployment, procurement trust, accumulated mission workflows, and Ontology embedding. Commercial switching costs become meaningful after production deployment, but pilots and standardized workflows are contestable. A 2046 governed operational context layer is specific and plausible, yet hyperscaler bundling, modular architectures, and foundation-model commoditization could compress pricing and reduce Palantir's role.

Stage 5 synthesis

The exact `Contribution` bridge and all 11 synthesis tie-outs clear. TTM revenue is USD 6.155941 billion, Contribution USD 4.391021 billion, statutory operating income USD 2.634652 billion, and diluted EPS USD 1.173334 on 2,571.043 million weighted shares. At 2031-06-30, Bear/Base/Bull revenue is USD 12.753092/20.896081/32.798045 billion and statutory operating income is USD 5.460195/12.026335/21.520588 billion. Required EPS at 15x and USD 179.94 is USD 11.996; Bear/Base/Bull EPS of USD 1.565904/3.724686/7.014825 all fail, with Base short by 68.95%. The P/E-15 gate is an entry-price diagnostic, not an exit-multiple instruction.

Red team

The strongest bear case is that incumbent AI/data/workflow bundles and Government portability mandates erode Palantir's control layer, NRR, pricing, and margins. The most optimistic Stage 5 number is the 57.55% Base terminal operating margin; holding it to 50% reduces Base gross IRR from -6.24% to -8.78%. A 30x rather than 35x Base exit reduces Base IRR to -9.09%. Conversely, the model is conservative through zero distributions and no separate terminal cash credit; even adding current net cash would improve Base IRR by only 0.51 percentage points. Reweighted scenario probabilities move the probability-weighted return by only 0.58 points. The red team independently reproduced Base and weighted gross/net-SGD returns and classifies the objection **IMMATERIAL**. No upstream repair is required. A force-exit signal is companywide NRR below 110% for two consecutive reported quarters.

Analyst flags

- The quote packet is valid through 2026-08-25T18:03:08+08:00 under its conservative fallback policy. A later price should update price-dependent actionability, not rewrite historical evidence.
- Platform-level revenue, margins, customer counts, churn, CAC, payback, billings, ARR, and cRPO are not disclosed. Conclusions are constrained rather than filled with estimates.
- Customer I represented 27% of Q2 receivables; concentration and Government procurement remain material even as total customer count expands.
- A ten-year mature-cycle margin record is unavailable; 2046 durability is therefore a reasoned but unproven claim.
- Investor-specific brokerage, tax wrapper, holdings, factor exposures, and risk budget were unavailable. Only portfolio fit and sizing remain unassessed.
- This is the first governed Production ER2 for PLTR, not a refresh; no prior-report comparison is required.

Appendix A - Company Primer

Palantir Technologies Inc. (PLTR) — Company Primer

Research date: 2026-08-24 SGT **Financial anchor:** TTM ended 2026-06-30 **Forward horizon used by the ER2:** 2031-06-30, exactly five calendar years from the TTM anchor **Trade execution authorized:** NO

Source discipline and reading guide

All financial figures in this primer come from the QFR-cleared [`quant_foundation.md`](./quant_foundation.md), whose SHA-256 is `980bbb1ad6377e8a361e0675982b89550e0d83a18c4f2d2c96c34453ec0a00e9`. The concurrent data package was still an in-progress placeholder when this primer was prepared and was not consumed. Qualitative orientation comes from Palantir's live [2025 Form 10-K](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm), [Q2 FY2026 Form 10-Q](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000041/pltr-20260630.htm), and current [product documentation](https://www.palantir.com/docs/foundry/architecture-center/platforms). The company's descriptions establish what its products are intended to do; they do not by themselves prove customer return on investment or competitive superiority.

Throughout the primer:

- **Fact** means a disclosed item or a direct calculation from the verified quantitative foundation.
- **Inference** means a reasonable economic interpretation of disclosed facts, not a company disclosure.
- **Judgment** means an investor conclusion that later ER2 stages must test rather than inherit.

1. Business overview: what Palantir actually sells

Palantir sells software that connects an organization's fragmented data to the decisions and operational actions that follow from it. The problem is not simply storing data or making a dashboard. Large institutions often have incompatible databases, legacy systems, security classifications, business rules, models, and human approval processes. Palantir attempts to create a governed operating layer across them. A hospital might connect capacity, staffing, patient-flow, and supply information; a manufacturer might connect orders, production constraints, inventory, and maintenance; and a defense organization might combine sensors, intelligence, logistics, and mission planning. The software is meant to turn those connections into repeatable workflows, applications, and increasingly AI agents.

Palantir calls the shared semantic representation of an institution the **Ontology**. In plain language, it maps raw data into recognizable operational objects—such as an aircraft, patient, supplier, order, or unit—then attaches permitted actions, rules, relationships, and access controls. This matters because a prediction that never reaches an operator's workflow has little economic value. Palantir's pitch is that the Ontology closes the gap between a data project and a production operating system. The [2025 10-K, Item 1](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm) describes a vertically integrated environment where data engineers, analysts, data scientists, frontline users, and leaders can work on the same governed representation.

The customer is usually a government agency or a commercial institution with complex, high-value operations. Palantir's own acquisition description targets large-scale, difficult opportunities at large institutions. The company does not fit the simplest self-serve SaaS archetype: deployment frequently requires detailed discovery, integration, ontology design, security configuration, and frontline adoption. Current documentation does offer an AIP Developer Tier and promotes bootcamps, so the top of the funnel is becoming more accessible, but the economic center remains institutional contracts and expansion after production use.

Fact: The company reports two operating segments—Government and Commercial—and four principal platforms—Gotham, Foundry, Apollo, and AIP. Government and Commercial are customer categories, not product lines. Customers can use multiple platforms together.

Inference: Palantir's unit of value is not a software seat alone. It is the operational workflow made possible by integrated data, software, and services. That can support high willingness to pay when the workflow is mission-critical, but it also makes selling

and implementation more labor-intensive than low-touch SaaS.

Judgment: The central underwriting question is whether Palantir can keep productizing deployment while preserving differentiated customer outcomes. If growth requires proportional growth in scarce forward-deployed talent, the headline software margins overstate long-run scalability. If bootcamps, reusable ontology components, and Apollo compress time-to-value, the opposite may be true.

2. Product architecture: Gotham, Foundry, AIP, and Apollo

The four product names are best understood as a stack rather than four isolated subscriptions.

Platform	Plain-English role	Primary users	Economic role
Gotham	Government-oriented operational and mission system that combines multi-source data, analysis, planning, and field action	Defense, intelligence, public-sector operators and analysts	Core Government workload and domain application layer
Foundry	Data operations and application platform that turns fragmented enterprise data into governed objects, logic, analytics, and workflows	Commercial teams and many government users	Foundational data and workflow layer across both segments
AIP	Generative-AI and agent layer that connects models to governed organizational data, tools, evaluations, and human approvals	Developers, analysts, business users, frontline operators	Accelerates new use cases and expands platform consumption
Apollo	Continuous-delivery and software-operations layer across cloud, on-premises, classified, disconnected, and edge environments	Palantir deployment teams and customers operating software fleets	Enables deployment breadth, upgrades, reliability, and cross-environment scale

Gotham

Gotham began in intelligence and counterterrorism and now serves broader defense and government missions. It integrates information from multiple domains and sensors, preserves access controls, supports investigation and analysis, and hands insights to operational users. The 10-K describes use from operations centers to the tactical edge, including planning and execution while maintaining an operational picture. The key point is the analyst-to-operator handoff: Gotham is intended to affect real decisions, not merely produce reports. It can be paired with Foundry, AIP, Apollo, and defense-specific applications.

Foundry

Foundry is the common data operations foundation. It connects sources, builds governed pipelines, represents the organization in the Ontology, runs models and analytics, and supports user-facing applications. Palantir says all commercial customers use Foundry and many government customers use it. A useful mental model is an integrated combination of data engineering, semantic modeling, application development, analytics, governance, and workflow orchestration. That breadth is a source of potential differentiation but also creates overlap with many incumbent categories.

AIP

AIP connects large language models and other AI systems to the organization's governed data and operations. Its value proposition is not that Palantir owns the best foundation model. Palantir's [AIP documentation](https://www.palantir.com/docs/foundry/aip) emphasizes connectivity to multiple models, tools for agents and automations, evaluations, security, auditability, and human review. The commercial thesis is therefore model-agnostic orchestration: customers can use evolving models while retaining permissions, lineage, operational context, and controls in the Palantir layer.

This architecture matters because general-purpose LLMs do not automatically know which data a user is permitted to see, what a "late order" means in a specific company, or which actions require approval. AIP uses the Ontology to supply that context. In a strong deployment, AI becomes another actor inside a governed workflow rather than a separate chatbot.

Apollo

Apollo manages software delivery and configuration across diverse environments. Palantir originally built it to update Gotham and Foundry in cloud, on-premises, restricted, and rugged environments. It now also supports customer software. Apollo's

economic importance is easy to underestimate: a platform cannot become an operating system if updates are unreliable or impossible in classified and disconnected settings. Current [architecture documentation](https://www.palantir.com/docs/foundry/architecture-center/platforms) describes Apollo as the infrastructure that manages the services underlying Foundry and AIP and orchestrates continuous upgrades.

Inference: The products reinforce one another. Foundry creates the governed data and workflow substrate; AIP adds AI-enabled agents and automation; Apollo makes the stack deployable; Gotham packages the architecture for government missions. The integration can increase switching costs because replacing one layer may require rebuilding semantic models, permissions, applications, and operating procedures—not just migrating data.

Judgment: AIP should not be valued as an independent model vendor. Its defensibility depends on Palantir remaining the trusted operational context and action layer as model capabilities commoditize.

3. Revenue architecture and segment economics

Palantir's exact issuer segments are **Government** and **Commercial**. The chief operating decision maker evaluates each using **Contribution**, defined as segment revenue less the segment's related cost of revenue and sales and marketing expense. Contribution excludes research and development, general and administrative expense, and all stock-based compensation. It is therefore not GAAP operating income and must never be labeled that way.

Segment snapshot

Period	Segment	Revenue ($m)	Share of total	Contribution ($m)	Contribution margin
FY2025	Government	2,402.3	53.7%	1,576.1	65.6%
FY2025	Commercial	2,073.2	46.3%	1,366.6	65.9%
TTM ended 2026-06-30	Government	3,210.8	52.2%	2,255.2	70.2%
TTM ended 2026-06-30	Commercial	2,945.2	47.8%	2,135.8	72.5%

Percentages above are direct calculations from the verified foundation. Segment revenue sums exactly to consolidated revenue in both periods.

Government: procurement, mission depth, and economics

Government includes U.S. and non-U.S. agencies. Gotham is especially relevant, but government customers may also use Foundry, AIP, Apollo, and domain applications. Buying differs from a normal commercial subscription. Budgets, appropriations, procurement rules, security requirements, contract options, and political priorities affect timing. U.S. federal customers cannot exercise some options more than one year in advance, and many contracts permit termination for convenience. A large announced award or total deal value is therefore not the same thing as enforceable backlog.

The Government segment's TTM contribution margin of 70.2% means that after its directly related cost of revenue and sales and marketing, roughly 70 cents of each revenue dollar remained to fund shared R&D, G&A, SBC, taxes, and shareholder returns. The margin rose from 65.6% in FY2025. That suggests strong operating leverage, mix, pricing, deployment efficiency, or some combination; the segment disclosure does not allocate enough detail to identify the exact driver.

Government economics can be attractive because mission-critical use, high security requirements, long-lived deployments, and accumulated workflows raise switching costs. The counterweight is concentration in public budgets, lengthy procurement, protest and compliance risk, potential use restrictions, and option uncertainty. Services and deployment needs can also be substantial in complex missions.

Inference: Government revenue durability is likely stronger after a platform becomes embedded across operational workflows, but contract duration alone overstates durability because convenience termination and unexercised options remain real.

Commercial: land, prove, and expand

Commercial includes customers in non-government industries. Foundry provides the common data and workflow layer; AIP can accelerate agent and automation use cases; Apollo supports deployment. Customers may start with one acute operational problem and expand across plants, business units, or workflows after proving value.

Commercial's TTM contribution margin of 72.5% exceeded Government's and rose from 65.9% in FY2025. The result is notable because Commercial includes customer acquisition and implementation activity that can be expensive. It suggests

that recent revenue has scaled faster than directly attributable cost of revenue and sales expense. However, Contribution excludes central R&D, G&A, and SBC, so it should not be treated as the segment's complete economic margin.

Commercial demand may be more discretionary than defense or intelligence demand. Prospective buyers can choose internal development, existing enterprise vendors, data platforms, cloud services, or consultants. Palantir must demonstrate a business outcome large enough to justify platform cost and organizational change. Once the platform is deeply integrated, expansion can be efficient; before that point, pilots and deployment work may consume meaningful resources.

Judgment: Commercial is the more important proof of product scalability. Sustained high contribution margin alongside rapid customer expansion would support the claim that bootcamps and productization have shortened the traditional high-touch sales motion. A reversal would suggest the current growth burst is more services-intensive than it appears.

4. How customers buy: direct sales, pilots, bootcamps, and expansion

Palantir primarily uses a direct, consultative selling motion. Sales teams and technical personnel identify a consequential workflow, connect a limited set of data, build an initial ontology and application, and demonstrate an operational result. The Q2 2026 10-Q says standard sales efforts involve considerable time and expense, often include no- or low-cost pilots and bootcamps, and may take six to nine months from demonstration to sale, sometimes a year or longer. Senior customer management is commonly involved because purchases can be financially and organizationally significant.

The **AIP bootcamp** is designed to compress the proof stage. Palantir's current [getting-started documentation](https://www.palantir.com/docs/foundry/getting-started/overview) describes working with Palantir engineers to move from zero to a use case in hours or days. In economic terms, the bootcamp is a product-led demonstration wrapped in forward-deployed engineering: real customer data and a real workflow make value tangible faster than a slide presentation. It is not automatically a durable deployment. The 10-Q explicitly warns that short-term pilots may not convert to longer-term revenue contracts.

After an initial use case, the expansion motion can proceed along several axes: more users, data sources, workflows, departments, geographic units, platform components, model consumption, or customer-built applications. This is why the Ontology matters commercially. A reusable operational model can make each additional workflow easier to build, while governance and access controls allow more teams to participate.

Contracts are generally one to five years, sometimes shorter, and revenue is generally recognized ratably across the term. Offerings include hosted Palantir Cloud subscriptions, software in a customer's environment with operating and maintenance services, and professional services. Many contracts contain termination-for-convenience provisions. Therefore three concepts must remain separate: stated contract value, noncancelable remaining performance obligations, and recognized revenue.

Channels and partners can assist implementation or distribution, and Palantir may collaborate with third-party service providers, but the company's own risk disclosure still emphasizes direct engagement and senior-management involvement. The product also has a Developer Tier, which broadens experimentation, yet there is insufficient evidence in the sealed foundation to treat self-service as a material revenue channel.

Inference: Bootcamps improve sales velocity only if four things happen: conversion to paid production, low incremental engineering effort, rapid customer adoption, and subsequent expansion. Bootcamp counts without those outcomes are a vanity metric.

Judgment: Investors should monitor cohort evidence—conversion, deployment duration, expansion, and support intensity—rather than treating a fast prototype as proof of scalable economics.

5. Unit economics: from a revenue dollar to shareholder economics

The cleanest segment lens is Contribution; the cleanest consolidated lens is GAAP operating income and free cash flow. The TTM bridge is:

TTM bridge to 2026-06-30	Amount ($m)	Share of revenue
Revenue	6,155.9	100.0%
Government contribution	2,255.2	36.6%
Commercial contribution	2,135.8	34.7%
Total contribution	4,391.0	71.3%
R&D excluding SBC	(473.6)	(7.7%)

TTM bridge to 2026-06-30	Amount ($m)	Share of revenue
G&A excluding SBC	(447.3)	(7.3%)
Total SBC	(835.5)	(13.6%)
GAAP operating income	2,634.7	42.8%

The two contributions sum to 71.3% of revenue. Shared R&D, G&A, and SBC then consume about 28.5% of revenue, leaving a 42.8% GAAP operating margin. This bridge is unusually important because an investor looking only at the roughly 70%-plus segment contribution margins would omit material central costs. Conversely, contribution margins show that direct delivery and sales costs have scaled well in the current period.

Cost of revenue is driven primarily by personnel performing operating, maintenance, and professional services; subcontractors and field-service representatives; third-party cloud hosting; and hardware or other direct deployment costs. Sales and marketing includes the direct sales force, commissions, pilots, customer-growth activity, bootcamp-related hosting, and travel. R&D supports the shared platform rather than a single segment, which is economically sensible but prevents precise product-level profitability analysis.

Palantir is physically capital-light. TTM capex was $42.0 million, only about 0.7% of revenue. TTM operating cash flow was $3,400.3 million and free cash flow was $3,358.3 million, a 54.6% margin. The free-cash-flow margin exceeding the GAAP operating margin reflects noncash expense, working capital and interest effects; it should not be interpreted as permanent profit superiority. Most importantly, SBC is added back in operating cash flow but is not free to shareholders.

Inference: The major resource constraint is human and organizational rather than physical plant. Product scalability depends on how effectively software and reusable deployment patterns substitute for forward-deployed labor, while cloud costs rise with platform usage.

Judgment: High FCF conversion is real at the corporate cash level, but per-share economics require a dilution charge. Any valuation that capitalizes FCF while ignoring share issuance would double-count the benefit of SBC.

6. Competitive landscape

Palantir competes across overlapping markets rather than within one clean category. The 10-K identifies large enterprise software companies, government contractors, systems integrators, emerging vendors, and internal customer development as competitors. A practical map is:

Competitive route	Why a customer may choose it	Palantir's claimed differentiation	Key risk to Palantir
Internal build	Control, existing talent, tailored architecture	Faster deployment and integrated data-to-action stack	Better AI coding tools reduce build cost
Cloud and data platforms	Existing spend, scale, broad ecosystems	Ontology, operational applications, cross-environment deployment	Incumbents bundle adjacent capabilities cheaply
Enterprise application suites	Installed workflows and vendor relationships	Cross-system integration and customization around unique operations	Suites extend AI and workflow features into installed base
Government primes and integrators	Procurement familiarity, program management, hardware integration	Software product depth and iterative deployment	Primes capture contract control or replicate functionality
Specialist analytics or AI vendors	Lower cost, narrower focus, faster point solution	Governance, security, model choice, and production operations	Point tools win simple use cases where Palantir is excessive
Consulting-led transformation	Executive access and organizational change expertise	Persistent software platform rather than one-off project	Consultants remain essential implementation gatekeepers

Palantir's strongest position should be in complex, heterogeneous, high-consequence environments where security, lineage, interoperability, and action matter simultaneously. Its weakest position should be a standardized, low-risk workflow adequately served by an incumbent application or a small point tool. The company itself says larger competitors may bundle products, use broader relationships, or price at very low margins.

Inference: AI is both accelerator and threat. Better models increase the value of connecting AI to real operations, benefiting AIP; they also lower the cost of writing software and may let incumbents or customers rebuild portions of Palantir's functionality.

Judgment: The moat cannot rest on model access or generic data integration. It must rest on operational ontology, deployment capability, security, accumulated applications, and customer-specific institutional learning that competitors cannot

cheaply reproduce.

7. Geography and concentration

The business is global but increasingly U.S.-weighted. The latest live filing reports a large majority of revenue from U.S.-headquartered customers or U.S. agencies, with no other country individually reaching the disclosure threshold in the latest quarter. Because the verified quantitative foundation does not include a geographic schedule, this primer deliberately does not import a new geographic revenue number. The qualitative conclusion is still important: U.S. government budgets, U.S. commercial AI adoption, regulation, and domestic macro conditions have an outsized influence.

Customer concentration also remains relevant. Palantir's Q2 filing states that its largest customers have long tenures and that its top customers represent a meaningful portion of revenue. The foundation does not contain the concentration schedule, so no concentration percentage is introduced here. Economically, concentration cuts both ways: large, long-lived deployments indicate stickiness and expansion; the loss, delay, or renegotiation of a major program can still move reported growth.

Government concentration is broader than customer concentration. Even if no single agency is dominant, U.S. appropriations, procurement policy, security requirements, and political priorities are common risk factors across many contracts. Commercial concentration can be more diversified by industry, but very large enterprise deployments may still create lumpy comparisons.

Judgment: Geographic diversification is not automatically desirable if U.S. unit economics and product-market fit are superior. The underwriting issue is whether reliance on one fiscal and regulatory ecosystem creates correlated downside that reported customer counts understate.

8. Balance sheet and capital intensity

At 2026-06-30, Palantir had $2,030.0 million of cash, $7,379.1 million of marketable securities, and no financial debt. Cash plus marketable securities totaled $9,409.1 million against $11,678.6 million of total assets. Noncurrent operating lease liabilities were $211.4 million. The company also had $110.7 million of noncontrolling interests, which must be respected in any later enterprise-to-equity bridge.

This balance sheet makes solvency risk remote under ordinary operating scenarios and gives Palantir substantial capacity to invest through a downturn. Interest income is already economically significant: the TTM statutory bridge included $266.4 million of interest income below operating profit. A later earnings model must not assume the full cash balance is both permanently surplus and simultaneously available to generate interest without an explicit capital-structure policy.

Physical investment is small relative to revenue, but cloud purchase commitments, software development, people, and customer deployment capacity are real operating requirements. Low capex does not mean zero reinvestment. Much of the reinvestment appears in the income statement as R&D, sales effort, and compensation rather than on the balance sheet.

Inference: Palantir can fund growth internally and need not issue equity to finance normal operations. Future issuance is therefore principally a compensation and capital-allocation choice rather than a liquidity necessity.

9. Stock-based compensation and dilution

SBC is the largest adjustment separating segment contribution from GAAP operating income. It was $835.5 million in the TTM, equal to 13.6% of revenue and nearly one-third of GAAP operating income. It is a noncash expense in the period of recognition, but existing shareholders pay through dilution unless repurchases offset issuance without consuming equivalent value.

At 2026-06-30, the company had 2,300.5 million Class A, 101.4 million Class B, and 1.0 million Class F shares outstanding, or about 2,402.9 million issued common shares across classes. The filed Q2 diluted weighted-average denominator was 2,568.7 million. Comparing a point-in-time share count with a weighted-average denominator is only a diagnostic, not a forecast, but the roughly 6.9% gap shows why a basic-share valuation would be inappropriate. Unrecognized RSU compensation was $916 million, and SARs remained outstanding.

The classes have the same economic rights but different voting and conversion rights. A financial common share maps one-for-one to a Class A economic share equivalent in the sealed quote bridge, while governance remains unequal. Future dilution depends on vesting, forfeitures, share price, award design, employee tax withholding, option exercises, SAR settlement, repurchases, and new grants.

Judgment: Downstream forecasts should model diluted shares explicitly in bear, base, and bull cases. Treating SBC only as an expense or only as dilution is incomplete: GAAP earnings already include the expense, while per-share value still requires

the correct share count.

10. What an investor should carry into the deeper ER2

Palantir is not simply a defense contractor, data warehouse, consulting firm, or AI model vendor. It is attempting to be the governed operating layer that connects data, models, decisions, and action across institutions. Government and Commercial share the underlying architecture, which allows product investment to serve both segments. The current financial profile shows exceptional direct contribution margins, strong GAAP operating profitability, high cash generation, negligible physical capital intensity, and a fortress balance sheet.

The unresolved question is durability per share. Growth and margin expansion must be tested against the costs hidden outside segment Contribution, customer acquisition intensity, convenience termination, public-budget correlation, powerful bundled competitors, and continuing SBC dilution. AIP increases the addressable set of workflows but does not by itself establish a moat; the moat depends on the Ontology, operational deployment, trust, security, and accumulated customer workflows remaining difficult to replicate.

The later ER2 should therefore focus on four falsifiable tests: whether commercial deployments can scale with less forward-deployed effort; whether Government growth survives procurement and budget cycles; whether contribution margin improvement persists after fully charging shared R&D, G&A, and SBC; and whether diluted per-share earnings compound fast enough to justify the security price at the exact 2031-06-30 horizon.

Glossary

Term	Meaning in this report
AIP	Artificial Intelligence Platform; Palantir's governed layer for connecting AI models to organizational data, tools, agents, evaluations, and workflows
Apollo	Continuous-delivery and software-operations platform for deploying and updating software across cloud, on-premises, classified, disconnected, and edge environments
Bootcamp	Short, hands-on deployment that uses customer data and Palantir engineers to build a real use case quickly; a pilot, not automatically recurring revenue
CODM	Chief operating decision maker; Palantir identifies its CEO in this role for segment reporting
Commercial	Reportable segment serving customers in non-government industries; a customer category, not a product
Contribution	Segment revenue less related cost of revenue and sales and marketing expense; excludes R&D, G&A, and all SBC
Contribution margin	Contribution divided by segment revenue; not GAAP operating margin
FCF	Free cash flow, calculated here as operating cash flow less capital expenditures
Forward-deployed engineer / FDE	Technical personnel working close to a customer to configure data, ontology, workflows, and applications for production use
Foundry	Palantir's foundational data operations, ontology, analytics, application, and workflow platform
Gotham	Government- and defense-oriented operational platform linking multi-source data, analysis, planning, and field action
Government	Reportable segment serving U.S. and non-U.S. government agencies; may use Gotham, Foundry, AIP, Apollo, and domain applications
LLM	Large language model; typically supplied by Palantir or a third party and governed through AIP
O&M	Operating and maintenance services that support hosted or customer-environment software over the contract term
Ontology	Governed semantic and operational representation of an institution's objects, relationships, logic, permissions, and possible actions
Palantir Cloud	Hosted subscription with continuous platform access and stand-ready O&M services
RPO	Remaining performance obligations; noncancelable contracted revenue not yet recognized, excluding cancelable amounts and certain short contracts
SAR	Stock appreciation right; an equity award whose settlement can create dilution depending on appreciation and award terms

Term	Meaning in this report
SBC	Stock-based compensation; noncash expense at recognition but an economic cost to shareholders through dilution or offsetting repurchase cash use
TTM	Trailing twelve months; here Q3 FY2025 + Q4 FY2025 + Q1 FY2026 + Q2 FY2026, ending 2026-06-30

Primary-source orientation links

- [Palantir 2025 Form 10-K, filed 2026-02-17](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm)
- [Palantir Q2 FY2026 Form 10-Q, filed 2026-08-04](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000041/pltr-20260630.htm)
- [Palantir platform architecture: AIP, Foundry, and Apollo](https://www.palantir.com/docs/foundry/architecture-center/platforms)
- [Palantir AIP documentation](https://www.palantir.com/docs/foundry/aip)
- [Palantir getting-started and bootcamp documentation](https://www.palantir.com/docs/foundry/getting-started/overview)

Appendix M - Management Assessment

Palantir Technologies Inc. (PLTR) — Management Assessment

Information cutoff: 2026-08-24 SGT **Assessment standard:** management capable of protecting and compounding *per-share* value for 20 years **Management gate: FAIL**

Palantir's operating leadership is unusually strong. The founder team made several non-consensus product and market choices that now show up in revenue growth, margins, cash generation, and customer expansion. The management gate nevertheless covers more than execution. It asks whether outside owners can entrust capital to the team through a full cycle with credible checks on control, compensation, dilution, succession, and related-party conduct. Palantir does not clear that broader standard. The Class F structure lets three founders control up to 49.999999% of votes while their economic ownership can fall, can survive their employment, and can concentrate control in as few as one remaining founder. Chronic equity issuance has also transferred a meaningful part of the operating achievement away from continuing Class A holders. These are structural minority-shareholder defects, not valuation complaints.

1. Who is running the company

Alexander C. Karp co-founded Palantir and has been CEO since 2005. He has therefore owned the entire modern strategy: prioritizing difficult government and commercial deployments, building Gotham/Foundry/Apollo/AIP as one operating stack, refusing certain markets on mission grounds, and commercializing AIP through hands-on bootcamps. He is not a hired manager inheriting someone else's plan. His tenure is a major strength for accountability and an equally important key-person risk.

The executive bench is internally developed and long-tenured. Co-founder Stephen Cohen has served as president, secretary, and a director since 2005. Shyam Sankar joined in 2006 and is CTO and executive vice president; he is the clearest product and technical deputy. Ryan Taylor joined in 2010 and is chief revenue officer and chief legal officer, combining commercial and legal responsibility. David Glazer joined in 2013 and is CFO and treasurer; the current proxy does not separately identify his initial CFO appointment date, so it is more precise to say he has been at Palantir since 2013 and is now CFO. Jeffrey Buckley returned as chief accounting officer in March 2025 after a brief interim period in which Glazer assumed principal-accounting-officer responsibilities. This history shows succession depth below the CEO, although no named CEO successor or formal external-facing succession plan is disclosed.

The seven-person board contains four Nasdaq-independent directors: Alexander Moore, Alexandra Schiff, Lauren Friedman Stat, and Eric Woersching. The other three are the founders Karp, Cohen, and Peter Thiel. Thiel is non-independent chairman. The proxy says an independent director *may* serve as lead independent director when the chair is non-independent, but it does not identify one. Thus the formal counterweight to a founder CEO is another founder, while the nominal independent majority is only 4-3. The board has two standing committees: Audit, and a combined Compensation/Nominating/Governance committee. This meets listing rules but is a thin accountability architecture for a company with exceptional founder voting rights. [2026 proxy, board and leadership](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000019/pltr-20260423.htm)

2. Tenure, strategic decisions, and track record

Karp has been CEO long enough to receive full credit and blame. Three decisions dominate the record.

First, management persisted with a product-led model for high-complexity operational software instead of turning Palantir into a conventional labor-heavy government contractor or broad horizontal data warehouse. From FY2021 to FY2025, revenue rose from $1.542 billion to $4.475 billion, while GAAP operating margin moved from negative 26.7% to positive 31.6%. On the DR-cleared TTM through June 2026, revenue reached $6.156 billion and operating margin 42.8%. The FY2025 ROIC diagnostic reached 15.0% after losses and low single-digit returns in earlier years. This is powerful evidence that the architecture can scale; it is not yet evidence of persistently greater-than-20% ROIC through a cycle.

Second, Karp, Cohen, and Sankar pushed AIP and bootcamp-led adoption beginning in 2023 rather than protecting the older Foundry/Gotham sales motion. Commercial revenue increased from $1.003 billion in FY2023 to $2.073 billion in FY2025 and $2.945 billion TTM. Commercial Contribution margin increased from 51.9% to 65.9% and then 72.5% TTM. The result supports management's claim that fast, production-oriented deployments can shorten time to value. It does not prove low customer-acquisition cost because Palantir still withholds ARR, billings, churn, CAC, payback, platform-level revenue, and

cohort conversion.

Third, management maintained a deliberate U.S. government and Western-alliance orientation while reusing those capabilities in commercial software. Government revenue rose from $1.222 billion in FY2023 to $2.402 billion in FY2025 and $3.211 billion TTM, with TTM Contribution margin of 70.2%. Management's willingness to keep investing in politically and operationally difficult deployments produced product credibility and a large growth engine. The associated concentration, procurement, and political risks remain real, but this was a coherent strategic bet with measurable outcomes.

The principal failed or mixed initiative was the 2021-2022 program of investments tied to strategic commercial contracts, including SPAC and private-company securities. Palantir stopped making new purchases under those agreements in 2024-2025, while related revenue declined from $87.3 million in 2023 to $52.3 million in 2024 and $15.3 million in 2025. The company now says it does not anticipate new SPAC investment agreements. That is a quiet but sensible retreat from a structure that complicated customer quality and capital allocation. Management disclosed the ongoing marks and reduced exposure, but it did not prominently frame the program on earnings calls as a capital-allocation error. [FY2025 10-K, Strategic Commercial Contracts](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm)

3. Capital allocation: strong reinvestment, weak per-share protection

The five-year allocation record is internally consistent with an asset-light software company.

Fiscal year	Capex ($B)	R&D ($B)	Buybacks ($B)	Dividends ($B)	Cash M&A ($B)
2021	0.013	0.387	0.000	0.000	0.000
2022	0.040	0.360	0.000	0.000	0.067
2023	0.015	0.405	0.000	0.000	0.000
2024	0.013	0.508	0.064	0.000	0.000
2025	0.034	0.558	0.075	0.000	0.000

R&D dollars grew 44% from FY2021 to FY2025 while R&D/revenue fell from 25.1% to 12.5%, evidence of operating leverage rather than underinvestment: revenue nearly tripled, product breadth expanded, and management continued hiring technical talent. Capex remained below 1% of revenue in four of five years, appropriate because human capital, cloud capacity, and expensed software development are the real reinvestment channels. M&A was negligible, and the absence of goodwill and acquired intangibles confirms that the results were overwhelmingly built rather than bought. The balance sheet reached $9.409 billion of cash and marketable securities at June 2026 with no financial debt, giving management genuine downturn capacity.

The weak point is per-share discipline. Buybacks totaled only $139 million in FY2024-2025, versus $1.376 billion of SBC expense in those two years. The $1.0 billion authorization retired just 0.6 million shares for $75 million in 2025 before termination in January 2026; that is neither opportunistic nor sufficient to offset issuance. Diluted weighted-average shares increased from 1,923.6 million in FY2021 to 2,565.2 million in FY2025, a 33% increase. TTM shares stabilized at 2,571.0 million, which is progress, but the company still carried substantial unrecognized option, SAR, and RSU compensation. Glazer has managed liquidity and guidance well, yet the buyback record shows no demonstrated price-sensitive capital return policy and no willingness to neutralize the accumulated dilution.

4. Compensation alignment and actual ownership

Founder ownership is real. At the April 2026 proxy record, Karp directly held 6.432 million Class A and 51.433 million Class B shares, excluding 34.5 million exercisable options, 0.975 million near-vesting RSUs, and his economic interest in 0.335 million Class F shares. Glazer held 353,553 Class A shares of record, plus exercisable options and near-vesting RSUs. These are material economic stakes, not token ownership.

The design of incremental executive pay is less satisfactory. There were no 2025 bonuses and no new 2025 equity grants to Karp, Cohen, or Sankar because legacy awards remained the retention mechanism. Glazer and Taylor each received 232,898 service-based RSUs, vesting mostly over five quarters with a 12-month post-vest sale restriction, and SARs tied to a $150 stock price with a limited 2033 exercise window and a capped benefit. Long duration and a high strike are positives. However, the awards use service and absolute stock price, not ROIC, free cash flow per diluted share, normalized operating profit, or relative multi-year returns. A rising market multiple can reward the SARs even if capital efficiency deteriorates. Palantir has no disclosed executive stock-ownership guideline, and the structure does not explicitly require retention of a multiple of salary.

SBC has improved but remains chronic: $778 million and 50.5% of revenue in FY2021; $565 million/29.6% in FY2022; $476 million/21.4% in FY2023; $692 million/24.1% in FY2024; $684 million/15.3% in FY2025; and $836 million/13.6% TTM. The percentage decline is meaningful. The absolute cost and historical dilution still violate a high bar for minority-owner protection. The proxy includes a compliant no-fault accounting-restatement clawback and prohibits hedging. Pledging is prohibited subject to limited exceptions for directors and the CEO; Karp had 3.0 million Class A shares pledged against personal debt. [2026 proxy, compensation and ownership](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000019/pltr-20260423.htm)

5. Disclosure quality, guidance credibility, and truthfulness

I reviewed eight live FMP transcripts, Q3 2024 through Q2 2026. Financial guidance is highly credible in the narrow forecasting sense. Each of the seven completed next-quarter revenue guides was exceeded: Q4 2024 actual $827.5 million versus $771 million high end; Q1 2025 $883.9 million versus $862 million; Q2 2025 $1.004 billion versus $938 million; Q3 2025 $1.181 billion versus $1.087 billion; Q4 2025 $1.407 billion versus $1.331 billion; Q1 2026 $1.633 billion versus $1.536 billion; and Q2 2026 $1.935 billion versus $1.801 billion. Management repeatedly raised full-year revenue and adjusted-profit guidance in prepared increments. This is conservative forecasting, not repeated scrambling or missed promises.

Candor is mixed. Glazer generally states the guide, segment result, and causal bridge precisely. When international commercial growth slowed, he quantified “continued headwinds in Europe” in Q3 2024 and again in Q1 2025. Those were the only two of eight calls where operationally hedged language appeared; both included concrete geography and timing explanations. Q2 2026 also explained the Government Contribution-margin dip as a customer transition to company-hosted cloud infrastructure, with higher current cost of revenue in exchange for faster deployment and future workflow expansion.

The weakness is voluntary disclosure and question selection. Palantir gives TCV, RDV, RPO, NRR, large-deal counts, U.S. commercial customer counts, and top-20 customer revenue. It does not disclose platform-level revenue, ARR, cRPO, billings, churn, CAC, payback, aggregate ACV, or a five-year NRR series; it also omitted total customer count in Q2 2026. RDV includes options and assumes no termination, so promotional use of RDV without an equally prominent cancelability caveat can overstate visibility. Q&A sessions are short and dominated by supportive strategic questions, often inviting expansive product or political narratives rather than testing SBC, dilution, international weakness, unit economics, or backlog quality. Karp frequently answers with conviction and strategic detail, but load-bearing financial questions are often left to Glazer and the company does not invite enough adversarial scrutiny to earn a top candor grade.

The conclusion is not that reported numbers are untruthful: there was no restatement, FY2025 disclosure controls and internal control were effective, and no material weakness was reported. It is that the communication format is optimized for narrative control. Forecast credibility is excellent; truthfulness of filed financial reporting is supported; completeness for outside underwriting is only average.

6. Strategic coherence

Management says Palantir exists to build software for high-stakes operational decisions while supporting “Western liberal democracy and its strategic allies.” The FY2025 10-K explicitly says culture prioritizes customer outcomes over short-term financial results and that Palantir may reject customers inconsistent with that mission. The Q2 2026 framing extended AIP into “sovereign AI”: customers retain control of data, logic, models, and operational workflows. The investment pattern supports the message—rising R&D dollars, very little acquisition spending, direct technical deployment, and a large liquidity reserve rather than a roll-up strategy. [FY2025 strategy and mission](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm)

This is a coherent 5-10 year destination, not a sequence of unrelated pivots. Foundry, Gotham, Apollo, and AIP share architecture; AIP is an extension of the governed operational layer, not a sudden abandonment of the old platform. Management names competitive categories—cloud/data platforms, enterprise software vendors, internal builds, systems integrators—and criticizes generic model access, but rarely names specific vendors on calls. The strategy is clearer than the disclosure needed to prove its unit economics.

7. Governance, insider sales, related parties, and succession

The governance defect is explicit. Class A has one vote per share, Class B ten, and 1.005 million Class F shares sit in a founder voting trust. Subject to a low 100 million-equity-security threshold, Class F plus founder-agreement shares can deliver 49.999999% voting power. The founders may sell substantial economic interests without equivalent loss of control; withdrawal, death, or disability can reduce the threshold and concentrate the vote among the remaining founders, potentially

one person. The arrangement can last until the death of the last living founder. The 10-K itself calls several provisions novel and warns that voting power may be far greater than economic interest. Limited protections exist: specified large founder transactions can require an independent committee, and a Rule 13e-3 founder take-private requires independent and majority-of-minority approval. These protections do not offset the absence of a conventional sunset, cumulative voting, or a named lead independent director. [FY2025 10-K, multi-class risks](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm)

Related-party transactions were disclosed and audit-committee reviewed but are not immaterial in optics. In 2025, Palantir received $17.66 million from Thiel-linked Anduril and invested $40 million in its preferred shares; it received $168,000 from another Thiel-linked company. Palantir incurred about $17.2 million using Karp's aircraft, under independently benchmarked rates, and sold Karp personal property for $358,275 using appraisals. Formal process reduces abuse risk, but a $17.2 million founder-aircraft arrangement is still a meaningful governance flag.

Live SEC Form 4 review for the 12 months ended August 24, 2026 found roughly 5.57 million shares sold across reported insiders for about $810 million in transaction value. The largest totals were Thiel's 2.0 million shares (about $290 million), Karp's 1.30 million (about $186 million), Cohen's 0.99 million (about $144 million), and Sankar's 0.72 million (about $102 million). Many executive sales were automatic tax-withholding transactions or pre-existing Rule 10b5-1 sales; Moore's monthly sales were also plan-based. This is not evidence of undisclosed adverse information, and Karp retained a vastly larger direct stake. It nevertheless reinforces that insiders are consistently monetizing while outside holders remain exposed to dilution and inferior votes. No cluster of open-market insider purchases offset that signal.

Succession is operationally deep but constitutionally weak. Cohen, Sankar, Taylor, and Glazer provide long-tenured continuity; the company is not a one-person operating bench. Yet a Karp departure could separate corporate leadership from voting control, because founder voting rights persist without employment and may concentrate. The board has not shown how it would choose and empower a successor against that structure.

8. Peer comparison

CEO / company	Capital allocation and execution	Disclosure and governance	Five-year stewardship versus Karp
Olivier Pomel / Datadog	Founder-led, organic product expansion and restrained M&A; strong but less explosive recent margins	More conventional investor communication; still founder influence and substantial SBC	Karp leads recent operating acceleration; Pomel offers a cleaner stewardship balance
Sridhar Ramaswamy / Snowflake	Shorter CEO record, active buybacks and product investment; transition risk remains	Publishes product revenue, RPO, customer cohorts and NRR; has a named lead independent director	Karp has the stronger in-seat record; Ramaswamy/Snowflake are materially easier for minorities to monitor
Bill McDermott / ServiceNow	Proven enterprise-sales scaler; greater M&A complexity and integration risk	Conventional board accountability and a broader recurring-revenue KPI set	Karp's per-dollar organic build is superior; McDermott's accountability architecture is superior

Karp ranks first on recent product conviction and organic operating execution, and last on minority-shareholder protection. Against realistic enterprise-software alternatives, that combination is not a passing management profile for a 20-year outside owner. [Datadog 2026 proxy](https://investors.datadoghq.com/static-files/737a728e-396b-4fd7-b932-c4bbf3106360), [Snowflake governance](https://investors.snowflake.com/governance/committee-composition/default.aspx), [ServiceNow leadership](https://www.servicenow.com/au/company/leadership/bill-mcdermott.html)

9. Management gate verdict

FAIL. The operating team is founder-led, deep, strategically coherent, and exceptionally credible on near-term financial guidance. Those strengths do not cure a Class F structure that separates voting power from economic ownership, can persist beyond employment, and can concentrate control in as few as one founder. Chronic SBC and a 33% increase in diluted shares from FY2021 to FY2025 show that management has not yet protected per-share value with the same discipline it has applied to product and margins; service/stock-price awards also omit ROIC and FCF-per-share hurdles. Related-party aircraft costs and the absence of a named lead independent director add weight, while the very short, supportive earnings-call Q&A format limits outside challenge. The specific signal that makes this a **FAIL rather than YELLOW** is the combination of perpetual/novel founder control and demonstrated dilution: both are present today, neither has a committed sunset or remediation plan, and either independently meets the framework's hard-fail criteria for minority-shareholder risk.

Appendix W - Moat and 2046 Durability

Palantir Technologies Inc. — Moat and 2046 Durability

Research date: 2026-08-24 SGT **Financial anchor:** TTM ended 2026-06-30 **Assessment:** Production ER2, based on the DR-cleared data package and independently reviewed Government and Commercial appendices **Authority boundary:** This appendix assesses business durability only. It authorizes no trade, ranking, sizing, portfolio action, outbound message, BUY queue, or ER2-process bypass.

Palantir has a real moat, but the evidence does not yet earn a 20-year PASS. The moat is not proprietary access to foundation models and it is not a conventional data network effect. It is the installed combination of customer-specific Ontology, governed data and permissions, operational applications, deployment across difficult environments, and accumulated institutional trust. That combination is strongest in Government. Commercial shows increasingly persuasive switching-cost and expansion evidence, but its current economics were established during an extraordinary AI adoption surge and have not been tested through a mature competitive cycle. Pricing power is also indirect rather than contractually visible. The appropriate calibration is therefore **YELLOW**, including for the separate 2046 test.

1. Moat taxonomy applied

Network effects — weak. Palantir does not have a classic cross-customer network in which every new customer directly raises the value of the platform for all existing customers. Customer data, ontologies, permissions, and mission logic are mostly private and often cannot be shared. There may be learning effects: deployment patterns, ontology templates, model-evaluation tools, and product features developed for one environment can improve later deployments. That is a scale or learning advantage, not a defensible network effect. It is strongest where reusable Government mission patterns or commercial industry modules reduce future implementation time, but no disclosed metric proves a self-reinforcing cross-customer network.

Switching costs — strong after production embedding, weak before it. A mature deployment represents an institution in an Ontology, connects data sources, implements security rules, embeds actions in workflows, trains users, and integrates with legacy systems. Replacement requires more than exporting tables: the customer must recreate semantics, permissions, applications, approval logic, validation, and operating procedures while avoiding mission or business disruption. These costs can be high in both segments and especially high in classified or safety-critical Government workflows. The caveat is phase dependence. A bootcamp, pilot, generic dashboard, or isolated agent is replaceable; the moat forms only when the deployment becomes a production operating layer. Palantir's [2025 Form 10-K](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm) also acknowledges that many contracts can be terminated for convenience, so technical embedding does not create legally guaranteed revenue.

Scale economies — moderate and strengthening. Palantir can amortize core R&D, security infrastructure, Apollo deployment tooling, and accumulated field knowledge across Government and Commercial. The clearest evidence is financial: sales and marketing fell from 39.9% of revenue in FY2021 to 20.1% TTM while GAAP gross margin rose from 78.0% to 84.8%. Both segment Contribution margins exceeded 70% TTM. Yet scale is not costless. Pilots and bootcamps are generally conducted at Palantir's expense without guaranteed conversion, and deployments still consume scarce technical talent, third-party cloud capacity, accreditation work, and customer-specific integration. Scale becomes a moat only if reusable software continues to replace labor per incremental workflow.

Brand — strong in high-consequence analytics, mixed elsewhere. Palantir's reputation in defense, intelligence, and difficult operational deployments reduces perceived execution risk and helps senior decision-makers justify adoption. Government references and large commercial outcomes are valuable distribution assets. The same brand can be polarizing in surveillance, defense, labor, privacy, and automated-decision contexts, creating procurement or employee resistance. Commercially, hyperscalers and incumbent enterprise vendors often have broader relationships and more familiar procurement brands. Brand helps Palantir reach the shortlist; it is not sufficient to prevent competition.

Regulatory advantage — conditional, not a protected license. Security clearances, classified deployment capability, facility approvals, export controls, privacy controls, and regulated-industry governance raise the entry barrier. Palantir has accumulated capabilities to operate across cloud, on-premise, edge, classified, and disconnected environments. These requirements favor proven vendors, but they do not grant an exclusive franchise. AWS, Microsoft, defense primes, systems integrators, and specialized vendors can achieve certifications. Regulation can also turn against Palantir through AI restrictions, data-rights rules, portability mandates, procurement reform, or political objections. In Government, the

commercial-item preference under FASA has helped the company, while Palantir itself warns that a statutory or judicial change could impair its position.

IP and patents — supportive, not the core moat. Palantir has patents, proprietary software, trade secrets, and contractual protections, but software functionality can be independently developed or assembled from cloud-native and open-source components. The 10-K says issued or pending patents may not be broad enough to prevent competition and that third parties may legitimately create similar technology. The defensible asset is therefore the integrated product plus accumulated deployment knowledge and customer embedding, not a patent wall around ontology, AI orchestration, or data integration.

High-trust procurement — strong in Government. Years of successful performance, security approvals, classified experience, procurement vehicles, and relationships with mission owners reduce the perceived risk of awarding additional work. An installed vendor knows the customer's architecture and mission context; a challenger must prove capability while navigating accreditation and bid processes. The July 2025 U.S. Army enterprise agreement can reduce ordering friction across many users. But its $10 billion ceiling is not an obligation, and volume discounts reveal concentrated buyer power. GAO procurement evidence also shows incumbency is an advantage rather than exclusivity: multiple vendors can bid for analytics requirements. This is the most defensible Government moat component, but it remains contestable at each new program or recompete.

2. Per-segment moat

Government — strongest segment; five-year strong, ten-year credible, twenty-year conditional. The protectable profit pool is mission software embedded across classified data, operational workflows, security boundaries, and procurement structures. A competitor must reproduce the working ontology, permissions, integrations, edge and classified deployment, user trust, and mission performance—not merely offer a better model. Palantir's reviewed Government appendix shows TTM revenue of $3.211 billion, 52.2% of company revenue, with 70.2% Contribution margin. Incumbency, procurement vehicles, and expanding workflows should protect the next five years. Ten-year durability is credible if Palantir remains the preferred integrator of models into operational systems. Twenty-year durability is conditional because open-architecture mandates, government-owned interfaces, multi-vendor acquisition vehicles, or hyperscaler maturity can deliberately reduce lock-in. The earliest break signal would be major recompete losses or renewal terms that require modular substitution while Government Contribution margin falls below the mid-60s despite continued revenue growth.

Commercial — real but less proven; five-year credible, ten-year uncertain, twenty-year unresolved. The moat is the customer-specific Ontology and the applications, agents, permissions, and operating routines built around it. Rising NRR, rapid U.S. customer additions, and expanding top relationships show that successful deployments can compound. TTM Commercial revenue was $2.945 billion and Contribution margin 72.5%. The attack is clearer than in Government: Microsoft, ServiceNow, Snowflake, Databricks, hyperscalers, integrators, and customer engineering teams can bundle or assemble data, model, workflow, and agent capabilities. A generic prompt-to-output workflow is vulnerable; a governed system that can safely execute actions across production operations is harder to replace. The earliest break signal is NRR normalizing below roughly 110%-115% while customer growth slows and Contribution margin retreats toward the low-60s, indicating bootcamps created pilots but not durable operating systems.

The blended moat is weaker than Government alone. The revenue mix is approximately 52% Government and 48% Commercial, so almost half the business sits in the segment more exposed to platform bundling and AI-native replication. However, Commercial is not unmoated: its TTM Contribution margin exceeds Government's, and current expansion metrics are stronger. The blend deserves YELLOW because the strongest segment has credible twenty-year mechanics while the nearly coequal Commercial segment has not yet demonstrated them through a normalized cycle.

3. Margin persistence evidence

Only five completed annual years plus the TTM are available in the cleared package, not a full ten-year comparable series. That limitation matters: the margin record captures Palantir's transition from heavy investment and losses to hypergrowth profitability, not a decade of stable mature economics.

Margin metric	FY2021	FY2022	FY2023	FY2024	FY2025	TTM 2026-06-30
GAAP gross margin	78.0%	78.6%	80.6%	80.2%	82.4%	84.8%
GAAP operating margin	(26.7%)	(8.5%)	5.4%	10.8%	31.6%	42.8%
Government Contribution margin	60.4%	57.9%	59.3%	60.4%	65.6%	70.2%

Margin metric	FY2021	FY2022	FY2023	FY2024	FY2025	TTM 2026-06-30
Commercial Contribution margin	55.5%	49.7%	51.9%	59.5%	65.9%	72.5%

The pattern is strong expansion, not mere stability. Consolidated operating margin improved by 69.5 percentage points from FY2021 to TTM; gross margin gained 6.8 points. Government Contribution held near 58%-60% through FY2021-FY2024 before accelerating, supporting an argument that established workflows carried durable direct economics even before the current surge. Commercial compressed to 49.7% in FY2022 and then expanded by 22.8 points to the TTM as acquisition and delivery costs scaled. Q2 FY2026 Commercial Contribution margin reached 77.7%; Government was 70.8% after a hosting change reduced it from 73.3% in Q1.

These results are genuine evidence that Palantir can scale revenue faster than directly attributable delivery and selling expense despite visible competition. They do not prove steady-state margins. Contribution excludes R&D, G&A, and all SBC, while the operating-margin history contains only three profitable annual years. The Government hosting change demonstrates that customer cost certainty or compute mix can move economics quickly. The key durability question is whether margins remain high after growth normalizes, customers gain bargaining power, and bundled alternatives become adequate.

4. AI disruption stress test

Government is the most protected segment. Its value is not the foundation model. A model can summarize intelligence or generate plans, but it cannot by itself obtain authorization, reconcile data rights, enforce classification and need-to-know rules, maintain lineage, connect legacy and edge systems, survive disconnected operation, or become an accepted mission workflow. Palantir is platform-adjacent: better models increase the value of AIP when Palantir remains the governed context and action layer. Government is not a physical or licensed chokepoint, however. Hyperscalers now deliver models and infrastructure at high classification levels, and agencies can combine those services with integrators and modular applications. AI enlarges the market while lowering the cost of building rival components.

Commercial is most exposed. Simple analytics, document agents, data preparation, dashboards, code generation, and generic workflow automation are increasingly replicable by foundation models plus an incumbent cloud or application suite. A customer with clean data and standardized processes may not need Palantir's full stack. The moat survives only where the Ontology contains valuable operational semantics, permissions, actions, evaluation, and accumulated applications. In those cases a model is interchangeable but the governed operating context is not. If customers can reproduce that context cheaply using Microsoft, ServiceNow, Databricks, Snowflake, or open-source orchestration, AIP becomes a feature rather than a platform.

Foundation-model commoditization is more opportunity than threat only if Palantir remains model-agnostic. Palantir does not need to own the leading model; it must let customers change models without rebuilding their data, controls, evaluations, and workflows. This lowers dependence on any one model provider. The opposite outcome is that hyperscalers vertically integrate model, data, security, and agents and bundle the orchestration layer at low or negative incremental margin. Palantir's own 10-K identifies broad product bundles and zero- or negative-margin pricing by larger competitors as a material threat.

Bootcamps improve acquisition speed but are not themselves a moat. Palantir conducts pilots and bootcamps largely at its own expense and without guaranteed returns. The motion can compress demonstration time and create a working use case with real customer data, but production conversion still faces a sales cycle that often lasts six to nine months and sometimes longer. Customer-level conversion, payback, churn, and cohort gross retention are not disclosed. Bootcamps become moaty only if they repeatedly seed production ontologies, lower sales expense per converted dollar, and create expansion. Current sales-and-marketing leverage and 157% NRR are encouraging indirect evidence, not a disclosed conversion proof.

5. The 2046 test

Today's date is **2026-08-24**; the twenty-year test date is **2046-08-24**.

Plausible 2046 business description. Palantir provides governed operational software that connects sovereign and enterprise data, autonomous systems, third-party models, human approvals, and physical operations. Government customers buy a successor to Gotham/AIP/Apollo for defense, intelligence, logistics, health, space, cyber, and civilian missions across classified, edge, and allied environments. Commercial customers buy an evolved Foundry/AIP operating layer for factories, energy networks, supply chains, healthcare systems, financial risk, and other regulated or high-cost workflows. Model access

has commoditized and is no longer separately valuable; Palantir earns high returns by owning the durable ontology, policy, evaluation, deployment, and action layer. Government remains a major segment and may be broader across allied sovereigns. Commercial is much larger in dollars but concentrated in complex operations; generic analytics and simple agents have withered or migrated to bundled incumbents. Directionally, GAAP operating margin is stable to modestly lower than the current 42.8% TTM—roughly a 30%-40% mature range—because growth and pricing normalize while software reuse offsets continuing R&D, compliance, and deployment expense. Customers still pay because replacing the operational representation and revalidating thousands of actions is riskier and more expensive than switching models.

This description is specific enough to be credible, but not yet proven enough to PASS. The moat that must survive is institutional operating context plus trusted deployment—not the current product names, foundation models, or AI enthusiasm.

Moat-erosion 2046 bear case. Government mandates open data models, portable ontologies, government-owned interfaces, and modular acquisition. Hyperscalers operate approved AI/data stacks across every classification level, while primes and specialist vendors compete for applications through standardized APIs. Commercial customers consolidate data, productivity, security, workflow, and agents with existing strategic vendors; AI coding reduces the cost of internal builds. Palantir remains a capable integration and applications vendor, but its ontology is no longer the control plane. Contracts become narrower and more services-intensive, NRR approaches ordinary enterprise software, Contribution margins fall toward the high-50s or low-60s, and GAAP operating margin settles below 25%. This is not the most likely outcome today, but it is material—roughly a one-in-three risk—because the attack aligns powerful suppliers, customers seeking portability, and incumbents with superior distribution.

2046 test: YELLOW. A protectable future profit pool can be named, but its survival requires Palantir to preserve the ontology and action layer while models, infrastructure, and agent-building tools commoditize.

6. Competitive threat inventory

Threat	Mechanism	Segment(s) affected	Earliest observable signal	Plausibility
Hyperscaler and incumbent bundling	Microsoft, AWS, ServiceNow, Snowflake, Databricks, and others bundle data, models, governance, agents, and workflows into installed enterprise or classified-cloud relationships	Commercial primarily; Government new workloads	Commercial NRR and customer additions slow together; price concessions rise; Contribution margin declines despite revenue growth	High
Open architecture and procurement reform	Government-owned interfaces, portable data/ontology standards, modular awards, and multi-vendor vehicles reduce switching costs without replacing every capability at once	Government	Major recompetes require portability; Palantir loses application scope while retaining only integration work; Government margin falls below mid-60s	Medium-high
Deployment economics fail to productize	Bootcamps create demonstrations but conversion needs persistent forward-deployed labor, cloud subsidies, and custom integration; customer bargaining power captures scale benefit	Commercial and Government	Sales and marketing stops declining as a share of revenue; hosting drag recurs; Contribution margins retreat toward low-60s	Medium

7. Pricing power verification

Pricing power is **weakly to moderately observable, not directly verified**. The cleared record contains no disclosed companywide price increase, adoption response to a price increase, churn series, gross retention, standard tariff, minimum commitment schedule, or segment price-versus-volume decomposition. Palantir negotiates heterogeneous institutional contracts, so a list-price framework may not be meaningful. This absence prevents a PASS-level claim that customers accept repeated price increases without churn.

Indirect evidence is favorable. Total NRR rose from 134% in Q3 FY2025 to 157% in Q2 FY2026, U.S. Commercial customers increased from 530 to 653, top-20 average TTM revenue per customer reached $124 million, and segment Contribution margins rose while revenue accelerated. Customers are clearly spending more on successful deployments, consistent with workflow expansion and value-based pricing. Yet NRR blends volume, new workflows, mix, and price; it does not isolate price.

Contract structure cuts both ways. Agreements are generally one to five years and RPO reached $4.9 billion, providing some visibility. RDV was $13.1 billion but assumes options are exercised and contracts are not terminated. Many Government and Commercial contracts permit termination for convenience. Government enterprise agreements can lower procurement friction but include volume discounts, and competitive multi-award structures can impose pricing pressure. Therefore Palantir appears able to expand wallet share where it proves operational value, but it has not demonstrated unilateral, recurring price power comparable to an indispensable regulated asset or standardized subscription with published escalators.

8. Moat verdict

Moat gate: YELLOW 2046 test: YELLOW

Government has a concrete moat in high-trust procurement, classified deployment, installed ontology, and mission workflow switching costs. Commercial has real customer-level switching costs and compelling current expansion economics, but nearly half the revenue mix faces serious bundling and AI-native replication risk. Five years of improving margins and recent NRR support durability, yet the record is not a ten-year mature-cycle test and pricing power remains indirect. The 2046 company can be described specifically, but the description depends on Palantir retaining the governed context and action layer after model and infrastructure commoditization; that unresolved control-point question holds both gates at YELLOW.

Observable calibration conditions:

- **Upgrade toward PASS:** through a normalized growth period, Commercial NRR remains above roughly 120%, customer count continues to broaden, and Commercial Contribution margin stays at or above 70% without rising deployment intensity.
- **Upgrade toward PASS:** Government renewals and new mission wins continue without recurring hosting drag or material price concessions, while open-architecture requirements expand Palantir's addressable workflows rather than displacing its control layer.
- **Downgrade toward FAIL:** NRR falls below roughly 110% as customer growth slows, either segment's Contribution margin moves sustainably into the low-60s, or major Government recompetes demonstrate that portable ontologies and modular stacks can replace Palantir with limited disruption.

Appendix B1 - Government

Government — Segment Deep Dive

Company: Palantir Technologies Inc. **Segment:** Government **TTM anchor:** 2026-06-30 **Forecast horizon:** 2031-06-30, exactly 5.0 years after the TTM anchor **Segment profit metric:** `Contribution`

STATUTORY_BRIDGE_STATUS: VERIFIED — CLEARED FOR FORECASTING

This appendix treats the DR-cleared `data_package.md` and sealed `statutory_bridge.json` as authoritative for Palantir historical dollars, growth rates, management quotations, and accounting classifications. The final data-package SHA-256 is `740b2e36a294903da1e182fe6cc426d17220e6a2a025a1bfd9e5e47e7593cab2`. Government is modeled on the issuer's exact additive CODM metric, **Contribution**, never as operating income. R&D, G&A, stock-based compensation, interest, other income, tax, and noncontrolling interests remain outside this segment model and are forecast once in the consolidated statutory bridge.

0. Segment boundary and taxonomy note

Government has been a stable customer-type operating segment since at least FY2020. It serves U.S. and allied government customers across defense, intelligence, health, and civilian agencies. The three-year taxonomy investigation found no reportable-segment realignment, restating 8-K, or change that would invalidate comparisons. FY2024 added disclosure of segment revenue and directly attributable expenses alongside Contribution, but that was a disclosure enhancement rather than an economic reclassification. In Q1 FY2026, one customer program moved from U.S. Commercial to U.S. Government; management quantified the effect on U.S. Commercial growth, but the transfer does not change the Government segment's underlying taxonomy. Historical comparisons use the published Government series and explicitly note the assignment effect where relevant. [data_package §1, paragraphs 1–3]

1. What this segment actually is

Palantir Government sells an operational software layer to public-sector organizations whose data, security, deployment, and mission requirements are difficult to satisfy with an ordinary cloud application. The economic buyer may be a defense program office, intelligence organization, combatant command, military service, health agency, civilian department, or an allied sovereign institution. End users range from analysts and logisticians to acquisition teams, commanders, investigators, clinicians, and planners. The common need is to combine fragmented data, impose fine-grained permissions, express real-world objects and relationships in an ontology, and turn governed data into decisions and workflows.

The product surface includes Gotham, Foundry, AIP, Apollo, and mission-specific defense capabilities. Gotham is most closely associated with intelligence, defense, targeting, and operational decision support. Foundry provides data integration, ontology, analytics, and operational applications. AIP connects models and agents to governed data and workflows. Apollo manages continuous delivery across public cloud, government cloud, on-premise, edge, and disconnected environments. A government deployment can use several capabilities together. Palantir does **not** disclose revenue by platform, and this appendix does not invent or back-solve AIP, Gotham, Foundry, Apollo, or Maven dollars. [data_package §4, platform inventory rows 1–5]

The commercial model mixes software subscriptions, platform access, implementation, ongoing support, and mission-specific expansion. Contracts may span multiple years, include funded task orders or delivery orders, contain options that are not yet obligated, and permit termination for convenience. Remaining deal value, or RDV, assumes options are exercised and contracts are not terminated; RPO is narrower and is described as primarily commercial because many government arrangements have short initial terms or termination rights. Government revenue visibility is therefore better understood through a portfolio of funded work, execution in existing programs, new awards, option conversion, and workflow expansion than through RPO alone. [data_package §0a, KPI discussion and §5, quarterly KPI table]

Go-to-market is mission-led and relationship-intensive. Palantir must earn security approvals, integrate with legacy and classified systems, demonstrate performance in operational settings, and expand from an initial use case into adjacent workflows, commands, or agencies. Procurement vehicles matter nearly as much as product fit. The U.S. Army's July 2025 enterprise agreement consolidated 75 contracts into one framework, applies volume discounts, and lets Army and other Department of Defense users acquire Palantir products for up to ten years. Its $10 billion ceiling is a maximum, not an obligation or revenue backlog. The framework can reduce procurement friction and reseller fees, but pricing concessions and

usage conversion determine realized economics. [U.S. Army enterprise agreement announcement](https://www.army.mil/article-amp/287506/u_s_army_awards_enterprise_service_agreement_to_enhance_military_readiness_and_drive_operational_efficiency)

For forecasting, the most defensible disclosed driver split is **U.S. Government** versus **allied and international Government**. This is not a product split and does not imply different software stacks. It reflects materially different procurement cycles, defense budgets, sovereignty rules, data-sharing restrictions, and current growth rates. Within each geography, the model considers existing-program execution, new awards, option and task-order conversion, workflow expansion, and hosting cost. These are drivers, not fabricated product revenue lines.

2. Current economics — TTM ended 2026-06-30

Government generated $3.211 billion of TTM revenue, 52.2% of Palantir's $6.156 billion consolidated revenue. TTM Contribution was $2.255 billion, producing a 70.2% Contribution margin. Contribution is revenue less cost of revenue and sales and marketing expense, excluding stock-based compensation; it is not statutory operating income. Government revenue accelerated throughout the last four quarters: 54.9% in Q3 FY2025, 60.3% in Q4, 76.3% in Q1 FY2026, and 79.0% in Q2. Contribution margin rose from 66.0% in Q3 to 73.3% in Q1, then eased to 70.8% in Q2 because of a specific cloud-hosting change. [data_package §3, TTM segment table; §5, quarterly segment table]

Current-economics metric	TTM or latest value	Data package source	Interpretation
Government revenue	$3.211B	data_package §3, TTM Government row	52.2% of consolidated revenue
Contribution	$2.255B	data_package §3a, TTM bridge Government line	Exact sealed segment profit basis
Contribution margin	70.2%	data_package §3, TTM Government row	$2.255B divided by $3.211B
Latest quarterly revenue	$0.990B	data_package §5, Q2 FY2026 Government row	79.0% YoY and 15% sequential growth
Latest quarterly Contribution margin	70.8%	data_package §5, Q2 FY2026 Government row	Down from 73.3% in Q1 on hosting mix
U.S. Government TTM revenue	$2.552B	data_package §5, U.S. subset rows Q3–Q2	Sum of $486M, $570M, $687M, and $809M
Allied and international Government TTM revenue	$0.659B	Calculated from data_package §3 and §5	$3.211B segment less $2.552B U.S. subset
Company capex	$0.042B	data_package §2, TTM metric table	0.7% of consolidated revenue; segment capex n/d

The U.S. Government driver represented approximately 79.5% of TTM Government revenue on the disclosed quarterly subset series. Allied and international Government represented the residual 20.5%. This residual is not a reported operating segment, and no Contribution split is disclosed. It is used only to make the forecast reconstructable from disclosed geography drivers. The latest Q2 figures show U.S. Government revenue of $809 million, up 90%, while total Government revenue was $990 million, up 79%. The implied international residual was about $181 million versus about $127 million a year earlier, or roughly 42% growth using rounded disclosed inputs. [data_package §5, Q2 U.S. subset and total Government rows]

Platform mix and customer-program mix are not disclosed in dollars. That prevents a credible statement such as "AIP is X% of Government revenue." It also prevents product-level margin claims. What is observable is the direction of mix: management attributes growth to continued execution in existing programs and new awards, says Government demand is tied to its AI platform, and is adapting hosting in a way intended to accelerate time to value and expand future workflows. Those facts support a driver model but not a platform revenue allocation. [data_package §6, Government rows 1, 4, and 6]

Capital intensity is low on Palantir's balance sheet because it generally buys cloud and hosting services rather than owning data centers. Consolidated TTM capex was $42 million, about 0.7% of revenue. Segment capex is not disclosed. Economic reinvestment occurs through product development, forward-deployed engineering, security and deployment work, and cloud capacity rather than property and equipment. R&D, G&A, and SBC are outside Contribution and must not be silently assigned to Government. Thus, the 70.2% Contribution margin is a measure of post-delivery and post-selling economics, but it is not a full-cost operating margin.

3. Unit economics

Government revenue becomes Contribution through four linked mechanisms: software value and pricing; delivery and hosting cost; implementation intensity; and sales or capture cost. First is the recurring value of a deployed operational platform. Once Palantir is embedded in a mission workflow, additional users, data domains, models, applications, and commands can expand contract value without a proportional increase in deployment cost. This is the source of potential incremental Contribution margins above the current segment average.

Second is hosting and compute. Some deployments are customer-hosted or run in government-controlled clouds; others require Palantir to bear more hosting cost. Q2 FY2026 illustrates that cost architecture can change faster than pricing. Management said a Government cloud-hosting change raised cost of revenue in the quarter, while arguing it should later improve time to value, efficiency, customer cost certainty, and workflow expansion. The near-term margin drag is real even if the long-term rationale is sound. Extrapolating Q1's 73.3% margin would overstate current economics; treating Q2's 70.8% as permanent would ignore the intended efficiency benefit. [data_package §5, Government quarterly margins; §6, Government cloud-hosting quote]

Third is implementation intensity. Government work can require security accreditation, bespoke connectors, deployment to classified or disconnected networks, data-rights negotiations, model evaluation, and integration with systems not designed for modern AI. Palantir's software is configurable, but the last mile is not frictionless. The U.S. Space Force's 2025 description of its Palantir-backed Space C2 Data Platform emphasizes secure environments, disparate sources, multiple networks, multiple security levels, and cross-service use. Those are barriers, but they also consume skilled labor. [U.S. Space Systems Command, May 2025](https://www.ssc.spaceforce.mil/Newsroom/Article-Display/Article/4188469/space-systems-command-advances-space-c2-dominance-decision-making-through-data)

Fourth is sales and capture cost. Government selling cycles can be long and lumpy, but enterprise procurement vehicles can amortize capture work across many task orders. The Army agreement is a useful example: consolidation can remove reseller fees and reduce procurement time, yet volume discounts share savings with the customer. The net Contribution effect depends on whether incremental adoption and lower capture cost outweigh discounting. Base assumes they do, but not immediately or one-for-one.

Historical margins show scale already appearing. Government Contribution margin was 60.4% in FY2021, 57.9% in FY2022, 59.3% in FY2023, 60.4% in FY2024, 65.6% in FY2025, and 70.2% on the TTM anchor. Improvement since FY2023 reflects revenue growth outpacing directly attributable cost of revenue and sales and marketing. It should not be credited to R&D or G&A leverage because those costs are excluded from Contribution. [data_package §3, Government historical table and TTM row]

What breaks unit economics upward is a mix of larger software-led renewals, expansion across workflows, reusable deployment assets, lower procurement friction, and better hosting efficiency. What breaks it downward is fixed-price or capped work with higher-than-expected integration effort, cloud compute absorbed by Palantir, pricing concessions, delayed funding, or services-heavy deployments. Competition can force modular pricing or portability commitments that reduce lock-in and increase sales effort. Because product-level gross margin is not disclosed, the model expresses these mechanisms through scenario-level U.S. and international Contribution margins rather than fabricated AIP or Gotham margins.

The segment's working-capital economics are harder to isolate. Government contracts can create milestone billing, customer deposits, receivables concentration, and timing swings. The package notes no customer was 10% or more of revenue at Q2 FY2026, but Customer I was 27% of accounts receivable. That does not prove Government concentration, yet it shows billing exposure can be more concentrated than revenue. RPO also understates option-heavy Government arrangements, so cash conversion and backlog should not be inferred mechanically from segment revenue.

4. Competitive dynamics

The competitive field is wider than "AI software" companies. Palantir competes against hyperscale cloud and AI platforms, data-platform vendors, defense primes and systems integrators, specialized application vendors, and internal government development. AWS and Microsoft provide classified cloud, model access, data services, security, and agent tooling. Microsoft states Azure OpenAI is available across U.S. Government classification levels including Secret and Top Secret. AWS made Bedrock available in its Top Secret cloud and has FedRAMP High and DoD IL4 and IL5 approvals for selected models in GovCloud. These offerings lower the infrastructure and model-access barrier for agencies that want to assemble a solution without adopting Palantir's full operating layer. [Microsoft Azure for U.S. Government](https://azure.microsoft.com/en-us/explore/global-infrastructure/government/); [AWS Bedrock Top Secret](https://aws.amazon.com/blogs/publicsector/amazon-bedrock-launches-with-claude-3-5-sonnet-in-the-aws-top-secret-cloud/)

Data and application competitors include Databricks, C3.ai, cloud-native services, and mission-specific software firms. Systems integrators and defense primes such as Booz Allen, Leidos, SAIC, Accenture Federal, and traditional program contractors can assemble or operate alternatives. Anduril and other defense-technology entrants compete for the application and command layer in selected missions. Internal government platforms and open architectures are also substitutes. The relevant choice is often Palantir's integrated commercial platform versus a consortium of cloud infrastructure, data tools, model providers, integrators, and custom code.

No reliable audited market-share denominator exists for the precise category Palantir addresses, and this appendix does not fabricate one. Procurement evidence shows incumbency advantage and real contestability. In a Veterans Affairs procurement protest, GAO recorded that market research identified several vendors capable of competing for a commercial data-analytics SaaS requirement and that five quotations were submitted, even though Palantir's incumbent experience conferred a normal advantage. [GAO decision B-419801](https://www.gao.gov/products/b-419801)

Palantir's pricing power is strongest when the workflow is mission-critical, deployment spans complex data and classification boundaries, and switching requires rebuilding ontologies, permissions, applications, and processes. It is weaker in generic analytics, stand-alone model access, simple data warehousing, or projects where government enforces modular interfaces and owns the data layer. The Army agreement's volume discounts show that even a deeply embedded vendor shares economics when buyers aggregate demand.

Entry barriers include accreditation, classified deployment capability, trust earned in live missions, integration with legacy systems, an installed ontology, and a workforce able to translate mission problems into software. These barriers can take years to reproduce. They are not permanent. Hyperscalers already operate across classification levels, and policy emphasizes open standards, data rights, modular contracting, and continuous evaluation. GAO's 2025 command-and-control work warned that reliance on one software provider can create vendor lock-in and described OPEN DAGIR as an effort to onboard infrastructure and applications through standardized licensing and a new acquisition process. [GAO-25-106454](https://files.gao.gov/reports/GAO-25-106454/index.html)

Substitution risk is asymmetric. A new vendor is unlikely to replace a mature deployment overnight, but new workloads can be competed, architectures modularized, and agencies can require portability at renewal. GAO's 2026 AI acquisition review highlights concern over IP and data rights, the mismatch between procurement timelines and AI development, and the need for performance metrics and continuous evaluation. Those reforms can accelerate commercial AI adoption while reducing any single vendor's lock-in. [GAO-26-107859](https://files.gao.gov/reports/GAO-26-107859/index.html)

5. Moat at the segment level

Government has a real but conditional moat. Its strongest component is cumulative deployment knowledge. Palantir is not merely hosting models; it has spent years learning to integrate operational data, enforce permissions, deploy across network boundaries, and maintain software in high-consequence environments. Each successful deployment creates reusable product capability, customer references, cleared relationships, and lower perceived execution risk. The Space C2 award is illustrative: the customer described a platform working across disparate sources, security levels, networks, and functional communities. That breadth is harder to copy than a dashboard or model endpoint.

The second moat component is workflow and ontology embedding. When operational decisions, objects, permissions, and applications are represented in a shared ontology, switching entails more than migrating tables. The agency must recreate semantics, controls, integrations, testing, and training. Incumbency can improve retention and expansion. GAO's VA decision recognized prior performance can create a normal competitive advantage, while still finding other vendors capable of bidding. The moat is meaningful, not absolute.

The third component is procurement leverage. Vehicles such as the Army agreement can turn many bilateral purchases into a reusable route to market, lowering friction and making Palantir a default commercial option for incremental workflows. Yet the vehicle imposes volume discounts and a visible ceiling; it is not a guaranteed $10 billion stream. A moat based on convenience remains vulnerable if agencies develop equally convenient multi-vendor vehicles.

The fourth component is breadth across cloud, edge, and classified settings. Apollo's deployment layer, Gotham's mission orientation, Foundry's ontology, and AIP's model orchestration can be combined without forcing the customer to integrate a vendor at every layer. The architecture can deliver faster time to value. However, hyperscalers are closing the classified AI gap, and open architectures can separate infrastructure, data, models, and applications. Palantir must prove integration value exceeds the price and flexibility cost of a modular stack.

Over five years, the moat should support above-market growth and Contribution margins around or above 70% if Palantir keeps winning workflows. Over ten years, durability depends on avoiding three failure modes: becoming a high-priced incumbent procurement reform targets; allowing hyperscaler-native tools to become good enough; and overreaching on data or IP rights in ways that provoke mandated portability. GAO's lock-in concerns show customer dependence can produce political and procurement backlash. The moat can protect economics, but pricing power must be exercised with restraint.

The base case treats Government as a high-quality, expanding platform franchise rather than a monopoly. It assumes deceleration from extraordinary current growth as comparisons and timing normalize, gradual margin improvement from scale and hosting efficiency, and continued competition for each new mission. Bull assumes procurement vehicles and mission proof create a reinforcing flywheel. Bear assumes open architecture, budget and award delays, and pricing concessions erode expansion economics.

6. Historical growth rates

Palantir does not disclose Government revenue by AIP, Gotham, Foundry, Apollo, or defense application. The finest recurring disclosed breakdown available for a bottom-up bridge is total Government plus the U.S. Government subset. The allied and international line is a calculated residual, not an issuer-reported segment. Every historical cell below identifies the controlling package row or the explicit non-disclosure row.

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 rev	3-yr CAGR	Most recent YoY	Data package cite
Government total	$1.222B [data_package §3, Government Revenue row FY2023]	$1.570B [data_package §3, Government Revenue row FY2024]	$2.402B [data_package §3, Government Revenue row FY2025]	30.9% [FY2022–FY2025; calculated from data_package §3]	79.0% [data_package §5, Q2 FY2026 Government row]	data_package §3 Government table; §5 quarterly segment table
U.S. Government	n/d [data_package §4, sub-segment non-disclosure]	$1.197B estimated [data_package §6, Government Q4 FY2025 row]	$1.855B [data_package §6, Government Q4 FY2025 row]	n/d [data_package §4, no three-year subset series]	90% [data_package §5, Q2 FY2026 U.S. subset row]	data_package §6 U.S. Government quote; §5 U.S. subset table
Allied and international Government	n/d [data_package §4, no geography series]	$0.373B calculated [data_package §3 FY2024 total; §6 Q4 U.S. row]	$0.547B calculated [data_package §3 FY2025 total; §6 Q4 U.S. row]	n/d [data_package §4, no three-year subset series]	about 42% calculated [data_package §5, Q2 total and U.S. rows]	Total Government less disclosed U.S. Government subset

Internal arithmetic checks. Government total: the three-year FY2022–FY2025 CAGR is ($2.402B ÷ $1.072B)^(1/3) – 1 = 30.9%; $1.222B × 1.284 = $1.569B, within 0.1% of the published $1.570B; $1.570B × 1.531 = $2.403B, within 0.1% of the published $2.402B. U.S. Government back-solve: $1.855B ÷ 1.55 = $1.197B, and $1.197B × 1.55 = $1.855B ✓. Allied and international residual: FY2024 $1.570B – $1.197B = $0.373B; FY2025 $2.402B – $1.855B = $0.547B; $0.373B × 1.466 = $0.547B ✓. Latest-quarter residual: current $0.990B – $0.809B = $0.181B; prior total $0.990B ÷ 1.79 less prior U.S. $0.809B ÷ 1.90 equals approximately $0.127B; $0.127B × 1.422 ≈ $0.181B ✓. Rounded source figures account for small differences.

The TTM forecast anchor is reconstructed from the same geography series. U.S. Government TTM revenue is $2.552 billion, the sum of $486 million, $570 million, $687 million, and $809 million for Q3 FY2025 through Q2 FY2026. Allied and international Government is $0.659 billion, calculated as $3.211 billion total Government less $2.552 billion U.S. Government. The two lines tie exactly to the displayed segment total.

TTM sub-driver bridge	TTM revenue	Share of Government	Construction	Platform-dollar treatment
U.S. Government	$2.552B	79.5%	data_package §5 U.S. rows Q3 FY2025–Q2 FY2026	No AIP, Gotham, Foundry, Apollo, or Maven dollars assigned
Allied and international Government	$0.659B	20.5%	$3.211B total less $2.552B U.S. subset	Residual geography driver only
Government total	$3.211B	100.0%	data_package §3 TTM Government row	Exact reportable-segment anchor

[DATA GAP: Palantir does not disclose AIP, Gotham, Foundry, Apollo, Maven-related, defense-versus-civilian, existing-program-versus-new-award, or agency-level Government revenue. No platform or contract-driver dollars are invented.]

7. Management guidance and commentary

Palantir provides no standalone FY2026 Government revenue or Contribution-margin guide in the cleared source set. The quotations below reproduce every Government-specific row in data-package §6. They are operational anchors, not

permission to treat contract ceilings or qualitative demand as guaranteed revenue.

- **Q3 FY2025, David Glazer — data_package §6, Government row 1.** “This growth was driven by continued execution in existing programs and new awards reflecting the growing demand for AI and our government software offerings.” Interpretation: installed-program expansion and award conversion are both active engines. The model tests both without assigning undisclosed dollars.

- **Q4 FY2025, David Glazer — data_package §6, Government row 2.** “Full year U.S. government revenue grew 55% year-over-year to $1.855 billion.” Interpretation: this is the only annual U.S. Government dollar anchor and supports the geography bridge. It is historical, not guidance, and 55% is not extrapolated flat.

- **Q1 FY2026, David A. Glazer — data_package §6, Government row 3.** “First quarter Government revenue grew 76% year-over-year and 18% sequentially to $858 million.” Interpretation: acceleration was broad at the reportable-segment level. Sequential growth establishes a strong TTM start but raises the hurdle for comparisons.

- **Q1 FY2026, David A. Glazer — data_package §6, Government row 4.** “This growth was driven by continued execution in existing programs and new awards reflecting the growing demand for our AI platform in government.” Interpretation: repeated driver language increases confidence growth is not one isolated award. It does not disclose AIP revenue.

- **Q2 FY2026, David Glazer — data_package §6, Government row 5.** “Second quarter government revenue grew 79% year-over-year and 15% sequentially to $990 million.” Interpretation: the latest anchor shows further acceleration. Base nevertheless decelerates because 79% sustained for five years would create an implausibly large base relative to observable procurement channels.

- **Q2 FY2026, David Glazer — data_package §6, Government row 6.** “While this change led to higher cost of revenue in Q2, going forward, we believe it will power faster time to value, drive greater efficiency, provide greater cost certainty to the customer and enable us to expand their future workflows.” Interpretation: the Q2 margin decline has a specific hosting cause. Base and bull allow later efficiency and expansion; bear assumes cost certainty is achieved partly at Palantir’s expense.

- **Q2 FY2026, Alexander Karp — data_package §6, Government row 7.** “And the story really does begin at the beginning when we dedicated ourselves to our most important partners in the U.S. government, and we built products to deliver value for them.” Interpretation: Government remains a strategic product-development source, but this is qualitative and does not justify a number by itself.

- **Q2 FY2026, David Glazer — data_package §6, Government row 8.** “We are raising our adjusted free cash flow guidance to between $4.5 billion and $4.7 billion, and we continue to expect GAAP operating income and net income in each quarter of this year.” Interpretation: this is a consolidated profitability anchor, not segment guidance. It does not change the Contribution basis or allocate corporate costs here.

Relevant company-wide context is the current FY2026 revenue guide of $8.150–$8.158 billion and Q3 FY2026 revenue guide of $2.160–$2.164 billion. Those guide consolidated synthesis, not an independent Government forecast. Management also describes AIP as the place where sovereign AI is built and says it is extending the stack for sovereign AI. Those statements support allied demand qualitatively, but base remains more conservative internationally because security, procurement, and data-sharing constraints delay conversion. [data_package §6, company-wide rows 7–12]

[DATA GAP: No standalone Government FY2026 revenue, Contribution, Contribution-margin, platform-revenue, or booking guide is disclosed.]

8. Forward growth assumptions — TTM 2026-06-30 to 2031-06-30

The model compounds from the $3.211 billion TTM anchor for exactly five calendar years. It does not compound from FY2025 for six fiscal years. Each scenario uses the disclosed U.S. Government subset and the allied and international residual. Segment total is the arithmetic sum, not a selected top-down number.

Base case growth

Product / sub-driver	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Base CAGR TTM→TTM+5y	Justification
U.S. Government	n/d; FY2025 +55%	90%	No segment guide	24.6%	Decelerates as comparisons normalize; existing programs, new awards, and the Army vehicle sustain expansion, offset by budget and procurement friction.
Allied and international Government	n/d; FY2025 residual +46.6%	about 42%	No segment guide	15.2%	Sovereign-AI demand grows, but classification, data sovereignty, partner access, and slower procurement keep growth below the U.S. path.
Government total	30.9% FY2022–FY2025	79%	No segment guide	22.9%	Revenue-weighted result; the pace fades from hypergrowth without assuming demand collapse.

The base U.S. path is 42%, 28%, 22%, 18%, and 15% across the five twelve-month periods. The first year remains elevated because Q2 exited at 90% growth, the TTM includes only partial benefit from recent awards, and the Army vehicle can accelerate ordering. The path then normalizes as the base expands from $2.552 billion to $7.679 billion. A fall directly from 90% to low teens would be inconsistent with disclosed acceleration absent a contract shock; a permanently high rate would ignore procurement capacity and budgets.

The base international path is 20%, 17%, 15%, 13%, and 11%. It starts below the latest implied 42% because the quarterly residual is small and affected by timing, while allied procurement and data-sharing constraints are structural. A 15.2% five-year CAGR is still strong because sovereign AI, defense modernization, and proven U.S. deployments can create reference-driven expansion.

Bull case growth

Product / sub-driver	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Bull CAGR TTM→TTM+5y	Justification
U.S. Government	n/d; FY2025 +55%	90%	No segment guide	34.6%	Enterprise procurement converts quickly, mission wins expand across agencies, and AIP-enabled workflows compound with limited displacement.
Allied and international Government	n/d; FY2025 residual +46.6%	about 42%	No segment guide	26.1%	U.S. mission proof transfers to allies and sovereign deployments overcome classification and procurement bottlenecks.
Government total	30.9% FY2022–FY2025	79%	No segment guide	33.1%	Both geographies sustain unusually strong conversion while international mix expands modestly.

The bull U.S. path is 52%, 39%, 33%, 28%, and 23%. This is not a flat extrapolation of 90%; it is controlled deceleration with repeated program expansion. The Army agreement's $10 billion cap is **not** inserted as revenue. The scenario instead assumes the vehicle reduces friction and that other defense, intelligence, health, and civilian customers create additional routes to scale.

The bull international path is 33%, 29%, 26%, 23%, and 20%. It requires several conditions together: sovereign-AI deployments progress from pilots to production, allied budget cycles accelerate, Palantir's security architecture is accepted across jurisdictions, and reference wins outweigh local-vendor preferences. It is less probable than base because classification and sovereignty cannot be solved only through product performance.

Bear case growth

Product / sub-driver	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Bear CAGR TTM→TTM+5y	Justification
U.S. Government	n/d; FY2025 +55%	90%	No segment guide	13.3%	Award timing, funding constraints, open-architecture mandates, and concessions cause rapid normalization after one still-strong year.

Product / sub-driver	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Bear CAGR TTM→TTM+5y	Justification
Allied and international Government	n/d; FY2025 residual +46.6%	about 42%	No segment guide	5.2%	Sovereignty, local procurement, export controls, and data-sharing limits constrain conversion and renewals.
Government total	30.9% FY2022–FY2025	79%	No segment guide	11.8%	Current work supports growth, but new-award and workflow-expansion engines weaken materially.

The bear U.S. path is 30%, 16%, 10%, 7%, and 5%. The first year does not collapse because the TTM exits with high quarterly growth and current programs continue. The sharper deceleration afterward has specific causes: large-award comparisons roll off, enterprise discounts reduce dollar growth, agencies require portability, and new workloads go to hyperscalers, integrators, or internal platforms. GAO's focus on lock-in, data rights, open standards, and continuous evaluation provides the structural basis; this is not generic maturation.

The bear international path is 8%, 6%, 5%, 4%, and 3%. Geopolitical demand remains high but revenue conversion disappoints because allies impose local hosting, domestic industrial participation, or data restrictions that make deployments slower and more services-intensive. A complete contraction is not central because Palantir has a growing installed base, but low-single-digit terminal growth reflects a weakened franchise.

Explicit year-by-year revenue bridge

Case and sub-driver	TTM 2026-06-30	2027-06-30	2028-06-30	2029-06-30	2030-06-30	2031-06-30
Bear U.S. Government	$2.552B	$3.318B	$3.848B	$4.233B	$4.530B	$4.756B
Bear allied and international	$0.659B	$0.712B	$0.754B	$0.792B	$0.824B	$0.849B
Bear Government total	$3.211B	$4.029B	$4.603B	$5.025B	$5.353B	$5.605B
Base U.S. Government	$2.552B	$3.624B	$4.639B	$5.659B	$6.678B	$7.679B
Base allied and international	$0.659B	$0.791B	$0.925B	$1.064B	$1.202B	$1.335B
Base Government total	$3.211B	$4.415B	$5.564B	$6.723B	$7.880B	$9.014B
Bull U.S. Government	$2.552B	$3.879B	$5.392B	$7.171B	$9.179B	$11.290B
Bull allied and international	$0.659B	$0.876B	$1.131B	$1.425B	$1.752B	$2.103B
Bull Government total	$3.211B	$4.756B	$6.523B	$8.596B	$10.931B	$13.393B

The bridge is geography-based. Product capability enters through new AI workflows, ontology expansion, deployment breadth, and mission conversion, but no product revenue is assigned. Existing programs support early years; new awards and procurement vehicles determine slope; workflow expansion and competitive retention determine duration. The forecast can be checked through quarterly U.S. Government revenue, total Government revenue, awards, disclosed margin, and commentary without relying on an undisclosed AIP mix.

9. Forward Contribution margin assumptions

The exact segment profit metric throughout this section is **Contribution**. The TTM Government Contribution margin is 70.2%. Palantir does not disclose Contribution by U.S. versus international customer, so sub-driver margins below are explicit analyst assumptions used only to construct the segment total. They are not historical facts and do not allocate R&D, G&A, SBC, or any statutory-bridge line into the segment.

Base Contribution margin

Product / sub-driver	TTM Contribution margin	2031-06-30 Contribution margin	2031-06-30 revenue	2031-06-30 Contribution	Mechanism
U.S. Government	n/d	75.0%	$7.679B	$5.759B	Scale, reusable deployments, and lower procurement friction outweigh hosting and volume discounts.
Allied and international Government	n/d	68.0%	$1.335B	$0.908B	Sovereignty, local deployment, and support intensity keep margin below U.S. assumptions.

Product / sub-driver	TTM Contribution margin	2031-06-30 Contribution margin	2031-06-30 revenue	2031-06-30 Contribution	Mechanism
Government total	70.2%	74.0%	$9.014B	$6.667B	Revenue-weighted Contribution; 380 bps above current.

The base gain exceeds 100 basis points, so its mechanism is concrete. Direct sales and marketing should grow slower than revenue as enterprise vehicles, references, and installed-program expansion reduce capture cost per dollar. Reusable product and deployment tooling should lower labor intensity. Management's hosting change should progress from a Q2 headwind toward efficiency. These benefits are offset by customer cost-certainty commitments, cloud compute, security work, and enterprise discounts. A 74.0% segment margin is meaningful leverage, not an extrapolation to a pure-license margin.

Bull Contribution margin

Product / sub-driver	TTM Contribution margin	2031-06-30 Contribution margin	2031-06-30 revenue	2031-06-30 Contribution	Mechanism
U.S. Government	n/d	79.0%	$11.290B	$8.919B	Software-led workflow expansion and efficient hosting create very high incremental Contribution.
Allied and international Government	n/d	73.0%	$2.103B	$1.535B	Larger sovereign deployments amortize accreditation and local support.
Government total	70.2%	78.1%	$13.393B	$10.454B	Revenue-weighted Contribution; about 790 bps above current.

Bull requires favorable mix, not only growth. Expansion must occur mainly through software capacity and workflows on deployed platforms, while implementation assets are reused. Hosting optimization must deliver management's expected efficiency. Procurement vehicles must cut capture and reseller cost by more than volume discounts reduce price. International scale must spread local security and deployment expenses. The U.S. sub-driver is capped at 79%, recognizing Government still consumes cloud, support, and mission engineering.

Bear Contribution margin

Product / sub-driver	TTM Contribution margin	2031-06-30 Contribution margin	2031-06-30 revenue	2031-06-30 Contribution	Mechanism
U.S. Government	n/d	67.0%	$4.756B	$3.187B	Volume discounts, competition, hosting cost, and services-heavy work overwhelm scale benefit.
Allied and international Government	n/d	60.0%	$0.849B	$0.509B	Local hosting, sovereignty requirements, and integration effort depress economics.
Government total	70.2%	65.9%	$5.605B	$3.696B	Revenue-weighted Contribution; about 430 bps below current.

Bear decline exceeds 100 basis points and has an explicit mechanism. Agencies demand modularity, portability, and pricing transparency; Palantir absorbs more compute or guarantees customer cost; mix tilts toward integration and support; and slower growth leaves capture and deployment teams underutilized. International projects remain bespoke and locally constrained. Bear does not include R&D, G&A, or SBC in Contribution—those remain outside the metric.

Product-level margin variation cannot be modeled honestly because the issuer discloses neither product dollars nor product Contribution. Assigning an 80% margin to AIP or a lower margin to Gotham would create false precision. Geography margins are scenario constructs tied to delivery conditions, not claims about named platforms.

10. Segment rollup — bottom-up reconciliation

Each case starts from $2.552 billion of U.S. Government TTM revenue and $0.659 billion of allied and international Government TTM revenue. Endpoints in section 8 are multiplied by sub-driver Contribution margins in section 9. Totals reconcile within rounding to less than 0.1%; there is no top-down plug.

Case	U.S. revenue	Allied and international revenue	Government revenue	U.S. Contribution	Allied and international Contribution	Government Contribution	Blended Contribution margin
Bear	$4.756B	$0.849B	$5.605B	$3.187B	$0.509B	$3.696B	65.9%
Base	$7.679B	$1.335B	$9.014B	$5.759B	$0.908B	$6.667B	74.0%
Bull	$11.290B	$2.103B	$13.393B	$8.919B	$1.535B	$10.454B	78.1%

Arithmetic tie-out. Bear revenue: $4.756B + $0.849B = $5.605B; Contribution: $4.756B × 67.0% + $0.849B × 60.0% = $3.696B; blended margin = 65.9%. Base revenue: $7.679B + $1.335B = $9.014B; Contribution: $7.679B × 75.0% + $1.335B × 68.0% = $6.667B; blended margin = 74.0%. Bull revenue: $11.290B + $2.103B = $13.393B; Contribution: $11.290B × 79.0% + $2.103B × 73.0% = $10.454B; blended margin = 78.1%. Display differences are rounding only.

Case	Revenue CAGR	U.S. ending mix	International ending mix	Contribution-margin change	Probability	Evidence-weighting rationale
Bear	11.8%	84.9%	15.1%	minus 430 bps	27%	Momentum limits near-term downside, but procurement, competition, and open-architecture risks are material.
Base	22.9%	85.2%	14.8%	plus 380 bps	50%	Best balance of acceleration, existing execution, new awards, leverage, and normalizing comparisons.
Bull	33.1%	84.3%	15.7%	plus 790 bps	23%	Requires award conversion, international acceleration, efficient hosting, and limited pricing pressure together.

Probabilities are deliberately non-uniform. Base receives 50% because it respects the strong Q2 exit rate and repeated attribution to existing programs and new awards while imposing substantial deceleration. Bear receives 27%, above bull's 23%, because procurement timing, budgets, pricing transparency, open architecture, and lock-in response can affect growth and margin at once. Bull remains substantial because acceleration, procurement consolidation, and mission proof create upside, but it requires more correlated successes.

The probability-weighted 2031-06-30 outcome is approximately $9.101 billion of Government revenue and $6.736 billion of Contribution, or a 74.0% implied weighted margin when weighted profits are divided by weighted revenue. This expected value is not a fourth case. It is useful only for consolidated sensitivity and shows how the higher-magnitude bull offsets its lower probability.

No corporate adjustment is embedded here. In the sealed bridge, Government Contribution adds to Commercial Contribution; R&D excluding SBC, G&A excluding SBC, and total SBC then bridge Contribution to statutory operating income. Interest, other income, tax, and NCI sit below that. Allocating them here would double count them in synthesis.

11. Key risks and wildcards

[Bear] Budget, continuing-resolution, and obligation timing. Demand can be strong while revenue is delayed. Contract ceilings and options are not funded backlog. Appropriations, priority, or task-order timing can create large comparisons after a 79% growth rate. The first bear year stays positive to reflect execution, but risk hits years two through five.

[Bear] Procurement concentration and pricing. Enterprise agreements reduce friction but aggregate buyer power. The Army vehicle includes volume discounts and a $10 billion maximum rather than a commitment. If Palantir becomes default, officials may demand greater discounts, transparency, or modular alternatives, reducing revenue per workflow and Contribution margin.

[Bear] Open architecture, data rights, and lock-in response. GAO documented concern that reliance on one provider creates lock-in and higher license cost. OPEN DAGIR, standardized licensing, government-owned data rights, and interfaces could make workloads portable. This is the most important structural moat risk because it can lower switching cost without a competitor replicating the whole platform.

[Bear] Hyperscaler and stack competition. AWS and Microsoft offer model access and AI tooling in classified environments. Agencies can combine those with data platforms, integrators, and custom applications. Palantir retains integration advantages, but the build-versus-buy threshold shifts as native capabilities improve. Generic analytics are most exposed; mission ontologies and edge deployments are less exposed.

[Bear] Hosting and fixed-price delivery. Q2's hosting change shows rapid growth can carry cost consequences. Customer cost certainty can transfer usage or compute risk to Palantir. If inference volumes, data egress, secure-cloud pricing, or integration effort exceed assumptions, Contribution can fall while revenue grows.

[Base and bear] Assignment and disclosure opacity. The Q1 transition from U.S. Commercial to U.S. Government complicates subset growth. Platform dollars, agency concentration, new-versus-existing program dollars, and Government bookings are not disclosed. The model can be monitored through total and U.S. revenue, but causal attribution remains incomplete. This uncertainty supports wider cases.

[Bear] Ethical, legal, and political constraints. Defense, intelligence, surveillance, targeting, immigration, and health applications can attract litigation, protest, policy restrictions, or political turnover. Rules on automated decisions, model evaluation, privacy, civil liberties, and human oversight may raise compliance cost or limit deployment. Strong governance may instead favor Palantir; direction is program-specific.

[Bull] Enterprise-vehicle conversion. The Army agreement could consolidate demand, reduce procurement time, eliminate reseller pass-through, and expand workflows. Bull treats it as an adoption accelerator, not $10 billion booked. Funded task orders, broader agency use, and sustained U.S. growth would move probability toward bull.

[Bull] Sovereign-AI expansion. Management's strategy could translate U.S. credibility into allied defense and civilian deployments. Upside is meaningful because international TTM revenue is only $0.659 billion. The wildcard is whether sovereignty drives Palantir adoption or favors local champions and country-specific stacks.

[Bull] Productization of deployment knowledge. If AIP, Apollo, reusable ontologies, and automation reduce forward-deployed labor per workflow, Palantir can grow faster while raising Contribution margin. Bull's 78.1% margin depends on this. It should be tested against quarterly Government Contribution margin and whether Q2 hosting drag reverses.

[Bull and base] New mission categories. AI-enabled logistics, readiness, command and control, intelligence fusion, health, procurement, and industrial-base planning can expand workflows. Space C2 shows demand across networks and security levels. The opportunity is not one AI budget; it is the replacement or augmentation of fragmented systems.

[All cases] Contract accounting and backlog interpretation. RDV includes options and assumes no termination; RPO is narrower and commercially biased. Neither should be multiplied by a conversion rate to fabricate Government revenue. Awards must be classified as ceilings, obligations, funded orders, or recognized revenue before influencing the model.

Monitoring scorecard. Key indicators are reportable Government growth; U.S. Government dollars and growth; Government Contribution margin; language distinguishing existing programs from new awards; funded task orders; international production deployments; hosting-cost commentary; and procurement requirements for data rights or modularity. Base or bull requires sustained growth without persistent margin penalty. Bear would be signaled by deceleration beyond the modeled path, recurring hosting drag, major recompete losses, or discounts without volume.

Summary

Case	Segment TTM+5y Revenue	Exact profit metric	Segment TTM+5y Profit Margin	Segment TTM+5y Profit	Probability
Bear	$5.605B	Contribution	65.9%	$3.696B	27%
Base	$9.014B	Contribution	74.0%	$6.667B	50%
Bull	$13.393B	Contribution	78.1%	$10.454B	23%

The segment starts from exceptional momentum and attractive Contribution economics, but the five-year debate is whether Palantir converts demand through funded, scalable programs faster than procurement reform, hyperscaler capability, data-rights requirements, and buyer concentration erode pricing and lock-in. Base answers yes, with pronounced deceleration and measured margin expansion. Platform-dollar non-disclosure remains a hard boundary: update the model through disclosed geography and contract-conversion evidence, never invented AIP or Gotham revenue.

Appendix B2 - Commercial

Commercial — Segment Deep Dive

Company: Palantir Technologies Inc. **Ticker:** PLTR **Segment:** Commercial **Information cutoff:** 2026-08-24 SGT **TTM anchor:** 2026-06-30 **Forecast horizon:** 2031-06-30, exactly 5.0 calendar years **Reporting currency:** USD **Authoritative data package SHA-256:** `740b2e36a294903da1e182fe6cc426d17220e6a2a025a1bfd9e5e47e7593cab2` **Exact segment profit metric:** `Contribution` **Authority boundary:** This appendix is a segment operating forecast. It authorizes no trade, broker instruction, ranking, position sizing, portfolio mutation, BUY queue, or ER2-process bypass.

STATUTORY_BRIDGE_STATUS: VERIFIED — CLEARED FOR FORECASTING

0. Segment boundary and taxonomy note

Commercial serves nongovernment institutions across industries using Palantir's AIP, Foundry, Apollo and related offerings. The Government and Commercial reporting structure has been stable since at least FY2020; the three-year taxonomy investigation found no segment realignment, reportable-segment change or restating 8-K. FY2024 expanded the segment note to disclose directly attributable expenses alongside Contribution, but that was a disclosure enhancement rather than an economic reclassification. A Q1 FY2026 customer-program transition from U.S. Commercial to U.S. Government affected the assignment of a customer program—management said U.S. Commercial growth would have been 143% rather than 133% without it—but did not change the segment boundary. Dollar comparisons therefore use the stable Commercial taxonomy throughout [data_package §1, paragraphs 1-3].

The exact additive CODM profit measure is **Contribution**, defined as revenue less cost of revenue and sales and marketing expense, excluding stock-based compensation. It is not operating income. There are no corporate or unallocated Contribution items in the sealed bridge. R&D excluding SBC, G&A excluding SBC, total SBC and all below-line items are consolidated bridge lines to be projected once by synthesis, not allocated here [data_package §3 introductory paragraph; §3a FY2025 and TTM bridge tables; statutory_bridge.json `common_segment_profit_metric`].

1. What this segment actually is

Commercial is an enterprise-software business aimed at putting data and models into operational use. Customers are nongovernment organizations in manufacturing, energy, healthcare, financial services, telecommunications, transportation, retail and other industries. The practical proposition is not simply a database, dashboard or model API. Palantir attempts to connect fragmented source systems, encode business objects and permissions in an ontology, expose those objects to analytics and AI, and then place decisions back into operational workflows. The buyer pays for software and deployment that can move from a focused use case—inventory, maintenance, underwriting, clinical operations, supply-chain planning or fraud—to a broader operating layer.

Foundry is the data-integration, ontology, analytics and operational-application foundation. AIP connects models and agents to governed enterprise context, permissions, workflows and evaluation. Apollo provides continuous delivery and orchestration across public cloud, private cloud, on-premise and edge environments. These products work together: AIP can generate demand, Foundry supplies governed context and workflow semantics, and Apollo makes deployment repeatable across heterogeneous infrastructure. Palantir does not disclose revenue for AIP, Foundry or Apollo. They are therefore treated as **non-additive product overlays**, not forecast dollar lines [data_package §4, Product and sub-segment historicals tables].

Pricing is contract-based rather than a simple disclosed per-seat tariff. Contracts can be multi-year but may include options and termination-for-convenience provisions. Remaining deal value consequently includes exercisable options and assumes no termination; it is not equivalent to RPO. RPO is a firmer contracted-revenue measure but is primarily commercial because short initial terms and termination rights common in government arrangements can exclude government commitments. The model therefore does not treat the $13.1 billion of total RDV at Q2 FY2026 as guaranteed backlog, and it does not mechanically assign the $4.9 billion RPO entirely to Commercial even though the disclosure says it is primarily commercial [data_package §0a; §5 customer and backlog KPI table].

The go-to-market motion combines intensive initial engagement with expansion. A focused deployment must reach production, establish an ontology around a real operating problem and demonstrate measurable value. If it works, Palantir can expand across workflows, business units and model families. AIP bootcamps and similar hands-on programs reduce the time needed for prospects to see a working system, but those programs do not eliminate enterprise procurement, integration,

security or change-management friction. The model distinguishes new-customer acquisition from expansion of existing customers because they have different risk: new lands depend on sales capacity and conversion, while expansion depends on realized customer value, renewal behavior and continued workflow depth.

The disclosed evidence shows both motions accelerating. U.S. Commercial customers rose from 530 in Q3 FY2025 to 571 in Q4, 615 in Q1 FY2026 and 653 in Q2. Reported NRR increased from 134% to 139%, 150% and 157% across the same quarters. The NRR disclosure is not a Commercial-only cohort statistic and is therefore used as a directional cross-check, not inserted mechanically into Commercial revenue. Total customer count was not disclosed for Q2 FY2026, which prevents a fully reconciled companywide logo-and-ARPU model [data_package §5 customer and backlog KPI table; §10 open data gaps].

2. Current economics

TTM Commercial revenue was **$2.945 billion**, 47.8% of consolidated revenue, and Contribution was **$2.136 billion**, for a **72.5% Contribution margin**. The TTM comprises Q3 FY2025, Q4 FY2025, Q1 FY2026 and Q2 FY2026. It is the sole financial starting point for this five-year forecast [data_package §3, TTM segment table, Commercial row; source: Q3 FY2025 and Q1/Q2 FY2026 Forms 10-Q, FY2025 Form 10-K, and sealed `SEC_2026_Q2_10Q` bridge record].

Momentum accelerated markedly through the anchor period. Quarterly Commercial revenue rose from $548 million in Q3 FY2025 to $677 million in Q4, $774 million in Q1 FY2026 and $945 million in Q2. Year-over-year growth accelerated from 72.9% to 81.8%, 95.1% and 109.8%. Contribution margin moved from 65.7% to 70.5%, 72.7% and 77.7%. This combination is unusually strong: growth more than doubled while the directly attributable cost ratio fell. It is also an unstable extrapolation base because the latest quarterly rate embeds rapid U.S. adoption from a much smaller prior-year base [data_package §5 quarterly segment table; sources: earnings releases dated 2025-11-03, 2026-02-02, 2026-05-04 and 2026-08-03 plus the corresponding Forms 10-Q/FY2025 Form 10-K].

The finest dollar-disclosed forward-model split is geography. Summing the four disclosed U.S. Commercial quarters—$397 million, $507 million, $595 million and $764 million—produces a TTM U.S. Commercial anchor of **$2.263 billion**. Subtracting that from the disclosed $2.945 billion Commercial total gives an **international Commercial residual of $0.682 billion**. These are calculated forecast lines, not new issuer disclosures [data_package §5 U.S. subset table; §3 TTM Commercial row; sources: earnings releases dated 2025-11-03, 2026-02-02, 2026-05-04 and 2026-08-03].

The regional contrast is important. U.S. Commercial grew 149% year over year in Q2 and management guided FY2026 U.S. Commercial revenue to more than $3.424 billion, at least 134% growth. International Commercial is not directly reported, but the Q2 residual is approximately $181 million: $945 million total less $764 million U.S. The prior-year quarterly values implied by each disclosed YoY rate are about $451 million total and $307 million U.S., leaving roughly $144 million international; that implies about 26% international growth. This inference calibrates the disclosed regional residual and is not presented as a company KPI [data_package §5 quarterly tables; §6 Commercial row 9].

`[DATA GAP: AIP, Foundry and Apollo revenue, gross margin and customer counts are not separately disclosed after the upstream official-source search; no platform dollars are back-solved.]` The financial model forecasts U.S. and international Commercial dollars, then uses product adoption, customer additions and NRR as non-additive evidence for the growth cadence.

Commercial is asset-light in the accounting sense: consolidated FY2025 capex was $34 million, 0.8% of revenue, and TTM capex was $42 million. That does not mean the segment requires no investment. Its economic reinvestment occurs mainly through personnel, cloud and deployment costs, product engineering and sales capacity. Contribution includes cost of revenue and sales and marketing, but excludes R&D, G&A and all SBC; consequently the high Contribution margin is not a full operating margin or a measure of shareholder earnings [data_package §2 annual and TTM tables; source: live FMP cash-flow endpoint checked to FY2025 Form 10-K and Q2 FY2026 Form 10-Q; §3a sealed bridge].

3. Unit economics

Commercial revenue turns into Contribution after two directly attributable cost pools: cost of revenue and sales and marketing expense, both excluding SBC. At the TTM anchor, their combined share is 27.5% because Contribution margin is 72.5%. The segment disclosure does not split those cost pools by geography or product, so this appendix does not invent U.S., international, AIP or Foundry gross margins. It models only the disclosed aggregate Contribution margin.

The most attractive economic feature is expansion. A successful deployment can spread across additional workflows without rebuilding the platform from zero. Foundry's ontology can become shared infrastructure; AIP can add model-driven workflows

on the same governed context; Apollo can deliver the software across environments. This can raise revenue per customer faster than delivery costs, improving Contribution margin. The recent rise in disclosed NRR to 157% is consistent with strong expansion, but it is not proof of a permanent cohort curve. NRR can fall as large expansions anniversary, usage stabilizes or customers optimize spend.

New-logo economics are less mature. Palantir often needs high-value technical staff to scope, integrate and operationalize a use case. A bootcamp may shorten the initial sales cycle, yet production deployment can still require solution engineering, data preparation, security review and organizational adoption. In early periods, those activities can depress Contribution through cost of revenue and sales and marketing. Over time, reusable software and repeatable patterns should reduce effort per deployment. The bear assumes repeatability disappoints and service intensity remains high; the bull assumes AIP makes initial value demonstration and production conversion dramatically more repeatable.

The historical margin path shows both risk and opportunity. Commercial Contribution margin fell from 55.5% in FY2021 to 49.7% in FY2022, recovered to 51.9% in FY2023, reached 59.5% in FY2024 and 65.9% in FY2025, before 72.5% in the TTM and 77.7% in Q2 FY2026. The recovery coincided with growth acceleration and better sales productivity, but Contribution's exclusions mean part of the enterprise's cost of supporting growth sits below this segment measure [data_package §3 Commercial table and §5 quarterly segment table; sources: FY2021-FY2025 Forms 10-K Contribution tables and Q2 FY2026 Form 10-Q/2026-08-03 earnings release].

Four mechanisms can lift Contribution margin. First, expansion revenue generally requires less acquisition cost than a new customer. Second, higher deployment reuse can reduce delivery labor per dollar. Third, greater scale can improve utilization of technical teams and cloud commitments. Fourth, pricing can reflect the mission-critical value of workflow outcomes rather than commodity compute. Four mechanisms can reverse it: price competition from hyperscalers and data platforms; unusually services-heavy implementations; cloud inference and hosting costs that grow with usage; and rapid sales hiring ahead of conversion.

R&D, G&A and SBC are explicitly outside this segment margin. Management's plan to invest in product pipeline and elite technical talent matters to consolidated earnings, but not as a direct adjustment to Commercial Contribution. Synthesis must apply those lines once through the sealed bridge. This boundary prevents the common error of calling Contribution “operating income” or treating a 75% Contribution margin as a 75% statutory operating margin.

4. Competitive dynamics

Palantir competes across several overlapping layers rather than in one clean market. Data and analytics platforms such as Snowflake and Databricks compete for the governed data and AI workload. Microsoft competes through Fabric, Foundry, Azure and its distribution inside the enterprise stack. ServiceNow competes for workflow control and agent orchestration. Hyperscalers and systems integrators can assemble alternatives from cloud data services, foundation-model APIs, orchestration tools and custom applications. In-house engineering is a substitute for customers with sufficient talent and time.

There is no reliable single “market share” denominator for Palantir Commercial. A narrow AI-platform share omits Foundry's data and workflow role; a data-platform share omits operational applications; an application-software share is too broad. Vendor rankings can validate product capability but do not provide a clean revenue-share denominator. The defensible context is adoption and scale: Commercial TTM revenue is $2.945 billion, while much larger platforms already operate at double-digit billions of annual revenue. This leaves runway but also means Palantir confronts competitors with deeper installed bases, cloud credits, channel ecosystems and procurement relationships.

Snowflake reported Q1 FY2027 product revenue of $1.33 billion, up 34%, with 126% NRR and $9.21 billion RPO. ServiceNow reported Q2 2026 subscription revenue of $3.877 billion, up 24.5%, with cRPO up 21% and 658 customers above $5 million ACV. These are not directly comparable revenue categories, but they show that 20%-30% growth remains achievable at larger recurring-revenue scale while NRR normally sits well below Palantir's current 157% [Snowflake Q1 FY2027 results](https://investors.snowflake.com/financials/quarterly-results/default.aspx); [ServiceNow Q2 2026 results](https://investor.servicenow.com/news/news-details/2026/ServiceNow-Reports-Second-Quarter-2026-Financial-Results/default.aspx).

Microsoft illustrates the bundling threat. It disclosed 35,000 paid Fabric customers, up 60%, and more than 15,000 customers using both Foundry and Fabric, also up 60%. Microsoft can combine cloud infrastructure, data, models, productivity software, security and enterprise distribution. Palantir must demonstrate that its ontology, operational deployment speed and governed decision layer create outcomes worth buying above or alongside that bundle [Microsoft FY2026 Q3 earnings call](https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q3).

Palantir's differentiation is strongest where data are fragmented, permissions are complex, decisions matter operationally and deployment must span cloud, on-premise or edge environments. Its position is weaker where a customer needs only a commodity warehouse, generic model endpoint, dashboard or low-code workflow. Price power varies by use case. A system embedded in production planning, supply-chain allocation or regulated operations can be hard to remove; an experimental AI assistant is easier to replace or cancel.

International execution is a specific competitive weakness relative to the U.S. growth engine. The estimated TTM international Commercial base is only $0.682 billion, and the latest inferred growth is far below U.S. Commercial. Local data-sovereignty requirements can help Apollo and Palantir's deployment flexibility, but regional procurement practices, partner coverage and incumbent relationships can slow conversion. The base assumes international growth improves modestly but remains below U.S.; the bull requires repeatable channel and sovereign-AI execution; the bear assumes the regional gap persists.

5. Moat at the segment level

The Commercial moat is real but still being tested at scale. Its strongest component is accumulated customer-specific ontology and workflow integration. Once business objects, permissions, actions and decision logic are encoded and used in production, switching is more difficult than replacing a dashboard. The cost includes technical migration, revalidation, retraining and operational risk. The moat deepens when multiple teams and applications depend on the same ontology.

A second component is deployment know-how. Palantir has spent years building software and field practices for messy, high-stakes environments. Apollo's ability to deploy across heterogeneous infrastructure and AIP's emphasis on governed model use can matter where generic cloud-native architectures are insufficient. This is not an absolute technical monopoly—well-funded competitors can build comparable features—but the integrated product and implementation system can shorten time to value.

A third component is expansion evidence. Rising NRR, larger customer relationships and accelerating U.S. Commercial revenue indicate that successful deployments can compound. Top-20 average TTM revenue per customer rose from $43.6 million in FY2021 to $93.9 million in FY2025 and $124 million at Q2 FY2026, although that KPI is companywide and customer concentration makes it a double-edged indicator [data_package §0a industry KPI table]. High expansion suggests customers are adding workflows, but a few very large relationships can also create lumpiness and bargaining power.

The moat is not yet proven for twenty years. Enterprise AI architecture is changing rapidly, and foundation models are not exclusive. Microsoft, ServiceNow, Snowflake, Databricks, hyperscalers and open-source components can replicate parts of the stack. Customer-built systems can improve as tools become easier to use. Palantir's durable advantage must be the governed ontology and operational outcome layer, not access to a model. Management's framing that AIP is where sovereign AI is built, deployed and compounded is strategically coherent, but remains to be validated through durable retention and margin after the current adoption surge [data_package §6 Commercial row 11].

Over five to ten years, the moat is most durable if three conditions hold: NRR normalizes above conventional enterprise-software levels rather than collapsing after initial AI enthusiasm; customer additions remain broad enough that growth is not driven by a handful of whales; and Contribution margin stays high while implementation becomes repeatable. A sustained NRR below roughly 110%, flattening U.S. Commercial customer count or a return of Commercial Contribution margin toward the low-60s would indicate that the moat is less powerful than current numbers imply.

6. Historical growth rates

The only complete dollar-disclosed product or sub-segment history is the Commercial customer-type segment itself. Palantir does not disclose platform dollars. The three-year CAGR below uses FY2022 to FY2025, a three-year compounding interval; the FY2023-FY2025 display supplies the requested three fiscal-year view. Every historical cell is tied to a specific authoritative data-package row.

Product / sub-segment	FY2023 revenue	FY2024 revenue	FY2025 revenue	3-year CAGR	Most recent YoY	Data package cite
Commercial total	$1.003B [§3 Commercial Revenue row]	$1.296B [§3 Commercial Revenue row]	$2.073B [§3 Commercial Revenue row]	35.5% [§3 FY2022 $0.834B to FY2025 $2.073B]	109.8% [§5 Q2 FY2026 Commercial row]	data_package §3 Commercial table and §5 quarterly segment table; official FY2022-FY2025 Forms 10-K and 2026-08-03 Q2 release
U.S. Commercial	n/d [§0a historical-series gap]	n/d [§0a historical-series gap]	n/d annually [§0a; quarterly only]	n/d [§0a]	149% [§5 Q2 FY2026 U.S. row]	data_package §0a KPI inventory and §5 U.S. subset table; 2026-08-03 Q2 release

Product / sub-segment	FY2023 revenue	FY2024 revenue	FY2025 revenue	3-year CAGR	Most recent YoY	Data package cite
International Commercial residual	n/d [§0a and §4]	n/d [§0a and §4]	n/d annually [§0a and §4]	n/d [§0a and §4]	approximately 26% calculated [§5 Q2 total and U.S. rows]	data_package §4 non-disclosure and §5 quarterly tables; 2026-08-03 Q2 release
AIP / Foundry / Apollo overlays	n/d [§4 Product table]	n/d [§4 Product table]	n/d [§4 Product table]	n/d [§4 Product table]	n/d [§4 growth-disclosure table]	data_package §4 rows AIP, Foundry and Apollo; FY2023-FY2025 Forms 10-K and latest releases searched

Historical arithmetic checks pass within the required 2% tolerance. Commercial FY2023 revenue of $1.003 billion multiplied by 1.292 equals approximately $1.296 billion; FY2024 revenue of $1.296 billion multiplied by 1.600 equals approximately $2.074 billion versus the disclosed rounded $2.073 billion. The difference is less than 0.1%. The three-year CAGR is $(2.073 / 0.834)^(1/3) - 1 = 35.5%$ [data_package §3 Commercial table].

Recent regional tie-out: the four U.S. Commercial quarters sum to $2.263 billion. $2.945 billion of TTM Commercial less $2.263 billion equals $0.682 billion international residual. Q2 FY2026 also ties: $0.764 billion U.S. plus approximately $0.181 billion international equals $0.945 billion total Commercial. No AIP, Foundry or Apollo dollars are inferred.

7. Management guidance

The following rows preserve every Commercial-relevant or companywide forward anchor in data-package §6. Short verbatim excerpts identify the source language; interpretation distinguishes current guidance from superseded guidance.

Data-package row	Verbatim excerpt	Interpretation for Commercial
§6 Commercial row 1, Q3 FY2025 CFO	“raising our full year 2025 revenue guidance midpoint to $4.398 billion”	Superseded by FY2025 actual revenue; evidence of repeated upward revisions, not a forward control.
§6 Commercial row 2, Q3 FY2025 CFO	“U.S. commercial revenue guidance to an excess of $1.433 billion”	Superseded historical guide that documented acceleration into FY2025 year-end.
§6 Commercial row 3, Q4 FY2025 CFO	“guiding a full year 2026 revenue of $7.190 billion at the midpoint”	Superseded first FY2026 company guide; later upgrades show unusually fast demand conversion.
§6 Commercial row 4, Q4 FY2025 CFO	“expect expenses to increase in 2026”	Consolidated reinvestment prevents direct translation of Contribution leverage into statutory margin.
§6 Commercial row 5, Q1 FY2026 CFO	“raising our revenue guidance to between $7.65 billion and $7.662 billion”	Superseded after one quarter; useful as evidence of forecast momentum.
§6 Commercial row 6, Q1 FY2026 CFO	“U.S. Commercial revenue guidance to in excess of $3.224 billion”	Superseded by Q2; establishes the direction of U.S. Commercial upgrades.
§6 Commercial row 7, Q2 FY2026 CFO	“For Q3 2026, we expect revenue of between $2.16 billion and $2.164 billion”	Current near-term consolidated anchor, not a Commercial guide. It supports strong second-half demand but cannot be allocated mechanically.
§6 Commercial row 8, Q2 FY2026 CFO	“raising our revenue guidance to between $8.15 billion and $8.158 billion”	Current consolidated FY2026 control range. Commercial plus Government forecasts must ultimately fit it in synthesis.
§6 Commercial row 9, Q2 FY2026 CFO	“U.S. commercial revenue guidance to in excess of $3.424 billion”	Current hard subset floor and the most important near-term Commercial anchor. It supports a strong first year but not 134% growth for five years.
§6 Commercial row 10, Q2 FY2026 CFO	“adjusted free cash flow guidance to between $4.5 billion and $4.7 billion”	Consolidated cash-conversion evidence; not a segment Contribution target.
§6 Commercial row 11, Q2 FY2026 CTO	“AIP is where your sovereign AI is built, deployed and compounded”	AIP is positioned as an orchestration and control layer, supporting expansion and international sovereign deployments. No dollar allocation follows.
§6 Commercial row 12, Q2 FY2026 CEO	“taking the AIP stack and extending it for sovereign AI”	Supports continued product investment and a potential international driver; it is qualitative, not revenue guidance.

`[DATA GAP: Management provides no total Commercial FY2026 guide, no international Commercial guide, no platform-level revenue guide, no Commercial Contribution-margin guide and no multi-year segment target.]` The forecast must combine the disclosed U.S. floor with segment history, regional residuals, customer and NRR evidence, and competitive scale paths.

Three interpretation rules prevent misuse of guidance. First, only the Q2 FY2026 figures are current; the earlier guides are evidence of revisions, not parallel scenarios. Second, consolidated guidance cannot be allocated to Commercial without a

disclosed segment split. Third, adjusted operating income and adjusted FCF are not Contribution. They may discipline the later synthesis, but they cannot change the exact segment metric here.

8. Forward growth assumptions

Modeling architecture and sub-driver bridge

Only two dollar lines are additive: U.S. Commercial and international Commercial. Product overlays are causal evidence. Within each geography, growth is cross-checked through customer additions, existing-customer expansion and price/workflow mix. The weights below refer to modeled **incremental revenue from TTM to 2031**, not disclosed platform revenue. AIP can influence both landing and expansion, so it is not added again as a third revenue line. Foundry provides the governed ontology and application context, while Apollo makes deployment more repeatable across infrastructure. None receives invented historical or forecast revenue.

Case	Geography	TTM revenue	Terminal revenue	Increment	Existing-customer expansion weight	New-customer weight	Price / workflow-mix weight
Bear	U.S. Commercial	$2.263B	$6.202B	$3.939B	45%	35%	20%
Bear	International Commercial	$0.682B	$0.947B	$0.265B	35%	45%	20%
Base	U.S. Commercial	$2.263B	$10.302B	$8.039B	55%	30%	15%
Base	International Commercial	$0.682B	$1.581B	$0.899B	45%	40%	15%
Bull	U.S. Commercial	$2.263B	$16.850B	$14.587B	60%	25%	15%
Bull	International Commercial	$0.682B	$2.555B	$1.873B	50%	35%	15%

The weights are deliberately non-uniform. The bear assumes reduced NRR makes expansion less dominant and pushes the company toward costlier new-logo selling. The base assigns more than half of U.S. incremental dollars to expansion because current NRR is 157% and top relationships are growing rapidly. The bull puts the greatest weight on expansion because AIP, Foundry and Apollo become a reusable operating layer across workflows. International weights tilt more toward new customers because the installed base is smaller. These weights describe the economic source of growth; they are not platform revenue estimates.

Base case growth rates

Product / sub-segment	3-year historical CAGR	Latest YoY	Management FY+1 guide	Base CAGR TTM to TTM+5y	Justification
U.S. Commercial	n/d annually	149%	FY2026 above $3.424B; at least 134%	35.4%	Year one reflects the disclosed FY2026 floor; growth then decelerates as the base expands. NRR-led expansion and AIP conversion remain strong, but the terminal 15% rate falls below mature peer growth.
International Commercial	n/d annually	approximately 26% inferred	No segment guide	18.3%	Sovereign-AI and Apollo deployment flexibility support growth, but local procurement and weaker channel density keep this below U.S. Commercial.
Commercial total	35.5%	109.8%	No total segment guide	32.2%	The weighted sum of the two geography lines reaches $11.882B. Deceleration is explicit and reflects saturation, harder comparables and peer scale paths.
AIP / Foundry / Apollo overlays	Revenue n/d	Revenue n/d	Qualitative sovereign-AI commentary	Non-additive	AIP accelerates landing and expansion, Foundry anchors ontology/workflows, and Apollo supports deployment. No platform dollars or CAGRs are invented.

Base year-by-year cadence:

Period end	U.S. Commercial revenue	U.S. growth	International revenue	International growth	Commercial total
TTM 2026-06-30	$2.263B	Anchor	$0.682B	Anchor	$2.945B
2027-06-30	$3.892B	72%	$0.853B	25%	$4.745B
2028-06-30	$5.605B	44%	$1.040B	22%	$6.645B
2029-06-30	$7.343B	31%	$1.227B	18%	$8.570B
2030-06-30	$8.958B	22%	$1.411B	15%	$10.369B
2031-06-30	$10.302B	15%	$1.581B	12%	$11.882B

The first-year U.S. number is consistent with FY2026 guidance above $3.424 billion because the 2027-06-30 TTM includes a later six-month period than FY2026. The path does not carry the current 149% rate beyond its early high-growth phase. Snowflake's 34% product growth with 126% NRR and ServiceNow's 24.5% subscription growth at larger scale make the base's mid-period 44%-31% rates demanding but not unprecedented for a breakout platform; the terminal 15% and 12% rates assume material maturation.

The base case customer logic is an audit cross-check rather than a hidden third revenue formula. Starting with 653 U.S. Commercial customers at Q2, a high-teens five-year customer CAGR and mid-teens revenue-per-customer growth can broadly support a $10.3 billion U.S. terminal line. NRR would have to fade from 157% toward approximately 118%-125% by the horizon rather than remain at today's level. That band remains above conventional software retention, reflecting workflow expansion, but concedes cohort normalization.

Bull case growth rates

Product / sub-segment	3-year historical CAGR	Latest YoY	Management FY+1 guide	Bull CAGR TTM to TTM+5y	Justification
U.S. Commercial	n/d annually	149%	FY2026 above $3.424B; at least 134%	49.4%	AIP bootcamp-to-production conversion remains exceptional, NRR stays well above conventional SaaS levels and customer additions scale without proportional sales cost.
International Commercial	n/d annually	approximately 26% inferred	No segment guide	30.2%	Sovereign-AI demand, partner coverage and Apollo's multi-environment deployment close part of the regional gap.
Commercial total	35.5%	109.8%	No total segment guide	45.8%	U.S. expansion dominates the weighted outcome; the terminal growth rate remains above 20%, requiring durable category leadership.
AIP / Foundry / Apollo overlays	Revenue n/d	Revenue n/d	Qualitative sovereign-AI commentary	Non-additive	AIP becomes the main demand wedge, Foundry remains the context layer and Apollo scales deployment. These roles drive the geography lines but receive no dollar allocation.

Bull year-by-year cadence:

Period end	U.S. Commercial revenue	U.S. growth	International revenue	International growth	Commercial total
TTM 2026-06-30	$2.263B	Anchor	$0.682B	Anchor	$2.945B
2027-06-30	$4.413B	95%	$0.955B	40%	$5.368B
2028-06-30	$7.149B	62%	$1.289B	35%	$8.438B
2029-06-30	$10.294B	44%	$1.676B	30%	$11.970B
2030-06-30	$13.588B	32%	$2.095B	25%	$15.683B
2031-06-30	$16.850B	24%	$2.555B	22%	$19.405B

This case is intentionally hard to earn. It requires Palantir to sustain a terminal U.S. growth rate similar to ServiceNow's current subscription rate even after U.S. Commercial grows more than sevenfold. It also requires NRR to remain unusually high after the current 157% reading and international execution to accelerate. AIP adoption must convert into production workloads rather than pilots, and Microsoft or data-platform bundles must not force material price concessions.

The bull's customer cross-check requires something like a mid-20s five-year U.S. Commercial customer CAGR together with high-teens revenue-per-customer growth. NRR would need to remain around 125%-135% at the horizon. This is materially above Snowflake's recent 126% NRR even after five years of Palantir scaling, which is why bull probability is only 20%.

Bear case growth rates

Product / sub-segment	3-year historical CAGR	Latest YoY	Management FY+1 guide	Bear CAGR TTM to TTM+5y	Justification
U.S. Commercial	n/d annually	149%	FY2026 above $3.424B; at least 134%	22.3%	The near-term guide still creates a strong first year, but large expansions anniversary, NRR normalizes rapidly and competitive bundles slow new production wins.
International Commercial	n/d annually	approximately 26% inferred	No segment guide	6.8%	Procurement friction, local incumbents and slower partner scaling offset sovereign-AI opportunity.
Commercial total	35.5%	109.8%	No total segment guide	19.4%	U.S. momentum prevents an immediate collapse, but the terminal blend slows to high single digits as the initial AI deployment wave matures.
AIP / Foundry / Apollo overlays	Revenue n/d	Revenue n/d	Qualitative sovereign-AI commentary	Non-additive	AIP remains useful but is increasingly bundled or replicated; Foundry and Apollo retain customers without sustaining exceptional new-logo growth.

Bear year-by-year cadence:

Period end	U.S. Commercial revenue	U.S. growth	International revenue	International growth	Commercial total
TTM 2026-06-30	$2.263B	Anchor	$0.682B	Anchor	$2.945B
2027-06-30	$3.395B	50%	$0.750B	10%	$4.145B
2028-06-30	$4.345B	28%	$0.810B	8%	$5.155B
2029-06-30	$5.127B	18%	$0.859B	6%	$5.986B
2030-06-30	$5.742B	12%	$0.902B	5%	$6.644B
2031-06-30	$6.202B	8%	$0.947B	5%	$7.149B

The bear is not a near-term revenue collapse because the FY2026 U.S. guide and current contracted demand make that inconsistent with disclosed evidence. The downside is a much faster fade after the guide period. This is plausible if NRR falls toward conventional software levels, large deals are not replicated internationally and enterprise customers consolidate onto incumbent platforms. A low-teens U.S. customer CAGR and high-single-digit revenue-per-customer growth can support the terminal U.S. line; NRR would fade toward roughly 108%-112%.

The customer and NRR audit below makes the qualitative cross-checks explicit. It uses the disclosed 653 U.S. Commercial customers at Q2 FY2026 and the $2.263 billion U.S. TTM revenue line, which implies a rough $3.466 million of TTM revenue per Q2 ending customer. That ratio mixes a trailing revenue numerator with a point-in-time denominator and is therefore an analyst reasonableness check, not an issuer KPI. Terminal customer counts and revenue per customer are calculated assumptions. NRR remains companywide, not Commercial-only, so the terminal bands are directional constraints and are not multiplied into revenue.

Case	U.S. customer CAGR	Implied terminal U.S. customers	Implied terminal revenue / customer	Revenue / customer CAGR	Directional terminal NRR	U.S. revenue tie-out
Bear	12.0%	1,151	$5.389M	9.2%	108%-112%	Approximately 1,151 × $5.389M = $6.202B
Base	17.0%	1,432	$7.196M	15.7%	118%-125%	Approximately 1,432 × $7.196M = $10.302B
Bull	25.0%	1,993	$8.455M	19.5%	125%-135%	Approximately 1,993 × $8.455M = $16.850B

These calculations preserve the disclosed data gaps. They do not create Commercial-only NRR, churn, ARR, CAC, payback or platform revenue. They simply show the combination of logo growth and account expansion that each U.S. dollar forecast requires.

Thesis-consistency conclusion

All cases decelerate explicitly from current triple-digit Commercial growth. The bear reaches 8% U.S. and 5% international terminal growth, the base 15% and 12%, and the bull 24% and 22%. None assumes an overnight fall without a mechanism. The base deceleration is comparable to the path from breakout growth toward the 20%-30% range observed at scaled Snowflake and ServiceNow, while the bear adds faster NRR normalization and international weakness. The bull assumes Palantir remains above those scaled benchmarks because it captures a larger operational-AI category.

9. Forward Contribution margin assumptions

The exact metric is **Contribution margin**. Current TTM Commercial Contribution margin is 72.5%; FY2025 was 65.9% and Q2 FY2026 was 77.7% [data_package §3 Commercial and TTM tables; §5 quarterly segment table]. Because product and geography margins are not disclosed, all cases model a segment-level blended Contribution margin. No R&D, G&A, SBC, interest, other income, tax or NCI line is included.

Case	TTM Contribution margin	2027-06-30	2028-06-30	2029-06-30	2030-06-30	2031-06-30	Change vs TTM	Specific mechanism
Bear	72.5%	70.0%	69.0%	68.5%	68.0%	67.5%	(500) bps	Higher services intensity, price concessions, slower expansion mix and sales capacity ahead of conversion raise directly attributable costs.
Base	72.5%	73.5%	74.0%	74.4%	74.7%	75.0%	+250 bps	Expansion revenue, implementation reuse and sales productivity offset continued hiring and inference or hosting costs.
Bull	72.5%	75.5%	76.5%	77.5%	78.3%	79.0%	+650 bps	AIP-led reuse, high NRR and workflow density sharply reduce acquisition and deployment cost per incremental dollar.

The bear's 500-basis-point contraction is specific: customers use more compute-intensive or bespoke deployments, competition forces discounts, and a lower share of growth comes from expansion. Even then, terminal margin remains above FY2025's 65.9%, recognizing structural scale already achieved. The base reaches 75.0%, between the TTM and recent Q2 exit rate. Its 250-basis-point expansion requires sales and deployment reuse but assumes Q2's 77.7% is partly an unusually favorable quarterly mix. The bull reaches 79.0%, only 130 basis points above Q2, yet maintaining that level at $19.4 billion of revenue requires exceptional repeatability.

The directly attributable cost ratio is the inverse of the margin: 32.5% in bear, 25.0% in base and 21.0% in bull. This is an arithmetic consequence, not an invented cost-of-revenue versus sales-and-marketing split. The data package does not give enough segment detail to allocate those two components. What matters for synthesis is the correctly labeled aggregate Contribution.

Contribution must not be confused with gross margin. It subtracts sales and marketing as well as cost of revenue, but excludes R&D, G&A and SBC. Consolidated statutory profitability will be materially lower after the sealed operating-bridge adjustments. A high segment Contribution outcome therefore does not by itself establish high shareholder returns.

10. Segment rollup

All terminal revenue is reconstructable from the regional lines in Section 8. Terminal profit equals terminal revenue multiplied by the stated Contribution margin. Calculations use the unrounded year-by-year cadence; displayed regional values can create residuals below $0.001 billion and remain far below the 1% tolerance.

Case	U.S. Commercial terminal revenue	International terminal revenue	Segment terminal revenue	Contribution margin	Contribution	Revenue tie-out	Profit arithmetic
Bear	$6.202B	$0.947B	$7.149B	67.5%	$4.826B	$6.202B + $0.947B = $7.149B	$7.149B × 67.5% = $4.826B
Base	$10.302B	$1.581B	$11.882B	75.0%	$8.912B	Unrounded $10.3016B + $1.5807B = $11.8823B	$11.882B × 75.0% = $8.912B
Bull	$16.850B	$2.555B	$19.405B	79.0%	$15.330B	$16.850B + $2.555B = $19.405B	$19.405B × 79.0% = $15.330B

The regional weights also tie. The U.S. share rises from 76.8% at the TTM anchor to roughly 86.7%-86.8% because every scenario assumes faster U.S. than international growth. This concentration is a forecast consequence, not a disclosed management target.

Case	U.S. terminal revenue mix	International terminal revenue mix	Mix tie-out
Bear	86.8%	13.2%	100.0%
Base	86.7%	13.3%	100.0%
Bull	86.8%	13.2%	100.0%

Contribution is not allocated between U.S. and international because geography-level cost of revenue and sales and marketing are not disclosed. Doing so would invent geographic margins. Only the Commercial segment total uses the sealed Contribution basis.

Case	Revenue CAGR	Terminal directly attributable cost ratio	Terminal Contribution	Probability	Weighted revenue	Weighted Contribution
Bear	19.4%	32.5%	$4.826B	30%	$2.145B	$1.448B
Base	32.2%	25.0%	$8.912B	50%	$5.941B	$4.456B
Bull	45.8%	21.0%	$15.330B	20%	$3.881B	$3.066B
Probability-weighted	n/a	n/a	$8.969B	100%	$11.967B	$8.969B

The probability weights are **Bear 30%, Base 50%, Bull 20%**. Base receives the highest weight because current revenue, customer, NRR and guidance evidence support continued strength, but five-year extrapolation from triple-digit growth is not the modal outcome. Bear receives more weight than bull because competitive bundling, cohort normalization and international underperformance can impair the cadence even while near-term guidance is met. Bull retains 20% because the observed acceleration and margin leverage are real and could persist if AIP becomes a standard operating layer.

Probability-weighted terminal Commercial revenue is $11.967 billion: 30% of $7.149 billion plus 50% of $11.882 billion plus 20% of $19.405 billion. Probability-weighted Contribution is **$8.969 billion** on unrounded scenario arithmetic: 30% of $4.825266 billion plus 50% of $8.911728 billion plus 20% of $15.330035 billion. The displayed weighted-Contribution components sum to $8.970 billion only because each component is independently rounded to three decimals. The $8.969 billion total controls. These are analytical expectations, not a fourth scenario and not company guidance.

11. Key risks and wildcards

[Bear] NRR normalization. The current 157% reading is exceptional. Large expansion cohorts can anniversary, optimization can reduce spend and new customers can dilute cohort quality. If NRR falls below roughly 115% within two years without faster customer growth, the base expansion weights are too high.

[Bear] Microsoft and platform bundling. Fabric, Microsoft Foundry, Azure, security and productivity distribution can offer a sufficiently integrated alternative. Snowflake, Databricks and ServiceNow also compete from established data or workflow positions. Palantir may retain technical differentiation yet face lower price or narrower scope.

[Bear] Implementation intensity. AIP bootcamps may generate pilots faster than production-scale deployments. If customer-specific engineering stays heavy, growth can remain strong while Contribution margin falls because cost of revenue and sales support rise with deployments.

[Bear] International stagnation. The latest inferred international growth is materially below U.S. Commercial. Local procurement, sovereignty requirements, incumbents and partner limitations may prevent the U.S. playbook from transferring. The base requires only 18.3% five-year CAGR internationally, but even that can fail if the regional engine does not broaden.

[Bear/Base] Customer concentration and deal lumpiness. No customer represented 10% or more of Q2 revenue, but Customer I was 27% of accounts receivable. Top-20 average TTM revenue per customer reached $124 million. Large relationships validate expansion but can produce volatility, collection concentration and bargaining power [data_package §0a].

[Base] Product overlap and attribution. AIP, Foundry and Apollo dollars are not disclosed. Investors cannot tell whether AIP creates net-new platform demand, accelerates Foundry expansions or merely relabels existing work. The model keeps product overlays non-additive, but disclosure remains insufficient to test platform-specific durability.

[Bull] Sovereign AI. Apollo and AIP may be unusually well suited to regulated, on-premise and sovereign deployments. If international governments and commercial partners demand controlled model orchestration, international Commercial could accelerate beyond the base. This is the most important upside wildcard because it addresses the segment's weaker region.

[Bull] Repeatable ontology network effects. If prebuilt ontology patterns, partner solutions and agent workflows become reusable across industries, implementation cost could fall while value grows. That would justify the bull's 79% Contribution margin and high expansion weight.

[All cases] Model commoditization. Foundation-model access is unlikely to be the moat. Palantir must remain valuable as the governed context, evaluation, permission and operational-action layer. If customers can assemble that layer cheaply from cloud-native and open-source tools, AIP-driven demand can erode even while AI usage grows.

[All cases] Metric opacity. ARR, cRPO, billings, aggregate ACV, churn, GRR, CAC and payback are not disclosed. RDV includes options and termination assumptions. This makes it harder to distinguish contracted durability from enthusiastic pipeline. RPO, NRR and customer counts should be monitored together rather than used in isolation [data_package §0a and §10].

The highest-value monitoring dashboard is therefore operational rather than narrative: U.S. Commercial revenue growth versus the current $3.424 billion FY2026 floor; U.S. Commercial customer additions; NRR; international residual growth; total and commercial-biased RPO; and Commercial Contribution margin. A base-case miss is most likely to appear first as NRR and customer growth decelerating together, followed by margin pressure as sales and deployment costs fail to scale.

Commercial segment scenario summary

Case	Segment TTM+5y Revenue	Exact profit metric	Segment TTM+5y Profit Margin	Segment TTM+5y Profit	Probability
Bear	$7.149B	Contribution	67.5%	$4.826B	30%
Base	$11.882B	Contribution	75.0%	$8.912B	50%
Bull	$19.405B	Contribution	79.0%	$15.330B	20%

The forecast horizon is exactly 2031-06-30, 5.0 calendar years after the TTM anchor of 2026-06-30. All scenario dollars are segment operating estimates on the sealed Contribution basis. They are not statutory operating income, price targets, portfolio recommendations or execution instructions.

Appendix S - Synthesis (Combined Earnings Model)

Palantir Technologies Inc. — Synthesis (Combined Earnings Model)

This appendix mechanically consolidates the two reviewed segment forecasts into one statutory earnings model. It preserves the issuer's exact segment metric, **Contribution**, forecasts the corporate and statutory bridge only once, converts consolidated net income into diluted EPS on the sealed security-unit basis, and applies the separate binary P/E 15 entry-price gate. It does not choose a security-return method or exit multiple, calculate an IRR, score the company, rank it, size it, authorize a trade, or mutate a portfolio.

1. Anchors and governed inputs

- **Today and cutoff:** 2026-08-24, Singapore time.
- **Fiscal year-end:** December 31.
- **Latest completed FY for historical tie-out:** FY2025, ended 2025-12-31.
- **TTM anchor:** 2026-06-30, built from Q3 FY2025, Q4 FY2025, Q1 FY2026, and Q2 FY2026.
- **Horizon:** 2031-06-30, exactly 5.0 calendar years after the TTM end. No forecast row compounds from FY2025 or for six fiscal years.
- **Valuation price:** USD 179.94 per Class A common share, XNAS:PLTR, raw exchange close for the completed 2026-08-21 session. Actionability expiry is 2026-08-25T18:03:08+08:00.
- **Quote evidence:** raw-source SHA-256 `3d0abdd4c04dba6a52c094d9ddb8a932cf7b4cff252d24ea459be86ab0e04d0b`; exact `quote_packet.json` SHA-256 `1a87c7195d065cbbb11be9896951d0c60eda8a55e1e8f5eff53b437a9ff7010b`. The strict validator returned `QUOTE_PACKET_VALID` at 2026-08-24T20:01:41+08:00.
- **Historical bridge:** exact `statutory_bridge.json` SHA-256 `70965001151a9ac518f8ce4c11f797b1dad4280818161b657c1c0cebcd6edd3f`. The validator returned `STATUTORY_BRIDGE_VALID`; the gate manifest records the DR-cleared `VERIFIED — CLEARED FOR FORECASTING` status.
- **Governed research artifacts:** `data_package.md` SHA-256 `740b2e36a294903da1e182fe6cc426d17220e6a2a025a1bfd9e5e47e7593cab2`; Government appendix SHA-256 `4826215b4b15333c68dbbf57abc92715e0992981301f24005d945ca4df21265a`; Commercial appendix SHA-256 `7b35a45fd12341554336fc34cde90975f131d0cd5a15c8cd03282fdd90cff698`; management appendix SHA-256 `6436f3ec34aea0fe7ef76b761d349a3a8b76c744839c87640dc551ec2b9c15d0`; moat appendix SHA-256 `4fecd179a577df678de946b62c8281ace2b0b4b5d962020851c01104577dfb7d`.

The quoted instrument and financial reporting currency are both USD. The listed unit is one Class A common share. The financial statements combine Class A, B, and F common stock on an equal economic-share basis. The depositary ratio is 1.0, share-class conversion factor is 1.0, and financial shares per listed unit are therefore 1.0. The financial-basis diluted share anchor of 2,568.694 million at 2026-06-30 multiplied by the no-split factor of 1.0 equals 2,568.694 million quote-basis diluted financial shares at the quote date. Dividing by the 1.0 listed-unit factor yields 2,568.694 million listed-unit equivalents. The split endpoint found no intervening split through 2026-08-21. The FX factor is the same-currency identity, USD 1.0 per USD 1.0, dated 2026-08-21.

Historical TTM diluted shares are separately 2,571.043 million, the mean of the four official standalone-quarter denominators. Historical TTM diluted EPS is exactly USD 3,016.692 million divided by 2,571.043 million, or USD 1.173333934905017. The historical weighted-average denominator is used only for the TTM EPS anchor; the forward security-unit path starts with the packet's 2,568.694 million quote-basis financial shares, as required.

The management appendix gate is **FAIL** because founder voting control, chronic dilution, compensation design, and minority-shareholder protection do not meet the Gideon compounder bar despite strong operating execution. The moat appendix is **YELLOW**, and the 2046 test is also **YELLOW**: Government switching costs are tangible, while Commercial

durability and pricing power remain less cycle-tested. Those judgments inform scenario weighting and monitoring, but they do not alter the arithmetic of a segment forecast already cleared by review.

2. Segment-level forward assumptions

The exact common additive segment profit metric is **Contribution**, not operating income. Government and Commercial Contribution add directly; R&D excluding SBC, G&A excluding SBC, and total SBC then bridge once to statutory operating income.

Segment	Exact profit metric	TTM revenue USD B	Bear TTM+5y revenue USD B	Base TTM+5y revenue USD B	Bull TTM+5y revenue USD B	Bear profit margin	Base profit margin	Bull profit margin	Source
Government	Contribution	3.210783	5.604331	9.013410	13.392346	65.94%	73.96%	78.06%	`appendix_B_government.md` §§8-10; recomputed from annual cadence and sub-driver margins
Commercial	Contribution	2.945158	7.148761	11.882671	19.405699	67.50%	75.00%	79.00%	`appendix_B_commercial.md` §§8-10; recomputed from annual cadence and segment margin

Government was recomputed rather than copied from rounded summary cells. U.S. Government begins at USD 2.552 billion; allied and international Government is the exact residual USD 0.658783 billion. Bear annual growth is 30%, 16%, 10%, 7%, and 5% for U.S. Government and 8%, 6%, 5%, 4%, and 3% internationally. Base uses 42%, 28%, 22%, 18%, and 15% for the U.S. and 20%, 17%, 15%, 13%, and 11% internationally. Bull uses 52%, 39%, 33%, 28%, and 23% for the U.S. and 33%, 29%, 26%, 23%, and 20% internationally. Terminal U.S./international Contribution margins are 67%/60% in Bear, 75%/68% in Base, and 79%/73% in Bull.

Commercial was also recomputed. U.S. Commercial begins at USD 2.263 billion; international Commercial is the exact residual USD 0.682158 billion. Bear annual growth is 50%, 28%, 18%, 12%, and 8% in the U.S. and 10%, 8%, 6%, 5%, and 5% internationally. Base uses 72%, 44%, 31%, 22%, and 15% in the U.S. and 25%, 22%, 18%, 15%, and 12% internationally. Bull uses 95%, 62%, 44%, 32%, and 24% in the U.S. and 40%, 35%, 30%, 25%, and 22% internationally. Because Commercial geography-level margins are not disclosed, the reviewed appendix appropriately applies one segment-level Contribution margin rather than inventing geographic profit margins.

3. Combined revenue build

No corporate eliminations are disclosed in the sealed historical bridge or the reviewed segment taxonomy. Therefore the combined forecast is the exact sum of the two segment lines.

Line	Bear	Base	Bull
Government TTM+5y revenue, USD B	5.604330537	9.013410029	13.392346118
Commercial TTM+5y revenue, USD B	7.148761356	11.882670507	19.405699343
Sum of segment TTM+5y revenue, USD B	12.753091893	20.896080536	32.798045461
Corporate eliminations, USD B	0.000000000	0.000000000	0.000000000
Combined TTM+5y revenue, USD B	**12.753091893**	**20.896080536**	**32.798045461**
Implied TTM 2026-06-30 to 2031-06-30 CAGR	15.6816%	27.6892%	39.7368%

The base case reaches USD 20.896 billion from the exact USD 6.155941 billion TTM anchor. The first-year base combined revenue of USD 9.159437 billion is not an unexplained top-down step. It is the exact sum of Government USD 4.414380 billion and Commercial USD 4.745058 billion generated by the annual segment cadences. This is 12.3% above the USD 8.154 billion midpoint of current FY2026 guidance, while the forecast period ends six months later. The current Q2 revenue annualized is USD 7.742 billion and the Q3 guide annualized is about USD 8.648 billion. Those dated primary-source guide points support the step-up while still requiring sharp deceleration after the extraordinary current growth phase.

4. Combined operating-income build

The corporate and operating lines are forecast once below the Contribution subtotal. There are no corporate/unallocated Contribution adjustments. Each ratio is an explicit scenario assumption anchored to the TTM ratios and the Section 3a driver ledger; none is a balancing plug.

Line	Bear	Base	Bull
Government Contribution, USD B	3.695522866	6.666666142	10.453831132
Commercial Contribution, USD B	4.825413915	8.912002880	15.330502481
Sum of segment TTM+5y Contribution, USD B	8.520936781	15.578669022	25.784333612
Corporate / unallocated adjustments, USD B	0.000000000	0.000000000	0.000000000
Consolidated TTM+5y Contribution, USD B	8.520936781	15.578669022	25.784333612
R&D excluding SBC, USD B	(0.956481892)	(1.149284429)	(1.475912046)
G&A excluding SBC, USD B	(0.828950973)	(0.940323624)	(1.147931591)
Total SBC expense, USD B	(1.275309189)	(1.462725638)	(1.639902273)
Combined statutory TTM+5y operating income, USD B	**5.460194727**	**12.026335331**	**21.520587702**
Implied statutory operating margin	42.81%	57.55%	65.62%

Forecast bridge line	TTM anchor	Bear rule	Base rule	Bull rule	Controlling driver
R&D excluding SBC	USD 0.473564B; 7.69% of revenue	7.5% of terminal revenue; heavy technical investment and slower leverage	5.5%; sustained absolute investment with scale leverage	4.5%; fastest reuse and revenue scale, still USD 1.476B of spend	Data package §3a and technical-hiring commentary
G&A excluding SBC	USD 0.447281B; 7.27% of revenue	6.5% of terminal revenue	4.5% of terminal revenue	3.5% of terminal revenue	Data package §3a; scale-driven ratio decline
Total SBC expense	USD 0.835524B; 13.57% of revenue	10.0% of terminal revenue	7.0% of terminal revenue	5.0% of terminal revenue	Data package §3a and §7; elevated but normalizing expense

Bear does not assume operating leverage beyond today's statutory margin: lower segment margins and 24.0% of revenue for the three bridge costs leave the terminal margin almost exactly at the TTM 42.8%. Base allows Contribution leverage and declines in R&D, G&A, and SBC ratios, but SBC still rises in dollars to USD 1.463 billion. Bull requires extraordinary platform reuse and revenue density. Even there, R&D, G&A, and SBC collectively remain USD 4.264 billion, so the result is not built by removing investment.

The management appendix and data package agree on capex: FY2025 capex was USD 0.034 billion and TTM capex USD 0.042 billion; capex remained below 1% of revenue in four of five historical years. Capex is a cash-flow and return-method input, not an operating-expense bridge line, and is therefore not subtracted again from statutory operating income. This avoids double counting capital allocation inside the income statement.

5. Net income and EPS build

Interest, other income, tax, and NCI are forecast independently below operating income. Interest income uses a conservative terminal range around the TTM USD 0.266384 billion: USD 0.150 billion in Bear, USD 0.300 billion in Base, and USD 0.500 billion in Bull. This reflects lower rates or cash in Bear and stronger retained liquidity in Bull without treating current net cash as a separate equity-value addition. Other income normalizes the TTM USD 0.170794 billion of marketable and equity security gains: Bear assumes a USD 0.250 billion loss, Base assumes zero, and Bull assumes USD 0.250 billion of gains. Taxes are 23%, 21%, and 18% of pretax income. NCI attribution is separately deducted at 0.75%, 0.50%, and 0.35% of pretax income, reflecting a higher outside share in Bear and a lower share in Bull.

Line	Bear	Base	Bull
Combined TTM+5y statutory OI, USD B	5.460194727	12.026335331	21.520587702
Interest income, USD B	0.150000000	0.300000000	0.500000000
Other income, net, USD B	(0.250000000)	0.000000000	0.250000000
Pretax income, USD B	5.360194727	12.326335331	22.270587702

Line	Bear	Base	Bull
Tax rate assumed	23.00%	21.00%	18.00%
Provision for income taxes, USD B	(1.232844787)	(2.588530419)	(4.008705786)
NCI attribution, USD B	(0.040201460)	(0.061631677)	(0.077947057)
Net income attributable to common, USD B	**4.087148479**	**9.676173235**	**18.183934859**
Quote-basis diluted financial shares at TTM+5y, M	2,610.088747	2,597.849227	2,592.215058
EPS in USD per quote-basis financial share	1.565904027	3.724686226	7.014824948
Financial shares per listed unit	1.000000000	1.000000000	1.000000000
EPS in USD per listed unit before FX	1.565904027	3.724686226	7.014824948
FX, quote USD per reporting USD	1.000000000	1.000000000	1.000000000
EPS in quote currency per listed unit at TTM+5y	**USD 1.565904027**	**USD 3.724686226**	**USD 7.014824948**

The exact unit formula is `EPS_REPORTING_PER_QUOTE_BASIS_FINANCIAL_SHARE = NI_REPORTING / QUOTE_BASIS_DILUTED_FINANCIAL_SHARES`. The listed-unit result then equals that EPS multiplied by 1.0 financial share per listed unit and USD 1.0 per reporting USD. No issuer EPS is relabeled with a quote currency without this bridge.

5a. Return-method input bridge

This section supplies inputs without selecting a terminal return method. Normalized NOPAT is available from the statutory OI and scenario cash-tax rates. A complete terminal capital-structure path is not available in the governed upstream package, so EV/NOPAT cannot control without new preflight evidence. That limitation does not impair the reconciled statutory EPS bridge or the separate P/E 15 gate.

Input	Bear	Base	Bull	Source-bound status
Statutory operating income, USD B	5.460194727	12.026335331	21.520587702	This synthesis §4
Normalized cash tax rate	23.00%	21.00%	18.00%	Scenario tax schedule, data package §3a
Normalized NOPAT, USD B	4.204349940	9.500804911	17.646881916	OI multiplied by one minus tax rate
Terminal cash and non-operating financial assets	Unavailable	Unavailable	Unavailable	Current Q2 cash and marketable securities are USD 9.409B, but no upstream five-year path is sealed
Terminal gross debt	Unavailable	Unavailable	Unavailable	Current package combines debt and leases; no separate terminal path
Terminal lease liabilities	Unavailable	Unavailable	Unavailable	No separate source-bound scenario path
Terminal preferred equity	Unavailable	Unavailable	Unavailable	No source-bound terminal schedule; do not assume zero merely from current presentation
Terminal minority interest	Unavailable	Unavailable	Unavailable	NCI earnings are forecast, but terminal balance-sheet claim is not provided
Terminal other senior claims	Unavailable	Unavailable	Unavailable	No upstream scenario path
Separately valued non-operating assets	Unavailable	Unavailable	Unavailable	No item proved absent from statutory EPS and eligible for separate addition
Diluted listed-equivalent shares, M	2,610.088747	2,597.849227	2,592.215058	Section 6; listed-unit factor 1.0

No residual net-debt plug is created. Current Q2 cash, marketable securities, and debt or lease information remain balance-sheet context only. Under a future P/E return method, terminal net cash must not be added after EPS times P/E

because interest income and capital structure already affect EPS. Under an EV/NOPAT method, a full terminal cash, debt, lease, NCI, and senior-claims bridge would be mandatory. Method choice, denominator normalization, and exit multiples belong to valuation preflight and are outside this synthesis.

5b. P&C insurer lane

NOT TRIGGERED — Palantir Technologies Inc. is enterprise software, not a P&C insurer or reinsurer. No insurer-return bridge is created.

6. Share-count path

The forward path begins with 2,568.694 million quote-basis diluted financial shares from the sealed packet. The repurchase authorization was terminated in January 2026 after only USD 75 million of FY2025 repurchases; therefore all three scenarios assume zero buybacks rather than inventing a new authorization. SBC issuance is approximated as annual SBC expense divided by a disclosed scenario average share price. This is explicitly an issuance approximation, not a claim that accounting expense equals grant-date cash value. The approximation is reconciled against the recent stabilization in diluted shares and avoids the unsupported assumption that buybacks automatically offset awards.

Case / year ending June 30	Starting shares M	Buyback USD B	Average price USD	Shares retired M	SBC expense proxy USD B	SBC issuance M	Ending shares M
Bear 2027	2,568.694000	0.000	160	0.000000	0.950000	5.937500	2,574.631500
Bear 2028	2,574.631500	0.000	150	0.000000	1.050000	7.000000	2,581.631500
Bear 2029	2,581.631500	0.000	140	0.000000	1.150000	8.214286	2,589.845786
Bear 2030	2,589.845786	0.000	130	0.000000	1.250000	9.615385	2,599.461170
Bear 2031	2,599.461170	0.000	120	0.000000	1.275309	10.627577	2,610.088747
Base 2027	2,568.694000	0.000	180	0.000000	0.950000	5.277778	2,573.971778
Base 2028	2,573.971778	0.000	190	0.000000	1.050000	5.526316	2,579.498094
Base 2029	2,579.498094	0.000	200	0.000000	1.150000	5.750000	2,585.248094
Base 2030	2,585.248094	0.000	210	0.000000	1.250000	5.952381	2,591.200475
Base 2031	2,591.200475	0.000	220	0.000000	1.462726	6.648753	2,597.849227
Bull 2027	2,568.694000	0.000	200	0.000000	1.000000	5.000000	2,573.694000
Bull 2028	2,573.694000	0.000	240	0.000000	1.150000	4.791667	2,578.485667
Bull 2029	2,578.485667	0.000	280	0.000000	1.300000	4.642857	2,583.128524
Bull 2030	2,583.128524	0.000	320	0.000000	1.450000	4.531250	2,587.659774
Bull 2031	2,587.659774	0.000	360	0.000000	1.639902	4.555284	2,592.215058

The cumulative share arithmetic is exact: Bear starts at 2,568.694 million and adds 41.394747 million SBC shares to end at 2,610.088747 million; Base adds 29.155227 million to end at 2,597.849227 million; Bull adds 23.521058 million to end at 2,592.215058 million. Zero repurchases mean there is no unsupported cash or share reduction. The paths model economic dilution explicitly even though terminal SBC as a percentage of revenue declines.

7. Binary P/E 15 gate

The required EPS at USD 179.94 is exactly `179.94 / 15 = USD 11.996` per listed unit. Cushion equals `(EPS × 15 - 179.94) / 179.94`. The confident case is Base; Bull cannot rescue a Base failure.

Case	EPS in USD per listed unit at 2031-06-30	EPS × 15	Quote-packet price	Cushion	Pass?
Bear	1.565904027	23.488560	179.940000	(86.95%)	FAIL
Base	3.724686226	55.870293	179.940000	(68.95%)	**FAIL**
Bull	7.014824948	105.222374	179.940000	(41.52%)	FAIL

P/E 15 gate: FAIL — base-case cushion (68.95%). Even Bull earns only USD 7.0148 per diluted listed unit, far below the required USD 11.996. This is an entry-price gate only. It does not select a terminal multiple for a security-return model.

8. Reflexivity check against reported anchors

Historical anchor	Segment sum or bridge, USD B	Consolidated statutory amount, USD B	Residual	Status
FY2025 revenue	2.402287 + 2.073159 = 4.475446	4.475446	0.000000	PASS
TTM revenue	3.210783 + 2.945158 = 6.155941	6.155941	0.000000	PASS
FY2025 Contribution	1.576070 + 1.366627 = 2.942697	2.942697	0.000000	PASS
FY2025 Contribution to OI	2.942697 - 0.420838 - 0.423811 - 0.684033 = 1.414015	1.414015	0.000000	PASS
TTM Contribution	2.255215 + 2.135806 = 4.391021	4.391021	0.000000	PASS
TTM Contribution to OI	4.391021 - 0.473564 - 0.447281 - 0.835524 = 2.634652	2.634652	0.000000	PASS

Both revenue ties are exact and therefore within the ±2% reflexivity threshold. Both Contribution-to-OI bridges reproduce the sealed packet to USD 1 thousand, its disclosed tolerance. No plug, residual scaling, or segment allocation is used.

9. Latest interim bridge reconciliation

NOT TRIGGERED — four non-overlapping standalone quarters are used: Q3 FY2025, Q4 FY2025, Q1 FY2026, and Q2 FY2026. Revenue is USD 1.181092 billion + USD 1.406802 billion + USD 1.632583 billion + USD 1.935464 billion = USD 6.155941 billion. Statutory OI is USD 0.393256 billion + USD 0.575394 billion + USD 0.753998 billion + USD 0.912004 billion = USD 2.634652 billion. Common NI is USD 0.475599 billion + USD 0.608676 billion + USD 0.870527 billion + USD 1.061890 billion = USD 3.016692 billion. The same monetary results cross-check to FY2025 minus H1 FY2025 plus H1 FY2026. The diluted-share denominator is the mean of the four standalone denominators, not a sum and not a cumulative-YTD denominator. There is no period overlap.

10. Sensitivity — variables that drive the verdict

Two variables are held at Base while the named variable changes. Government CAGR sensitivity preserves the Base effective Government Contribution margin and all other corporate ratios; Commercial margin sensitivity changes only Commercial Contribution margin. Share sensitivity uses the exact Base NI. Tax sensitivity changes the tax rate while leaving pretax income and NCI assumptions fixed.

Variable	Down case	Base	Up case	EPS range and impact
Government five-year CAGR minus or plus 200 bps	20.93%; EPS 3.602518	22.93%; EPS 3.724686	24.93%; EPS 3.855069	Minus USD 0.122168 to plus USD 0.130383
Commercial Contribution margin minus or plus 200 bps	73.0%; EPS 3.652874	75.0%; EPS 3.724686	77.0%; EPS 3.796499	Minus or plus USD 0.071812
Terminal diluted shares minus or plus 5%	2,467.957M; EPS 3.920722	2,597.849M; EPS 3.724686	2,727.742M; EPS 3.547320	Plus USD 0.196036 to minus USD 0.177366
Tax rate minus or plus 300 bps	18.0%; EPS 3.867031	21.0%; EPS 3.724686	24.0%; EPS 3.582342	Plus or minus USD 0.142345

The most powerful displayed single perturbation is the 5% share-count change, followed by tax and Government growth. Yet none comes remotely close to closing the USD 8.271314 gap between Base EPS and the USD 11.996 gate requirement. Even combining the favorable ends of these local sensitivities would not transform the verdict. The kill variables are therefore not one marginal assumption; they are the current price and the requirement for an exceptional five-year earnings scale far beyond the already demanding Base case. Monitoring should focus on whether U.S. Commercial growth, Government contract conversion, segment Contribution margins, and actual dilution jointly outperform, rather than treating one quarter's adjusted margin as sufficient evidence.

11. Combined earnings summary

Scenario weights are blended consistently across both segments. Government's reviewed weights are Bear 27%, Base 50%, Bull 23%; Commercial's are Bear 30%, Base 50%, Bull 20%. Weighting those by exact TTM revenue mix, Government 3.210783 / 6.155941 and Commercial 2.945158 / 6.155941, produces consolidated weights of **Bear 28.435275939%, Base**

50.000000000%, Bull 21.564724061%. This retains each segment review's evidence judgment without allowing a different probability set for revenue, Contribution, OI, NI, EPS, or valuation arithmetic. Bear remains more likely than Bull because management is FAIL and moat/2046 are YELLOW despite strong recent momentum.

Case	Revenue USD B	OI USD B	OI margin	NI USD B	EPS USD per listed unit	Probability	EPS × 15	Versus USD 179.94
Bear	12.753092	5.460195	42.81%	4.087148	1.565904	28.4353%	23.488560	(86.95%)
Base	20.896081	12.026335	57.55%	9.676173	3.724686	50.0000%	55.870293	(68.95%)
Bull	32.798045	21.520588	65.62%	18.183935	7.014825	21.5647%	105.222374	(41.52%)
Probability-weighted	21.147225	12.206644	57.72%	9.921594	3.820340	100.0000%	57.305098	(68.15%)

The probability-weighted EPS is the direct sum of each scenario EPS times the shared scenario probability. It is not a fourth operating scenario, and the probability-weighted OI margin is weighted OI divided by weighted revenue. Because the binary gate is controlled by Base, the probability-weighted row does not override the FAIL.

Tie-out checks

1. Reported-financial and statutory-bridge tie-out — PASS

FY2025 Government plus Commercial revenue equals USD 4.475446 billion, exactly the consolidated amount. TTM segment revenue equals USD 3.210783 billion + USD 2.945158 billion = USD 6.155941 billion, exactly the consolidated TTM. For FY2025, Contribution of USD 2.942697 billion less R&D excluding SBC USD 0.420838 billion, G&A excluding SBC USD 0.423811 billion, and SBC USD 0.684033 billion equals statutory OI USD 1.414015 billion. For TTM, USD 4.391021 billion less USD 0.473564 billion, USD 0.447281 billion, and USD 0.835524 billion equals USD 2.634652 billion. Residuals are zero, satisfying both the packet's USD 1 thousand tolerance and the ±2% reflexivity test.

2. NI, shares, listed unit, and FX equal EPS — PASS

Case	NI USD M	Diluted financial shares M	Reporting EPS	Unit factor	FX factor	Quote EPS	Residual
Bear	4,087.148479	2,610.088747	1.565904027	1.0	1.0	1.565904027	0.000000000
Base	9,676.173235	2,597.849227	3.724686226	1.0	1.0	3.724686226	0.000000000
Bull	18,183.934859	2,592.215058	7.014824948	1.0	1.0	7.014824948	0.000000000

Every case follows the exact mandatory two-step security-unit formula. No ADR, class, split, or FX factor is omitted merely because each happens to equal one.

3. Segment sum at horizon — PASS

Bear segment revenue sums to USD 12.753091893 billion, Base to USD 20.896080536 billion, and Bull to USD 32.798045461 billion. Corporate eliminations are zero. Combined revenue divided by the segment sum is 100.0000% in every case, with no unnamed gap.

4. Product-to-segment rollup — PASS

Government Bear equals USD 4.756064914 billion U.S. plus USD 0.848265624 billion allied/international, or USD 5.604330537 billion. Base equals USD 7.679247454 billion + USD 1.334162575 billion = USD 9.013410029 billion. Bull equals USD 11.290307757 billion + USD 2.102038361 billion = USD 13.392346118 billion. Commercial Bear equals USD 6.201683067 billion U.S. plus USD 0.947078289 billion international, or USD 7.148761356 billion. Base equals USD 10.301594709 billion + USD 1.581075797 billion = USD 11.882670507 billion. Bull equals USD 16.849704478 billion + USD 2.555994864 billion = USD 19.405699343 billion. All residuals are zero at calculation precision and within 1%.

5. Weighted margin check — PASS

Government Bear Contribution is USD 4.756064914 billion times 67% plus USD 0.848265624 billion times 60%, or USD 3.695522866 billion; divided by Government revenue, the margin is 65.94%. Base is USD 7.679247454 billion times 75% plus USD 1.334162575 billion times 68% = USD 6.666666142 billion, or 73.96%. Bull is USD 11.290307757 billion times 79% plus USD 2.102038361 billion times 73% = USD 10.453831132 billion, or 78.06%. Commercial product or geography margins are not disclosed, so its reviewed segment-level blended margin is applied once; a weighted sub-product margin test is not

triggered for that segment. The absence of invented Commercial sub-margins is a control, not a missing calculation.

6. Share-count and security-unit path — PASS

The 2026-06-30 financial-basis share count of 2,568.694 million multiplied by the 1.0 split factor equals 2,568.694 million quote-basis diluted financial shares. Dividing by the depositary ratio 1.0 and class factor 1.0 gives 2,568.694 million listed-unit equivalents. From that starting point, five years of explicit SBC issuance and zero retired shares reproduce every scenario endpoint: 2,610.088747 million in Bear, 2,597.849227 million in Base, and 2,592.215058 million in Bull. The share path never runs in ADR units, and no unsupported buyback offsets dilution.

7. Growth plausibility — PASS

Government Base CAGR is 22.93%, below 1.2 times its 30.9% FY2022-FY2025 disclosed CAGR. Commercial Base CAGR is 32.18%, below 1.2 times its 35.5% historical CAGR. Government Bull at 33.06% remains below the Government historical threshold. Commercial Bull at 45.80% exceeds 1.2 times 35.5%, so explicit justification is required and present in `appendix_B_commercial.md` §§8 and 10: current Commercial growth is 109.8%, U.S. Commercial is 149%, companywide NRR is 157%, and the bull case nevertheless decelerates annual U.S. growth to 24% and international growth to 22% by the horizon. The appendix also benchmarks the demanded trajectory against Snowflake and ServiceNow and assigns only a 20% segment probability. Therefore the rule is satisfied with explicit evidence rather than ignored.

8. Historical arithmetic and five-year compounding — PASS

Segment case	TTM revenue USD B	Recomputed CAGR	TTM × one plus CAGR to the fifth	Published horizon USD B	Difference
Government Bear	3.210783	11.784749%	5.604330537	5.604330537	0.000000000
Government Base	3.210783	22.929364%	9.013410029	9.013410029	0.000000000
Government Bull	3.210783	33.060497%	13.392346118	13.392346118	0.000000000
Commercial Bear	2.945158	19.405530%	7.148761356	7.148761356	0.000000000
Commercial Base	2.945158	32.178583%	11.882670507	11.882670507	0.000000000
Commercial Bull	2.945158	45.802221%	19.405699343	19.405699343	0.000000000

Every row uses exactly five calendar years from 2026-06-30 to 2031-06-30. The constant CAGR is a mathematical summary of the explicit annual cadence, not a replacement for it.

9. Recency anchor — PASS with quantified step-change evidence

The Base first-forward-year revenue is USD 9.159437 billion versus TTM USD 6.155941 billion, a 48.79% increase and therefore outside the mechanical 10% band. The step is explicitly supported, not silently introduced. Q2 FY2026 revenue was USD 1.935464 billion and grew 92.8%; Q3 guidance is USD 2.160-2.164 billion; FY2026 guidance is USD 8.150-8.158 billion; U.S. Commercial FY2026 guidance is above USD 3.424 billion and at least 134% growth. The 2027-06-30 TTM model is only 12.33% above the FY2026 guide midpoint and covers a period ending six months later. Both segment appendices quantify the first annual cadence and then decelerate. That primary-source evidence meets the exception; there is no unsupported run-rate reset.

10. Period and horizon integrity — PASS

The TTM monetary anchor uses four standalone, non-overlapping quarters. The equivalent FY2025 minus H1 FY2025 plus H1 FY2026 construction is only a cross-check and uses matching January-June periods. Cumulative FMP share denominators are quarantined and do not enter EPS. Every displayed forecast begins on 2026-06-30 and ends on 2031-06-30. The horizon is exactly 5.0 years; FY2025 is used only for the historical tie-out. No six-year compounding or period overlap appears.

11. Valuation, unit bridge, packet hashes, and probabilities — PASS

The exact quote packet and statutory bridge byte hashes match the governed hashes. The quote validator and statutory validator both cleared. XNAS:PLTR Class A common stock, raw completed close USD 179.94, split factor 1.0, depositary ratio 1.0, share-class factor 1.0, and FX 1.0 reproduce the listed-unit basis. Bear, Base, and Bull EPS multiplied by 15 equal USD 23.488560, USD 55.870293, and USD 105.222374. Their cushions are negative 86.9464%, 68.9506%, and 41.5236%. Consolidated probabilities sum exactly to 100%; probability-weighted EPS equals 28.435275939% × 1.565904027 + 50% × 3.724686226 + 21.564724061% × 7.014824948 = USD 3.820339887. The Base case controls the binary gate and fails.

Synthesis conclusion

The combined model is arithmetically coherent and deliberately demanding. Base revenue compounds at 27.69% for five years, statutory operating margin expands to 57.55%, and common NI reaches USD 9.676 billion. Even after that performance, modeled dilution leaves Base EPS at only USD 3.7247. The P/E 15 entry-price value is USD 55.87, 68.95% below the sealed USD 179.94 quote. Bull requires USD 32.798 billion of revenue and USD 21.521 billion of statutory OI, yet its USD 105.22 value at 15 times EPS also fails the gate by 41.52%.

Management FAIL and moat/2046 YELLOW strengthen the caution already visible in the price test: the investor is being asked to underwrite exceptional operating outcomes while accepting unresolved governance, dilution, competitive, and duration risks. This synthesis makes no verdict or return claim beyond the required binary gate. Subsequent valuation work must independently choose and validate its primary method, denominator, capital-structure treatment, and scenario exit multiples.

SYNTHESIS CLEARED — 11/11 TIE-OUTS PASSED

Appendix CS - Compounder Score

Palantir Technologies Inc. (PLTR) — Compounder Score and Security-Return Underwriting

Analysis date: 2026-08-24 SGT **Security:** XNAS:PLTR Class A common stock **Reporting and quote currency:** USD **Price anchor:** USD 179.94, completed 2026-08-21 session **Operating anchor:** TTM ended 2026-06-30 **Exact horizon:** 5.0 years, ending 2031-06-30 **Primary valuation method:** P/E **Authority boundary:** diagnostic research only. Trade execution, ranking, sizing, broker action, messages, portfolio mutation, BUY-queue placement, and ER2 bypass are not authorized.

1. Executive Verdict

Compounder Score: 55/100. Business Quality Score: 53/80. 20% Return Underwriting Score: 2/20. Score-implied classification: Reject. Confidence: Medium.

Palantir is an exceptional current operating performer but not an acceptable security at the governed price under this rubric. The evidence supports a large reinvestment runway, unusually strong recent organic growth, expanding segment economics, a debt-free balance sheet, and genuine switching costs in high-consequence deployments. It does not support dream-investment status because normalized returns on capital have only recently become positive, a complete acquisition-inclusive ROIIC series is unavailable, Commercial durability and direct pricing power remain unproven through a normal cycle, chronic stock issuance diluted outside owners, and the Management gate is a mandatory **FAIL**. That FAIL is not softened: perpetual founder voting control, separation of control from economic ownership, weak buyback discipline, and compensation without ROIC or per-share FCF hurdles are structural minority-owner defects.

The security return is worse than the operating story. The validator-cleared preflight locks one P/E method and exit P/Es of 25x / 35x / 45x for Bear / Base / Bull. Those multiples were chosen before return calculation and are far below Palantir's comparable normalized current P/E of 203.106162x. Even so, zero-distribution terminal shareholder wealth is only USD 39.147601, USD 130.364018, and USD 315.667123 per share. The resulting exact five-year gross USD IRRs are **-26.29%, -6.24%, and 11.90%**. Probability-weighted terminal wealth is USD 144.386481 per share and its CAGR is **-4.31%**. This is not an average of scenario IRRs.

The 20% annualized hurdle requires USD 447.748301 of terminal wealth per share. At the locked Base 35x exit P/E and zero distributions, that requires terminal EPS of USD 12.792809, versus governed Base EPS of USD 3.724686. The Base-case maximum purchase price for a 20% IRR is USD 52.390375, 70.88% below USD 179.94. Even Bull reaches only 11.90% gross USD. The price therefore capitalizes a degree of continued operating and multiple perfection that the governed scenarios do not earn.

Preflight and evidence gate

The controlling artifact is `C:\Users\Admin\OneDrive\GPT\ER 2.0\Palantir Technologies Inc.\valuation_preflight.json`, SHA-256 `a49f6d6fc7dbdceec4c7cb6bcadb2a78c9faf24c20f6f63e9e443aca27731da1`. The ER2 validator returned `VALUATION_PREFLIGHT_VALID` with evidence-file checks. Status is `CLEARED`; all five foundation flags pass; blockers are empty. The primary method is `PE`; denominator integrity, capital-structure inventory, dilution treatment, independent exit-multiple selection, excessive-optimism challenge, excessive-bearishness challenge, and double-count checks all pass.

The price is hash-bound to `quote_packet.json`, SHA-256 `1a87c7195d065cbbb11be9896951d0c60eda8a55e1e8f5eff53b437a9ff7010b`; the raw quote SHA-256 is `3d0abdd4c04dba6a52c094d9ddb8a932cf7b4cff252d24ea459be86ab0e04d0b`. The governed synthesis SHA-256 is `de2463238d5b093b8400180d633c54963257af0a24684dc9b92b161a7e89fa5b`.

2. Data Source Map

Evidence need	Controlling source	Period / currency	Use and confidence
Price and listed security	`quote_packet.json`; live FMP historical-price raw file	Completed 2026-08-21; USD	XNAS:PLTR Class A at USD 179.94; High
Market capitalization	Price times 2,568.694M quote-basis diluted shares	2026-08-21 price / 2026-06-30 shares; USD	USD 462.210798B diluted-equity diagnostic; High

Evidence need	Controlling source	Period / currency	Use and confidence
Enterprise value context	Q2 balance sheet plus calculated market capitalization	2026-06-30; USD	Approx. USD 453.013099B after cash, marketable securities and leases; secondary only because P/E controls
Income statement and statutory EPS	FY2021-FY2025 10-Ks; Q3 2025, Q1/Q2 2026 10-Qs; Q4 release; `statutory_bridge.json`	FY2021-TTM 2026-06-30; USD	SEC controls; eight QF reconciliations and statutory bridge pass; High
Balance sheet	Q2 FY2026 10-Q; `quant_foundation.md`	2026-06-30; USD	Cash, securities, leases, equity, NCI; High
Cash flow and capital allocation	FY2021-FY2025 10-Ks and live FMP cash-flow rows checked to filings	FY2021-TTM; USD	CFO, capex, FCF, SBC, repurchases, M&A; High
Share count and dilution	SEC standalone-quarter diluted denominators, quote packet, synthesis share schedule	TTM and 2031; millions	Official share counts override FMP cumulative-denominator defects; High
SBC	SEC cash-flow and expense bridge; option/SAR/RSU schedules	FY2021-TTM and modeled 2031; USD	Treated as expense and explicit dilution, once each; High
Segment data	SEC Contribution tables; reviewed Government and Commercial appendices	FY2021-TTM and 2031; USD	Government / Commercial Contribution, never relabeled operating income; High
Management commentary	Eight live FMP transcripts reviewed in `appendix_management.md`	Q3 FY2024-Q2 FY2026	Guidance, disclosure and execution record; Medium-High
Compensation and governance	2026 proxy, 10-K multi-class disclosures, live Form 4 review	2025-2026	Management FAIL; High
TAM / runway	Segment appendices, company filings, agency procurement sources, larger-software peer scale	Current and five-year	No single fabricated TAM; assessed from addressable workflows and scale; Medium
Competitors and moat	SEC risk disclosures, GAO procurement evidence, Microsoft/AWS/Snowflake/ServiceNow primary materials	Current	Bundling, open architecture, internal build and workflow competition; Medium-High
USD/SGD overlay	Live FMP USDSGD EOD file	2026-08-21 close 1.26976 SGD/USD	Base FX anchor; SHA-256 `3edb508d65f09df8c169f3579de4bd07ab4a8d2d4f7b0260cd0e260b471b8302`; High for spot, Low for five-year stress

Five completed annual years plus TTM are available in the cleared package. A comparable ten-year public-company series is not available, so all ten-year cells are explicitly `Insufficient evidence`; no stale cache or invented pre-listing series is used.

3. Key Financial and Quality Metrics

Metric	1-year	3-year	5-year / longest available	10-year	Interpretation
Revenue CAGR	56.15% FY2024-25	32.91% FY2022-25	30.52% FY2021-25 over four years	Insufficient evidence	Strong organic scale, but current 92.8% quarterly growth is not a mature anchor
NOPAT CAGR	355.92% FY2024-25	Not meaningful from negative FY2022 NOPAT	Not meaningful from negative FY2021 NOPAT	Insufficient evidence	Profit inflection makes CAGR unstable
FCF CAGR	84.06% FY2024-25	125.29% FY2022-25	59.91% FY2021-25 over four years	Insufficient evidence	Cash generation scaled rapidly
Simple ROIC	15.0% FY2025	Rising from negative to 15.0%	2.6% FY2021-25 median	Insufficient evidence	Uses NOPAT / year-end invested capital; not an adjusted average-capital ROIC
R&D-capitalized / lease-adjusted ROIC	Insufficient evidence	Insufficient evidence	Insufficient evidence	Insufficient evidence	Full capitalization schedule unavailable; score capped
ROIIC proxy	n/m	26.40% FY2022-25	28.93% FY2021-25 four-year proxy	Insufficient evidence	Change in NOPAT / change in invested capital; directional, not a complete reinvestment ledger
Acquisition-inclusive ROIIC	Similar direction because M&A was negligible	Incomplete	Incomplete	Insufficient evidence	FY2022 cash M&A USD 0.067B; exact working-capital and capitalized-R&D ledger missing

Metric	1-year	3-year	5-year / longest available	10-year	Interpretation
FCF / normalized NOPAT	188.06% FY2025	Positive and rising	161.35% TTM	Insufficient evidence	Working-capital timing and SBC add-back mean this is not pure owner conversion
CFO / net income	131.35% FY2025	Positive	112.72% TTM	Insufficient evidence	Statutory earnings are cash-supported
Gross margin	82.4% FY2025	80.2%-82.4% FY2023-25	78.0%-82.4% FY2021-25; 84.8% TTM	Insufficient evidence	Strong scale economics
Statutory operating margin	31.6% FY2025	5.4%-31.6% FY2023-25	(26.7%)-31.6% FY2021-25; 42.8% TTM	Insufficient evidence	Only three profitable annual years
Net debt / EBITDA	(4.13x) FY2025	Net cash	USD 9.198B cash less leases at Q2 FY2026	Insufficient evidence	Balance sheet is exceptionally strong
Diluted-share CAGR	4.67% FY2024-25	7.52% FY2022-25	7.46% FY2021-25 over four years	Insufficient evidence	33% cumulative rise from FY2021 to FY2025
SBC / revenue	15.29% FY2025	Elevated	13.57% TTM	Insufficient evidence	Improving ratio, large absolute cost
SBC / FCF	32.56% FY2025	Elevated	24.88% TTM	Insufficient evidence	FCF still includes noncash SBC add-back
EV / normalized NOPAT	n/a	n/a	Approx. 217.65x current diagnostic	n/a	Secondary only; P/E is locked primary method
Market cap / FCF	n/a	n/a	137.63x TTM	n/a	FCF yield only 0.73%
Dividend yield	0.00%	0.00%	0.00%	n/a	No dividend is currently paid or underwritten
Net buyback yield	Negative economic direction	Negative economic direction	Not precisely measurable from disclosed issuance proceeds	n/a	Repurchases far below SBC and were terminated

The most important normalization is SBC. It remains in GAAP operating expenses and therefore reduces statutory EPS. The forward share schedule also captures aggregate future award issuance. That is not double counting: one line measures compensation cost in earnings, while the other measures ownership dilution. What would be double counting is separately subtracting the SBC expense again from terminal equity value or adding award dilution on top of the sealed share schedule; neither occurs.

4. Rubric Scorecard

Category and test	Score	Maximum	Evidence-based judgment
1A Normalized ROIC high	1	5	Five-year simple median is only 2.6%; adjusted series unavailable
1B ROIIC high	6	8	3-year invested-capital proxy is 26.40%, but incomplete reinvestment ledger caps confidence
1C ROIIC not deteriorating sharply	3	3	Profit and incremental returns improved materially
1D Unit economics support returns	3	3	Both segments exceed 70% TTM Contribution margin
1E Cash conversion is real	2	3	Strong CFO/FCF, tempered by SBC add-back and working-capital timing
1F ROIC not an accounting artifact	1	3	GAAP profit is real, but low capital intensity, excess cash and no R&D capitalization limit comparability
Incremental returns subtotal	**16**	**25**	Good incremental economics; insufficient long-cycle proof
2A Large remaining opportunity	5	5	Government and complex commercial workflows are large and underpenetrated
2B Proven reinvestment at scale	4	5	Organic revenue, margin and product expansion are strong
2C Organic growth opportunity	4	4	M&A is negligible; growth is built internally
2D Repeatable expansion playbook	2	3	NRR and customers support it; bootcamp cohort economics are undisclosed

Category and test	Score	Maximum	Evidence-based judgment
2E Market structure supports runway	1	2	Demand is strong, but bundling and buyer power are material
2F Growth without ROIC collapse	1	1	Recent growth coincided with improving profitability
Reinvestment runway subtotal	**17**	**20**	The strongest part of the score
3A Clear source of advantage	3	4	Ontology, governance, deployment and trust are identifiable
3B Pricing power	1	3	Indirect expansion evidence; no price/volume or churn proof
3C Competitive evidence	2	3	Strong wins and margins, but powerful bundled alternatives
3D Customer stickiness	2	2	Production embedding creates real switching costs
3E Resilience through stress	1	2	Limited mature-cycle evidence; termination rights remain
3F Low disruption risk	0	1	AI lowers rival component-build costs; Commercial is exposed
Durability subtotal	**9**	**15**	Consistent with Moat YELLOW and 2046 YELLOW
4A Clean revenue / earnings quality	2	2	SEC bridges reconcile and one-offs are normalized
4B Cash flow supports earnings	2	2	CFO and FCF exceed net income/NOPAT
4C Conservative SBC / dilution	0	2	33% diluted-share increase and chronic award overhang
4D Balance sheet strength	2	2	USD 9.409B cash/securities, no financial debt
4E Transparent disclosures	1	1	Filings are strong, though key SaaS cohort metrics are absent
4F No governance/accounting red flags	0	1	Founder control and related-party optics are material
Accounting / financial quality subtotal	**7**	**10**	Accounting is credible; owner-quality caveats remain
5A Rational reinvestment priority	2	2	Organic R&D and platform build are coherent
5B Acquisition discipline	2	2	Minimal M&A; SPAC-linked program was discontinued
5C Buyback discipline	0	2	Buybacks did not offset issuance and authorization ended
5D Incentive alignment	0	2	Awards lack ROIC, FCF/share and relative-return hurdles
5E Owner mentality	0	1	Control/economics separation fails minority-owner standard
5F Past decisions	0	1	Strong operations do not cure dilution and strategic-investment error
Capital allocation subtotal	**4**	**10**	Management FAIL is fully reflected, not averaged away
Business Quality Score	**53**	**80**	Promising operating quality, below dream-business threshold
6A Base-case IRR	0	6	-6.24%, below 10%
6B Conservative exit multiple	2	4	Exits are independently justified and compress sharply, but 20% still does not clear
6C Bear case acceptable	0	3	-26.29% annualized is unacceptable
6D Reverse-return hurdle beatable	0	3	Base EPS must be 3.43x the governed Base estimate
6E Multiple does not price perfection	0	2	Current normalized P/E is 203.11x

Category and test	Score	Maximum	Evidence-based judgment
6F Cash returns help	0	2	Zero distributions and zero buybacks underwritten
20% Return Underwriting Score	**2**	**20**	Current price does not support the stretch return hurdle
Compounder Score	**55**	**100**	Reject under this diagnostic rubric

5. ROIIC and Reinvestment Analysis

Palantir's conventional ROIC history is not yet dream-company evidence. Using statutory operating income after a normalized 21% tax divided by year-end invested capital, ROIC moved from -13.2% in FY2021 and -4.7% in FY2022 to 2.6% in FY2023, 4.7% in FY2024 and 15.0% in FY2025. The five-year median is only 2.6%. The latest TTM operating-profit acceleration suggests further improvement, but a point estimate during hypergrowth is not a through-cycle median.

The directional incremental-return evidence is better. NOPAT rose from negative USD 0.127B in FY2022 to USD 1.117B in FY2025 while invested capital increased from USD 2.723B to USD 7.435B. The resulting change-in-NOPAT / change-in-invested-capital proxy is 26.40%. From FY2021 to FY2025 the four-year proxy is 28.93%. Those values justify a 6/8 ROIIC score, but they are not mislabeled as the rubric's ideal cumulative-reinvestment calculation. A complete series for non-cash working capital, capitalized R&D amortization, growth opex, dispositions and acquisitions is not available in the cleared package. Acquisition-inclusive ROIIC should be similar because cash M&A was only USD 0.067B in FY2022 and otherwise negligible, but exactness would be false precision.

Economic reinvestment is overwhelmingly organic. Capex was below 1% of revenue in four of five years; R&D rose from USD 0.387B to USD 0.558B while falling from 25.1% to 12.5% of revenue. Customer deployment, cloud capacity, product engineering and forward-deployed talent are the real growth inputs. Revenue nearly tripled from FY2021 to FY2025 and the TTM statutory operating margin reached 42.8%. This is credible evidence that incremental spending has created value. The cap is that some reinvestment is expensed, some labor is paid in equity, and the five-year record starts in losses, making standard ratios unusually sensitive to accounting perimeter and endpoint choice.

6. Runway Analysis

The runway case is strong without relying on one management TAM. Government can expand across defense, intelligence, logistics, health, space, cyber, allied sovereign and civilian workflows. Commercial can expand across manufacturing, energy, healthcare, finance, telecom, transport and supply-chain operations. The forward model begins from only USD 6.156B of TTM revenue, while much larger enterprise vendors demonstrate that durable software categories can support tens of billions of dollars of annual revenue.

Evidence of the playbook is current rather than hypothetical: TTM Government and Commercial grew to USD 3.211B and USD 2.945B; Q2 growth was 79.0% and 109.8%; NRR rose to 157%; U.S. Commercial customers rose from 530 in Q3 FY2025 to 653 in Q2 FY2026; RPO reached USD 4.9B. The caveats matter. NRR is companywide, RDV includes options and assumes no termination, RPO is commercially biased, no bootcamp conversion cohort or payback is disclosed, and the customer can terminate many contracts for convenience. The Base forecast already assumes 27.69% consolidated revenue CAGR for five years. That is demanding, not a conservative extrapolation.

7. Moat Analysis

The moat is the installed operating context: customer-specific Ontology, permissions, governed data, applications, approval logic, deployment across cloud/on-premise/edge/classified environments, and institutional trust. Switching a mature production deployment means reconstructing semantics and validating operating behavior, not just exporting tables. Government adds classified capability, mission experience, procurement vehicles and high-trust references.

The moat is nevertheless YELLOW. Commercial is nearly half of revenue and faces Microsoft, ServiceNow, Snowflake, Databricks, hyperscalers, integrators and internal-build alternatives with deeper distribution. Foundation models commoditize generic analytics and agent construction. Pricing power is inferred from NRR and margin rather than proven with price/volume, gross retention or churn data. The 2046 case depends on Palantir retaining the ontology and action layer after models and infrastructure commoditize. A credible future profit pool exists, but a 20-year protected control point is not yet proven.

8. Accounting and Capital Allocation Analysis

Accounting quality is better than governance quality. The FY2025 and TTM segment-to-statutory bridges reconcile at zero residual. Current statutory EPS of USD 1.173334 is explicitly normalized to USD 0.885941 by zeroing volatile other-income marks and applying a 21% tax rate. FCF is strong, the balance sheet has USD 9.409B of cash and marketable securities and no financial debt, and goodwill/intangibles are zero. These facts reduce insolvency and acquisition-accounting risk.

The owner-economics weaknesses are material. SBC was USD 0.836B TTM, 13.57% of revenue and 24.88% of FCF. Diluted shares increased 33% from FY2021 to FY2025. FY2024-FY2025 buybacks of USD 139M were far smaller than USD 1.376B of SBC, and the authorization was terminated. Founder control can persist as economic ownership declines and may concentrate in one surviving founder. Compensation is tied mainly to service and absolute share-price outcomes rather than ROIC or per-share FCF. Management's operational achievement is real; outside-owner stewardship does not meet the Gideon bar.

9. 20% Return Underwriting

Valuation denominator integrity

Reported TTM statutory diluted EPS is USD 1.173333934905017, making reported current P/E 153.35787591839x. Normalized TTM statutory diluted EPS is USD 0.885940624096913, making normalized comparable current P/E 203.106162090064x. The normalization removes USD 170.794M of volatile other income and substitutes a 21% tax rate for the anomalous 1.34% TTM effective tax rate. The terminal denominators are statutory diluted EPS after scenario tax, other income, NCI and explicit dilution. The preflight concludes current normalized and terminal denominators are comparable; the raw reported multiple is disclosure only and was not used as a cap.

Security return underwriting (mandatory even though the classification is Reject):

Measure	Result	Notes
Primary valuation method	P/E	Same audited method for all scenarios, probability weighting, reverse return and net-SGD
Current price / horizon	USD 179.94 / 5.0 years to 2031-06-30	XNAS:PLTR completed 2026-08-21 session
Scenario probabilities	28.435275939% / 50.000000000% / 21.564724061%	Bear / Base / Bull; sum 100%
Terminal statutory diluted EPS	USD 1.565904027 / 3.724686226 / 7.014824948	Governed synthesis
Terminal diluted shares	2,610.088747M / 2,597.849227M / 2,592.215058M	Explicit SBC issuance, zero buybacks
Exit P/E	25.00x / 35.00x / 45.00x	Independently justified and preselected
Terminal equity value	USD 102.178712B / 338.666063B / 818.277069B	Common NI times exit P/E
Terminal equity value per share	USD 39.147601 / 130.364018 / 315.667123	Terminal EPS times exit P/E
Distributions per share	USD 0.000000 / 0.000000 / 0.000000	No dividend paid or source-bound distribution underwritten
Total terminal shareholder wealth	USD 39.147601 / 130.364018 / 315.667123	Equity value/share plus distributions once
Gross USD Bear / Base / Bull IRR	**-26.29% / -6.24% / 11.90%**	Exact five-year CAGR of terminal wealth / price
Expected terminal shareholder wealth	USD 144.386481	Probability-weighted wealth, not a fourth scenario
Probability-weighted gross USD IRR	**-4.31%**	CAGR of expected wealth using exact scenario weights
Net cash/debt count-once treatment	Current net cash after leases USD 9.197699B; terminal standalone adjustment USD 0	Interest/lease effects remain in EPS; no cash addition or debt subtraction after EPS x P/E

All per-share, percentage and multiple values below are reader-facing two-decimal displays; exact values remain in the companion JSON. Operating dollars are USD billions and shares are millions. Net-SGD adjustments are separate.

Bear-case operating and return bridge

Bear case	Start (2026-06-30)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.16B	USD 12.75B	15.68%	Government and Commercial decelerate materially; no top-down plug

Bear case	Start (2026-06-30)	Horizon (2031-06-30)	CAGR / change	Justification
Operating margin	42.80%	42.81%	+0.01 pp	Lower segment margins and 24% corporate-cost ratio prevent leverage
Operating earnings	USD 2.63B	USD 5.46B	15.69%	Statutory operating income after R&D, G&A and SBC
Diluted share count	2,568.69M	2,610.09M	0.32%	41.39M aggregate SBC shares; zero buybacks
EPS	USD 1.17	USD 1.57	5.94%	Statutory diluted EPS; start uses official TTM weighted denominator
Valuation denominator	Reported EPS 1.17; normalized EPS 0.89	Statutory EPS 1.57	Comparable after normalization	Other-income mark and anomalous tax removed from current denominator
Valuation	153.36x reported / 203.11x normalized P/E	25.00x P/E	(178.11x) vs normalized	Peer lower end, governance/moat/dilution haircut; not set by target IRR
Share value	USD 179.94	USD 39.15	(26.29%) CAGR	EPS 1.565904027 x 25x; no separate net-cash add
Annual dividend	USD 0.00	USD 0.00	0.00%	No current dividend and zero source-bound distribution path
Dividend yield	0.00%	0.00%	0.00 pp	Annual dividend divided by endpoint value
Five-year dividends	-	USD 0.00	cumulative	Counted once at zero
Total ending value	USD 179.94	USD 39.15	**-26.29% gross USD IRR**	Share value plus zero distributions; terminal equity USD 102.18B

Base-case operating and return bridge

Base case	Start (2026-06-30)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.16B	USD 20.90B	27.69%	U.S. Commercial and Government compound rapidly but decelerate annually
Operating margin	42.80%	57.55%	+14.75 pp	Segment Contribution leverage and lower R&D/G&A/SBC ratios
Operating earnings	USD 2.63B	USD 12.03B	35.48%	Statutory operating income after all corporate bridge costs
Diluted share count	2,568.69M	2,597.85M	0.23%	29.16M aggregate SBC shares; zero buybacks
EPS	USD 1.17	USD 3.72	25.99%	Statutory diluted EPS after 21% tax and NCI attribution
Valuation denominator	Reported EPS 1.17; normalized EPS 0.89	Statutory EPS 3.72	Comparable after normalization	Same earnings perimeter and listed-unit basis
Valuation	153.36x reported / 203.11x normalized P/E	35.00x P/E	(168.11x) vs normalized	Premium to mature peers, large maturity discount to current PLTR
Share value	USD 179.94	USD 130.36	(6.24%) CAGR	EPS 3.724686226 x 35x; no separate net-cash add
Annual dividend	USD 0.00	USD 0.00	0.00%	No distribution is source-bound
Dividend yield	0.00%	0.00%	0.00 pp	Annual dividend divided by endpoint value
Five-year dividends	-	USD 0.00	cumulative	Counted once at zero
Total ending value	USD 179.94	USD 130.36	**-6.24% gross USD IRR**	Share value plus zero distributions; terminal equity USD 338.67B

Bull-case operating and return bridge

Bull case	Start (2026-06-30)	Horizon (2031-06-30)	CAGR / change	Justification
Revenue	USD 6.16B	USD 32.80B	39.74%	Category leadership persists; terminal segment growth remains low-20s
Operating margin	42.80%	65.62%	+22.82 pp	Exceptional reuse, Contribution margins and corporate-cost leverage
Operating earnings	USD 2.63B	USD 21.52B	52.20%	Statutory operating income with R&D, G&A and SBC retained
Diluted share count	2,568.69M	2,592.22M	0.18%	23.52M aggregate SBC shares; zero buybacks
EPS	USD 1.17	USD 7.01	42.99%	Statutory diluted EPS after 18% tax and NCI attribution
Valuation denominator	Reported EPS 1.17; normalized EPS 0.89	Statutory EPS 7.01	Comparable after normalization	Same common-share and accounting perimeter
Valuation	153.36x reported / 203.11x normalized P/E	45.00x P/E	(158.11x) vs normalized	Premium for 23%-like terminal growth, below NOW/current PLTR extremes
Share value	USD 179.94	USD 315.67	11.90% CAGR	EPS 7.014824948 x 45x; no separate net-cash add
Annual dividend	USD 0.00	USD 0.00	0.00%	No distribution is source-bound
Dividend yield	0.00%	0.00%	0.00 pp	Annual dividend divided by endpoint value
Five-year dividends	-	USD 0.00	cumulative	Counted once at zero
Total ending value	USD 179.94	USD 315.67	**11.90% gross USD IRR**	Share value plus zero distributions; terminal equity USD 818.28B

Reverse-return and maximum-price diagnostic

Scenario exit P/E	Required terminal wealth for 20%	Required terminal EPS	Governed terminal EPS	Maximum current price for 20% IRR
Bear 25x	USD 447.748301	USD 17.909932	USD 1.565904	USD 15.732543
Base 35x	USD 447.748301	USD 12.792809	USD 3.724686	USD 52.390375
Bull 45x	USD 447.748301	USD 9.949962	USD 7.014825	USD 126.859537

The reverse hurdle is not beatable within the governed cases. Even Bull terminal EPS is 29.50% below the USD 9.949962 required for 20% at 45x. Base needs 3.43 times its modeled EPS. The problem is not that the locked exits are punitive: 35x Base remains a premium multiple for mature software and still compresses 83% from the normalized current P/E.

CIO actionability: net-SGD diagnostic only

The live FMP USDSGD close on the security price date, 2026-08-21, is 1.26976 SGD per USD. To avoid false precision about a five-year currency forecast, the diagnostic applies a transparent symmetric terminal stress: Bear USD weakens 10% to 1.142784; Base stays at 1.26976; Bull strengthens 10% to 1.396736. It applies a 0.50% one-time terminal-wealth haircut for aggregate brokerage and FX conversion costs. Withholding is zero because distributions are explicitly zero. Investor-specific capital-gains tax, wrapper, custody, financing and portfolio facts are unavailable; the calculation is therefore pre-personal-tax and constrains only actionability, not the governed USD return.

Measure	Result	Notes
Status	PASS	Management FAIL and negative Base/PW returns prevent actionability
Bear / Base / Bull net SGD IRR	**-27.90% / -6.34% / 13.94%**	Same P/E wealth, FX stress and 0.50% cost haircut
Probability-weighted net SGD IRR	**-3.66%**	CAGR of probability-weighted net terminal wealth; same weights
Probability-weighted net terminal wealth	SGD-equivalent USD 149.330180 per starting share	Wealth translated relative to current FX, not averaged IRRs
10% mandate hurdle	FAIL	Base and probability-weighted both fail
15% stretch hurdle	FAIL	Even net-SGD Bull is below 15%

Measure	Result	Notes
Downside and portfolio fit	UNAVAILABLE	Security downside fails; portfolio context is absent
Proposed sizing ceiling	none	No sizing, ranking, or portfolio action authorized

10. Hard Veto Checklist

Hard veto	Result	Evidence
1. ROIIC below cost with no path	Pass	Directional 3-year ROIIC proxy 26.40%; operating margins improving
2. ROIC high only from artificial capital	Watch	Current economics are asset-light; adjusted R&D/average-capital series unavailable
3. Growth needs external financing	Pass	USD 9.409B cash/securities, no financial debt, strong FCF
4. Material share growth without owner protection	**Fail**	33% diluted-share rise FY2021-FY2025; buybacks immaterial and terminated
5. Large acquisitions with undisclosed returns	Pass	M&A negligible; failed strategic-investment program discontinued
6. Structural FCF shortfall	Pass	FCF exceeds earnings and normalized NOPAT, though SBC is an add-back
7. Concentration/regulation/platform risk	Watch	Customer I 27% of receivables; government procurement and bundling risks
8. 20% requires multiple expansion	Pass mechanically	Locked exits are far below normalized current P/E; return still fails
9. Flawless execution needed at price	**Fail**	Base growth/margin excellence still produces negative return; Bull only 11.90%
10. Moat cannot be explained	Pass	Governed ontology, deployment, workflow and trust form an identifiable moat
11. Governance/accounting trust impaired	**Fail**	Perpetual founder control, control/economics separation and related-party optics
12. Balance sheet can wipe out equity	Pass	Net-cash balance sheet and no financial debt

11. Final Decision

Score-implied classification: Reject. The Business Quality Score of 53/80 recognizes exceptional current execution and runway but remains below the 55-point promising-quality band because normalized ROIC history is short, ROIIC evidence is incomplete, moat durability is YELLOW, SBC/dilution are material, and Management is a hard FAIL. The 20% Return Underwriting Score of 2/20 is decisive: Base and probability-weighted returns are negative and even Bull does not clear 15%, let alone 20%.

Compounder Score note: 55/100; the evidence does not support dream-investment status, and hard vetoes make the score-implied classification Reject.

Business Quality Score note: 53/80; operating quality and runway are strong, but normalized capital-return history, moat duration, governance and dilution are below the dream-business standard.

20% Return UW Score note: 2/20; governed Base IRR is -6.24%, probability-weighted gross USD IRR is -4.31%, and the 20% reverse hurdle requires Base EPS far beyond the model.

The ER2 decision gates remain Management **FAIL**, Moat **YELLOW**, 2046 **YELLOW**, and P/E-15 **FAIL**. Under the exhaustive verdict rule, any gate FAIL implies **PASS** for the overall ER2 verdict. This diagnostic score does not authorize or queue a purchase.

12. Machine-Readable JSON

The complete full-precision artifact is `C:\Users\Admin\OneDrive\GPT\ER 2.0\Palantir Technologies Inc.\compounder_score.json`. It binds the price, horizon, source hashes, validated preflight, method consistency, scenario probabilities, terminal EPS and shares, P/E equity bridge, zero distributions, gross USD IRRs, probability-weighted wealth CAGR, net-SGD overlay, reverse-return hurdle, score detail, hard vetoes, and trade-authority flags. All scenario returns are finite. P/E is used consistently; net cash/debt is counted through statutory EPS and never added or subtracted

separately; distributions are counted exactly once at zero.

Appendix CSA - Scoring Audit

Palantir Technologies Inc. (PLTR) — Independent Scoring Audit

Analysis date: 2026-08-24 **Audit status: SCORING AUDIT CLEARED Materiality conclusion: NO MATERIAL ISSUE.** No repair to the Compounder Scorer outputs is required. **Authority boundary:** This audit authorizes no ranking, sizing, trade, broker action, outbound message, portfolio mutation, BUY queue, or ER2 bypass.

The audit independently reconstructed the P/E denominator, locked exits, terminal equity values, per-share values, gross local-currency returns, probability-weighted wealth CAGR, net-SGD bridge, rubric score, hard vetoes, ER2 gates, and authority flags. It did not merely accept the scorer's JSON. The five required upstream hashes match the governed artifacts byte for byte, and the valuation-preflight validator returned `VALUATION_PREFLIGHT_VALID` with the stipulated SHA-256. The original score is appropriately skeptical and is preserved without revision.

1. Original vs Audited Scores

Item	Original	Audited	Notes
Business Quality Score / 80	53/80	53/80	Evidence supports the component scores; missing adjusted-return history is already penalized
Valuation / 20% Return UW Score / 20	2/20	2/20	Base and probability-weighted returns are negative; only independent multiple compression earns limited credit
Total Score / 100	55/100	55/100	Arithmetic is 53 + 2 = 55
Classification	Reject	Reject	Score is below 70 and three hard vetoes remain

The diagnostic classification is distinct from the exhaustive ER2 verdict. The score rubric says **Reject**. Separately, Management is **FAIL**, Moat is **YELLOW**, the 2046 durability test is **YELLOW**, and the P/E-at-15 entry-price gate is **FAIL**; under the ER2 gate rule, any FAIL produces an overall ER2 **PASS**, meaning do not buy. Neither label creates action authority.

Hash-bound inputs

Artifact	Recomputed SHA-256	Audit result
`appendix_compounder_score.md`	`443aa09dc4625a84410cbf4106f39370a42798198aa39a909d771153896867b7`	Exact match
`compounder_score.json`	`a6a733b3379a10d73e3ddf8d1e84fcac553f7caa88d30a1a61493c05d0109d01`	Exact match and valid JSON
`valuation_preflight.json`	`a49f6d6fc7dbdceec4c7cb6bcadb2a78c9faf24c20f6f63e9e443aca27731da1`	Exact match; status CLEARED; five foundations passed; zero blockers
`appendix_synthesis.md`	`de2463238d5b093b8400180d633c54963257af0a24684dc9b92b161a7e89fa5b`	Exact match
`quote_packet.json`	`1a87c7195d065cbbb11be9896951d0c60eda8a55e1e8f5eff53b437a9ff7010b`	Exact match; price USD 179.94

2. Score Reductions and Reasons

No further score reduction is justified. The scorer already reduced the business-quality result for the important evidentiary weaknesses rather than filling gaps with optimistic estimates. The 5-year median simple ROIC is only 2.6%, while a full R&D-capitalized, lease-adjusted, average-capital ROIC history is unavailable; this leaves 1A at 1/5 and 1F at 1/3. The reconstructed 3-year directional ROIIC is approximately 26.4%, but it is based on change in NOPAT divided by change in year-end invested capital, not a complete cumulative reinvestment ledger. The 1B award of 6/8 fits the rubric's 25%-35% band and is tempered by the missing ledger rather than presented as an ideal calculation.

The runway score of 17/20 is supported by organic growth, negligible acquisitions, rising NRR, customer expansion, and both reportable segments exceeding 70% TTM Contribution margin. It does not receive full credit because bootcamp cohort economics, CAC, payback, churn, gross retention, and platform-level financials are undisclosed. Moat receives 9/15, consistent with YELLOW: Government has meaningful installed-workflow and trust barriers, while Commercial pricing power

and durability have not been tested through a normal competitive cycle. Accounting quality receives 7/10 after zero points for conservative SBC/dilution treatment and governance cleanliness. Capital allocation receives only 4/10 because buybacks failed to offset issuance, the repurchase program ended, incentives omit ROIC and per-share FCF hurdles, and founder control is not aligned with outside owners.

The 2/20 return-underwriting score is also fair. Category 6A is zero because Base gross USD IRR is negative. Category 6C is zero because Bear loses 26.29% annually. Category 6D is zero because the Base 20% hurdle needs USD 12.792809 EPS versus USD 3.724686 modeled. Category 6E is zero at a 203.11x normalized current P/E, and 6F is zero because distributions and buybacks are explicitly zero. The limited 2/4 in 6B recognizes that all exits were independently selected before returns and are sharply below a comparable normalized current multiple. It does not imply that the 20% hurdle clears.

3. Unsupported Claims

The evidence does not support a complete 5-year adjusted ROIC, 5-year ROIIC, acquisition-inclusive ROIIC, 5-year reinvestment-rate series, or ten-year public operating record. It also does not support platform-level revenue or margins for AIP, Foundry, Gotham, or Apollo, customer-level CAC/payback/churn cohorts, or a normal-cycle proof of Commercial pricing power. These claims are not used to inflate the score: the structured artifact records the unavailable metrics as null, and the Markdown expressly calls the ROIIC calculation a directional proxy.

The statement that acquisition-inclusive ROIIC should point in a similar direction is reasonable but not a precise metric. Cash acquisition spend was only USD 0.067B in FY2022 and otherwise negligible in the reviewed period, so the missing acquisition-inclusive calculation is not material to Category 1 or Category 5. The scorer properly refrains from publishing a fabricated percentage.

4. Calculation Errors

No material calculation error was found. The visible recomputations below use the locked P/E method throughout.

Denominator and P/E replay

Item	Formula	Recomputed result	Published result
Reported TTM diluted EPS	USD 3,016.692M / 2,571.043M shares	USD 1.173333934905017	USD 1.173333934905017
Reported current P/E	USD 179.94 / USD 1.173333934905017	153.357875918390x	153.357875918390x
Normalized current EPS	((2,634.652 + 266.384) × 79% – 14.027) / 2,571.043	USD 0.885940624096913	USD 0.885940624096913
Normalized current P/E	USD 179.94 / USD 0.885940624096913	203.106162090064x	203.106162090064x
20% terminal wealth hurdle	USD 179.94 × 1.20^5	USD 447.748300800000	USD 447.748300800000
Base required EPS	USD 447.7483008 / 35x	USD 12.792808594286	USD 12.792808594286

The normalized denominator removes USD 170.794M of volatile other income and replaces the anomalous 1.34% TTM effective tax rate with 21%. It is therefore the relevant current comparator to terminal statutory diluted EPS. The preflight locked Bear/Base/Bull P/Es at 25x/35x/45x before IRR or scoring. The raw 153.36x reported multiple was not used as a cap. The evidence anchors—CRM/MSFT at the lower end, ServiceNow at the upper end, plus growth, margin, governance, moat, maturity, and dilution adjustments—support the band. The optimism challenge prevents preservation of Palantir's extreme present/historical premium; the bearishness challenge prevents forcing a still-profitable, growing software franchise to the separate 15x entry gate. Both directional tests pass.

P/E terminal-equity and gross-return replay

Scenario	Weight	Terminal EPS	Exit P/E	Terminal value/share	Terminal diluted shares	Terminal equity value	Gross USD IRR
Bear	28.435275939%	1.565904027	25x	USD 39.147600675	2,610.088747M	USD 102.178711994B	-26.2918541213075%
Base	50.000000000%	3.724686226	35x	USD 130.364017910	2,597.849227M	USD 338.666063156B	-6.242503096994723%
Bull	21.564724061%	7.014824948	45x	USD 315.667122660	2,592.215058M	USD 818.277068675B	11.897485420845944%

For each row, value per share equals terminal EPS times exit P/E; terminal equity value equals that per-share value times terminal diluted shares; and gross IRR equals `(terminal value per share / 179.94)^(1/5) - 1`. Cumulative distributions are zero in every scenario and are included exactly once at zero. The P/E count-once rule is respected: current cash and marketable securities are USD 9.409099B, financial debt is zero, lease liabilities are USD 0.2114B, and current net cash after leases is USD 9.197699B. No terminal cash is added, and no debt or lease claim is subtracted after EPS times P/E. Interest, lease expense, NCI attribution, SBC expense, and dilution already enter the statutory per-share earnings bridge.

Probability-weighted wealth replay

Component	Full-precision calculation	Result
Expected USD terminal wealth/share	0.28435275939 × 39.147600675 + 0.50 × 130.36401791 + 0.21564724061 × 315.66712266	USD 144.3864811833615
Probability-weighted gross USD CAGR	(144.3864811833615 / 179.94)^(1/5) – 1	-4.307089631320182%
Arithmetic mean of scenario IRRs	Not used	Not applicable

The probability-weighted return is correctly the CAGR of probability-weighted terminal wealth, never an arithmetic average of scenario IRRs.

Net-SGD and cost bridge

Starting wealth is USD 179.94 translated at spot USDSGD 1.26976, or SGD 228.4806144. Each terminal USD wealth figure is translated at the locked scenario FX rate and multiplied once by the 0.995 aggregate transaction/FX-cost factor. With zero distributions, withholding is zero. No personal capital-gains tax is invented because the investor wrapper and tax status are unavailable.

Scenario	Terminal USDSGD	Net terminal wealth SGD	Recomputed net-SGD IRR
Bear	1.142784	44.513565431330306	-27.90110565751335%
Base	1.269760	164.7033603044946	-6.336448673412931%
Bull	1.396736	438.6991160644596	13.936658980179972%
Probability weighted	Scenario weighted	189.6134891502802	-3.660591751132525%

The net-SGD overlay uses the same P/E terminal wealth and does not change exit multiples, share counts, distributions, or capital structure. The method-consistency, same-method net-SGD, and double-count flags are all correctly true.

5. Missing Data

The most important missing evidence is a complete capitalized-R&D, non-cash-working-capital, growth-opex, and acquisition-inclusive reinvestment ledger. Other material gaps are the short public history, absent platform-level revenue and margin disclosure, missing customer cohort economics, limited long-run NRR history, and no investor-specific tax, broker, wrapper, or portfolio context. These gaps support Medium—not High—overall confidence. They do not invalidate the finite security-level gross or net-SGD diagnostics.

6. Hard-Veto and Threshold Test Audit

Test	Audited outcome	Reason
Share count rises materially	FAIL	FY2021-FY2025 diluted shares rose about 33%; four-year CAGR is 7.46%
Flawless execution needed at today's price	FAIL	Base operating excellence still produces -6.24%; Bull is only 11.90% gross USD
Governance or related-party issues impair trust	FAIL	Perpetual founder control, control/economics separation, chronic dilution, and related-party optics
Total score at least 80/100	FAIL	55/100
Business Quality at least 65/80	FAIL	53/80
Return UW at least 15/20	FAIL	2/20
ROIIC plus runway at least 35/45	FAIL	33/45
Base IRR at least 20%	FAIL	-6.242503096994723%
Bear IRR at least 8%-10%	FAIL	-26.2918541213075%

Test	Audited outcome	Reason
No hard vetoes	FAIL	Three hard vetoes remain
No terminal multiple expansion	PASS	25x/35x/45x are below 203.106162x normalized current P/E

7. Final 20% Hurdle Pass/Fail

FAIL. The Base maximum buy price for a 20% five-year IRR is USD 52.39037499598124, versus USD 179.94. The 20% case requires Base EPS of USD 12.792808594285711 at 35x, 3.43 times the governed Base terminal EPS. Even the Bull gross USD return is only 11.897485420845944%. The evidence cannot support a dream investment at the locked price.

8. Dream Business / Dream Investment Classification

Neither. Palantir is an exceptional current operating performer with substantial runway, but 53/80 remains below the rubric's dream-business threshold. The investment is decisively not a dream investment because the starting valuation overwhelms strong operating growth. The single biggest reason is the USD 179.94 price against terminal EPS and independently supported exit P/Es: even demanding Base execution destroys shareholder wealth on an annualized basis.

9. Stage 7 Scorecard Notes

- **Compounder Score note:** 55/100; strong recent execution does not overcome incomplete through-cycle quality evidence, hard vetoes, and an unacceptable starting price.
- **Business Quality Score note:** 53/80; runway and current margins are strong, but adjusted return history, moat duration, governance, and dilution remain below dream-business quality.
- **20% Return UW Score note:** 2/20; Base is -6.24%, probability-weighted gross USD is -4.31%, and Bull is only 11.90%.

10. Auditor JSON

The full-precision machine-readable audit is in `compounder_score_audit.json`. It preserves the audited scores, return fields, source hashes, preflight status, method-consistency and double-count checks, exit-multiple integrity, hard-veto outcomes, net-SGD bridge, and all negative authority flags.

SCORING AUDIT CLEARED — ORIGINAL 55/100 SCORE, P/E RETURN BRIDGE, AND AUTHORITY FLAGS INDEPENDENTLY VERIFIED.

Appendix RT - Red Team

Palantir Technologies Inc. (PLTR) — Independent Red Team

Analysis date: 2026-08-24 **Information cutoff:** 2026-08-24 SGT **Primary valuation method challenged:** P/E **Audited current price:** USD 179.94 **Operating horizon:** 2026-06-30 to 2031-06-30, exactly 5.0 years **Mandate boundary:** This appendix authorizes no ranking, sizing, trade, broker action, outbound message, portfolio mutation, BUY queue, or ER2 bypass.

The governed data package, reviewed Government and Commercial appendices, synthesis, valuation preflight, scorer, and scoring audit all hash-match the supplied SHA-256 values. The valuation-preflight validator independently returned `VALUATION_PREFLIGHT_VALID` at SHA-256 `a49f6d6fc7dbdceec4c7cb6bcadb2a78c9faf24c20f6f63e9e443aca27731da1`. The red team therefore attacks the sealed assumptions and their economic meaning; it does not silently replace audited return fields.

Executive attack

The strongest bear thesis is not that Palantir suddenly stops growing. It is that the current AI adoption shock pulls several years of demand forward while hyperscalers and workflow incumbents make the model/data/context/governance stack sufficiently complete, portable, and bundle-priced that Palantir loses control of the action layer. Commercial NRR then normalizes faster than modeled, bootcamps require persistent forward-deployed effort, and Contribution margin loses its assumed leverage. In Government, modular-open-system rules, government-owned data rights, volume discounts, and multi-vendor procurement reduce lock-in and turn a portion of Palantir's expanding footprint into lower-priced integration scope. Revenue can still compound, yet per-share value can fail because the market currently prices far more than even the governed Bull economics deliver.

This attack reinforces, rather than overturns, the published conclusion. The model is already a do-not-buy result: Base gross USD IRR is negative 6.24%, probability-weighted gross USD IRR is negative 4.31%, and even Bull produces only 11.90%. The most plausible adverse corrections make those returns worse. The strongest conservative shortcut—zero distributions and no separate net-cash addition under P/E—has only an immaterial quantified effect using the evidence available. No supportable defect was found that makes the published conclusion insufficiently conservative.

`RETURN_MODEL_OBJECTION_STATUS: IMMATERIAL`

1. The most fragile assumption

The most fragile assumption is Commercial's expansion-led Base path: USD 2.945 billion of TTM Commercial revenue becomes USD 11.883 billion in five years, a 32.18% CAGR, while Commercial Contribution margin rises from 72.5% to 75.0%. U.S. Commercial alone rises from USD 2.263 billion to USD 10.302 billion. The appendix uses companywide NRR of 157%, U.S. Commercial customer additions, and current guidance to support that trajectory, but the critical evidence is not segment-pure. Palantir does not disclose Commercial-only NRR, gross retention, churn, ARR, cRPO, billings, CAC, payback, bootcamp conversion, or cohort expansion. Companywide NRR may benefit from exceptionally strong Government expansion and cannot prove the Commercial cohort economics required by the Base model.

The fragility is magnified by competition at the exact layer Palantir claims as its durable control point. Microsoft reported that more than 15,000 customers use both Foundry and Fabric, and it explicitly describes Fabric/Foundry/Agent 365 as data-context, orchestration, evaluation, governance, identity, security, and agent-control layers. ServiceNow is likewise pushing AI Control Tower, Workflow Data Fabric, forward-deployed engineering, and autonomous workflows into an installed enterprise base. These are not claims that the products are identical to AIP. They are evidence that competitors do not need full feature parity: a bundled "good enough" context-and-action layer can lower Palantir's price, scope, or expansion rate. Sources: [Microsoft FY2026 Q3 call](https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q3) and [ServiceNow Q2 2026 results](https://investor.servicenow.com/news/news-details/2026/ServiceNow-Reports-Second-Quarter-2026-Financial-Results/default.aspx).

The synthesis sensitivity shows that a 200-basis-point change in Government CAGR moves Base EPS by only about USD 0.12, but it does not publish a Commercial growth sensitivity even though Commercial is the more fragile segment. A mechanical 200-basis-point Commercial CAGR reduction, before considering margin interaction, would not alone close the valuation gap; the real risk is correlation. Lower expansion reduces revenue while also keeping sales, deployment, hosting,

and technical labor elevated. The Base path assumes the opposite joint outcome: rapid growth and better margins.

2. The most-likely-to-be-optimistic Stage 5 number

The single Stage 5 number most likely to be revised down is the **57.55% Base statutory operating margin in 2031**. It requires Government and Commercial Contribution margins of 73.96% and 75.00%, R&D excluding SBC to fall from 7.69% to 5.5% of revenue, G&A excluding SBC from 7.27% to 4.5%, and SBC from 13.57% to 7.0%. Every component moves favorably while revenue compounds at 27.69%. That is possible, but it assumes productization overwhelms cloud compute, customer cost-certainty commitments, security/accreditation work, implementation labor, and incumbent price competition.

The margin is also inconsistent with the moat appendix's own long-run caution unless Palantir retains the ontology/action control point. That appendix describes a plausible mature GAAP margin of roughly 30%-40% and a moat-erosion outcome below 25%. The five-year Base is earlier and faster-growing than the 2046 steady state, so it need not equal those values. Still, 57.55% is an aggressive peak-like terminal margin, not a conservative normalized endpoint.

As a non-controlling sensitivity, holding Base revenue, interest, tax, NCI, shares, and 35x P/E fixed but lowering terminal statutory operating margin to 50% produces about USD 3.2478 of EPS, USD 113.67 of terminal value, and a negative 8.78% gross USD IRR. That is 2.53 percentage points below the audited Base return. This is a material outcome sensitivity, but not a model correction: the locked Base margin is explicitly disclosed, arithmetically reconciled, and supported by the current scaling evidence. The sensitivity shows where optimism resides; it does not establish that the audited return bridge is invalid.

3. The strongest bear case, stated in its strongest form

Palantir is priced as if it will own the governed operating layer of enterprise and sovereign AI, but its competitors are building the same economic control points—data semantics, orchestration, evaluation, identity, policy, security, and action—inside platforms customers already buy. Within two years, Commercial NRR can fall from 157% toward 110%, customer additions can slow, and Contribution margin can retreat because production conversion remains labor- and cloud-intensive. At the same time, DoD open-architecture and data-rights initiatives can force portable interfaces and multi-vendor scope, while large procurement vehicles exchange faster ordering for volume discounts and buyer power. Palantir would remain a strategically relevant, profitable software vendor, but not the scarce control plane implied by a current normalized P/E above 200x; multiple compression then overwhelms even excellent EPS growth.

The observable anchors are not press-release deal ceilings. They are companywide NRR, U.S. Commercial customer growth, segment Contribution margins, funded Government task orders, and recompete scope. GAO has specifically warned that dependence on a single software provider complicates open standards and can create licensing-cost lock-in, while OPEN DAGIR and modular-open-system efforts aim to improve portability and competition. Source: [GAO-25-106454](https://files.gao.gov/reports/GAO-25-106454/index.html).

4. Cross-examination of management commentary

Management statement	Skeptical reading	Earliest disconfirming evidence
David Glazer, Q2 FY2026: the Government hosting change should "power faster time to value, drive greater efficiency, provide greater cost certainty to the customer and enable us to expand their future workflows."	Cost certainty for the customer can mean Palantir absorbs compute, egress, security, and usage risk. Faster deployment does not guarantee better unit economics.	Government Contribution margin remains below 70% or cost of revenue rises faster than Government revenue for the next two quarters.
Alexander Karp, Q2 FY2026: Palantir is extending AIP into a sovereign stack that can orchestrate and fine-tune models.	Sovereignty is a requirement, not an exclusive franchise. Microsoft already markets national-partner, public, private, and sovereign-cloud choices; local procurement can also favor domestic vendors.	International Government and Commercial residual growth stays below 10% while local-hosting and partner costs rise.
David Glazer, Q2 FY2026: U.S. Commercial FY2026 revenue will exceed USD 3.424 billion, at least 134% growth.	Near-term guidance has been highly credible, but the valuation depends on post-guide durability, not merely hitting FY2026. Pulled-forward AI projects and a few large expansions can make the guide real while the 2028-2031 slope fails.	NRR falls below 120% and U.S. Commercial customer growth decelerates together as the FY2026 cohorts anniversary.

Management deserves credit for repeated guide beats. That history makes an immediate FY2026 miss less likely; it does not validate a five-year cohort, margin, or multiple assumption. The skeptical distinction is between **forecast credibility for the next quarter** and **durability of the economic control point through 2031**.

5. Stage 5 tie-out reverification and model pressure points

The recency anchor is valid. The TTM uses four non-overlapping standalone quarters ending 2026-06-30, and the first forward year is supported by Q3 and FY2026 guidance. The price and share bridge is XNAS:PLTR Class A on a one-for-one economic-share basis. The model's temporal convention is exactly five years from the TTM operating anchor, as instructed. The August 21 quote date is 1,774 days before June 30, 2031, approximately 4.857 years; using that shorter price-to-terminal interval would make Base about negative 6.42% and probability-weighted return about negative 4.43%, differences of only 0.18 and 0.12 percentage points. That presentation nuance is immaterial to the conclusion and does not warrant reopening the exact-five-year governed convention.

The share path is more debatable than the arithmetic. Base adds only 29.155 million shares over five years, about 1.14% cumulatively, even though diluted weighted-average shares increased 33% from FY2021 to FY2025. The justification is the current 0.2% stabilization and falling SBC/revenue, while future issuance is approximated as SBC expense divided by scenario share price. That approximation can understate grants if compensation targets remain high or if the stock price falls. Yet the model also deducts GAAP SBC expense from earnings, appropriately capturing both compensation cost and dilution rather than treating one as a substitute for the other. A 5% higher terminal share count lowers Base EPS to roughly USD 3.5473 and Base gross IRR to negative 7.15%, only 0.91 percentage point worse. Dilution remains a major stewardship and moat signal, but the displayed local stress is not a material return-bridge defect.

No product CAGR clearly exceeds a competitor pattern without a named reason. The risk is instead that the named reason—AIP becoming a reusable operating layer—cannot be verified from product revenue or cohort metrics. Base Commercial's 32.18% five-year CAGR and Base Government's 22.93% CAGR are demanding but both decelerate sharply from current growth. The arithmetic ties; the economic uncertainty is correlation between growth and margin.

6. Score-derived 20% return claim

The audited score does **not** claim that a 20% return is plausible. It gives the return-underwriting category 2/20, sets Base gross USD IRR at negative 6.24%, probability-weighted gross USD IRR at negative 4.31%, and Bull at only 11.90%. The red team agrees. The Base 20% hurdle requires USD 12.7928 of terminal EPS at 35x versus USD 3.7247 modeled, and the Base maximum buy price is USD 52.39 versus USD 179.94.

The assumption that would break a hypothetical 20% case first is the exit multiple. Even the extraordinarily favorable Bull EPS of USD 7.0148 would need about 63.83x P/E to produce USD 447.75 of required terminal wealth with zero distributions. That would demand both Bull operating execution and a premium terminal valuation despite governance FAIL, moat/2046 YELLOW, chronic dilution, and greater platform competition. The score correctly rejects that stack.

7. Two-sided valuation challenge and independent return replay

Audited arithmetic replay

Item	Independent formula	Recomputed result	Audited result
Base terminal value	USD 3.724686226 EPS × 35x	USD 130.364017910	USD 130.364017910
Base gross USD IRR	(130.364017910 / 179.94)^(1/5) – 1	-6.242503096995%	-6.242503096995%
Probability-weighted terminal wealth	28.435275939% × 39.147600675 + 50% × 130.364017910 + 21.564724061% × 315.667122660	USD 144.386481183361	USD 144.386481183362
Probability-weighted gross USD IRR	(144.386481183361 / 179.94)^(1/5) – 1	-4.307089631320%	-4.307089631320%
Base net-SGD IRR	((130.364017910 × 1.26976 × 0.995) / (179.94 × 1.26976))^(1/5) – 1	-6.336448673413%	-6.336448673413%
Probability-weighted net-SGD IRR	(SGD 189.613489150280 / SGD 228.480614400)^(1/5) – 1	-3.660591751133%	-3.660591751133%

Strongest reason the model is too optimistic. The 35x Base P/E assumes Palantir still deserves a material premium to Microsoft and Salesforce's cited raw TTM P/Es of about 27x and 24x after five years, despite governance FAIL and unresolved moat durability. A clearly labelled 30x sensitivity produces USD 111.7406 of terminal value and a negative 9.09% Base gross IRR, 2.85 percentage points worse. This is the most important valuation attack. It is not an unresolved defect because the preflight selected 35x before returns, explains the premium through mid-teens terminal growth and 57.55% margin, and applies a vast haircut from PLTR's normalized current 203.11x P/E. Thirty-five times is contestable, but not denominator-invalid or reverse-engineered.

Scenario probability construction is the next attack. Revenue-weighting the two segment probability sets into 28.435%/50%/21.565% is a disclosed heuristic, not a fully specified joint distribution of Government and Commercial outcomes. Because the consolidated Bull endpoint requires both segments' Bull paths, its 21.565% weight slightly exceeds Commercial's 20% Bull marginal. Moving that 1.565 percentage points from Bull to Bear lowers probability-weighted IRR to negative 4.89%, a 0.58-point change. A much larger 10-point Bull-to-Bear stress lowers it to negative 8.29%, but that is a new probability judgment, not a correction demonstrated by the evidence. Future work would be cleaner with directly underwritten consolidated probabilities or a nine-cell joint segment grid; the present issue is immaterial under the 2-point rule.

Strongest reason the model is too bearish. The P/E bridge assigns zero distributions, gives no standalone terminal value to the current USD 9.409 billion of cash and marketable securities, and assumes only USD 0.300 billion of Base terminal interest income despite substantial retained earnings capacity. Economically, that is punitive. Mechanically, it is correct: adding net cash after statutory EPS × P/E would double count a capital structure already reflected in interest income and the equity multiple. As an upper-bound, non-controlling sensitivity only, current net cash after leases is about USD 3.58 per diluted share. Adding that amount to Base terminal wealth would improve Base IRR to negative 5.73%, just 0.51 percentage point. The available source-bound cash evidence therefore does not support a material bearishness objection, and no terminal-cash path may be fabricated.

Challenge	Audited case	Sensitivity	IRR effect	Classification
Base exit multiple	35x / -6.24%	30x / -9.09%	-2.85 pp	Material sensitivity, not a proved correction
Base operating margin	57.55% / -6.24%	50.0% / -8.78%	-2.53 pp	Material sensitivity, not a proved correction
Consolidated Bull probability	21.565%	20.0%, shifted to Bear	PW IRR -4.89%	-0.58 pp, immaterial
Current net cash upper bound	No separate addition	USD 3.58/share added	Base IRR -5.73%	+0.51 pp, immaterial and method-inconsistent if controlling

The P/E count-once treatment is sound. Cash, leases, NCI attribution, SBC expense, and diluted shares enter once; distributions are explicitly zero; no unsupported terminal balance-sheet plug appears. The audited Bear/Base/Bull gross USD returns of negative 26.29%, negative 6.24%, and positive 11.90%, and the probability-weighted negative 4.31%, all reproduce at full precision.

8. Governance, KPI gaps, concentration, procurement, and moat durability

Governance deserves to affect both multiple and probability. The Class F trust can preserve 49.999999% founder voting control while economic ownership falls, can survive employment, and can concentrate among remaining founders. The board's nominal independent majority is only 4-3 under a founder chairman; no named lead independent director is disclosed. Diluted shares rose 33% from FY2021 to FY2025, buybacks were negligible relative to SBC, the authorization was terminated, and substantial insider sales continue. Strong execution does not cure the outside-owner asymmetry. Management FAIL is correctly hard, not cosmetic.

Visibility is weaker than the headline KPIs imply. RPO of USD 4.9 billion is narrower and commercially biased. RDV of USD 13.1 billion includes options and assumes no termination, so it is not funded backlog. NRR is companywide and has only a short disclosed series. The absence of ARR, cRPO, billings, GRR, churn, CAC, payback, aggregate ACV, and platform-level revenue prevents a reliable separation of new production demand from expansion, pricing, mix, and contract timing. AIP may be creating a new category, accelerating Foundry, or relabelling existing deployments; outsiders cannot quantify which.

Concentration is economically meaningful even without a customer above 10% of Q2 revenue. Customer I represented 27% of accounts receivable, and the top-20 average TTM revenue per customer reached USD 124 million. Large relationships validate product depth but also create billing lumpiness, buyer power, and discount risk. In Government, the USD 10 billion Army vehicle is a ceiling, not an obligation; contracts can terminate for convenience; funding, task-order timing, appropriation, and option exercise govern conversion. Faster procurement can coexist with worse unit economics if enterprise purchasing consolidates demand and extracts volume discounts.

The moat survives only if Palantir remains the governed context-and-action layer after model access commoditizes. Government's classified deployments, installed ontology, mission trust, and accreditation make displacement difficult. Commercial's moat is less proven: Microsoft, ServiceNow, Snowflake, Databricks, cloud providers, integrators, and customer engineering teams can bundle enough of the stack to narrow Palantir's scope. Open interfaces do not need to recreate Palantir; they need only make selected modules substitutable. The moat appendix's YELLOW and 2046 YELLOW verdicts are

therefore appropriate.

Palantir is not a P&C insurer, reinsurer, or comparable risk-bearing financial business. The insurer-return bridge test is not applicable, and no `insurer_return_bridge.json` is required.

9. The signal that forces an exit

Forced-exit signal: companywide net revenue retention below 110% for two consecutive reported quarters within the next four quarters. Palantir does not disclose Commercial-only NRR, so the companywide series is the best available expansion metric. Two sub-110% readings would mean the expansion engine has fallen to ordinary enterprise-software levels despite an unusually favorable AI demand environment. It would directly invalidate Commercial's Base assumption that NRR fades only toward roughly 118%-125%, weaken the reuse-based margin thesis, and indicate that incumbent bundles or cohort anniversary effects are arriving faster than modeled. One quarter is insufficient because large contracts are lumpy; two consecutive quarters are a structural signal.

10. Conviction at the end of red team

Thesis does not survive this red team as an investment at USD 179.94. The break is not a hidden arithmetic error; it is the valuation's dependence on exceptional, correlated operating outcomes while governance, durability, dilution, and KPI opacity remain unresolved. Even the audited Bull case does not reach the owner's return target. The conclusion should remain organized around price overwhelming operating excellence, with Commercial control-layer commoditization and Government portability/procurement reform as the mechanisms that can make the already-negative Base return worse.

The return model itself survives. Its P/E method is consistent across scenarios and net-SGD, the denominator normalization is explicit, terminal EPS and share paths reconcile, current cash is counted once through statutory earnings/multiple economics, and the probability-weighted return is correctly the CAGR of weighted terminal wealth rather than an average of scenario IRRs. The valuation can be debated, but no supportable correction was found that would move the published conclusion in a less-conservative direction or expose an invalid bridge. The material sensitivities to a 30x Base multiple and 50% Base operating margin both make the conclusion more conservative; they do not rescue it.

Appendix D - Data Package (Raw Historicals)

Palantir Technologies Inc. (PLTR) — Production ER2 Data Package

Information cutoff: 2026-08-24 SGT. Fiscal year-end: December 31. Latest completed quarter: Q2 FY2026 ended 2026-06-30, reported 2026-08-03 and filed 2026-08-04. Forecast horizon required downstream: exactly 2031-06-30. Reporting currency is USD. Monetary tables use billions unless explicitly labeled USD thousands or USD millions. Facts are identified as reported or calculated; forecast treatments are analyst inferences, not company guidance. No trade execution is authorized.

0. Controlling live source set and sealed inputs

This package uses live pulls made for this stage and archived under `C:\Users\Admin\OneDrive\GPT\ER 2.0\Palantir Technologies Inc.\raw\pulse-cio-20260824T180308+0800\data_package`. FMP endpoints were fetched live without exposing the API key. Controlling filing links are the [FY2025 10-K](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm), [Q3 FY2025 10-Q](https://www.sec.gov/Archives/edgar/data/1321655/000132165525000131/pltr-20250930.htm), [Q1 FY2026 10-Q](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000028/pltr-20260331.htm), and [Q2 FY2026 10-Q](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000041/pltr-20260630.htm). Current releases are [Q3 FY2025](https://www.sec.gov/Archives/edgar/data/1321655/000132165525000130/a2025q3ex991earningsrelease.htm), [Q4 FY2025](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000004/a2025q4ex991earningsrelease.htm), [Q1 FY2026](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000026/a2026q1ex991pressrelease.htm), and [Q2 FY2026](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000039/a2026q2ex991pressrelease.htm).

The mandatory input hashes were recomputed before drafting: `quant_foundation.md` = `980bbb1ad6377e8a361e0675982b89550e0d83a18c4f2d2c96c34453ec0a00e9`; `quant_foundation_review.md` = `27ac3ceb8a3fb85bb0f1a7c34a04919b7d7afe44b9a26a19590bccba5636f18e`; `quote_packet.json` = `1a87c7195d065cbbb11be9896951d0c60eda8a55e1e8f5eff53b437a9ff7010b`; `statutory_bridge.json` = `70965001151a9ac518f8ce4c11f797b1dad4280818161b657c1c0cebcd6edd3f`. QFR gate: `QUANT FOUNDATION VERIFIED — CLEARED FOR WAVE 1`.

0a. Industry classification and disclosure inventory

Palantir is enterprise software with SaaS-like subscription and cloud economics, but contracts can contain multi-year options and termination-for-convenience clauses. It reports under U.S. GAAP, has a calendar fiscal year, and discloses two customer-type operating segments. FMP annual and quarterly product-segmentation endpoints contain Government and Commercial revenue, but no platform-level dollars or Contribution; official filings control Contribution and industry KPIs. The live FMP FY2021 Government row is defective at $8.973356 billion and is quarantined; the FY2021 10-K controls at $897.356 million.

Official FY2021-FY2025 10-Ks, Q3 FY2025/Q1 FY2026/Q2 FY2026 10-Qs, four earnings releases, four FMP transcripts, SEC submissions, FMP statements, segment and geography endpoints were searched. Search terms: `ARR`, `ending ARR`, `NRR`, `net dollar retention`, `RPO`, `cRPO`, `remaining performance obligations`, `billings`, `deferred revenue`, `remaining deal value`, `TCV`, `ACV`, `customer count`, `top twenty customers`, `Contribution`, `Government`, `Commercial`, `stock-based compensation`, `segment realignment`, and `reclassification`.

Dataset / table family	Official source searched	Search terms or table names	Years available	Included?	If excluded, why	Downstream relevance
Segment economics	FY2021-FY2025 10-Ks; latest 10-Qs	Segment and Geographic Information; Contribution	FY2021-FY2025; last four quarters; TTM	Yes	None	Segment forecasts and statutory bridge
ARR and net new ARR	Same filings and releases	ARR; ending ARR; net new ARR	None	No	True non-disclosure	SaaS growth durability
NRR / DBNR	Calls and releases	Net dollar retention; NRR; DBNR	Q3 FY2025-Q2 FY2026	Yes	Annual five-year series not disclosed	Expansion economics
RPO	Note 3 or Note 12; releases	Remaining performance obligations	FY2021-FY2025; latest quarters	Yes	cRPO not separately disclosed	Contracted revenue visibility
RDV / TCV / ACV	MD&A; releases; calls	Remaining deal value; total contract value; annual contract value	FY2021-FY2025; last four quarters	Yes	Aggregate ACV not disclosed; definition only	Bookings and option-weighted backlog

Dataset / table family	Official source searched	Search terms or table names	Years available	Included?	If excluded, why	Downstream relevance
Billings	Statements and MD&A	Billings; customer billings	None as a reported KPI	No	True non-disclosure; only descriptive timing references	Cash conversion
Deferred revenue	Balance sheets and contract-liability notes	Deferred revenue; contract liabilities	FY2021-FY2025; latest quarter	Yes	None	Noncancelable billed backlog
Customers / large customers	Business overview; calls	Customer count; top twenty; U.S. commercial customer count	FY2021-FY2025; latest quarters	Yes	Q2 FY2026 total customer count not stated	Land-and-expand and concentration
CAC / payback / magic number / GRR / churn	Same official set	CAC; payback; magic number; GRR; churn	None	No	True non-disclosure	Unit economics
Platform-level revenue	Same official set	AIP; Foundry; Gotham; Apollo; revenue	None	No	True non-disclosure	Product mix
SBC and dilution	10-Ks, 10-Qs, FMP statements	Stock-based compensation; weighted-average shares	FY2021-FY2025; TTM	Yes	FMP H1/YTD share denominators quarantined	Per-share economics

Industry KPI history follows. RDV includes exercisable options and presumes no termination, so it is not equivalent to RPO. Deferred revenue is current plus noncurrent and excludes customer deposits. Q2 FY2026 values are the latest available unless marked otherwise.

KPI family	FY2021	FY2022	FY2023	FY2024	FY2025	Latest interim / TTM	Source
Total customers	237	367	497	711	954	1,007 at Q1 FY2026; Q2 total not disclosed	10-K business overview; Q1 transcript
Top-20 average TTM revenue per customer ($M)	43.6	49.4	54.6	64.6	93.9	124 at Q2 FY2026	10-Ks; Q2 transcript dated 2026-08-03
Total RDV ($B)	3.8	3.7	3.9	5.4	11.2	13.1 at Q2 FY2026	10-K MD&A; Q2 transcript
RPO ($B)	1.1	0.973	1.2	1.7	4.1	4.9 at Q2 FY2026	Contract-liability notes; latest 10-Q
Deferred revenue ($B)	0.268	0.193	0.275	0.300	0.455	0.613 at Q2 FY2026	Balance sheets; calculated current plus noncurrent
GAAP gross margin	78.0%	78.6%	80.6%	80.2%	82.4%	84.8% TTM	Live FMP income statements; official filings; calculated
Sales and marketing expense ($B)	0.615	0.703	0.745	0.888	1.057	1.235 TTM	Live FMP income statements; official filings
Sales and marketing / revenue	39.9%	36.9%	33.5%	31.0%	23.6%	20.1% TTM	Calculated from the preceding sourced dollars
NRR / net dollar retention	n/d	n/d	n/d	n/d	139% at Q4	157% at Q2 FY2026	Calls; quarterly KPI only
U.S. commercial customers	n/d	n/d	n/d	n/d	571 at Q4	653 at Q2 FY2026	Q4 and Q2 calls
ARR / cRPO / billings / aggregate ACV	n/d	n/d	n/d	n/d	n/d	n/d	True non-disclosure after stated source/term search

Contrary evidence preserved: the Q4 FY2025 call rounded RPO to $4.2 billion, while the later filed FY2025 10-K reports $4.1 billion; the controlling statutory filing value is used. RPO is primarily commercial because short initial terms and termination-for-convenience constructs common in government contracts are excluded. At Q2 FY2026, no customer was 10% or more of revenue, but `Customer I` was 27% of accounts receivable, up from 25% at FY2025.

1. Reporting structure and taxonomy

Government serves U.S. and allied government customers, including defense, intelligence, health, and civilian agencies. Commercial serves nongovernment institutions across industries using Foundry, AIP, Apollo, and related offerings. FY2025 Government revenue was $2.402287 billion, 53.7% of total; Commercial was $2.073159 billion, 46.3%.

Reporting structure has been stable as Government and Commercial since at least FY2020. No segment realignment, reportable-segment change, or restating 8-K was found in the three-year taxonomy investigation. In FY2024, expanded

segment disclosures began showing segment revenue and directly attributable expenses alongside Contribution; this was disclosure enhancement, not an economic reclassification. Q1 FY2026 included a customer-program transition from U.S. Commercial to U.S. Government; management said U.S. Commercial growth would have been 143% rather than 133% absent that transition. This is a customer assignment effect, not a segment taxonomy change.

FMP taxonomy flags: `Government Operating Segment` maps to issuer label `Government`; `Commercial` maps directly. FMP FY2021 Government revenue is off by a factor of ten and is not used. FMP geography adds a duplicate FY2025 `Geographic Concentration Risk` total and separately identified France through FY2021; the table below aggregates France into rest of world for consistency.

Geography revenue ($B)	FY2021	FY2022	FY2023	FY2024	FY2025	FY2025 mix	Source
United States	0.879	1.161	1.378	1.900	3.320	74.2%	FMP live geography endpoint; 10-K Note 13
United Kingdom	0.173	0.221	0.235	0.305	0.427	9.5%	Same
Rest of world, including France	0.489	0.524	0.612	0.661	0.728	16.3%	Same; calculated aggregation in FY2021
Total	1.542	1.906	2.225	2.866	4.475	100.0%	Tie to consolidated revenue

2. Consolidated historicals

Live FMP sources: [income statement](https://financialmodelingprep.com/stable/income-statement?symbol=PLTR&period=annual&limit=8), [cash flow](https://financialmodelingprep.com/stable/cash-flow-statement?symbol=PLTR&period=annual&limit=8), and [key metrics](https://financialmodelingprep.com/stable/key-metrics?symbol=PLTR&period=annual&limit=8), checked against FY2021-FY2025 10-Ks.

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	1.542	1.906	2.225	2.866	4.475	FMP live; 10-Ks
Operating income ($B)	(0.411)	(0.161)	0.120	0.310	1.414	FMP live; 10-Ks
Operating margin	(26.7%)	(8.5%)	5.4%	10.8%	31.6%	Calculated
Net income attributable to common ($B)	(0.520)	(0.374)	0.210	0.462	1.625	FMP live; 10-Ks
Diluted EPS ($)	(0.27)	(0.18)	0.09	0.19	0.63	FMP live; 10-Ks
Diluted weighted-average shares (M)	1,923.617	2,063.793	2,297.927	2,450.818	2,565.197	FMP live; 10-Ks
Capex ($B)	0.013	0.040	0.015	0.013	0.034	FMP live cash flow; absolute cash outflow
Free cash flow ($B)	0.321	0.184	0.697	1.141	2.101	FMP live cash flow; CFO less capex
NOPAT at 21% tax ($B)	(0.325)	(0.127)	0.095	0.245	1.117	Calculated
FMP year-end invested capital ($B)	2.451	2.723	3.649	5.197	7.435	FMP live key metrics
ROIC on stated basis	(13.2%)	(4.7%)	2.6%	4.7%	15.0%	NOPAT divided by year-end invested capital

ROIC definition is explicit and deliberately simple: statutory operating income multiplied by 79%, divided by FMP year-end invested capital. The 21% U.S. federal statutory rate is used consistently so loss-year tax benefits and very low reported cash-tax rates do not create false comparability. This is an analyst calculation, not a company KPI; it does not use average invested capital and therefore should not be mixed with FMP's precomputed ROIC field.

TTM metric, four official quarters	TTM ended 2026-06-30	Construction / source
Revenue ($B)	6.156	FY2025 less H1 FY2025 plus H1 FY2026
Operating income ($B)	2.635	Same statutory construction
Operating margin	42.8%	Calculated
Net income attributable to common ($B)	3.017	Same statutory construction

TTM metric, four official quarters	TTM ended 2026-06-30	Construction / source
Diluted shares (M)	2,571.043	Mean of official standalone Q3, Q4, Q1, Q2 denominators
Diluted EPS ($)	1.173333934905017	$3,016.692M divided by 2,571.043M
Capex ($B)	0.042	FY2025 less H1 FY2025 plus H1 FY2026
Free cash flow ($B)	3.358	CFO less capex on same construction
SBC ($B)	0.836	Sealed bridge / official statements

Official standalone diluted denominators are Q3 FY2025 2,571.057M, Q4 FY2025 2,573.497M, Q1 FY2026 2,570.924M, and Q2 FY2026 2,568.694M. FMP nine-month/H1 denominators are cumulative and quarantined; they do not control TTM per-share math.

3. Segment historicals

The exact common additive CODM profit basis is `Contribution`. Contribution is revenue less cost of revenue and sales and marketing expense, excluding stock-based compensation; it is not statutory operating income. Annual revenue comes from the live FMP product endpoint and official 10-Ks; Contribution comes from the official segment tables.

Government	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	0.897	1.072	1.222	1.570	2.402	Official 10-Ks; FMP except quarantined FY2021
YoY growth	47.1%	19.4%	14.0%	28.4%	53.1%	Calculated from official dollars
Contribution ($B)	0.542	0.621	0.725	0.948	1.576	Official 10-K Contribution tables
Contribution margin	60.4%	57.9%	59.3%	60.4%	65.6%	Calculated

Commercial	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	0.645	0.834	1.003	1.296	2.073	Official 10-Ks; FMP live
YoY growth	33.6%	29.4%	20.2%	29.2%	60.0%	Calculated from official dollars
Contribution ($B)	0.358	0.414	0.521	0.772	1.367	Official 10-K Contribution tables
Contribution margin	55.5%	49.7%	51.9%	59.5%	65.9%	Calculated

TTM segment	Revenue ($B)	TTM mix	Contribution ($B)	Contribution margin	Construction
Government	3.211	52.2%	2.255	70.2%	Q3 FY2025 + Q4 FY2025 + Q1 FY2026 + Q2 FY2026
Commercial	2.945	47.8%	2.136	72.5%	Same four official quarters
Total	6.156	100.0%	4.391	71.3%	Ties sealed TTM bridge

3a. Historical statutory-profit bridge and forecast-driver ledger

Sealed packet: `statutory_bridge.json`, SHA-256 `70965001151a9ac518f8ce4c11f797b1dad4280818161b657c1c0cebcd6edd3f`. Schema 1.1.0; basis `issuer_disclosed`; reporting currency USD; scale USD thousands. There are no corporate/unallocated adjustments. Labels, signs, classifications, and values below reproduce the packet exactly.

FY2025 bridge

Bridge line	Classification	Value ($000)	Source ID	Running subtotal ($000)
Government contribution	segment profit	1,576,070	SEC_2025_10K	1,576,070
Commercial contribution	segment profit	1,366,627	SEC_2025_10K	2,942,697
Sum of reportable segment profit	segment profit subtotal	2,942,697	SEC_2025_10K	2,942,697
Contribution	subtotal	2,942,697	SEC_2025_10K	2,942,697

Bridge line	Classification	Value ($000)	Source ID	Running subtotal ($000)
Research and development expenses, excluding stock-based compensation	operating	(420,838)	SEC_2025_10K	2,521,859
General and administrative expenses, excluding stock-based compensation	operating	(423,811)	SEC_2025_10K	2,098,048
Total stock-based compensation expense	operating	(684,033)	SEC_2025_10K	1,414,015
Statutory operating income	statutory subtotal	1,414,015	SEC_2025_10K	1,414,015
Interest income	interest	229,181	SEC_2025_10K	1,643,196
Other income, net	other_non_operating	14,172	SEC_2025_10K	1,657,368
Pretax income	statutory subtotal	1,657,368	SEC_2025_10K	1,657,368
Provision for income taxes	tax	(22,724)	SEC_2025_10K	1,634,644
Net income attributable to noncontrolling interests	noncontrolling_interests	(9,611)	SEC_2025_10K	1,625,033
Net income	statutory subtotal	1,625,033	SEC_2025_10K	1,625,033

TTM ended 2026-06-30 bridge

Bridge line	Classification	Value ($000)	Source ID	Running subtotal ($000)
Government contribution	segment profit	2,255,215	SEC_2026_Q2_10Q	2,255,215
Commercial contribution	segment profit	2,135,806	SEC_2026_Q2_10Q	4,391,021
Sum of reportable segment profit	segment profit subtotal	4,391,021	SEC_2026_Q2_10Q	4,391,021
Contribution	subtotal	4,391,021	SEC_2026_Q2_10Q	4,391,021
Research and development expenses, excluding stock-based compensation	operating	(473,564)	SEC_2026_Q2_10Q	3,917,457
General and administrative expenses, excluding stock-based compensation	operating	(447,281)	SEC_2026_Q2_10Q	3,470,176
Total stock-based compensation expense	operating	(835,524)	SEC_2026_Q2_10Q	2,634,652
Statutory operating income	statutory subtotal	2,634,652	SEC_2026_Q2_10Q	2,634,652
Interest income	interest	266,384	SEC_2026_Q2_10Q	2,901,036
Other income, net	other_non_operating	170,794	SEC_2026_Q2_10Q	3,071,830
Pretax income	statutory subtotal	3,071,830	SEC_2026_Q2_10Q	3,071,830
Provision for income taxes	tax	(41,111)	SEC_2026_Q2_10Q	3,030,719
Net income attributable to noncontrolling interests	noncontrolling_interests	(14,027)	SEC_2026_Q2_10Q	3,016,692
Net income	statutory subtotal	3,016,692	SEC_2026_Q2_10Q	3,016,692

Source integrity: SEC_2025_10K local path is `C:\Users\Admin\OneDrive\GPT\ER 2.0\Palantir Technologies Inc\raw\pulse-cio-20260824T180308+0800\official\SEC_PLTR_2025_10K.htm`, SHA-256 `a4fef9542c4d1a99a9265df88948e5a115223940db01a0bd01f1d8b6c00acd46`. SEC_2026_Q2_10Q local path is `C:\Users\Admin\OneDrive\GPT\ER 2.0\Palantir Technologies Inc\raw\pulse-cio-20260824T180308+0800\official\SEC_PLTR_2026_Q2_10Q.htm`, SHA-256 `a36c04dc261071a706cc84f0e94ed1bff7ce7745b73ac2af03a0e07118a5e741`. Each of the five sealed residuals is zero for both anchors; disclosed rounding unit and tolerance are both one USD thousand.

The following schedule is analyst inference for downstream modeling. It does not alter the sealed history.

Historical bridge line	FY2025 ($000)	TTM ($000)	Forecast treatment	Driver and source	Bear / base / bull rule
R&D excluding SBC	(420,838)	(473,564)	Percent of revenue	Official bridge; technical hiring commentary	Bear higher ratio; base gradual leverage; bull faster leverage

Historical bridge line	FY2025 ($000)	TTM ($000)	Forecast treatment	Driver and source	Bear / base / bull rule
G&A excluding SBC	(423,811)	(447,281)	Percent of revenue	Official bridge	Bear slower leverage; base declining ratio; bull faster leverage
Total SBC expense	(684,033)	(835,524)	Independently forecast expense	10-K/10-Q award schedules and dilution	Bear elevated; base declines as percent of revenue; bull faster normalization
Interest income	229,181	266,384	Cash and interest-rate schedule	Balance sheet cash and U.S. Treasuries	Bear lower rates/cash; base forward-rate path; bull higher cash yield
Other income, net	14,172	170,794	Independently forecast, not a margin plug	Marketable/equity security gains in 10-Q	Bear losses; base zero normalized; bull modest gains
Provision for income taxes	(22,724)	(41,111)	Tax-rate schedule	Statutory filings; valuation allowance context	Bear 23%; base ramps toward 21%; bull lower transition rate
NCI attribution	(9,611)	(14,027)	Percentage of JV income	Japan/Korea controlled entities	Bear higher NCI; base recent share; bull lower NCI

STATUTORY_BRIDGE_STATUS: VERIFIED — CLEARED FOR FORECASTING

4. Product and sub-segment historicals

Palantir names AIP, Foundry, Gotham, Apollo, and related defense offerings, but it does not disclose platform revenue in dollars. Government and Commercial are customer-type segments, not products. No platform dollars are back-solved.

Segment	Product	What it is	Dollar disclosure?	Last three years revenue and source
Both	AIP	Generative-AI platform connecting models, agents, workflows, data, security, and evaluations	No	True non-disclosure in FY2023-FY2025 10-Ks and latest releases
Primarily Commercial; also Government	Foundry	Data integration, ontology, analytics, and operational applications	No	True non-disclosure in same source set
Primarily Government	Gotham	Mission, defense, intelligence, and operational decision software	No	True non-disclosure in same source set
Both	Apollo	Continuous delivery and orchestration across cloud, on-premise, and edge environments	No	True non-disclosure in same source set
Government	Broader defense offerings / Maven-related capabilities	Mission-specific defense and AI capabilities integrated with Gotham and AIP	No	True non-disclosure in same source set

Segment	Product	Disclosure method	Last four quarters YoY growth	Source
Both	AIP	Strategic commentary and deal examples only	Not separately disclosed	Four live FMP transcripts and SEC releases
Both	Foundry	Strategic commentary only	Not separately disclosed	Same
Government	Gotham / defense offerings	Strategic commentary and contract announcements only	Not separately disclosed	Same
Both	Apollo	Strategic commentary only	Not separately disclosed	Same

5. Most recent quarterly trend

Palantir's fiscal year ends December 31. Most recently completed fiscal quarter at time of writing: FY2026 Q2, covering April-June 2026, reported on 2026-08-03 via the [Q2 FY2026 earnings release](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000039/a2026q2ex991pressrelease.htm).

Segment margin below means Contribution margin, not operating margin. Consolidated margin is statutory operating margin.

Quarter	Total revenue ($B)	YoY	Statutory operating margin	Source
Q3 FY2025	1.181	62.8%	33.3%	2025-11-03 release and 10-Q
Q4 FY2025	1.407	70.0%	40.9%	2026-02-02 release and FY2025 10-K
Q1 FY2026	1.633	84.7%	46.2%	2026-05-04 release and 10-Q
Q2 FY2026	1.935	92.8%	47.1%	2026-08-03 release and 10-Q

Quarter	Government revenue ($B)	YoY	Contribution margin	Commercial revenue ($B)	YoY	Contribution margin
Q3 FY2025	0.633	54.9%	66.0%	0.548	72.9%	65.7%
Q4 FY2025	0.730	60.3%	69.5%	0.677	81.8%	70.5%
Q1 FY2026	0.858	76.3%	73.3%	0.774	95.1%	72.7%
Q2 FY2026	0.990	79.0%	70.8%	0.945	109.8%	77.7%

Quarter	U.S. commercial revenue / YoY	U.S. government revenue / YoY	Official headline distinction
Q3 FY2025	$397M / 121%	$486M / 52%	U.S. subsets, not total segments
Q4 FY2025	$507M / 137%	$570M / 66%	U.S. subsets, not total segments
Q1 FY2026	$595M / 133%	$687M / 84%	U.S. Commercial would have grown 143% absent customer transition
Q2 FY2026	$764M / 149%	$809M / 90%	U.S. subsets; total Government grew 79%

Quarter	TCV ($B)	Total RDV ($B)	RPO ($B)	Total customers	U.S. commercial customers	NRR
Q3 FY2025	2.760	8.6	2.6	911	530	134%
Q4 FY2025	4.262	11.2	4.1 filed	954	571	139%
Q1 FY2026	2.410	11.8	4.5	1,007	615	150%
Q2 FY2026	3.373	13.1	4.9	Not disclosed	653	157%

6. Management commentary on forward growth

All quotes below are verbatim from live FMP transcript pulls archived for this stage. Dates match the fiscal periods: Q3 FY2025 call 2025-11-03; Q4 FY2025 call 2026-02-02; Q1 FY2026 call 2026-05-04; Q2 FY2026 call 2026-08-03. Endpoint pattern: `https://financialmodelingprep.com/stable/earning-call-transcript?symbol=PLTR&year={year}&quarter={quarter}`. No separate recent investor day or capital-markets-day source was found in the SEC/IR source inventory; therefore none is fabricated.

Commercial and company-wide

Period / speaker	Verbatim quote	Topic	Modeling interpretation
Q3 FY2025, David Glazer, CFO	“We’re also raising our full year 2025 revenue guidance midpoint to $4.398 billion, representing a 53% year-over-year growth rate and 8-point or $252 million increase over our full year 2025 revenue guidance last quarter.”	Company outlook	Then-current floor; superseded by reported FY2025 revenue
Q3 FY2025, David Glazer	“In addition, we’re raising our full year U.S. commercial revenue guidance to an excess of $1.433 billion, representing a growth rate of at least 104% year-over-year, a 19-point increase over the guidance we gave just last quarter.”	U.S. Commercial	Demonstrates acceleration into year-end
Q4 FY2025, David Glazer	“On the strength of our 2025 results, we are guiding a full year 2026 revenue of $7.190 billion at the midpoint, representing 61% growth year-over-year.”	Initial FY2026 outlook	Superseded upward by Q1 and Q2 guidance

Period / speaker	Verbatim quote	Topic	Modeling interpretation
Q4 FY2025, David Glazer	“We continue to expect expenses to increase in 2026 as we remain committed to investing in the product pipeline and the most elite technical talent, all while delivering on our goals of sustained GAAP profitability.”	Expense reinvestment	Do not extrapolate margin solely from revenue growth
Q1 FY2026, David A. Glazer	“For full year 2026, we are raising our revenue guidance to between $7.65 billion and $7.662 billion.”	Revised FY2026 outlook	Superseded by Q2 guidance
Q1 FY2026, David A. Glazer	“We are raising our U.S. Commercial revenue guidance to in excess of $3.224 billion, representing a growth rate of at least 120%.”	U.S. Commercial	Superseded upward one quarter later
Q2 FY2026, David Glazer	“For Q3 2026, we expect revenue of between $2.16 billion and $2.164 billion and adjusted income from operations of between $1.292 billion and $1.296 billion.”	Near-term outlook	Current quarter anchor at cutoff
Q2 FY2026, David Glazer	“For full year 2026, we are raising our revenue guidance to between $8.15 billion and $8.158 billion.”	Current FY2026 outlook	Current consolidated revenue control range
Q2 FY2026, David Glazer	“We are raising our U.S. commercial revenue guidance to in excess of $3.424 billion, representing a growth rate of at least 134%.”	Current U.S. Commercial outlook	Current disclosed commercial subset floor
Q2 FY2026, David Glazer	“We are raising our adjusted free cash flow guidance to between $4.5 billion and $4.7 billion, and we continue to expect GAAP operating income and net income in each quarter of this year.”	Cash conversion	Current FCF anchor; adjusted, not GAAP FCF
Q2 FY2026, Shyam Sankar, CTO	“AIP is where your sovereign AI is built, deployed and compounded.”	Product strategy	Management frames AIP as the control/orchestration layer, not model resale
Q2 FY2026, Alexander Karp, CEO	“Now we’re taking the AIP stack and extending it for sovereign AI, which requires us to be able to orchestrate and fine-tune models to provide a completely sovereign stack to our partners.”	Product roadmap	Supports continued product investment; qualitative, not a dollar forecast

Government

Period / speaker	Verbatim quote	Topic	Modeling interpretation
Q3 FY2025, David Glazer	“This growth was driven by continued execution in existing programs and new awards reflecting the growing demand for AI and our government software offerings.”	Government growth	Both expansion and new awards contributed
Q4 FY2025, David Glazer	“Full year U.S. government revenue grew 55% year-over-year to $1.855 billion.”	U.S. Government	Strong base entering FY2026; historical, not guidance
Q1 FY2026, David A. Glazer	“First quarter Government revenue grew 76% year-over-year and 18% sequentially to $858 million.”	Government acceleration	Confirms broad segment momentum
Q1 FY2026, David A. Glazer	“This growth was driven by continued execution in existing programs and new awards reflecting the growing demand for our AI platform in government.”	Government driver	Contract execution and award conversion are model drivers
Q2 FY2026, David Glazer	“Second quarter government revenue grew 79% year-over-year and 15% sequentially to $990 million.”	Government acceleration	Latest reportable-segment growth anchor
Q2 FY2026, David Glazer	“While this change led to higher cost of revenue in Q2, going forward, we believe it will power faster time to value, drive greater efficiency, provide greater cost certainty to the customer and enable us to expand their future workflows.”	Government cloud hosting	Q2 Government Contribution margin dip has a specific hosting-cost cause

Period / speaker	Verbatim quote	Topic	Modeling interpretation
Q2 FY2026, Alexander Karp	“And the story really does begin at the beginning when we dedicated ourselves to our most important partners in the U.S. government, and we built products to deliver value for them.”	Strategic origin	Government R&D creates capabilities reused across the platform
Q2 FY2026, David Glazer	“We are raising our adjusted free cash flow guidance to between $4.5 billion and $4.7 billion, and we continue to expect GAAP operating income and net income in each quarter of this year.”	Profitability	Current profitability floor is consolidated, not segment-specific

7. Capital allocation and dilution history

Capex, R&D, repurchases, dividends, and acquisition cash flows are from live FMP statements and controlling 10-Ks. Cash outflows are shown positive. M&A is cash acquisition outflow; the FY2022 amount relates to consolidation/business-combination activity and should not be treated as recurring.

Year	Capex ($B)	Capex / revenue	R&D ($B)	R&D / revenue	Buybacks ($B)	Dividends ($B)	M&A ($B)	Source
FY2021	0.013	0.8%	0.387	25.1%	0.000	0.000	0.000	FMP live; 10-K
FY2022	0.040	2.1%	0.360	18.9%	0.000	0.000	0.067	FMP live; 10-K
FY2023	0.015	0.7%	0.405	18.2%	0.000	0.000	0.000	FMP live; 10-K
FY2024	0.013	0.4%	0.508	17.7%	0.064	0.000	0.000	FMP live; 10-K
FY2025	0.034	0.8%	0.558	12.5%	0.075	0.000	0.000	FMP live; 10-K

Dilution metric	FY2021	FY2022	FY2023	FY2024	FY2025	TTM / latest	Source
SBC expense ($B)	0.778	0.565	0.476	0.692	0.684	0.836 TTM	Cash-flow statements; sealed bridge
SBC / revenue	50.5%	29.6%	21.4%	24.1%	15.3%	13.6% TTM	Calculated
Diluted weighted-average shares (M)	1,923.617	2,063.793	2,297.927	2,450.818	2,565.197	2,571.043 TTM	Official/FMP; TTM official-only
YoY diluted-share growth	n/a	7.3%	11.3%	6.7%	4.7%	0.2% vs FY2025	Calculated

The August 2023 authorization allowed up to $1.0 billion of Class A repurchases. Palantir repurchased and retired 0.6 million shares for $75.0 million in FY2025, then terminated the program in January 2026. Repurchases have been far smaller than cumulative SBC and option/RSU issuance. At FY2025, unrecognized compensation included $406.0 million for options, $144.6 million for SARs, and $914.9 million for RSUs over multi-year service periods. This overhang is a fact; any future expense recognition schedule is inference.

8. Balance sheet snapshot

Item	FY2025 ($B)	Source
Cash and equivalents	1.424	FMP live balance sheet; 10-K
Short-term investments / marketable U.S. Treasuries	5.753	Same
Total debt, including lease obligations	0.229	FMP live
Net debt / (net cash)	(6.948)	Debt less cash and short-term investments
EBITDA	1.684	FMP live income statement
Net debt / EBITDA	(4.13x)	Calculated; net cash position
Goodwill plus intangibles	0.000	FMP live; 10-K balance sheet
Stockholders' equity	7.387	FMP live; 10-K
Tangible book value	7.387	Equity less goodwill/intangibles

At Q2 FY2026, cash and marketable securities were $9.409 billion, total stockholders' equity was $9.774 billion, and debt/lease obligations were $0.211 billion. Deferred revenue was $0.613 billion, but customer deposits added another $0.453 billion; those categories must not be conflated.

9. Recent corporate actions, last 12 months

- 2026-08-03: Q2 results raised FY2026 revenue guidance to $8.150-$8.158 billion, U.S. Commercial revenue to more than $3.424 billion, adjusted operating income to $4.889-$4.897 billion, and adjusted FCF to $4.5-$4.7 billion. Source: Q2 release.
- 2026-06-03 meeting, filed 2026-06-09: shareholders re-elected all directors, ratified EY, approved executive compensation on an advisory basis, and rejected three shareholder proposals. Source: [June 2026 8-K](https://www.sec.gov/Archives/edgar/data/1321655/000132165526000033/pltr-20260603.htm).
- January 2026: Palantir terminated its $1.0 billion share-repurchase program after only $75.0 million of FY2025 repurchases. Source: FY2025 10-K Note 9.
- No material acquisition, divestiture, CEO/CFO transition, financial restatement, large legal settlement, or new dividend was found in the live SEC filing inventory for the period. This is a source-inventory conclusion, not a claim that immaterial events did not occur.

10. Open data gaps, caveats, and live self-checks

Gap / issue	Classification	Exact source search	Treatment
AIP, Foundry, Gotham, Apollo revenue	True non-disclosure	FY2023-FY2025 10-Ks; latest 10-Qs/releases; platform names plus revenue	No platform dollars or back-solves
ARR, cRPO, billings, aggregate ACV	True non-disclosure	Same official set; exact KPI terms in Section 0a	ARR/cRPO/billings absent; ACV defined but no aggregate value
GRR, churn, CAC, payback, magic number	True non-disclosure	Same official set; exact KPI terms in Section 0a	Not modeled as historical fact
Q2 FY2026 total customer count	True non-disclosure	Q2 release, 10-Q, and 2026-08-03 transcript	U.S. Commercial count 653 disclosed; total omitted
Five-year NRR history	True non-disclosure	FY2021-FY2025 10-Ks and calls	Only recent quarterly values included
FY2021 FMP Government revenue	Source defect	Live FMP annual product endpoint says $8.973356B	Quarantined; FY2021 10-K $897.356M controls
Q4 FY2025 RPO rounding	Contrary official evidence	Q4 call $4.2B; FY2025 10-K $4.1B	Filed 10-K controls; discrepancy disclosed
FMP H1/YTD diluted denominators	Source-basis mismatch	Live quarterly FMP statements vs standalone SEC quarters	Quarantined; official four-quarter 2,571.043M controls
Other income, net forecast	Nonrecurring / volatile	Q2 10-Q marketable-security disclosures	Normalize separately; never use as operating plug
Investor or capital markets day	Not fetched / omitted only if one exists outside searched IR/SEC inventory	SEC submissions, four current releases/transcripts	No event-day quote included

Live self-checks completed in this stage: (1) the Q2 FY2026 release contains the U.S. Government headline of 90% YoY; the reportable Government segment grew 79.0% from official $990.032 million and $552.983 million quarterly values. (2) The Q2 transcript was fetched live, dated 2026-08-03, and contains the current quote: “For full year 2026, we are raising our revenue guidance to between $8.15 billion and $8.158 billion.” (3) The live FMP annual income-statement endpoint reports FY2025 revenue of $4,475,446,000, an exact tie to the FY2025 10-K.

The most important analytical caveats are the option-weighted nature of RDV, commercial bias in RPO, continued SBC/share dilution despite strong per-share earnings growth, Customer I's 27% receivables concentration, the volatile below-line contribution of marketable securities, and the difference between Contribution and statutory operating income. None is allowed to rewrite the sealed statutory bridge.