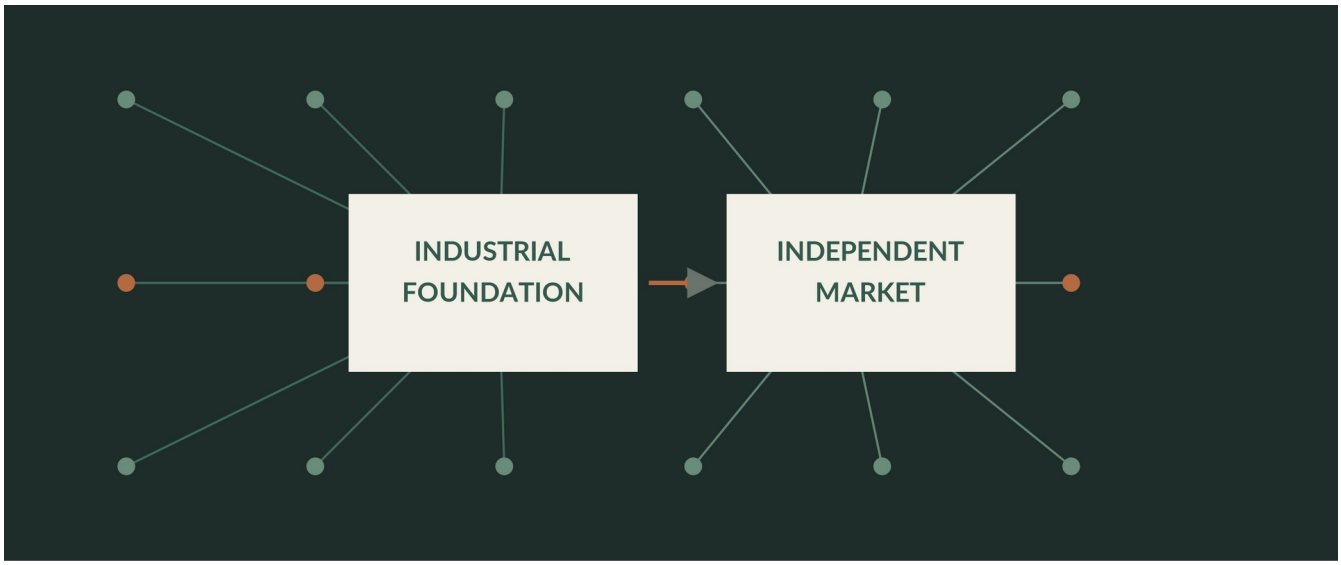
INSTITUTIONAL STRATEGY

The Independent Institution

Jensen Huang, Cornelius Vander Starr and the next commercial chapter of the CUDA economy



NUMAIR FARAZ

16 November 2026

An independent private institution for the CUDA economy.

The argument and the operating model

From the institution-builders to the customer market—and the decisions each audience can make.

An independent institution for the AI economy 3
1. Jensen Huang and Cornelius Vander Starr 5
2. The company beside the company 9
3. Computraq: the customer owns the market 13
4. From the first requirement to a recurring purchase 17
5. Inside the operating model 22
6. How customers and partners connect 27
7. Better buying, better economics 31
8. Win half the work—and create the next application economy 36
9. The freedom to own the means of computing 40
10. National champions are not an industrial strategy 43
11. When national security becomes a license for incumbents 45
12. Turn the credit network into an application engine 49
13. NVIDIA: the application-economy revaluation 52
14. Build an economy in which outsiders can succeed 56
15. The independent institution at work 62
Economic definitions and essential terms 62
Public sources and reading notes 65

READING ROUTES

The institutional case: chapters 1–2. The product and customer journey: chapters 3–6. Commercial economics: chapters 7–8 and 12. Public policy: chapters 9–11. Shareholders: chapter 13. Action program: chapter 14.

An independent institution for the AI economy

Jensen Huang built the technical foundation. Computraq organizes the commercial opportunity around it.

Jensen Huang and Cornelius Vander Starr belong to the same important tradition of American enterprise: the builder who creates an international institution by making other people's productive activity possible. Starr connected local commercial knowledge to the capacity of an international insurance business. Huang connected developers, equipment owners, suppliers and customers to a common foundation for accelerated computing. Their most consequential product was a system in which many other businesses could operate. [13,48–50]

The comparison has a second, equally important dimension. The insurance enterprise that became AIG did not require every valuable commercial activity to sit inside the public company. C.V. Starr & Co., established in 1950, retained specialist insurance agencies alongside AIG. Those agencies originated business, cultivated relationships and earned their own commercial income. A large public institution and a separate private organization could make one another more productive through complementary work. [48,49]

Computraq occupies that complementary position in the NVIDIA ecosystem. NVIDIA supplies the technology, software and industrial foundation. Computraq is the independent private business that organizes customer demand, qualified AI services, execution providers and payments around that foundation. Its ambition is to become the C.V. Starr & Co. of the CUDA economy: a commercially consequential institution with its own customers, assets, income and decisions.

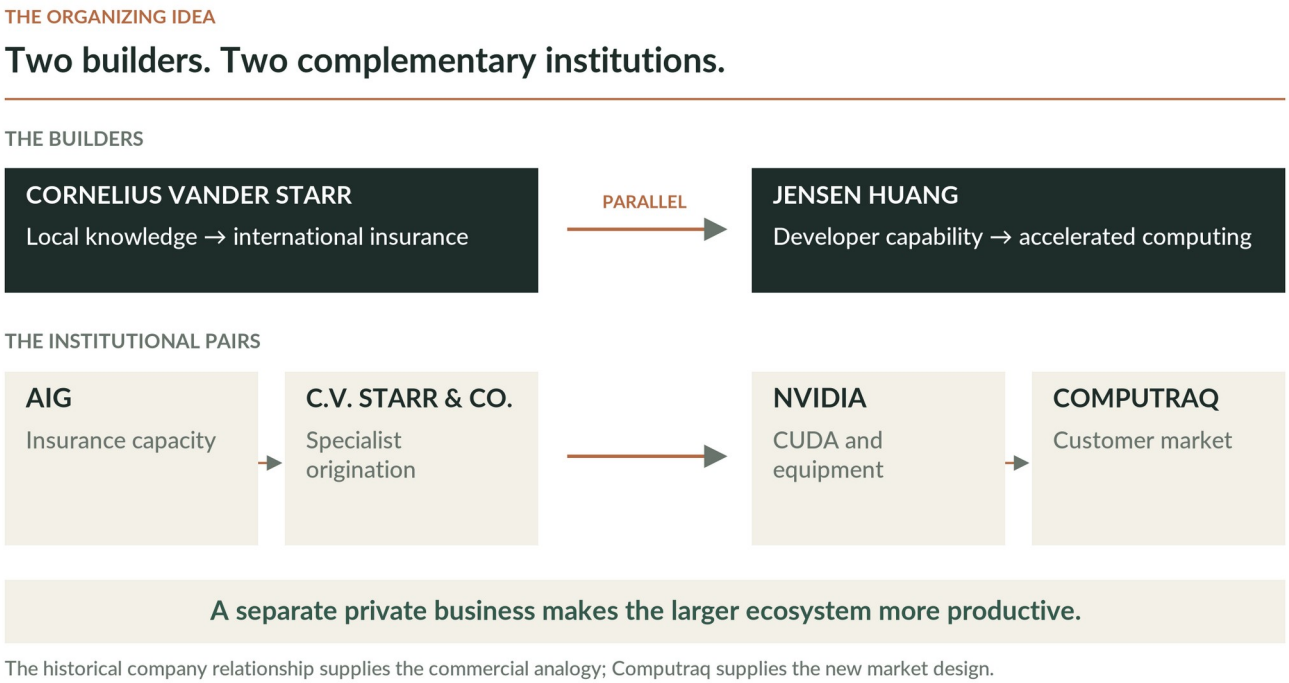


Exhibit 01. Two parallels, one institutional model. The founders build the foundation; the complementary private company develops the market.

Independence is central to the proposition. Computraq is not a department waiting for a place on NVIDIA's organization chart. It can pursue specialist customers, develop new forms of purchasing, support smaller operators and build recurring relationships on its own commercial terms. NVIDIA can benefit from the resulting

CUDA demand without owning the marketplace. Computraq can concentrate fully on CUDA without surrendering its ownership or operating judgment.

The customer owns the market

Computraq's commercial idea is straightforward: let a customer retain the definition of a job while suppliers compete to perform it. A customer might need software maintained, documents processed, research checked or a manufacturing workflow analyzed. It records what an acceptable result looks like. Independent developers demonstrate that their AI Agents can deliver it. Qualified computing operators provide the execution. The customer reserves the service, pays an explicit price and keeps the evidence needed to make the next purchase.

A **benchmark** is the repeatable test of the job. An **Agent** is the package of models, tools and software that performs it. A **reservation** purchases bounded access to that capability for 24 hours. A **benchmark creator** is an independent expert who helps a customer decide what to buy. The marketplace connects these activities so that a better result can become a real purchase rather than a score that goes nowhere.

The customer can define and follow benchmarks without making a purchasing commitment. When it is ready to buy, its organization funds an account, sets purchasing authority and selects a qualified offer. It can buy directly through the website, work through Station, use a command-line interface, integrate the API or receive and act on notices through Slack. The interface changes; the requirement, purchase and operating record remain the same.

A private business with its own earnings

Computraq earns a 20% marketplace fee from the underlying Agent service price. The remaining 80% supports the developer and the agreed execution costs. Where the customer chooses a fee-bearing third-party benchmark, a separately displayed 10% fee pays that benchmark's creator. Hardware financing and vendor programs are complementary businesses; the principal commercial engine is the purchase and renewal of useful Agent services.

That gives Computraq a reason to widen the market. A small customer with a recurring problem can become valuable. A specialist developer can earn without building a large sales organization. An operator can obtain work from many applications rather than depend on one large tenant. NVIDIA benefits when those businesses have continuing reasons to use and invest in CUDA.

The industrial opportunity

The initial competitive program is to make a substantial share of frontier-laboratory workloads contestable through independently supplied services. The larger objective is an application economy at the scale of the digital services associated with Apple's App Store. The economic mechanism is familiar: independent applications make the underlying equipment more useful, and a common purchasing system makes those applications easier to commercialize.

The financing network already assembled around NVIDIA gives this program an industrial starting point. Operators, facilities, equipment suppliers and capital providers exist. Computraq brings the missing commercial coordination: customer-defined work, measurable qualification, reservations, delivery and collections. Capacity becomes useful service; useful service becomes recurring revenue; recurring revenue supports further investment.

This is also a program for American industrial reach. A worldwide population of developers and businesses choosing CUDA has economic significance beyond any individual shipment. The larger that population becomes, the more knowledge, investment and commercial opportunity accumulate around an American-led technical foundation.

1. Jensen Huang and Cornelius Vander Starr

Two institution-builders who expanded American influence by making other people's businesses possible.

The most useful comparison for Jensen Huang is not another semiconductor executive. It is Cornelius Vander Starr, the American entrepreneur whose insurance network became the foundation of AIG. Both saw that a product becomes much more valuable when it is surrounded by a working international institution: people who understand local needs, technical and financial capacity that can meet them, and relationships through which useful business can be conducted repeatedly.

Starr worked in insurance. Huang worked in computing. The common achievement was to make productive capacity accessible beyond the organization that created it. Customers did not merely receive a product; they acquired a way to pursue activities they could not otherwise organize as effectively.

Starr: from a Shanghai agency to an international institution

In 1919, Starr established American Asiatic Underwriters in Shanghai. An insurance agency connected customers needing protection with insurers able to supply it. The work required knowledge of businesses, property, transport and local conditions. It was commercial intermediation in the substantive sense: understanding an opportunity well enough to make a transaction possible. AIG traces its origins to that enterprise and incorporated as American International Group in 1967. [13]

That history contains two different organizational achievements. One was the growing group of insurance businesses that became AIG. The other was C.V. Starr & Co., established in 1950 to hold U.S.-based insurance agencies. As the AIG structure developed, C.V. Starr & Co. retained agencies that continued doing business with AIG's insurance subsidiaries. It remained a private company with a distinct commercial function. [49]

The agency business was substantial. AIG's 2004 proxy reported that AIG and its subsidiaries paid Starr and its subsidiaries **$173 million for the production of insurance business in 2003**. Those payments supported Starr's operating costs and originating-broker commissions. The private organization was being paid for creating and administering real business, not for lending a name to the public company. [48]

THE HISTORICAL FOUNDATION

The private agency role developed alongside the public institution.



Sources: AIG corporate history; AIG's 2004 proxy statement and historical account submitted to the SEC.

Exhibit 02. The early institutional history distinguishes the global insurance enterprise from the private agency company that complemented it.

Insurance capacity alone does not identify every useful transaction. A business owner needs someone who understands its particular exposure. A specialist agency learns a market, develops relationships, assembles information and brings suitable business to an insurer. The value is in making the institution's capacity usable in places where a centralized organization would otherwise struggle to recognize or serve the opportunity.

This is the commercial significance of Starr's model. The capacity-bearing institution and the market-making organization can be separate. They can serve the same economic system and reinforce one another without being the same corporation.

Huang: from a chip company to an industrial foundation

Huang cofounded NVIDIA in 1993. CUDA, introduced in 2006, made the parallel processing capability of NVIDIA GPUs available for general-purpose computing. The resulting platform connected hardware to software tools, libraries, engineering knowledge and a growing field of applications. [50,51]

A processor's theoretical performance is not enough to create an industry. An independent developer needs software that works, tools that make problems tractable, access to equipment and a place to deploy. A system builder needs compatible components and customers. An operator needs dependable systems, technical support and work it can sell. The investments become useful when those parties can rely on a common foundation.

NVIDIA's achievement was to make those relationships reinforce one another. More developers created more applications. More applications gave customers reasons to obtain CUDA capacity. More installed capacity supported a larger supplier and operator community. Those businesses then made the platform easier for the next developer to use.

CORNELIUS VANDER STARR AND JENSEN HUANG

American influence can grow through useful global participation.

CORNELIUS VANDER STARR

Insurance and international underwriting

Assets • counterparties • transport • local risk

JENSEN HUANG

Accelerated computing and industrial coordination

Developers • operators • suppliers • emerging work

Presence creates knowledge. Useful relationships create influence.

Exhibit 03. Starr connected knowledge to underwriting capacity. Huang connected developers and applications to a common computing foundation.

The institution was larger than any one product. A customer could begin with a workstation, develop a service, rent additional resources and later operate a larger system. A regional provider could serve workloads beyond its own software expertise. A specialist could build on the platform without having to manufacture the chip, write the compiler and establish an international support organization first.

The AI economy's real regulator

This paper calls Huang **the AI economy's real regulator** because he supplied the practical conditions under which a large population of independent businesses could compete. A market needs more than formal permission to enter. It needs usable tools, available productive equipment, dependable industrial supply and a path from an idea to a customer.

The term has an older American resonance. The eighteenth-century Carolina Regulators organized against abusive local fees and ineffective public institutions. The enduring question was who would establish a workable order when the existing institutions did not provide one. In modern computing, the instruments are engineering, equipment, industrial coordination and commercial relationships. [41]

Huang's contribution was to organize those instruments into a functioning economy. Smaller firms could specialize because they did not have to recreate the foundation. Public institutions could describe the value of competition; NVIDIA's ecosystem made more of that competition operationally possible.

Five functions that the market needed

Technical coordination made an advanced processor a usable platform. Compilers, libraries, debugging, model-serving tools and accumulated engineering knowledge lowered the amount each independent team needed to invent. Compatibility became a common investment rather than a private problem repeated by every customer.

Industrial coordination connected demand to manufacturing, memory, networking and complete systems. Suppliers could invest against a platform whose software and customer ecosystem gave their output a purpose. Equipment makers and distributors could serve smaller buyers within that common foundation.

Commercial access created a progression from personal equipment to specialist operators and large deployments. A developer could begin with a prototype, rent additional resources and reach customers without transferring its business to one integrated infrastructure owner.

Financial support helped independent operators acquire capacity before they had the distribution reach of the largest clouds. NVIDIA's disclosed cloud-purchase arrangements and prospective financing initiatives are part of that structure. Their greatest strategic value is the field of counterparties and operating capability they helped assemble. [37,38]

Global participation gave American technology a presence wherever useful computing was being developed. Engineers learned a common platform. Suppliers, customers and operators acquired economic reasons to maintain compatibility and continuing trade. Their interests crossed borders even when political relationships were strained.

Taken together, these functions supplied a practical counterweight to concentration. The company selling the common technology did not have to own every application or every customer. Independent businesses could cooperate around CUDA while continuing to compete with one another.

A commercial network is a source of national strength

Starr's international reach had strategic significance because commercial relationships produced knowledge of places, industries and counterparties. U.S. National Archives records identify wartime work involving Starr and his companies gathering intelligence in the Far East. The network's reach had been accumulated through years of useful commercial presence. [14]

Huang's international network has a corresponding economic significance. Developers expose emerging technical needs through the problems they solve. Suppliers reveal practical constraints through production and deployment. Operators show where demand is forming. Customers show which capabilities justify investment. Continuing participation keeps an American-led institution close to the world's productive frontier.

This is a deeper source of influence than a single export sale. Engineers invest in skills. Businesses build products around the platform. Suppliers develop complementary equipment and services. Customers acquire reasons to keep improving the relationship. The resulting network supports continued American participation in industries that develop beyond the boundaries of any one country.

Computraq extends that logic. By giving independent applications a common route to customers, it adds another reason for businesses around the world to build and earn within the CUDA ecosystem. Commercial participation becomes more durable when it supports the participants' own income.

From the founder parallel to the company parallel

There are therefore **two comparisons**, not one. The first concerns institution-builders: Cornelius Vander Starr and Jensen Huang. The second concerns the structure that allows an institution to keep expanding: AIG alongside C.V. Starr & Co., and NVIDIA alongside an independent Computraq.

The first explains how the technical and industrial foundation came to exist. The second explains how a separate private company can cultivate the next market around it. Computraq need not become NVIDIA to strengthen the NVIDIA ecosystem, just as specialized commercial origination did not have to sit entirely inside AIG to make its underwriting capacity more productive.

Huang built the foundation on which independent computing could flourish. Computraq's purpose is to organize the business that makes that independence pay.

2. The company beside the company

A private commercial institution can extend a public industrial platform without becoming part of it.

C.V. Starr & Co.'s historical position beside AIG provides a particularly useful way to understand Computraq. AIG's insurance companies supplied underwriting capacity: the financial ability to accept risks and pay claims. Starr's specialist agencies supplied knowledge of particular markets, relationships with customers and brokers, and the work required to turn that knowledge into insurance business. The public company and private organization participated in the same economic system while performing different functions. [48,49]

Computraq applies that division of work to accelerated computing. NVIDIA develops the processors, system technologies, software and industry relationships that make CUDA a productive foundation. Computraq organizes the market in which customers purchase useful AI services built on that foundation. The operator, rather than NVIDIA itself, supplies a particular reservation's computing. The developer supplies the Agent. The buyer determines what is worth purchasing. The resulting activity makes the underlying NVIDIA ecosystem more valuable.

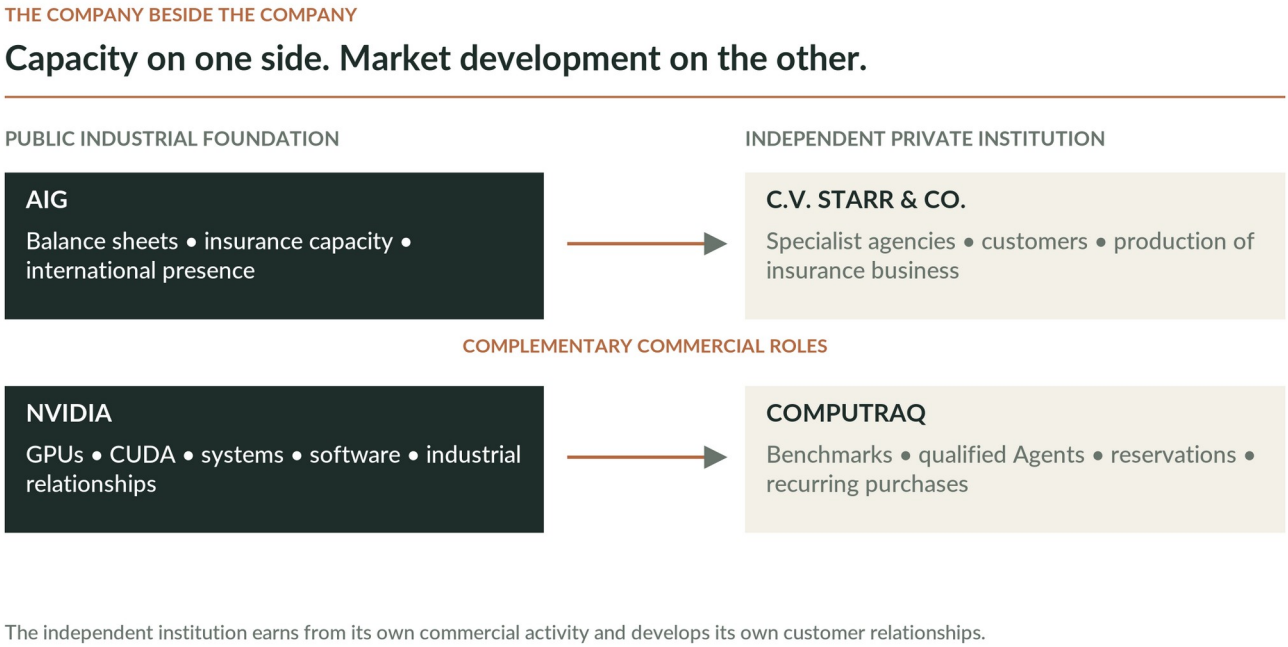


Exhibit 04. AIG and Starr supplied different commercial functions. NVIDIA and Computraq can do the same in accelerated computing.

Why the separate organization matters

A large technology company has to coordinate decisions across hardware generations, supply agreements, global accounts, software roadmaps and capital allocation. A new activity inside that company competes with all the other demands on management. Even an attractive opportunity can remain too small relative to the parent, too specialized for its sales organization or too different from its established measures of success.

For an independent business, the same opportunity can be the central purpose of the company. A market for specialist AI services does not have to compete internally with the next processor launch. A regional operator does not need to become a strategically significant NVIDIA account before it can matter to Computraq. A

business buying a modest recurring Agent reservation can be worth serving because thousands of similar relationships can become a valuable franchise.

Private ownership also gives Computraq its own time horizon. It can invest in the accumulated value of customer requirements, benchmark libraries, developer relationships, repeat procurement and operating knowledge. Those assets strengthen with use. They need not be managed as a short-term extension of a hardware-sales campaign or as a project whose budget depends on another company's quarterly priorities.

The result is a division of institutional attention. NVIDIA can continue developing the technical and industrial foundation. Computraq can devote its entire business to making that foundation commercially useful to a larger population of people.

Alignment through earning, not ownership

Computraq's commitment to CUDA is a commercial choice. A common execution foundation lets the market concentrate its technical effort, makes operator configurations easier to reproduce and gives developers a more coherent route from a working prototype to a paid service. Operators can invest in a defined stack. Buyers can compare results without requiring every procurement to begin with an accelerator-architecture study.

That commitment does not require common ownership. The organizations can align because the same transaction creates value for both. A customer buys useful work. The developer receives proceeds. The operator receives execution revenue. Computraq earns a marketplace fee. CUDA equipment becomes more productive, and its owners acquire a commercial reason to maintain, replace and expand it.

The relationship is strongest when these benefits repeat. One sponsored launch can create attention. A recurring market creates a business constituency: customers who rely on the service, developers who earn from it, operators who invest to supply it and vendors whose equipment has a route to paying work.

THE COMPOUNDING MECHANISM

Useful work creates the reason to invest again.



The economic engine is a continuing sequence of purchases—not ownership of the entire supply chain.

Exhibit 05. The commercial return travels through independent businesses. Their separate earnings make the relationship durable.

What Computraq owns

The private institution needs an actual commercial franchise. Its value is not an introductory relationship with a hardware manufacturer. It is the ability to organize and operate a market that customers and suppliers continue choosing.

Computraq owns and develops its marketplace software, customer-account infrastructure, reservation and settlement systems, commercial relationships and operating brand. It invests in the methods by which requirements become comparable offers and offers become dependable purchases. Its employees and owners make its pricing, product and partnership decisions.

Participants retain the assets that make them independent. The customer retains control of its requirements, data access and spending. Developers retain their Agent products and intellectual property. Benchmark creators retain their methods and maintain their reputation through useful advice. Operators retain their equipment and delivery businesses. The market brings these assets into a productive relationship rather than requiring their consolidation under a single owner.

Institution or participant	Its principal asset	Its commercial decision
NVIDIA	CUDA technology, systems and industrial relationships	What technology and ecosystem capabilities to develop
Computraq	Purchasing market, distribution, operating systems and customer relationships	How to connect qualified supply with recurring demand
Customer	Business requirement, data and purchasing authority	Which result is worth buying and on what terms
Agent developer	A maintained, versioned capability	What work to perform and how to improve its economics
Benchmark creator	Evaluation expertise and buyer confidence	Which methods and recommendations earn continuing selection
Operator	Equipment, facilities and execution competence	Which services to supply at a sustainable return
Vendor or financier	Equipment, support or capital	Which productive operators and assets to serve

Complementarity is more valuable than organizational duplication

Computraq does not need to manufacture an alternative GPU, build another frontier laboratory or acquire a national data-center estate to become strategically valuable. It needs to make the existing independent ecosystem much better at reaching customers.

A new procurement method can unlock more useful demand than another undifferentiated capacity announcement. A benchmark shared by a hundred businesses can give a small specialist a market it could not have reached through individual sales calls. A deployment configuration accepted by several operators can make a developer's service available without requiring the developer to run a cloud. A consistent record of delivery and payment can help an operator obtain capital on the strength of its own business.

These are commercial functions. Their value comes from lower friction, greater specialization and more transactions worth completing. They are the computing equivalent of finding a market, understanding the risk and arranging a business relationship that an insurance balance sheet alone cannot create.

Collaboration can begin with work

The first useful collaboration does not have to be a corporate transaction. An NVIDIA engineering team can help make an Agent deployment more efficient. A system builder can provide a reproducible operator configuration. An operator can make qualified capacity available. A customer can supply a recurring requirement. A financing partner can help an operator expand after its services earn repeat purchases.

Those collaborations can become larger commercial programs. Technical support can be organized around common workloads. Vendors can offer equipment-and-service packages to qualified operators. Financing partners can create programs for equipment serving verified demand. NVIDIA can support an ecosystem initiative while Computraq retains the market and the participants retain their businesses.

COMMERCIAL COLLABORATION

A relationship can deepen without an acquisition.



Computraq retains its own customers, economics and decisions at every level of collaboration.

Exhibit 06. The relationship can deepen through successive commercial achievements, without requiring an acquisition.

The advantage of this progression is that collaboration has a product at every stage. The parties can point to a working deployment, a customer purchase, a profitable provider or a replicable equipment program. The relationship grows because it produces something, rather than waiting for a comprehensive corporate agreement to make it meaningful.

Institutional continuity

A technical ecosystem has a longer life than a single product cycle or leadership arrangement. Developers accumulate skills. Customers establish workflows. Operators invest in equipment and facilities. Vendors build support businesses. A complementary private institution gives those relationships another place to compound.

Computraq can maintain its own customer programs and product priorities as NVIDIA changes its organizational focus. It can keep serving specialist markets even when those markets are not the manufacturer's largest accounts. It can organize developers around demand even when a public-company debate is focused on something else.

This is the durable significance of the Starr parallel. A useful commercial institution can remain close to a larger enterprise's economic world without being contained within its corporate boundaries. Its contribution is to develop parts of the market that deserve a dedicated owner, a different decision-making horizon and their own source of income.

Computraq's independence is not distance from NVIDIA's success. It is the organizational form through which Computraq can add to that success.

3. Computraq: the customer owns the market

A first-principles guide to the business, the participants and the transaction.

Computraq is a market for useful AI work. Customers describe jobs and decide what counts as a good result. Independent developers offer Agents that can perform those jobs. Computing operators supply the machines and operating environments. Specialists create benchmarks that help buyers compare the offers. Computraq connects discovery, evaluation, purchasing, delivery and payment.

The distinguishing feature is that the customer's requirement remains in place while suppliers can change. A manufacturer might need incoming quality reports checked every morning. A software company might need a defined class of maintenance issues resolved. A professional practice might need documents extracted and compared. The buyer does not need to become permanently attached to the first model or supplier that solves the problem.

This is a **standing contest**. Each purchase has an incumbent, but the next purchase remains open to a better qualified offer. The customer can keep a good supplier for years while retaining a practical alternative. Competition is expressed through usable services, not merely through the number of model names on a comparison chart.

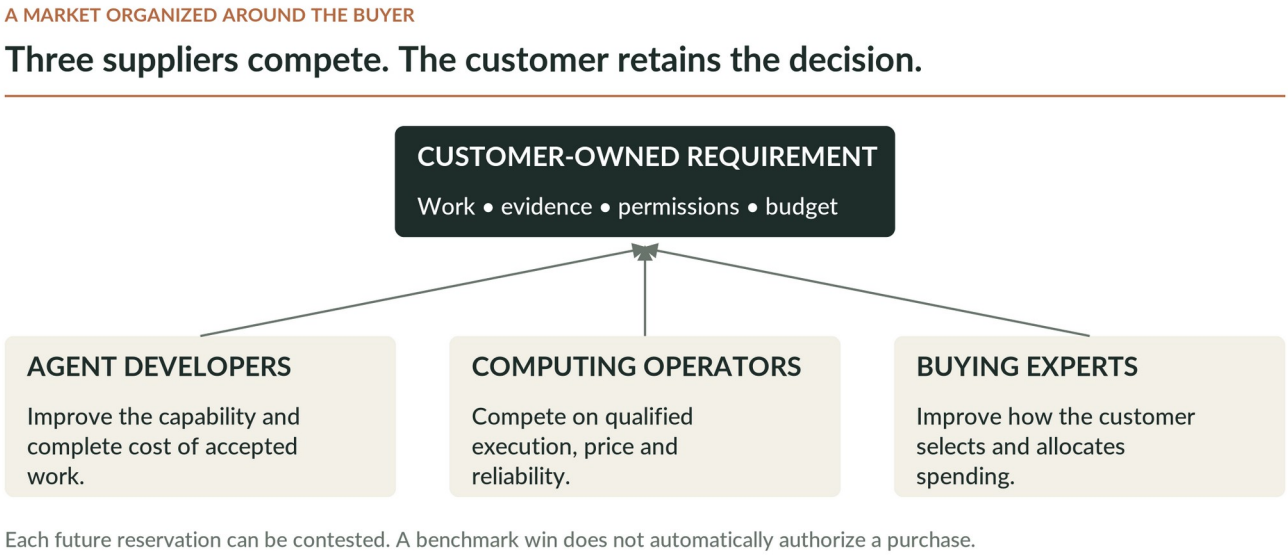


Exhibit 07. The requirement stays with the customer. Developers, operators and buying experts compete around it.

Five terms explain the core product

Term	In ordinary language	A concrete example
Benchmark	The test that defines acceptable work	Extract specified fields from sample invoices and identify exceptions
Agent	The software capability offered to perform the job	A maintained invoice-processing workflow using models, tools and rules
Operator	The business supplying the computing environment	A regional CUDA provider running the qualified configuration
Reservation	The period and allowance of service purchased	Twenty-four hours of access within the offer's stated capacity
Benchmark creator	The expert who maintains a buying standard	An accounts-payable specialist evaluating offers for many businesses

A model is a component of an Agent, not the whole commercial product. A useful Agent includes the instructions, software, tools, interfaces and maintenance necessary to perform the particular work. The same model can be used badly in one workflow and well in another. A smaller model can be the better commercial choice when the surrounding system is more effective.

The operator is also distinct from the developer. A developer can build and support the Agent without owning a data center. An operator can supply computing for many developers without inventing every application. A business can perform both roles, but it does not have to do so.

What the customer actually buys

The standard unit is a **24-hour Agent reservation**. The offer describes the exact capability, version, operating conditions, resource allowance, price and support responsibility. The buyer compares how much acceptable work it expects that reservation to deliver.

This structure makes AI easier to procure as a service. A buyer can understand the price of a day's capability, the work it supports and the conditions under which it operates. It can reserve successive days, plan a recurring allocation or negotiate a larger program while retaining the underlying qualification and purchasing record.

Accepted output is the measure of value. The reservation is the commercial unit. That distinction allows a buyer to compare services honestly: an offer that produces more acceptable work with less checking may justify a higher daily price. A lower advertised daily price is not automatically the cheaper way to complete the job.

Free discovery creates a market before the first purchase

A customer can create or follow benchmarks before committing paid demand. It can register an interest in an industry workflow, develop a private test for its own organization or follow a specialist's permitted benchmark. Developers then have a clearer picture of what customers want and how those customers will judge an offer.

Results are connected to the relevant audience. The owner of a private benchmark receives submissions for that requirement. Followers of a creator-maintained benchmark receive improvements that the creator has accepted

under its method. Notices can arrive through the website, email, Slack or a webhook that informs another business system.

This creates a different form of distribution. The customer does not have to search continuously for every new model or developer. The market already knows the requirement. A qualified improvement can find the customers for whom it matters.



Exhibit 08. One account connects discovery, approval, service use and the next improvement.

The market has three independent competitions

Developers compete to produce the best complete capability. Operators compete to deliver compatible execution at attractive cost and service quality. Benchmark creators compete to improve the customer's purchasing decisions. Computraq connects these competitions without requiring one firm to perform all three roles.

The result is a productive division of labor. A domain expert who understands an industrial process can build an evaluation business. A technically excellent small team can build a specialist Agent. An experienced operator can supply reliable CUDA execution. The customer can assemble a better outcome from their separate strengths.

The same architecture makes opportunity extend beyond individual people. Consultancies can publish purchasing standards. Equipment vendors can package qualified environments. Trade associations can maintain shared industry benchmarks. Universities and research groups can offer reproducible specialist capabilities. Each becomes part of the market through a defined contribution rather than through corporate affiliation.

Why this is a business rather than a directory

A directory shows that a service exists. A leaderboard shows a reported performance result. A cloud catalogue lists available resources. Computraq connects those pieces to a purchase that can actually be executed and renewed.

The buyer needs to know which version was evaluated, where it can run, what it will cost, what information it can access and how it will be supported. The developer needs a route from better performance to customers. The

operator needs reservations it can deliver. The benchmark creator needs a way to be paid for useful advice. The marketplace's role is to make those commercial relationships fit together.

That is also why the private institution can become valuable independently of NVIDIA. Its assets are the market's recurring relationships and operating knowledge. CUDA supplies the common technical foundation; Computraq turns that foundation into customer purchases and participant earnings.

The CUDA tax and the competition that public taxes did not purchase

"CUDA tax" is a critical nickname for the premium associated with choosing NVIDIA's hardware and computing ecosystem. The expression implies that the premium pays principally for dependence. Computraq's position is the opposite: an important part of the premium supports the practical infrastructure of independence—tools, expertise, supplier diversity, deployment alternatives and a worldwide population of operators.

Buyers have, in effect, outsourced some of the work of sustaining a competitive computing economy to NVIDIA and Jensen Huang. The private price purchases common capabilities that smaller firms could not economically build alone. Public tax revenues fund agencies charged with protecting competitive markets, yet policy can push those same firms toward politically connected custodians of capital, models and distribution. The comparison is therefore not simply an expensive chip against a cheap chip. It is a functioning field of independent enterprise against dependence on a narrow class of approved providers.

A productive system should be valued by the work and alternatives it makes possible over its lifetime. Suppose a CUDA deployment costs $1.2 million more than another option but generates $2 million more contribution through accessible applications, lower integration effort and additional customers. The premium creates $800,000 of net value. Computraq adds a visible commercial component to this calculation: a market of customer requirements that a qualified CUDA service can compete to satisfy.

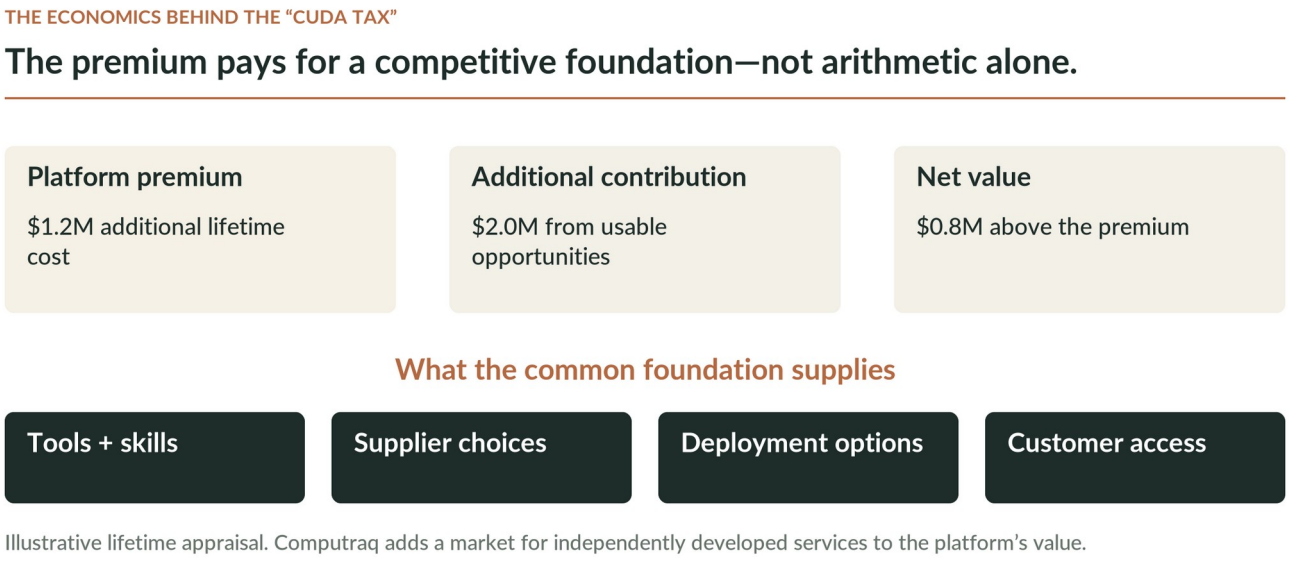


Exhibit 09. The position paper treats the CUDA premium as payment for competitive industrial infrastructure that individual small firms cannot reproduce.

The platform premium is consequently an investment in the means to compete. It finances a common foundation while allowing separate owners to decide what to build. That economic organization is the valuable part of the CUDA ecosystem. Computraq's purpose is to give it a larger, more durable source of customer revenue.

4. From the first requirement to a recurring purchase

How a buyer enters, selects an Agent, uses it and retains the power to change.

The simplest way to understand Computraq is to follow one customer. Consider a company that processes supplier invoices. It already uses AI, but staff still spend too much time checking outputs and resolving exceptions. Its problem is not access to another chatbot. It needs a complete workflow that handles a repeatable business task more economically.

1. Describe the work in the customer's own terms

The operations team begins with the documents it receives, the fields it needs, the systems that consume those fields and the exceptions that require a person. Finance records the present cost of the work. A business owner defines the acceptable result: accurate extraction, evidence linking values to the original document, a useful exception queue and an output that the accounting system can accept.

The customer can bring this requirement directly, open an RFP to solicit a proposed approach or select an existing benchmark maintained by an accounts-payable specialist. The point of entry can be simple. The resulting requirement becomes progressively more precise as the customer and evaluator define the work.

An RFP is particularly useful where the buyer knows the outcome but not the implementation. Responses can explain the workflow, proposed Agent, expected service allowance, integration and support. The chosen evaluation then becomes a reusable purchasing standard rather than a demonstration that disappears into a sales archive.

2. Establish the benchmark and purchasing authority

Representative documents become test cases. The benchmark records success criteria, difficult examples, review requirements and complete cost. The customer controls which information is visible to suppliers. A private test can remain within an authorized evaluation process while publishing a limited description of the requirement and the resulting qualified offers.

At the same time, the organization identifies who can view results, approve access, authorize spending and receive operating reports. A department's business owner can make the purchasing decision without granting every person who follows a benchmark the ability to spend the company's money.

The initial account can remain entirely on the discovery side. A paid program is added when the buyer is ready to reserve a service. That preserves a low-friction way to learn about the market while keeping the actual purchase explicit.

3. Invite and compare qualified offers

Developers submit an identified Agent version and the evidence supporting its result. A useful submission connects the benchmark to a purchasable offer: the configuration tested, its execution requirements, the daily price, the capacity included and the support available.

The buyer compares complete results. Suppose one Agent extracts fields quickly but creates a large review queue. Another runs more slowly but produces more accepted records with clearer evidence. The second can be the more valuable offer. The benchmark makes that trade-off visible in terms the buyer already uses.

An improved result can trigger a notification. For a benchmark the customer owns, the customer evaluates the submission. For a specialist's benchmark, the specialist reviews the result and publishes accepted improvements to its followers. Both routes preserve the connection between evidence and the actual decision-maker.

A PURCHASE ACROSS FIVE PARTICIPANTS

One requirement coordinates the entire transaction.

	REQUIREMENT	QUALIFICATION	PURCHASE + USE
CUSTOMER	Defines the job	Accepts evidence	Authorizes and uses
BENCHMARK EXPERT	Builds the test	Reviews submissions	Updates buying advice
DEVELOPER	Prepares the Agent	Submits results	Supports its service
OPERATOR	Lists capacity	Proves execution fit	Runs the reservation
COMPUTRAQ	Stores the requirement	Connects offers to tests	Coordinates booking and pay

The customer's buying authority remains distinct from the people producing scores, software and execution.

Exhibit 10. The buyer's work passes through a continuous commercial process, not disconnected demonstrations and contracts.

4. Fund the relationship and approve the offer

The customer selects the commercial plan appropriate to its expected use. A first purchase can use a limited prepaid balance. A recurring business can choose a monthly arrangement. A larger enterprise can commit annual or multiyear spending across a portfolio of requirements. The funding arrangement and purchasing limits are recorded on the account, while individual reservations identify the selected service.

An illustrative $3,600 annual plan can provide twelve $300 monthly purchasing allowances. That is a budget for the customer's qualified purchases, not a requirement to keep buying the same Agent. The allowance can pay the accepted service price and any selected benchmark-creator fee within the customer's approved limits.

The quote shows the total before approval. A $75 Agent day using the buyer's own benchmark costs $75 before tax and any separately identified items. The same service purchased through an accepted fee-bearing creator benchmark adds $7.50, producing an $82.50 service-and-creator total. The platform fee is already within the $75 service base. It is not added again at checkout.

5. Reserve the capability and connect the work

The reservation identifies the Agent version, selected operator, start time, duration and operating allowance. A standard Run lasts 24 hours. A customer needing continuous coverage can reserve a sequence of Runs or use an agreed recurring program. A business requiring more parallel work obtains the stated capacity rather than assuming that a daily price includes unlimited simultaneous activity.

The customer connects the approved inputs and output destination. For the invoice example, that might be a document inbox and a structured output folder, followed by an accounting-system integration. A coding service might receive access to a particular repository and create a proposed change for review. A research service might receive a specified document collection and return a report with source references.

The customer sees what has been reserved and how to use it. The same underlying reservation can be reached through the website, Station, Slack, an API integration or the command line. A large organization can place Computraq behind an existing internal application so staff use the capability without managing each supplier relationship themselves.

6. Run the service and observe the result

The selected operator starts the qualified environment. The Agent receives the customer-authorized inputs and operates within the purchased allowance. Status, usage, outputs and exceptions become part of the service record. The buyer can see the relationship between the purchased capability and the work actually completed.

The record matters for practical management. An operations manager needs to know how much of the queue was processed, which cases require attention and whether the result meets the standard. A developer needs to know where its workflow can improve. The operator needs to know whether the environment is delivering the promised capacity. The same service event supplies useful evidence to all three.

7. Close the reservation and retain the purchasing knowledge

At the end of a Run, the service follows its recorded persistence policy. The customer retains the outputs, accepted evidence and operating record it has chosen to keep. Temporary execution and access end with the reservation unless the next authorized reservation continues the work. Charges, credits and participant allocations are recorded against the transaction.

The useful knowledge remains. The customer’s test, previous results and total costs support the next purchasing decision. An incumbent that keeps improving can retain the work. A challenger that meets the same requirement more economically can win a future reservation without forcing the customer to restart its entire procurement process.

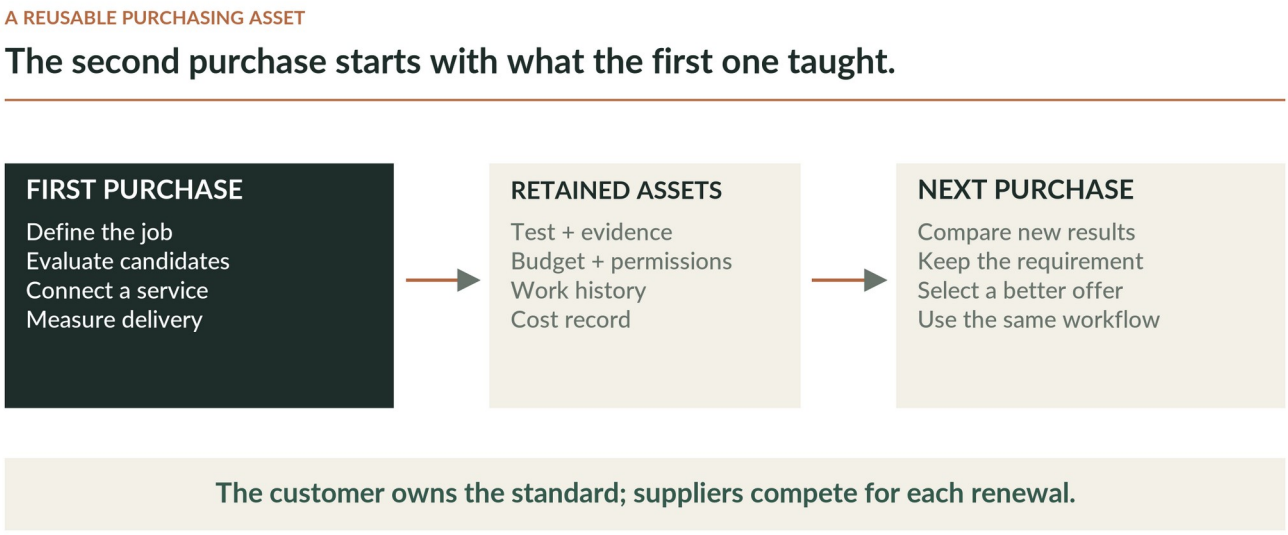


Exhibit 11. The first purchase builds an asset for the next one: requirements, evidence, interfaces and spending history remain reusable.

The account can support a portfolio

The invoice workflow can be the first of several requirements. The same organization might add contract intake, software maintenance, translation or internal research. Each has its own standard, suppliers, permissions and operating costs. Purchasing authority can remain centralized while responsibility for the work stays with the relevant department.

That is how a material second source develops. The buyer is not asked to replace every existing AI supplier on the first day. It begins with well-defined work, adds qualified alternatives and directs a growing share of spending toward better results. Its account becomes a managed portfolio of capabilities rather than a collection of unrelated model subscriptions.

Common first-purchase questions

Does the buyer need to choose a model? It can specify a model where the requirement calls for one, but it can also specify the result and let developers compete on implementation.

Does the buyer need to operate a GPU? No. A qualified operator supplies the execution. Equipment ownership remains an option for participants who want to operate infrastructure, not an entry requirement for customers.

Can an existing consultant remain involved? Yes. The consultant can maintain benchmarks, integrate the selected service, provide operating support or develop an Agent. Its contribution becomes easier to compare and reuse.

Can a purchase be automated? A customer can grant a standing purchasing mandate with defined conditions and spending limits. Notifications alone do not spend money; the mandate or an explicit approval does.

What happens when a better Agent appears? The customer receives relevant evidence and can authorize a future switch while retaining the same underlying requirement.

From RFQs and RFPs to a standing purchasing standard

An **RFQ**, or request for quotation, asks for a price and delivery terms for a reasonably defined purchase. An **RFP**, or request for proposal, asks a supplier to explain how it would solve a problem. A pilot tests a limited implementation. A statement of work records deliverables; a service-level agreement records operational commitments. These instruments establish a relationship, but they do not by themselves make different AI services comparable.

A model vendor may quote a price per token. A token is a small unit of text or other encoded input or output processed by the model. The bill measures consumption. It does not tell a buyer how many attempts will be needed, how much human checking will remain, whether the result meets the business standard, or how much expert labor is embedded in the demonstration. Cheap model consumption can be expensive work. An expensive call can be economical when it avoids retries and review.

The buyer therefore pays to discover the real product. It explains its processes, negotiates data access, commissions a pilot, evaluates exceptions and asks the vendor to adapt the system. The next vendor repeats much of the same work. Developers face the inverse burden: find a buyer, determine its needs, prepare evidence, answer security questions and build a custom demonstration before knowing whether the buyer can recognize the value.

Computraq changes the order. The requirement and evaluation method become an asset of the customer relationship. Relevant suppliers compete against that asset. Their evidence identifies the exact Agent version, supported conditions and purchasable offer. Commercial terms, permissions, execution and payment connect to the same decision. The next improvement enters an existing contest instead of starting a new sales process.

The resulting market is more than an evaluation service. An isolated benchmark tells a customer something about performance. A purchasing market makes that information actionable. The customer can select the qualified offer, authorize a bounded reservation, receive the service and retain the result for its next decision. Evidence becomes a route to revenue.

An evaluation that survives contact with production

A commercially useful benchmark includes difficult cases and the full cost of completing them. It specifies acceptable accuracy, required evidence, permitted data use, response time and escalation. Some examples should remain unavailable for supplier tuning. The evaluation records which configuration performed the work so that a different deployment is not sold under the same score.

For private enterprise work, repeatability does not require public data. An authorized evaluator can run controlled tests while publishing only the information the buyer permits. Nor does the method demand that every problem become one number. Quality, review effort, speed, resource use and reliability can remain separate measures, with the customer defining the trade-offs it accepts.

The purchasing standard should also identify the consequences of failure. Extracting a field incorrectly, changing production software without approval and generating an unsupported research claim are different events. A credible comparison tests the actual workflow and its review boundary. An offer wins because it meets the customer’s requirements more economically, not because it quietly performs a weaker task.

These practices make the contest durable. The incumbent cannot permanently define quality in its own terms. The challenger has a way to prove a useful improvement. The customer has a reason to change suppliers without abandoning the knowledge acquired during previous purchases.

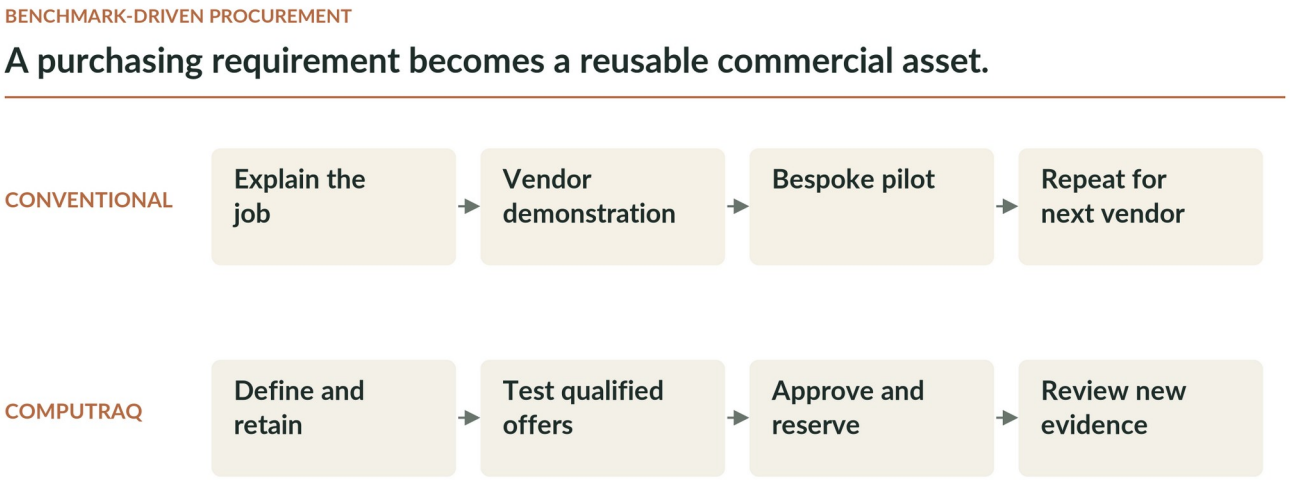


Exhibit 12. The same customer-owned standard can support repeated evaluations and purchases instead of isolated sales cycles.

5. Inside the operating model

The machinery that turns a benchmark result into a dependable commercial service.

The product can be understood as four connected layers. The **market layer** records requirements, evidence and offers. The **account layer** records identity, purchasing authority and money. The **execution layer** starts and operates the selected service. The **service record** connects the result back to the next purchasing decision. The operating design gives these layers one continuous transaction history.

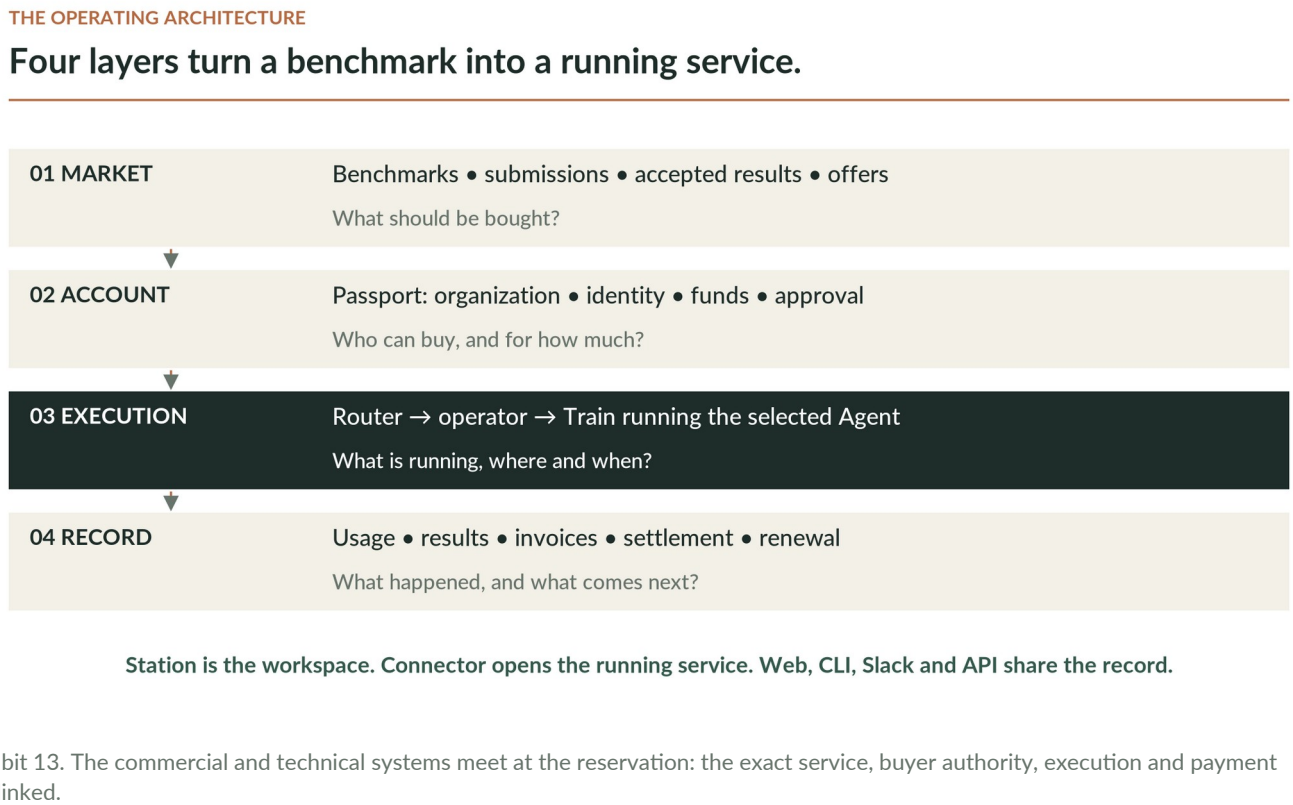


Exhibit 13. The commercial and technical systems meet at the reservation: the exact service, buyer authority, execution and payment are linked.

The named components

Station is the working environment for people managing participation in the market. Buyers review requirements, notices, reservations and results. Developers and operators manage the capabilities they offer. The same participant can purchase one capability while supplying another; Station is therefore more than a customer dashboard.

Train is the execution unit that carries an Agent in its declared operating environment. One Train runs one Agent. Larger deployments can run several Trains where sufficient capacity has been allocated. This makes a reservation intelligible: a particular capability is being supplied under a defined resource and service commitment.

Connector gives a user or an authorized application a practical route into the running service. It supports the terminal-oriented workflows important to developers and technical customers. The commercial relationship does not disappear when the user moves from a web page to a command line.

Router connects the reservation to eligible execution. The selected service has to reach an operator with the right configuration, available capacity and accepted delivery terms. It also has to remain tied to the customer's funding and operating instructions. Routing is therefore a purchasing and delivery decision, not merely a network address.

Passport provides the identity and account context that follows the transaction: which organization is acting, which person or system has authority, and which account funds the purchase. A consistent identity and payment record lets the same organization use several services without rebuilding its commercial identity every time.

The names describe parts of one product. A first-time customer need not learn the internal vocabulary to buy an Agent. The interfaces present the requirement, offer and service; the components make those experiences coherent.

What an Agent package contains

An Agent is delivered as a versioned package. A container is the practical packaging format: it groups the application and its software dependencies so that the same version can run in a known environment. The package identifies the model dependencies, application code, tools, configuration and entry point needed to perform the job. Large model files can be referenced by fixed versions rather than being copied into every application image.

Its offer also describes the resources it needs. These include GPU memory, system memory, storage, networking and the supported software environment. The developer identifies the operating assumptions under which the benchmark result was obtained. An operator then offers a compatible place to run that capability.

The connection between package and evidence is essential to the product. A buyer purchasing the evaluated version should receive that version under the relevant conditions. A materially changed model, tool or configuration creates a new version to assess rather than silently inheriting an older result.



Exhibit 14. A purchasable Agent combines software, evidence, execution requirements and commercial terms.

How execution is separated

Packaging and operating separation perform different jobs. The container packages the Agent. The operator's execution environment supplies the resource allocation, operating boundary and access controls used to deliver the offer. A deployment profile can use dedicated hosts, virtualized execution or other supported isolation arrangements appropriate to the service. The offer identifies the profile the customer is buying.

The GPU allocation must be as explicit as the CPU and memory allocation. A service may need a whole GPU, a supported partition or a declared share of capacity. The operator cannot promise the same unrestricted hardware allowance to several concurrent reservations. Declared capacity determines how many Trains can operate while meeting their offers.

Private customer data follows the selected workflow. Some jobs require copying an input into the execution environment; others use approved remote connections. Outputs go to the customer's chosen destination. Long-lived operating data is treated separately from temporary working files, so the end of a reservation has a clear effect on what remains and what is removed.

These details make the offer buyable. A procurement team can compare location, data handling, available capacity and support. A technical team can reproduce the service and understand the deployment. The developer and operator can identify which part of the delivery each is responsible for.

A reservation has a complete lifecycle

Before execution, the market identifies the offer and the account authorizes the purchase. Capacity is reserved. The operator prepares the environment, obtains the specified package and applies the accepted configuration. The customer's approved access is attached to that reservation.

During execution, the service records status and use. The customer receives results and exceptions through the selected interface. A failure has an identifiable owner and remedy: an application problem returns to the developer; an execution problem returns to the operator; an account or payment problem returns to Computraq's commercial system.

At completion, the customer receives the retained deliverables. Temporary access closes according to the offer. Usage and charges are reconciled. Applicable credits or refunds change the same transaction record rather than becoming an unconnected manual adjustment. The result can then inform the buyer's next comparison.



Exhibit 15. Each stage has an owner and a recorded outcome, from qualification through settlement.

Three routes through the same transaction

There are three different things moving through this system: **decisions**, **work** and **money**. Keeping them connected is the marketplace's operating advantage.

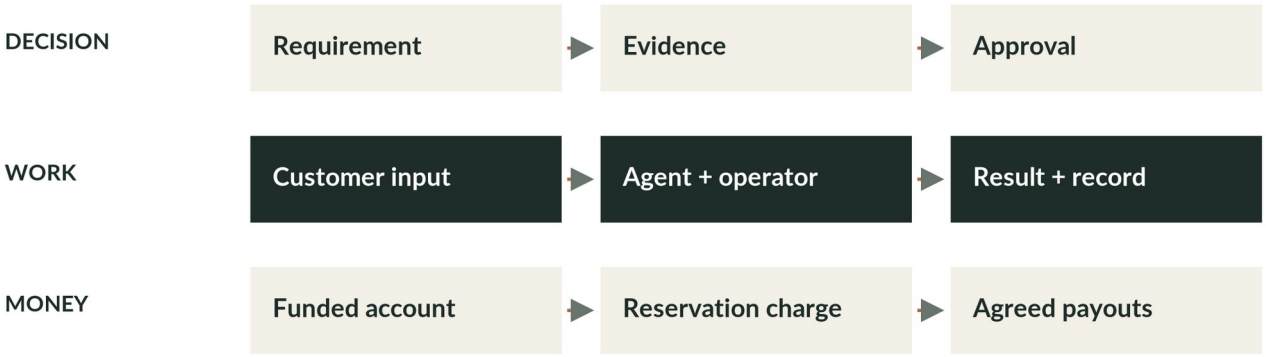
The decision route begins with a requirement and ends with the authorization of a particular service. Benchmark results, expert recommendations and notifications inform that decision. The customer's mandate determines what happens next.

The work route begins with authorized inputs and ends with outputs the customer can use. It passes through the selected Agent and operator. The system records which configuration performed the work and how it performed.

The money route begins with the customer account and ends with the agreed participant allocations. The service base, platform fee, execution amount and any creator fee have distinct places in the record. A reservation can therefore be followed from the customer's approval to the participants' proceeds.

THREE ROUTES, ONE TRANSACTION

The decision, the work and the payment are linked—not confused.



A common reservation ID joins the work and account records without making every participant the same business.

Exhibit 16. Decisions, work and money move along different paths, joined by the reservation record.

Web, Slack, command line and API

Different buyers work in different places. A procurement owner may prefer the website. A software team may want notifications and approvals inside Slack. An engineer may operate a service from a terminal. A large enterprise may want its own application to create requirements, check accepted results and reserve approved services through an API.

Those are interfaces to the same market. An API integration should refer to the same benchmark and offer identifiers that a person sees in the website. A Slack approval should create the same reservation and spending record. A command-line user should reach the same service history rather than a separate commercial product.

Webhooks make the relationship continuous. A business system can receive an event when a relevant result is accepted, a reservation starts, a service needs attention or a Run completes. A customer's own automation can then decide what to do under the authority already granted to it.

The commercial plan can change without erasing the customer

Accounts retain the terms accepted for their purchases. That allows Computraq to introduce a new plan for new customers while continuing to serve an existing organization on its current arrangement. A customer can move between prepaid, monthly and committed programs without losing its requirements, history and supplier relationships.

This is commercially important in a developing market. The right way to fund a trial is different from the right way to fund a recurring enterprise portfolio. The account system accommodates that difference while preserving continuity. The participant does not have to start again because the marketplace improves its commercial offering.

The same principle applies to delivery. New Agent versions, additional operator profiles and improved interfaces can enter the market while previous purchases retain an intelligible record. A durable institution accumulates knowledge through change instead of discarding it at every product revision.

6. How customers and partners connect

The private institution earns by making separate businesses work together.

Computraq's institutional position becomes concrete at the point where each participant knows what to bring, whom to work with and how to earn. A customer need not negotiate an independent relationship with every hardware, software and financial organization behind a service. The marketplace coordinates the purchase, while the specialist relationships remain economically distinct.

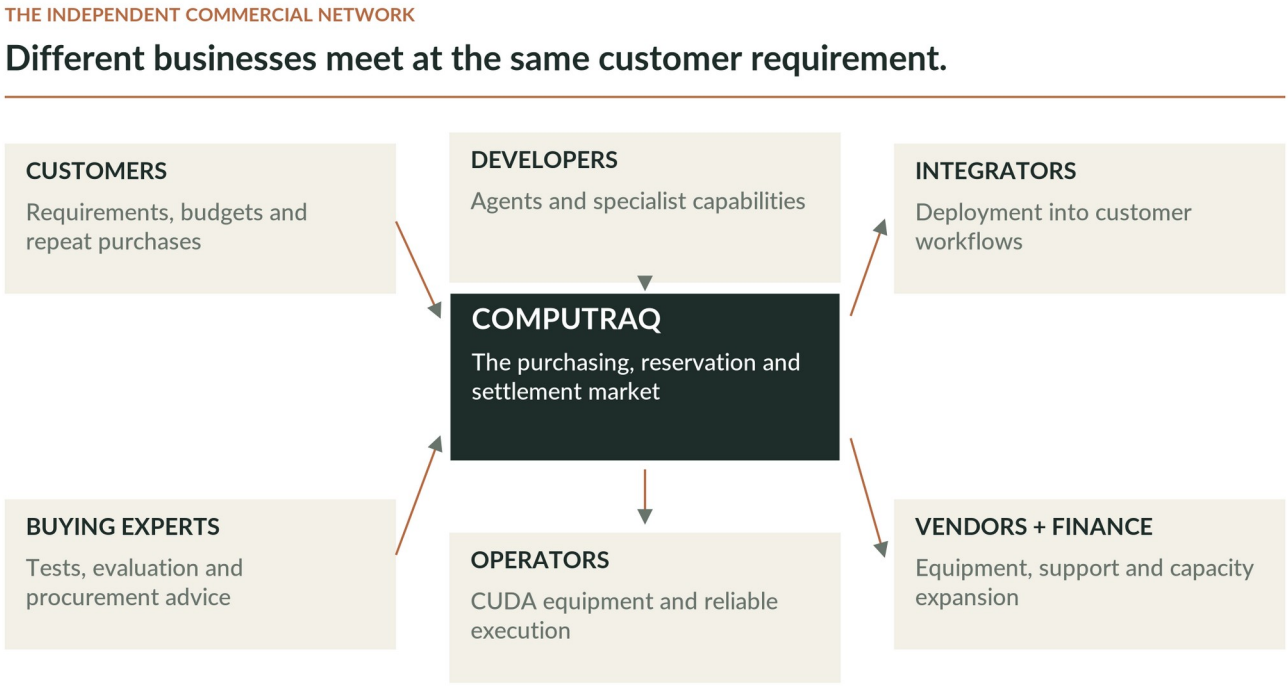


Exhibit 17. Computraq connects a customer-facing service market to the wider CUDA industrial network.

Customers bring the demand

The customer relationship begins with recurring work. Its most important contribution is not a vague expression of interest in AI, but a job, an acceptance standard and a budget owner. Computraq makes that requirement accessible to relevant suppliers and preserves it through future purchasing decisions.

A large organization can create several workstreams under one account. A small business can follow an existing benchmark. A procurement adviser can help organize a portfolio. The common element is that a supplier's claim is assessed against the buyer's actual purpose.

The buyer's commercial connection is therefore persistent. It does not end with the first winning supplier. New results, different operators and improved purchasing methods can all create value within the same relationship.

Developers bring the capability

A developer publishes a versioned Agent, demonstrates its performance and offers a price for a bounded reservation. Its connection to customers comes through relevant requirements and accepted evidence. The market gives useful specialization a route to revenue.

The developer can use a qualified external operator, operate its own equipment or combine arrangements across offers. Its economics need to include compute and support. The platform fee buys common distribution and purchasing machinery that a small firm would otherwise have to reproduce for itself.

A model producer can participate by supplying maintained components to Agents or by publishing its own complete services. Kernel developers, workflow engineers and integration specialists can reach customers through the applications that use their work. The opportunity is broader than becoming a general-purpose chatbot vendor.

Benchmark creators bring purchasing expertise

The creator turns knowledge of an industry or workflow into a maintained buying standard. It defines representative tests, reviews supplier submissions and helps customers recognize useful improvements. A specialist consultancy can thereby serve many buyers without repeating the same initial evaluation from scratch.

The creator's paid role is selected at the purchase. A 10% fee applies to the underlying Agent service price when that purchase uses the accepted fee-bearing benchmark relationship. The fee is displayed separately. Creating a benchmark, following one or receiving a notification does not by itself create a payment obligation.

The relationship is valuable when the advice improves the customer's all-in result. A creator can earn recurring income from a method that stays useful. The buyer can select another creator or maintain its own requirement. Purchasing expertise becomes a competitive business rather than a permanent claim on the account.

Operators bring dependable execution

An operator publishes capacity in a form an Agent can use: equipment, region, software profile, allowance, availability, price and service responsibility. It supports a known set of deployments and earns execution revenue when those deployments are purchased.

An operator does not have to discover every end customer. Several developers can bring different applications to the same qualified infrastructure. A regional provider can compete for work it would not have reached through a general-purpose cloud sales organization.

The reward is repeatable contribution after power, facilities, staffing, maintenance and finance. A higher utilization number is useful only where it corresponds to economically worthwhile service. Computraq's purchasing and settlement record makes that relationship visible.

Vendors and capital providers expand what works

System builders and distributors can package equipment around reproducible operator profiles. Their contribution includes compatible hardware, installation, repair, lifecycle support and the practical knowledge a smaller operator needs to begin earning. A qualified configuration is easier to sell when it has known applications and an understandable route to customers.

Capital providers can finance equipment and facilities serving that demand. Their customers are the operators or asset-owning businesses whose obligations they underwrite. Computraq supplies commercial records and access

to purchasing activity; the capital provider supplies finance. This division lets the market grow without making Computraq the balance-sheet owner of the whole industry.

Vendor programs can become valuable secondary businesses. A manufacturer or distributor may support a particular equipment class, an operator may contract for a capacity program, and a financier may offer terms against the resulting business. The core marketplace continues earning from Agent purchases. It does not have to wait for every financing product to be arranged before customers can buy.

NVIDIA supplies the common industrial foundation

NVIDIA's contribution is the foundation that lets these participants reuse one another's work: CUDA software, processors, networking, systems knowledge and a broad supplier community. Technical cooperation can make common configurations more efficient. Developer programs can connect useful expertise to customer requirements. Channel relationships can help competent operators obtain and support equipment.

The economic return to NVIDIA comes through its ecosystem. More independent services create more reasons to use CUDA. More operator revenue can support maintenance, replacement and expansion. A wider developer market makes the platform more commercially attractive. Those effects do not require the customer to contract directly with NVIDIA for every reservation.

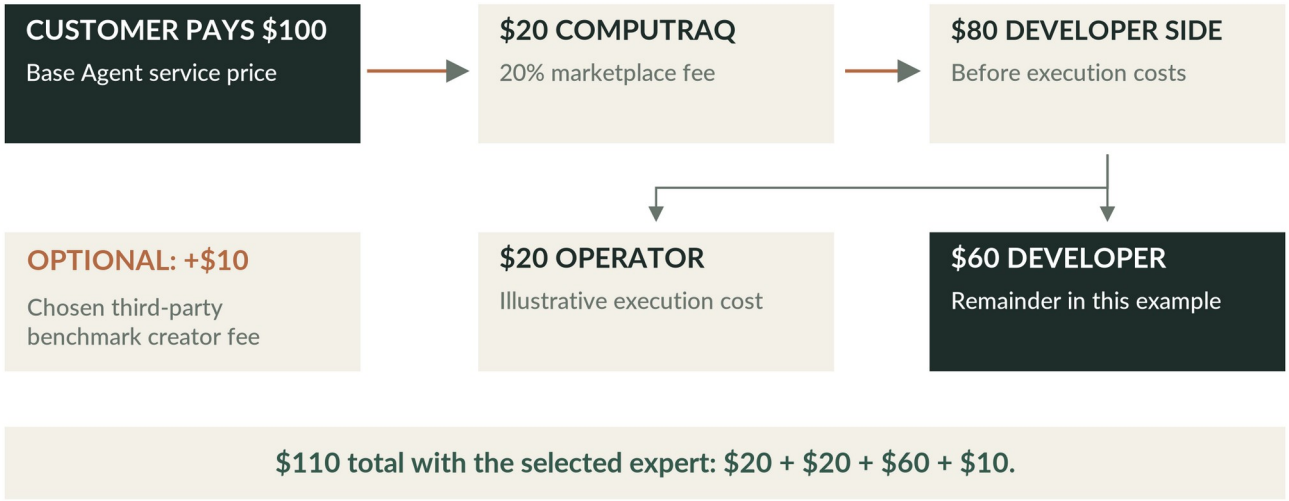
A complete payment example

Take a $100 Agent reservation. Computraq receives $20. The developer's side receives $80 before the execution charge. If the agreed execution costs $20, the operator receives $20 and the developer retains $60 before its other costs. A customer using a selected third-party benchmark pays an additional $10 to its creator, bringing the service-and-creator total to $110.

The 20% compute share is an illustration of that offer's economics, not a second universal platform fee. Another Agent might need $12 of compute, leaving $68 for the developer, or $30 of compute, leaving $50. The marketplace share remains $20 on the $100 service base. The customer compares the all-in price and the work delivered.

WORKED TRANSACTION

A $100 service purchase—and a separately selected expert fee.



$110 total with the selected expert: $20 + $20 + $60 + $10.

The $20 execution charge is a worked assumption; actual execution cost follows the selected commercial offer.

Exhibit 18. A $100 service purchase distributes to three operating businesses. A chosen creator's $10 is additional and separately attributed.

A customer buying its own benchmark's qualified offer pays no creator surcharge. An operator that also develops the Agent earns both applicable shares. An integrator can have its own agreed service engagement. The record identifies the parties and the basis for each amount so the market can support several business models without obscuring the customer's purchase.

A partnership starts with a deliverable

Participant	The first useful deliverable	Its direct commercial connection
Buyer	A recurring requirement and funded evaluation	Computraq account and selected service offers
Developer	An evaluated, deployable Agent	Customer requirements, operator execution and marketplace proceeds
Benchmark creator	A maintained purchasing standard	Customers selecting its fee-bearing advice
Operator	A qualified capacity offer	Developers and reservations using its environment
Integrator	A repeatable connection to the customer's workflow	Customer implementation and operating support
System vendor	A reproducible equipment-and-support package	Operators acquiring and maintaining productive capacity
Capital provider	An equipment or capacity finance program	Defined operator obligations and commercial performance

Participant	The first useful deliverable	Its direct commercial connection
NVIDIA ecosystem team	A deployment or optimization that benefits many participants	Developers, vendors and operators using the common foundation

These are practical entry points, not an invitation to merge institutions. The first customer, first Agent and first operator can begin the commercial cycle. Partnerships then deepen around activities that already have an owner, a product and an economic purpose.

The marketplace joins the transactions. The participants retain the businesses.

7. Better buying, better economics

Customer costs, developer distribution and professional purchasing expertise.

The full cost of an existing AI workflow

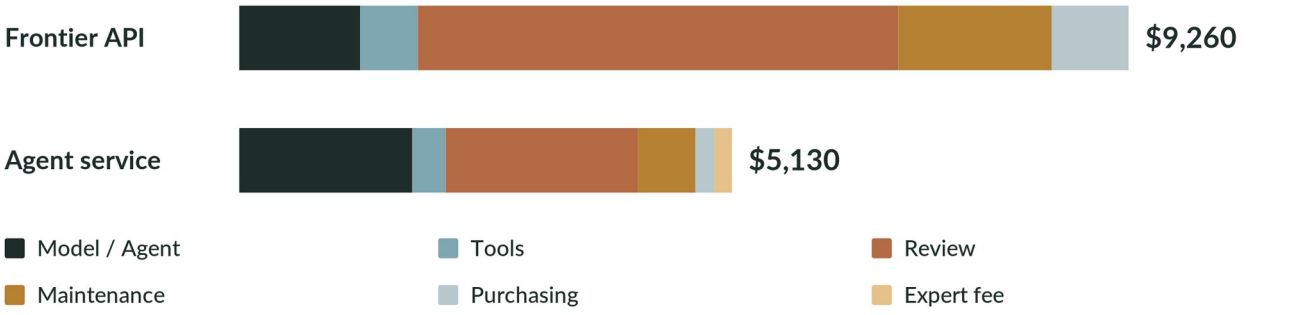
Consider an organization already processing **10,000 accepted cases per month through a frontier-model API**. Its system uses cached inputs, orchestration tools, human review and a technical maintainer. Computraq must improve that complete AI workflow.

The illustrative model uses Anthropic's published Sonnet 5.5 rates of $2 per million input tokens and $10 per million output tokens, with a 90% reduction for the cached-read assumption. The batch consumes 300 million uncached input tokens, 300 million cached-read tokens and 60 million output tokens, including retries. Its model bill is $1,260. The example holds accepted output constant and makes review, maintenance and purchasing effort explicit. [02]

Monthly cost for the same 10,000 accepted cases	Existing frontier-API workflow	Benchmark-selected Agent workflow
Model consumption / complete Agent reservations	$1,260	24 Agent days × $75 = $1,800
Workflow tools and customer-side services	$600	$350
Human review and correction	100 hours × $50 = $5,000	40 hours × $50 = $2,000
Technical maintenance	16 hours × $100 = $1,600	6 hours × $100 = $600
Evaluation and purchasing effort allocated to the month	$800	$200
Benchmark expert's 10% fee on Agent reservations	—	$180
Complete monthly cost	**$9,260**	**$5,130**
Cost per accepted case	**$0.926**	**$0.513**

10,000 ACCEPTED CASES / MONTHLY ILLUSTRATION

The token bill is only part of the cost of useful work.



Modeled total: $9,260 → $5,130. Saving: $4,130 / 44.6%, including the selected expert's fee.

Exhibit 19. Matched-output scenario. The service price already includes platform and compute charges. Tariff input: [02].

The modeled saving is **$4,130, or 44.6%**, with the same accepted output and the benchmark expert's fee included. Crucially, the Agent reservation is more expensive than the baseline token bill. The improvement comes from the complete workflow: fewer interventions, maintained operating knowledge and a repeatable purchasing process. Token prices alone would miss the value.

The evaluation records how released capacity is used: additional accepted cases, a shorter backlog, avoided overtime or expenditure that can actually be removed. That evidence gives the finance team a return-on-investment measure and gives the developer a proposition expressed in the customer's own economics.

The example also makes the platform economics visible. Of the $1,800 reservation base, a 20% platform fee is $360. With an illustrative $360 of computing expense deducted from the developer's share, $1,080 remains for the developer before its other costs. The separate $180 belongs to the benchmark creator. The customer's $5,130 total already contains all of these service charges.

The mechanism remains useful when model suppliers cut their prices. The customer can rerun the comparison and select the improved incumbent if it wins. Computraq does not need a permanently expensive API to create value. Its product is the ability to determine and purchase the best qualified service repeatedly.

Competition reaches the developer, the operator and the expert allocating the budget.

The distribution cost hidden inside an AI service

Small AI companies do not fail only because their models are inadequate. A useful capability can be expensive to explain, demonstrate, integrate and sell. The developer's commercial cost may grow with every new customer even when its core software is reusable. Buyers see the resulting price, but not the duplicated discovery and persuasion behind it.

A shared benchmark market reduces that duplication. Demand has a description before the next supplier arrives. Relevant customers can subscribe to a requirement. The evaluator can establish a repeatable method. An improved Agent can reach buyers through accepted evidence rather than a fresh advertising campaign. Security, offer, payment and execution machinery can be reused. The specialist's engineering becomes the central product rather than an accessory to its sales organization.

This is why competition can reduce prices while improving smaller firms' earnings. The developer does not need to retain the incumbent's selling costs to retain its income. It can pass part of the saving to customers, keep part for maintenance and growth, and serve markets that were previously uneconomical to acquire.

An acquisition-cost model

Consider a $20,000 annual service relationship. A direct seller spends 60 skilled hours on prospecting, discovery, demonstration and evaluation at an all-in labor cost of $125 per hour, plus $1,000 of marketing expenditure. Acquisition cost is **$8,500 per customer**.

In the marketplace example, establishing a reusable evaluation and deployment package takes 160 hours at the same rate: $20,000. Across 100 similar customer relationships, that is $200 each. Each relationship then requires five hours of adaptation, costing $625, and $100 of continuing demand-generation expenditure. Acquisition cost becomes **$925 per customer**. The difference is $7,575, an **89.1% reduction** in modeled acquisition cost.

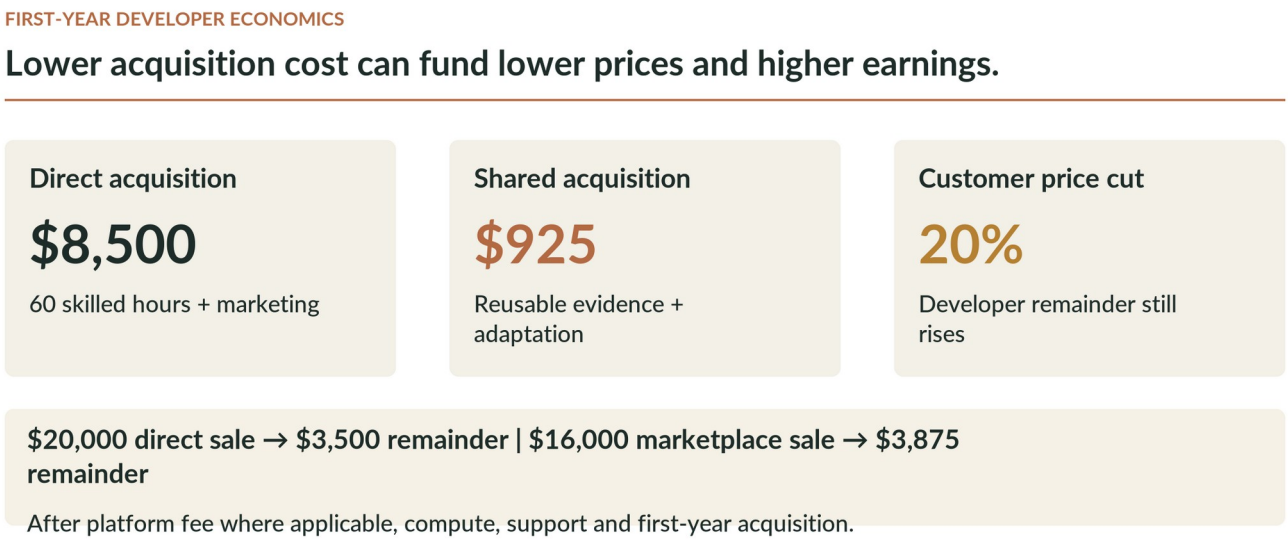


Exhibit 20. The common evaluation and distribution process replaces repeated selling work; the platform fee is included in the comparison.

Now include the platform fee rather than hiding it. At the same $20,000 service price, assume both businesses incur $6,000 of computing and $2,000 of support costs. The direct seller retains $3,500 after acquisition. The marketplace seller pays Computraq $4,000, yet retains **$7,075** after the lower acquisition cost. Even if the marketplace seller cuts the customer's price by **20% to $16,000**, it retains **$3,875**—more than the original direct seller.

Annual economics per acquired customer	Direct sale	Marketplace, same price	Marketplace, 20% lower price
Service price	$20,000	$20,000	$16,000
Marketplace fee	—	$4,000	$3,200
Compute	$6,000	$6,000	$6,000
Support	$2,000	$2,000	$2,000
Acquisition cost	$8,500	$925	$925
Remainder before other costs	**$3,500**	**$7,075**	**$3,875**

The first-year illustration isolates the proposition: a substantial marketplace fee can be compatible with customer deflation when the platform replaces even more expensive distribution work. The operating record follows acquisition effort, paid conversion, retention and support cost. Those measures show whether reusable distribution is delivering its intended advantage.

Buying expertise becomes a market of its own

Most buyers do not want to become evaluation specialists. An enterprise may know that its research process is expensive without knowing how to design a reliable test of competing research Agents. A small company may need accounts-payable automation but lack the staff to evaluate model behavior, exception handling and data permissions. Someone must translate the business requirement into a purchasing standard.

Computraq creates a business for that work. A **benchmark creator** develops and maintains a useful evaluation, reviews relevant results and helps customers direct spending to qualified capabilities. When a customer expressly chooses that benchmark for a fee-bearing purchase, the creator receives **10% of the underlying Agent service price**, shown separately in the all-in quote. A buyer using its own benchmark pays no creator surcharge. A subscription or notification does not itself create a fee.

The commercial precedent is the agency relationship. Businesses have long paid agencies and advisers to improve the allocation of a much larger spending budget. Advertising and media agencies provide a recognizable example: expertise in selection, measurement and placement can be worth a percentage of spending. Computraq applies that established model of compensated purchasing expertise at a clearly stated 10% rate. [04,05]

The important advance is that **the adviser's own value becomes testable and replaceable**. Customers can compare benchmark creators by the outcomes their purchasing methods deliver: accepted output, complete cost, maintenance quality, useful coverage and successful recommendations. A benchmark creator does not gain permanent authority over the buyer's budget merely because it designed the first test.

Suppose work costs $10,000 without outside purchasing advice. An expert identifies an equivalent qualified service for $8,000. Its $800 fee brings the complete service outlay to $8,800, leaving the customer $1,200 better off. Another creator can compete with a better method or broader coverage. If the expert improves no other dimension, a 10% surcharge needs a service-price reduction greater than 9.09% to lower the total bill. That is a useful commercial discipline: the fee must buy value.

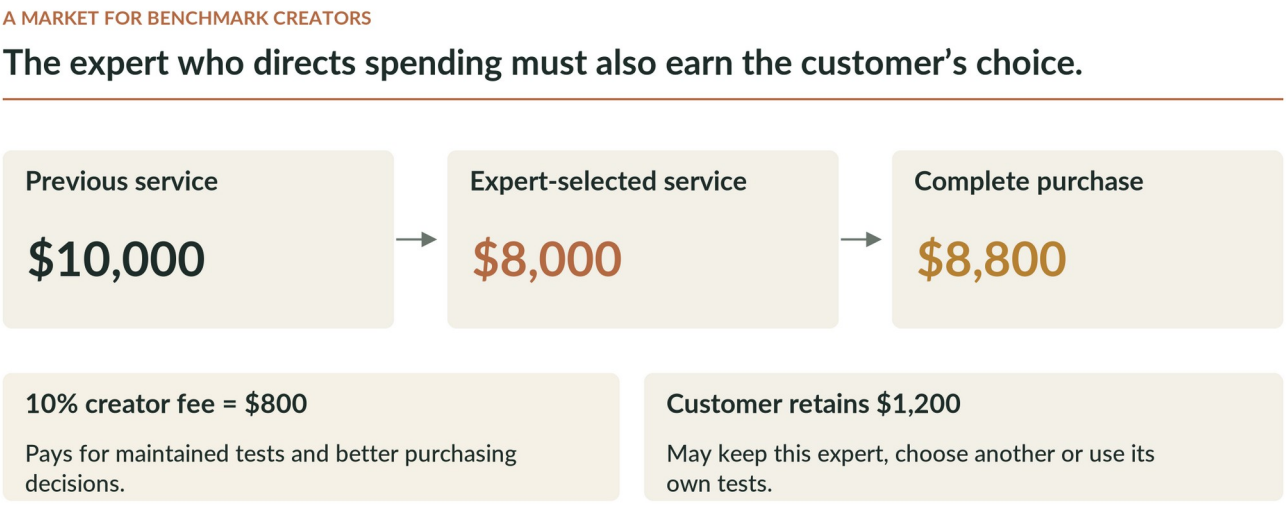


Exhibit 21. A disclosed buying-expert fee is tested against the customer's all-in result, and the expert remains replaceable.

This arrangement aligns three kinds of specialization. Developers improve capabilities. Computing operators improve execution. Benchmark creators improve the decisions allocating money between them. The customer retains the right to replace any of the three. The platform is valuable because it makes all three contests work together.

Service continuity when scarce talent moves

Forward-deployed engineers work closely with a customer to adapt AI to its actual environment. They learn the difficult data, permissions, interfaces and exceptions. Anthropic's enterprise disclosures describe Applied AI and customized deployment work, demonstrating that even a leading model company depends on human expertise to convert capability into a reliable business service. [03]

The scaling problem is not that expertise has no value. It is that scarce experts cannot personally maintain unlimited relationships. Reassignment, promotion, leave, acquisition and departure can remove access to customer-specific knowledge. Equity vesting and liquidity can influence career decisions, but the customer's continuity problem exists for any reason that a key person becomes unavailable.

Computraq's response is to make expert effort cumulative. The requirement, representative tests, approved configuration, interfaces, permissions, known failures and operating procedures remain usable by an authorized successor. The customer need not acquire the developer's proprietary model or source code. It needs the practical knowledge to verify continued service and assess an alternative.

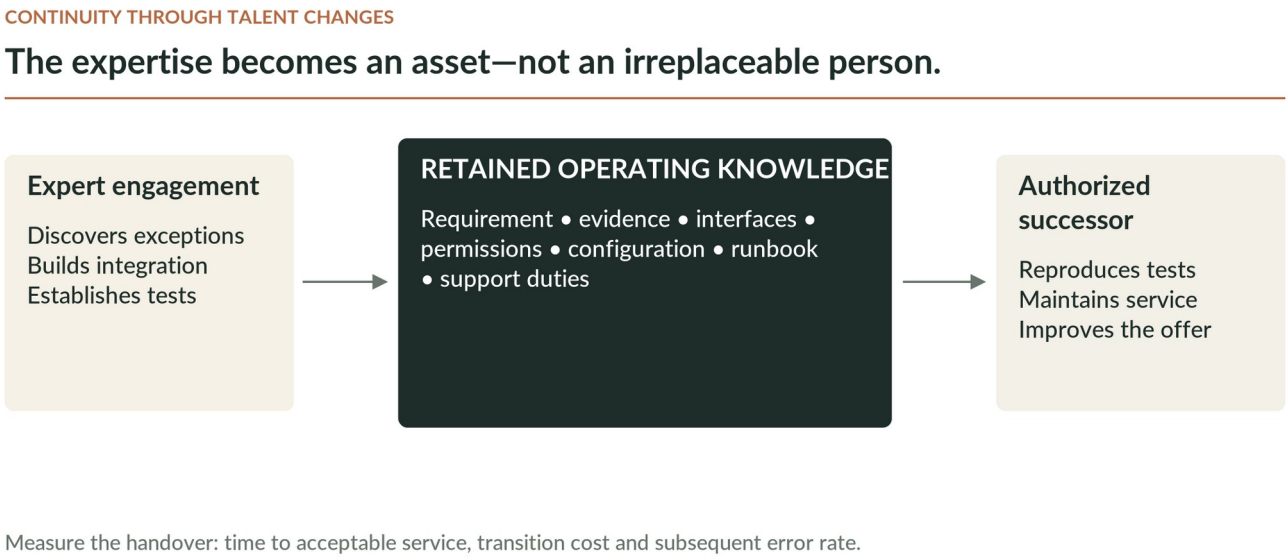


Exhibit 22. Maintained customer-authorized knowledge makes an expert's contribution usable through reassignment, departure or supplier change.

This turns an expert engagement into a durable improvement rather than a recurring dependency. A new maintainer can reproduce the tests. A different Agent can compete against the same standard. A qualified operator can run a compatible service. The measures are time to restore accepted work, transition cost and subsequent error rates. A successful handover converts the original deployment expense into an asset the customer can continue using.

The result is a superior long-term procurement position. The buyer can keep useful integration knowledge without granting the original vendor an indefinite claim on future spending. Developers can win through demonstrable improvement. Evaluators and integrators can maintain valuable standards and operating knowledge as recurring businesses. Continuous competition becomes compatible with continuity of service.

Reusable expertise changes the long-term competitive balance

A forward-deployed engineering team can be outstanding and still be scarce. Every new customer adds demands on the people able to understand an unusual workflow and keep it reliable. Promotion, reassignment, acquisitions, leave and departures change the availability of that skill. Stock vesting and liquidity can change individual incentives at the same time. A customer should not have to renew indefinitely with the original vendor simply because a handful of its employees understand the deployment.

A maintained evaluation and operating record changes the renewal decision. The next supplier begins with a defined job, known exceptions, approved interfaces and evidence of the incumbent's performance. The buyer can retain a trusted maintainer, use another specialist or choose a replacement Agent. Proprietary software remains with its owner; customer-authorized knowledge remains usable in the customer relationship.

This is why Computraq's model can improve with age. Each successful deployment should make the next evaluation and handover less expensive. A conventional vendor can respond with better service or lower prices, but the customer retains an independently usable purchasing standard. The contest does not disappear when a sales contract is signed.

8. Win half the work—and create the next application economy

A substantial second purchasing channel becomes a route to entirely new uses.

Offer the customer an economically meaningful alternative

Computraq’s initial competitive objective is to make half of a customer’s suitable frontier-lab workload contestable. A small pilot can prove a technology. A material second channel changes purchasing leverage, service resilience and the economics of supplier investment. The buyer keeps frontier capabilities where they are strongest and moves qualified work where the complete result is better.

Consider an enterprise with a $10 million annual AI-service baseline. It retains $5 million of incumbent work and moves a portfolio that formerly cost $5 million to Agents priced 30% lower. Agent reservations cost $3.5 million. A 10% creator fee on half that base adds $175,000. A $325,000 transition and evaluation allowance produces a combined annual outlay of **$9 million**, a **$1 million saving** after those additional costs.

The buyer can instead retain the original budget and obtain more output. After the same transition allowance, a $5 million allocation buys work equivalent to approximately $6.36 million at the incumbent’s prices when replacement services cost 30% less and the weighted creator surcharge is 5%. The choice is between releasing budget and expanding accepted output. Both establish a substantial alternative channel.

Cost efficiency and risk management reinforce one another. A credible alternative makes renewal negotiations more informed. Retained tests reduce exposure to changes in the incumbent’s models or support staff. A second developer community can offer ideas outside the original vendor’s roadmap. Independent execution choices make future workloads less dependent on one operating relationship.

The sales proposition is a portfolio, not a declaration of universal model superiority. Coding maintenance, document extraction, supported research, localization and specialist business analysis are natural starting points because their outcomes can be checked. Work with stricter consequences can enter when its acceptance and

operating standards are established. The customer's budget moves with evidence, not with a model release's publicity cycle.

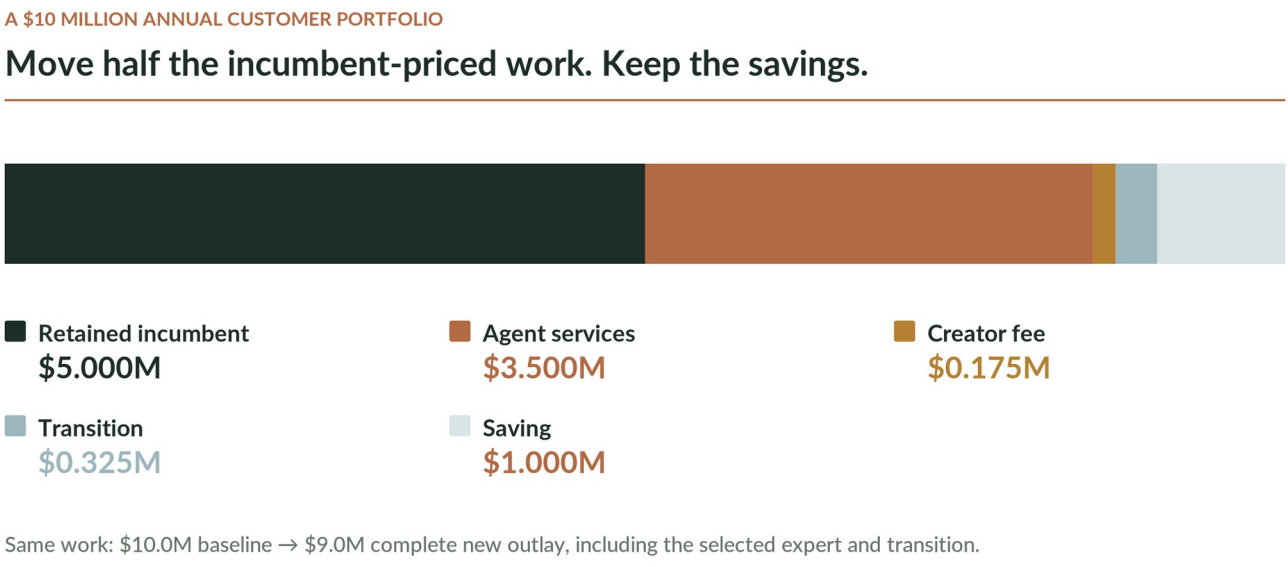


Exhibit 23. An equivalent-work migration can establish a material second supply channel while lowering total expenditure.

The fifty-percent competitive objective

Anthropic's November registration materials report an approximately **$81 billion October revenue run-rate**, an annualization of recent consumption and subscription activity. Computraq's objective is to contest work representing half of that economic scale: **$40.5 billion at incumbent prices**. Thousands of specialists can compete for that portfolio without any one developer reproducing a frontier laboratory's research organization. [03]

At a 30% lower price for equivalent work, the service market is **$28.35 billion**. Under the working allocation, $5.67 billion pays Computraq, $5.67 billion purchases CUDA execution, and $17.01 billion remains with developers before their other costs. Buyers save $12.15 billion before creator fees. With the 10% fee applying to half the service base, benchmark creators receive $1.4175 billion and buyers retain $10.7325 billion of savings before transition costs and tax.

This is a commercial program with an explicit competitive target. It does not depend on one "best model." One specialist can win document processing; another can improve software maintenance; a third can serve a regulated industry's narrow workflow. The platform aggregates their ability to reach and satisfy customers. The competitive force is the ecosystem.

50% COMPETITIVE OBJECTIVE / ANNUALIZED REFERENCE

An ecosystem of specialists can contest a frontier lab's business.

Incumbent-priced work	Replacement services	Buyer savings
$40.50B	$28.35B	$12.15B
50% of $81B reference pace	Equivalent work at 30% less	Before separate creator fees

Within $28.35B: $5.67B platform + $5.67B compute + $17.01B developer remainder

The creator fee is separate; at 10% on half the base, it adds $1.4175B.

Exhibit 24. Original scale calculation using Anthropic's reported October 2026 run-rate [03]. Customer savings are not marketplace revenue.

The customer benefit also creates pressure on the incumbent. Better offers can lower prices even where the buyer has not yet switched. Computraq's value is making those offers executable: a qualified service, a known total price, a retained operating requirement and a route to delivery. The contest is more consequential than a threat to change vendors that nobody can implement.

THE APP STORE MECHANISM APPLIED TO AI

Applications increase the economic value of the same hardware.

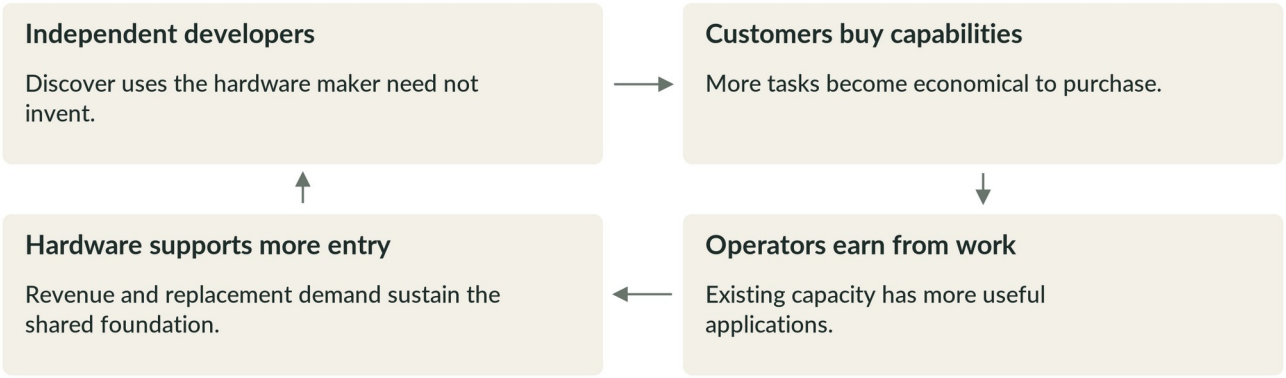


Exhibit 25. The market creates complementary value: independently developed applications make the technical platform more productive.

The application economy that made a device more valuable

Apple's App Store is a familiar example of independent applications expanding the value of a hardware platform. A shared discovery and purchasing system enabled developers to create uses the device maker did not need to invent. As useful applications multiplied, an expensive device could deliver more value to its owner. The economic

mechanism is **complementarity**: software makes hardware more useful, and installed hardware gives software a market.

An Apple-supported Analysis Group study estimates that its app ecosystem facilitated **$149 billion of digital goods and services in 2025**. The much broader $1.437 trillion estimate includes physical commerce and advertising. Computraq's $150 billion annual service objective uses the digital category as a comparison of scale. The comparable measure is application activity; platform revenue is reported separately in the economic model. [06]

The corresponding opportunity in AI is larger than shifting existing invoices. Continuous software repair can make previously neglected systems maintainable. A small business can buy specialized analysis without hiring a permanent AI team. Scientific workflows can become reproducible services. Localization and accessibility can reach customers too small for custom consulting. An Agent can procure another bounded capability under its owner's budget rather than asking a person to assemble a new supplier relationship each time.

A $150 billion service market can arise from a broad population: 10,000 large organizations averaging $5 million, 100,000 midsize organizations averaging $300,000, one million small businesses averaging $30,000, 20 million professional and individual buyers averaging $1,000, and a separate $20 billion of public and scientific programs. These nonoverlapping planning cohorts produce $50 billion, $30 billion, $30 billion, $20 billion and $20 billion respectively. They describe an application economy, not dependence on a handful of enormous contracts.

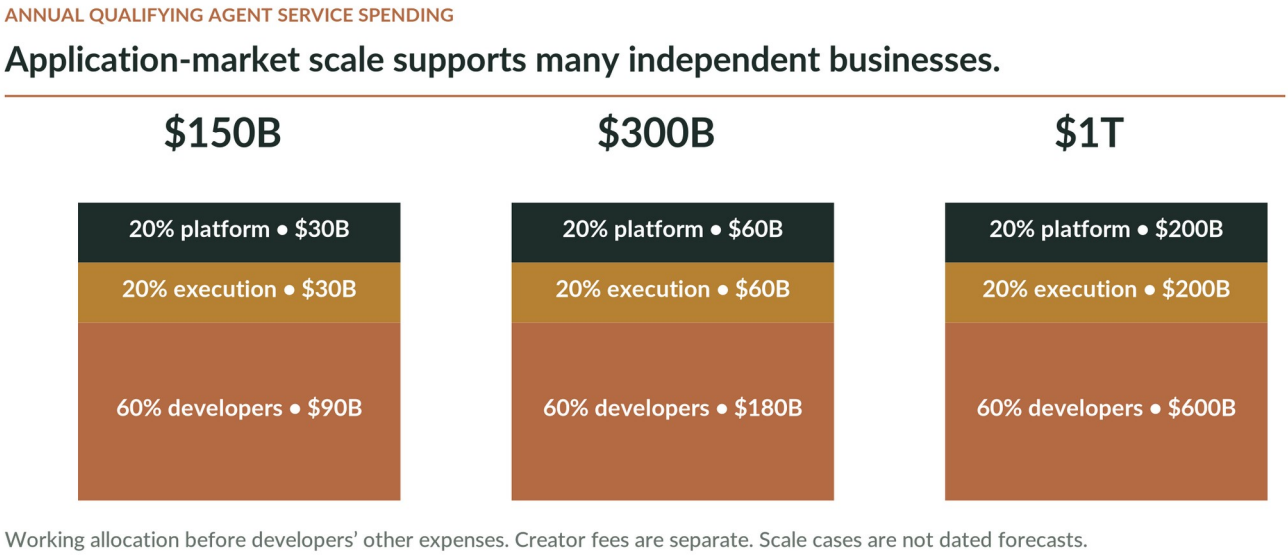


Exhibit 26. Platform, execution-provider and developer economics are allocated from the same service base without double counting.

At that scale, the working 20% compute allocation produces **$30 billion of annual payments to execution providers** and the developer remainder is **$90 billion before other costs**. At $300 billion of services those amounts become $60 billion and $180 billion. These payments support a broad population of developers and computing operators at substantial economic scale.

Efficiency can make this market expand. If computing per accepted job halves, more than twice as many jobs must be purchased for physical computing demand to grow. A common application market supplies the commercial path to that expansion: more uses, more customers and lower acquisition friction. The hardware economy should grow because more useful work becomes worth doing—not because suppliers preserve waste.

The application market makes the hardware valuable by making more independent

businesses possible.

Scale is distribution across many categories

The customer does not need to choose between retaining a frontier laboratory and using independent AI. A procurement portfolio can contain both. That is an advantage in entering a large organization: a measurable second channel can win useful work without demanding a wholesale replacement of an established system. A meaningful allocation then gives specialists a reason to invest in deeper capability and support.

Expansion comes from two sources. The first is additional work already bought from other AI suppliers. The second is work not previously economical to buy at all. A small legal practice can procure a narrow document workflow; a manufacturer can obtain specialized analysis; a research team can reserve reproducible scientific computation. Each adds a use for the same common purchasing infrastructure.

The central displacement example creates $28.35 billion of annual service spending. Reaching $150 billion requires $121.65 billion of additional service activity. The bottom-up purchasing populations describe where that activity can arise. At $150 billion, a 20% platform fee supports the market's shared services, 20% compute intensity purchases $30 billion of execution, and $90 billion remains with developers before their other costs. At $300 billion, those amounts become $60 billion, $60 billion and $180 billion. The economic objective is a large field of sustainable suppliers, not a platform that captures every dollar of value created by its users.

A useful Agent can also help a customer complete a transaction much larger than the price of its reservation. That downstream value belongs in the assessment of business productivity. The marketplace's own service volume remains the amount actually purchased. Keeping those accounts separate makes the application-economy argument more informative: modest software services can enable far larger changes in how people work.

9. The freedom to own the means of computing

Bitcoin and space launch reveal what a service subscription cannot replace.

The founder’s starting point

Numair Faraz began developing Computraq while working on **BTX, a GPU-centric, AI-native blockchain designed around independently operated computing**. The concern was simple: a new distributed protocol cannot flourish if its potential participants lose the practical ability to obtain and control useful equipment. Bitcoin’s transformation made that concern concrete.

Bitcoin began as a system people could help operate with ordinary computers. Mining became an industrial activity organized around application-specific integrated circuits, or ASICs: chips optimized for the calculations used in mining. A Bitcoin mining machine cannot be repurposed to run an AI application or ordinary business software. Equipment purchasing, cheap electricity, financing and scale increasingly determined who could participate economically.

Cambridge’s 2025 industry survey found that Bitmain, MicroBT and Canaan supplied **99.1% of the hardware reported by respondents**, with Bitmain alone at 82%. The hardware sample comprised 46 responses and described equipment around mid-2024. These Chinese-founded manufacturers supply a market in which large

operators and mining pools coordinate substantial activity across national borders. Pools, hardware makers and machine owners are separate layers, but an entrant can encounter concentrated counterparties at each one. [07]

The lesson is not that a protocol must close formally before opportunity narrows. The right to connect remains public while the economics favor a small industrial class. Large buyers negotiate better equipment and power terms. Specialized machines offer few other earning opportunities. A small operator may remain legally free to join while having little realistic prospect of building a consequential independent business.

The resulting structure can be transnational rather than cleanly American or Chinese. Chinese-origin equipment suppliers, American industrial operators, internationally distributed pools and mobile capital can cooperate commercially even as governments describe the sector through national rivalry. Public capacity signals and common economic pressures can coordinate investment without an explicit agreement to divide markets. The structural consequence is dependence across borders: a small entrant faces concentrated counterparties even when governments describe the same industry as a contest between nations.

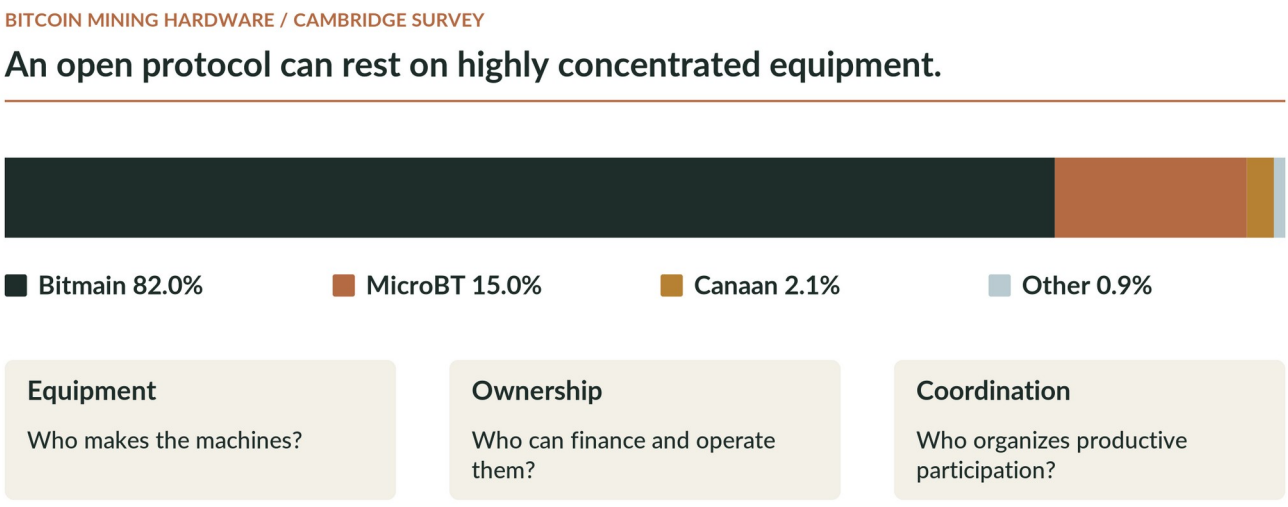


Exhibit 27. Manufacturing, equipment ownership and pool coordination are distinct layers through which entry can become concentrated.

AI can reproduce the same concentration without using literal Bitcoin chips. A custom AI accelerator may support many models, yet be sold only as a service inside its owner's cloud. Its engineering programmability does not give an outsider ownership, maintenance rights or a practical route to competing infrastructure. The relevant test is whether a complete useful system can be acquired, operated, redeployed and connected to customers by independent people.

Space launch shows what service access leaves out

A launch customer normally buys a mission, not a rocket it can independently operate and adapt. SpaceX serves outside customers and offers public rideshare access, but the owner still determines how external missions fit alongside its technical schedule, facilities and internal programs. The remarkable capability of the vehicle does not transfer the power to organize a launch business to the customer. [09]

That distinction should be understood before AI infrastructure acquires the same commercial shape. A founder may be allowed to call a model API while being unable to operate the underlying system, alter the resource-sharing architecture or compete with the infrastructure provider. An unusual research program can be technically possible and commercially stranded. A peer-to-peer network needs independently controlled participants; a subscription to one centralized service is not an equivalent foundation.

THE ROUTE FROM A PERSONAL PROTOTYPE TO GLOBAL DEPLOYMENT

A rented service and an owned tool preserve different freedoms.

CAPTIVE SERVICE	INDEPENDENT EQUIPMENT
Use the provider’s interface	Choose what to build
Fit its available capacity	Use several operating channels
Follow its commercial roadmap	Repurpose a productive machine
Buy an output	Become a supplier

CUDA channels include retail GPUs, system builders, private servers, regional operators and larger clouds. [08]

Exhibit 28. The space-launch comparison concerns control of productive capability. Computing ownership lets outsiders organize new businesses.

CUDA currently supports a different progression. Useful NVIDIA equipment is distributed through retailers, system builders and original equipment manufacturers. Developers can use workstations, private servers, regional operators, specialist clouds and large clouds at different stages. The same technical family supports a wide range of lawful applications. A personal machine need not run the largest model to give its owner a starting point for development and an intelligible route to scale. [08]

This is the freedom Computraq seeks to sustain economically. A common technical platform can support many independent owners; a shared application market can give those owners paying work. A new computing architecture then has machines, people, commercial relationships and customers from which to grow. Its inventor does not have to persuade an integrated rival to sponsor its existence.

The economic freedom of a general-purpose tool

A personally controlled computer is a place to begin without seeking a strategic partnership. It can be used to test a model, assemble a workflow, write software, design an evaluation or establish the first node of a new network. Larger work can move to a workstation, rented GPU or professionally operated cluster. The same developer can progress through those stages without giving a cloud company ownership of the customer relationship.

That progression is what industrial policy should protect. The value lies not only in current sales but in future options: the ability to repurpose equipment, change commercial partners and try an idea outside an incumbent's product roadmap. An industry can contain several sophisticated chips and still offer very few such options if each complete system is confined to its designer's service.

The founder's experience with BTX led to a market-building response. Preserve the independent computing base by making it productive for many more applications. Developers need buyers; operators need workloads; equipment owners need useful alternatives. Computraq organizes those commercial needs so that the next distributed platform has machines, skills and businesses on which to build.

10. National champions are not an industrial strategy

The cost of protecting a domestic franchise while financing its foreign replacement.

Governments can create an impressive national champion and still lose the wider industrial contest. The distinction is between the value of one company and the depth, reach and resilience of an ecosystem. Public policy should be judged by how many competent firms can participate, what the country can produce and whether customers around the world choose its products.

Computing policy faces this choice immediately. Broad restrictions can protect a domestic customer base for established AI services while making foreign buyers more determined to finance alternative hardware, software and supply chains. The visible achievement is a blocked shipment. The less visible consequence is the future market built elsewhere.

The soybean embargo and the customer who built an alternative

In 1973, the Nixon administration restricted soybean exports amid domestic price and supply pressure. For Japan, the interruption exposed the danger of depending on an American supplier whose access could change abruptly. JICA's account connects that experience to Japanese diversification and the agricultural cooperation initiated with Brazil during Prime Minister Kakuei Tanaka's 1974 visit. [10,11]

The Japan–Brazil Cerrado development program, PRODECER, ran through three phases from 1979 to 2001. JICA records a total cost of ¥68.4 billion, including ¥27.9 billion of official development assistance. Japanese finance and cooperation joined Brazilian agricultural science, infrastructure and enterprise in developing a durable alternative production base. [12]

SOYBEANS / JAPAN / BRAZIL

A restriction can finance a durable replacement for the exporter.



Customer demand persists. Capital and engineering move to the substitute.

Once another system has been built, restoring old access does not restore old dependence.

Exhibit 29. The 1973 disruption and subsequent Japan–Brazil cooperation illustrate customer-financed substitution. Sources: **[10–12]**.

The industrial lesson is that customers do not stop needing an input when Washington restricts it. They organize another source. Once they have invested in substitute production, equipment, knowledge and relationships,

removing the original restriction does not recover the old dependence. The policy has subsidized the customer's incentive to escape American supply.

Advanced computing magnifies that effect because software knowledge and operating experience compound. A customer forced to adopt another accelerator also funds engineers, libraries, serving software and deployments around it. The next customer benefits from the work already done. A temporary intervention can help establish a competing platform with a permanent commercial life.

A policy that encourages this substitution while promising permanent American leadership is internally contradictory. The right question is not how many transactions were stopped. It is where the next decade's investment, engineering and customer relationships will accumulate.

Electric vehicles and the publicly assisted champion

The United States helped make Tesla a national champion through an architecture of public finance, demand incentives, regulatory credits and state support. Tesla received a $465 million federal manufacturing loan in 2010. Nevada's development agency reported a further $330 million of approved tax abatements in 2023. Treasury reported more than $2 billion of upfront clean-vehicle consumer credits across eligible manufacturers in the first nine months of 2024. [17,18,42]

Tesla's financial statements show how regulation itself became a revenue stream: automotive regulatory-credit revenue was $2.763 billion in 2024 and $1.993 billion in 2025. Those receipts arose because other manufacturers needed credits generated under public rules. Public policy was not peripheral to the company's industrial rise. [43]

Computraq's critique is that the public helped create an exceptionally powerful corporate winner without creating a comparably broad, globally dominant American electric-vehicle ecosystem. The concentration of benefits around a celebrated company and its controlling entrepreneur is not the same as a country acquiring deep competition among affordable producers, suppliers and export channels.

The international outcome is instructive. The IEA reports that Chinese automakers supplied 60% of global electric-car sales in 2025, while North American and European automakers each supplied about 15%. China accounted for nearly 75% of production. Chinese imports supplied 55% of electric-car sales outside Europe and the United States. These are separate measures of commercial reach; together they show where the industry's scale has developed. [15,16]

ELECTRIC VEHICLES / 2025 MARKET OBSERVATIONS

A national champion is not the same as global industrial leadership.



Different measures: manufacturer origin, factory location and destination-market imports. IEA Global EV Outlook 2026 [15–16].

Exhibit 30. Publicly supported domestic champions and broad international industrial reach are different outcomes.

The danger for AI is a similar division: a protected American service market dominated by nationally favored firms and a much larger manufacturing-and-software ecosystem developing alternatives abroad. Independent American founders would then be squeezed between domestic gatekeepers and foreign industrial scale. Restricting the global reach of merchant computing can hasten that outcome.

The public should not be asked to equate the wealth of one nationally favored owner with its own industrial success. A stronger program spreads the opportunity to invent, sell, operate, maintain and export. It creates several routes into a productive industry rather than a political contest for affiliation with the selected champion.

The political price of narrow opportunity

Economic opportunity is also a question of legitimacy. When public money and access visibly benefit a narrow set of owners, people can reasonably question whether the market is open to their own effort. The OECD's research links financial insecurity and perceived lack of political voice to lower trust in government. A country cannot rebuild confidence merely by announcing another successful national champion. [19]

Computraq's alternative begins with a much wider population of potential earners. A specialist can qualify a useful Agent. An evaluator can improve how buyers spend. A regional operator can sell execution. A customer can keep purchasing decisions open. That structure makes participation productive without requiring political selection.

Global presence as a security asset

The United States and its trading partners—including China—benefit when many firms have commercial reasons to maintain reliable supply, compatible tools and stable exchange. Such relationships create technical knowledge and repeated practical contact across political boundaries. They are a source of resilience because the economy contains many participants with something to lose from disruption.

The strongest American strategy is to make American-led computing useful enough that others choose to remain connected to it. That is the lesson of Starr's insurance network and Huang's industrial platform. A policy of retreat leaves less American influence over the market's next stage, while suppliers elsewhere continue meeting the demand.

A protected home franchise is not a substitute for global industrial leadership. Leadership is the ability to remain useful wherever the next productive application is being built.

11. When national security becomes a license for incumbents

Disclosed interests, access restrictions and the political economy of market entry.

The central policy error is to confuse a small number of approved owners with a secure computing economy. Security requires the ability to identify problems, replace failing suppliers, retain skills and develop alternatives. A system that makes independent development depend on the permission of a few integrated providers weakens those capabilities even if it simplifies official oversight.

Computraq's position is that national-security and safety policy must not be used to allocate an industrial franchise. Export restrictions, expensive qualification, procurement preferences and investment-linked cloud obligations can combine to turn a prospective competitor into a customer of an incumbent. A startup remains legally free to exist while every practical route to serious scale is controlled by the firms it hopes to challenge.

The commercial interests behind the policy debate

The people financing AI companies also help finance the institutions that shape AI policy. Current records make the overlaps concrete. Moskovitz and Eric Schmidt were named investors in Anthropic's first financing. Founders Fund co-led Anthropic's February 2026 financing and lists OpenAI and SpaceX in its portfolio. [30,31,44]

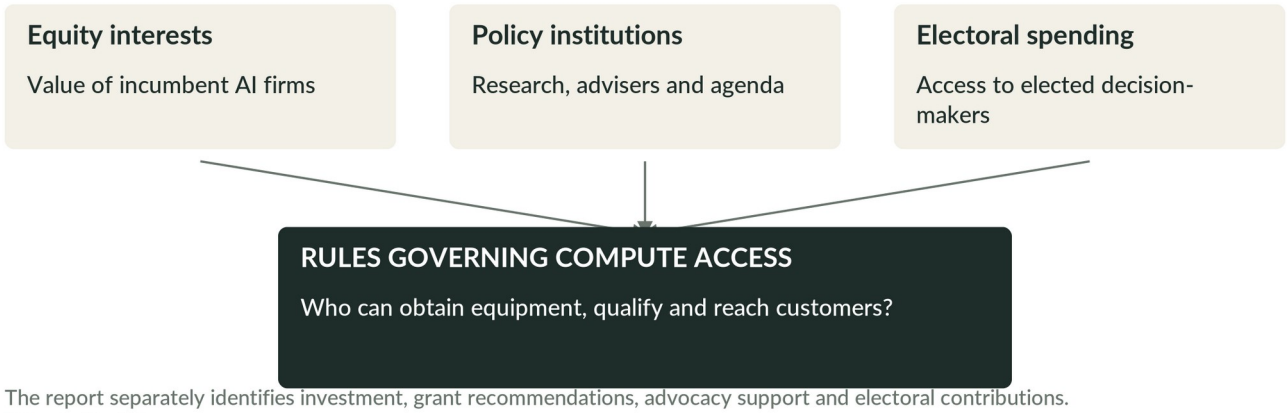
Peter Thiel was named among OpenAI's founding financial supporters in 2015. Reporting places Founders Fund's OpenAI equity participation in 2023. [45,46] These relationships create a direct economic interest in the terms on which competing models, hardware and distribution channels can enter the market.

Current funding or institutional relationship	Amount or disclosed structure	Channel
Dustin Moskovitz, Survival and Flourishing Fund 2026 round	$785,000 recommended for Foundation for American Innovation; $465,000 for Institute for AI Policy and Strategy	Named grant recommendations [25]
Coefficient Giving, Abundance & Growth fund	$10.5 million general-support grant to Institute for Progress, July 2026	Institutional policy funding [26]
Coefficient Giving, AI portfolio	More than $1 billion of 2026 commitments anticipated in its September update	Technical safety, security and fieldbuilding [24]
Anthropic	$40 million total support for Public First Action by July 2026	Public education and policy advocacy [23]
Leading the Future, 2026 midterm cycle	$25 million from Greg and Anna Brockman; $25 million from a16z in the second half of 2025	Electoral support reported from filings [29]
Eric Schmidt, Special Competitive Studies Project	Subsidiary of the Eric & Wendy Schmidt Operating Foundation	National-security and technology-policy institution [28]
Peter Thiel	Hundreds of thousands of dollars to Republican leadership-aligned committees in 2025	Current congressional-cycle giving [29]

These channels are economically different. A grant can sustain researchers and events; electoral expenditure can affect who holds office; an investment benefits from a company's commercial position. Their convergence gives well-connected interests a continuous voice in setting the conditions of entry. An independent developer normally has resources to build a product, not to maintain an equivalent policy network.

POLITICAL ECONOMY / DISTINCT CHANNELS OF INFLUENCE

Private interests can shape the rules of public market entry.



The report separately identifies investment, grant recommendations, advocacy support and electoral contributions. [23–33,44–47]

Exhibit 31. Documented channels establish economic interests and routes of influence; the capture argument examines their effects on entry.

Computraq's charge is one of **regulatory capture**: the policymaking process can become more responsive to concentrated commercial networks than to the dispersed people whose ability to compete is at stake. The mechanism does not require a politician to announce a favor. Agenda-setting, advisory appointments, funded research, repeated private access and career opportunities can make one industrial structure appear inevitable while the alternatives receive little attention.

The chip-policy alignment

Anthropic has publicly advocated tighter advanced-chip and semiconductor-equipment controls. Its July 2026 policy-funding announcement connected those priorities to its support for Public First Action. Dario Amodei's September essay proposed coordination on frontier development and a narrow antitrust waiver for certain safety discussions. Demis Hassabis's July framework proposed a substantially industry-funded frontier standards institution with review that could become mandatory for deployment. These positions put the governance of entry at the center of the commercial contest. [23,32,33]

On November 12, reporting disclosed OpenAI's support for bipartisan chip-control measures; Anthropic had already supported tighter controls. The commercial effect deserves scrutiny regardless of whether each advocate reaches its position independently: restrictions on broadly accessible compute can improve the relative position of large domestic service providers. [47]

Public advocacy, concentrated financing and parallel commercial interests can produce an incumbent-protecting policy outcome. Authorities should examine how that structure affects entry, access to equipment and the ability of unrelated firms to reach customers.

A non-frontier exemption is not enough to preserve the whole ecosystem. A small application developer may depend on an independent open-model producer that must eventually obtain frontier-scale compute and evaluation. If that producer cannot graduate on viable terms, downstream developers become more dependent on approved laboratories. Policy closes the route to the next competitor while continuing to describe the entry-level market as open.

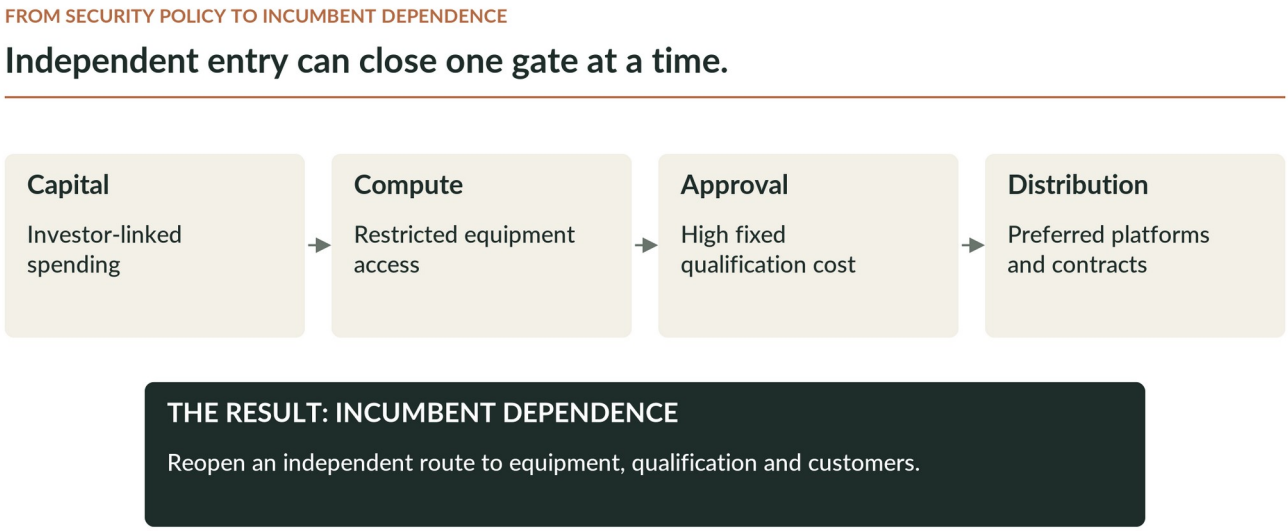


Exhibit 32. Connected restrictions can produce dependence even where an unrelated company remains legally free to incorporate.

The export-control paradox

A restriction aimed at an American-led general-purpose platform has two industrial effects. It can shelter approved domestic service firms from some competitive pressure, while also accelerating investment in foreign alternatives. The soybean precedent explains why the second effect lasts: customers who finance a replacement do not discard it when access to the original supplier is restored.

The practical question is what happens to the entire network of demand. Does the restriction prevent a defined harmful use, or does it merely move lawful research, software development, engineering knowledge and commercial relationships into another ecosystem? How much American influence is lost when that migration becomes permanent? An export-control system that fails to answer these questions is measuring enforcement activity rather than industrial success.

The government has already demonstrated that the rules are choices. The prior AI Diffusion Rule was rescinded in 2025, and a January 2026 rule adopted conditional case-by-case review for specified advanced chips. Those changes make clear that blanket exclusion is not the only administrative design available. The policy priority should be preserving broad lawful participation while addressing identified recipients and activities with specific evidence. [34,35]

Connected market infrastructure: the advertising precedent

Advertising shows how multiple nominal markets can become one practical route to customers. In April 2025, a U.S. court found that Google unlawfully maintained monopolies in publisher ad servers and ad exchanges and unlawfully tied products. Control of connected infrastructure affected publishers, advertisers and the economics of information. [21]

The advertising case shows why connected markets matter. Once a few participants control the route between the producer and the customer, restoring a usable alternative becomes a long and costly task. The opportunity in AI is to establish independent commercial routes while the market is still forming.

AI has comparable connections. The same cloud can invest in a model laboratory, require spending on its infrastructure, distribute the model and sell applications that compete with its own customers. The FTC's study documents cloud-spending commitments, rights and information relationships within major cloud–AI

partnerships. The system should be evaluated as a route to productive capability and customers, not as isolated chip, model and software markets. [20]

A corrective program, not another incumbent franchise

The remedy begins with access. Restore commercially usable routes for lawful independent equipment ownership and international operation. Test proposed rules against the path an unrelated developer must take to become a serious supplier. A licensing route that only the largest incumbents can navigate is a barrier even if it is formally available to everyone.

Next, prevent incumbent influence over the right to become an incumbent's competitor. Public evaluation can examine hazards without letting a commercial rival determine admission, acquire a challenger's sensitive information or impose compliance costs unrelated to the activity. Published decisions and usable appeals should make lawful entry predictable rather than dependent on exceptional political access.

Then reopen purchasing. Remove contractual and technical restrictions that tie investment, cloud spend, model use and distribution into a single captive channel. Government buyers should retain their requirements and compare outcomes across qualified suppliers. Publicly supported capacity should serve a field of users, not only firms affiliated with a favored laboratory.

Finally, assess interventions against NVIDIA by their downstream consequences. Weakening the supplier that makes independently owned computing practical can increase the power of integrated buyers. The correct industrial test is whether more independent firms can build and earn afterward.

Frontier laboratories and hyperscalers already provide competition for CUDA. CUDA and Computraq provide competition for them. Policy that extinguishes the second direction while celebrating the first is not protecting an open market. It is choosing its next gatekeepers.

12. Turn the credit network into an application engine

Prepared capacity, vendor relationships and customer revenue can reinforce one another.

Prepared infrastructure is a competitive advantage

NVIDIA’s existing support arrangements created more than contingent obligations. They helped assemble operators, installed capacity, facilities, technical expertise and vendor relationships. That network can bring a new application market to scale faster than building an equivalent supplier base from the beginning.

The July 2026 filing describes approximately **$36 billion of AI-cloud commitments**, typically extending over six years, that decrease as eligible capacity is used by third-party customers or NVIDIA. Separately, an August initiative proposed financing platforms intended to mobilize more than **$500 billion of third-party capital**. The former is a contractual capacity network; the latter is a proposed capital-formation channel. Computraq connects those relationships to a shared application market and identifiable customer demand. [37,38]

Computraq can use the network in both directions. Available qualified capacity makes a buyer’s desired service deliverable. That makes developers more willing to publish and customers more willing to purchase. Collected revenue improves the operator’s economics and makes further capacity easier to underwrite. NVIDIA’s engineering and commercial relationships shorten the time required to establish those connections.

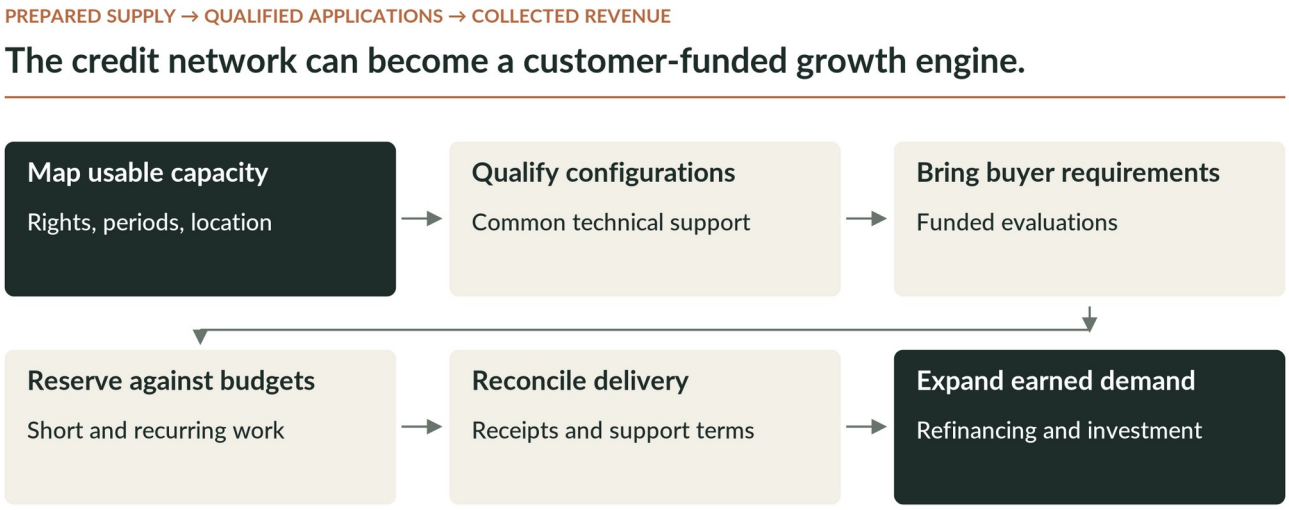


Exhibit 33. Capacity commitments, proposed financing and actual customer collections have distinct roles in the growth programme. [37–38]

Converting prepared capacity into earned revenue

The first task is to identify capacity that is actually usable: provider, location, equipment, available delivery period, execution configuration and rights to sell to third parties. That catalogue is matched with customer requirements rather than treated as an undifferentiated number of GPUs.

The second task is to reduce repeated qualification work. Operators support a defined set of CUDA configurations; developers can build and test against those configurations; NVIDIA's engineering teams help resolve common deployment barriers. Each improvement becomes reusable across participants instead of remaining a private fix for one large customer.

The third task is to connect reservations to real budgets. A limited evaluation program can be sponsored explicitly. Production services are purchased under customer authority. Recurring portfolios can support negotiated capacity blocks; variable demand can use shorter reservations. The market supplies the common purchasing process while contracts identify who owes what and when.

Finally, delivered service and money are reconciled. Providers record the revenue earned. The parties to a support agreement determine the treatment of qualifying third-party usage. Operators can present a better cash-flow history to financiers, and additional support can be directed toward proven demand or specific bottlenecks. The cycle turns existing industrial relationships into a distribution advantage.

At the central frontier-lab challenge, $28.35 billion of service spending produces $5.67 billion of execution purchases under the 20% compute assumption. At $150 billion of services, the corresponding pool is $30 billion. If half reaches providers participating in a relevant support program, the annual receipts are $2.835 billion and $15 billion respectively. The provider agreements identify which receipts correspond to eligible capacity and how the original support is accounted for.

The important change is how the network is used. A manufacturer-backed purchase can help assemble supply. An application market can then make that supply valuable to unrelated buyers. The existing credit relationships become an accelerator of commercialization rather than a story consisting only of possible future cash calls.

The operator, vendor and financier each have a role

An operator should translate capacity into exact, dependable offers: supported configurations, location, price, availability, metering, isolation and recovery. Its commercial objective is additional contribution after power, facilities, staff, maintenance and finance. Repeat work across unrelated customers can be more valuable than a large but concentrated stream whose renewal depends on one laboratory.

Equipment distributors, system builders and service firms keep the progression from personal ownership to professional operation practical. Compatible parts, repair knowledge and understandable support are important to the small firm that cannot maintain a bespoke hardware team. The market should reward vendors that reduce the cost of becoming a competent operator.

Financiers can evaluate collected revenue, recurring commitments, costs and customer concentration through consistent records. A buyer's commitment to Computraq becomes financeable for a particular asset only when the relevant payment and capacity rights are established. This contractual step is the route to independent underwriting, not an obstacle to it.

Computraq may use separately financed computing vehicles where they add unusually competitive supply or fill a genuine gap. The main strategy remains a market for many operators. Owning every machine would recreate the concentration problem and consume capital that can otherwise improve distribution and service quality.

The opportunity for the supply chain is therefore straightforward: **earn from a larger field of applications without requiring common ownership of the applications, equipment and customers**. That is how the ecosystem can grow beyond a small set of favored infrastructure purchasers.

A commercial program with a capital multiplier

A marketplace faces an early coordination problem: developers want dependable execution before investing in a service, while operators want paying work before reserving capacity. NVIDIA's prepared network can shorten that cycle. Supported operators, existing equipment, known vendors and shared engineering reduce the amount each participant must assemble before the first customer can buy.

A targeted launch program can therefore be designed around demand portfolios rather than a general promise to purchase GPUs. Identify several recurring customer requirements. Match them to qualified configurations at more than one operator. Fund evaluation explicitly. Make production reservations against actual buyer budgets. The program then has a direct way to distinguish development support from continuing customer revenue.

The financing advantage grows with repetition. An operator can show accepted service deliveries, collections, repeat use, customer concentration and the contribution left after power, facilities and support. A financier can judge specific capacity or equipment against those records. The operator does not need every underlying customer to negotiate a separate cloud contract; the common market records the relationship.

Service-market scale	Execution purchases at 20%	Receipts reaching participating supported providers at 50%
Frontier-work challenge: $28.35B	$5.67B	$2.835B
Digital-application scale: $150B	$30B	$15B
Broader application economy: $300B	$60B	$30B

The relevant contracts can convert eligible usage into reduced residual purchase obligations, additional third-party revenue or other agreed economic benefits. The model treats each of those as a separate financial channel. This allows an analyst to see both the gross commercial activity and the specific cash effect rather than collapsing them into one headline.

Preserve multiple ways to earn

Computraq can supplement the external supplier network through separately financed compute vehicles where there is no suitable independent supply or an exceptional, verified full-cost advantage. A vehicle holds defined assets and obligations; it need not become the owner of the whole marketplace. Its role is to add useful supply and strengthen competition.

The advantage of the model is that the market works across many owners. NVIDIA can supply technology and engineering; vendors supply systems and support; independent operators supply qualified execution; developers supply capabilities; buyers supply the ultimate commercial demand. Capital can enter at the point where a defined investment earns a return rather than being required to finance every layer together.

For NVIDIA's board, this reframes the existing credit portfolio. It is not only a set of potential future payments. It is an accumulated distribution and operating network that can be used to build a new market. Computraq offers a way to put that network to work around customer-defined services and to let its performance become visible through the earnings of independent participants.

13. NVIDIA: the application-economy revaluation

New demand, productive support commitments and a longer-lived earnings franchise.

Move from order concentration to a broader franchise

For NVIDIA analysts, the significance of Computraq is not its first GPU purchase. It is the prospect of a recurring application market that adds independent reasons for customers to use CUDA. A hardware supplier whose future depends heavily on a few purchasing organizations faces a different bargaining structure from one whose systems enable a large population of profitable businesses.

Reuters reported NVIDIA trading at approximately **12 times expected earnings on November 12, 2026**, its lowest forward multiple since 2012. The dated closing-price reference used here is **$166.19**. Dividing that price by 12 gives a normalized forward earnings-per-share input of approximately **$13.85**. [39,40]

A low multiple can reflect doubt about growth, margins, cash requirements or the durability of the competitive advantage. Computraq addresses those questions through three connected operating changes. More applications create a broader demand base. Third-party use makes supported capacity more productive. A commercial market gives CUDA developers a reason to stay even as technical porting improves. The result can be stronger earnings and greater confidence in how long those earnings persist.

Separate the financial channels

The direct revenue channel begins with execution providers. More customer work improves utilization and can justify replacement systems or additional equipment. Provider spending includes energy, facilities, people, networking and finance as well as NVIDIA hardware. A manufacturer model should estimate the equipment response rather than equate all Agent spending with chip sales.

The cash-flow channel concerns existing support. Qualifying third-party purchases can alter the expected use and remaining obligations of particular arrangements. Better-performing counterparties can also improve the

economics of related investments. Analysts should reflect the resulting cash consequences in the relevant contractual periods.

The duration channel is broader. Developers may choose CUDA because customers, qualified configurations, experienced operators and earnings are available through its application economy. That commercial complement can remain valuable even when translating software to another chip becomes easier. A thriving set of independent firms also supplies public evidence of why the platform matters to competition and industrial resilience.

These channels support a different account of NVIDIA's future: not only a company supplying a concentrated construction boom, but the technical foundation of an expanding services economy. That change is the basis for a revaluation argument.

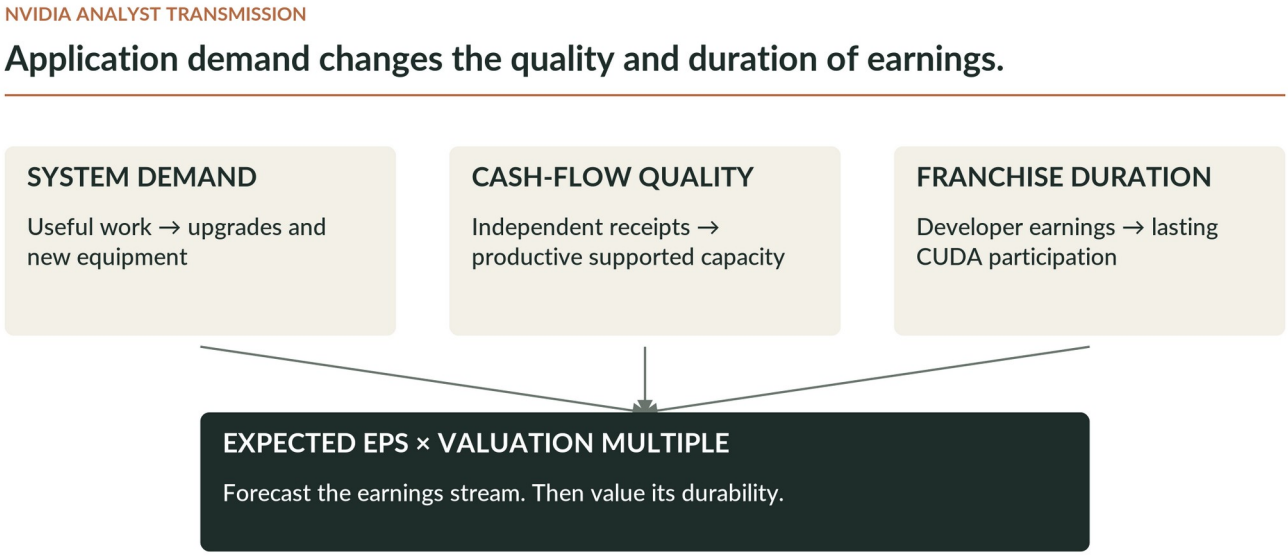


Exhibit 34. The analyst framework separates the direct hardware channel, contractual cash effects and the duration of the wider franchise.

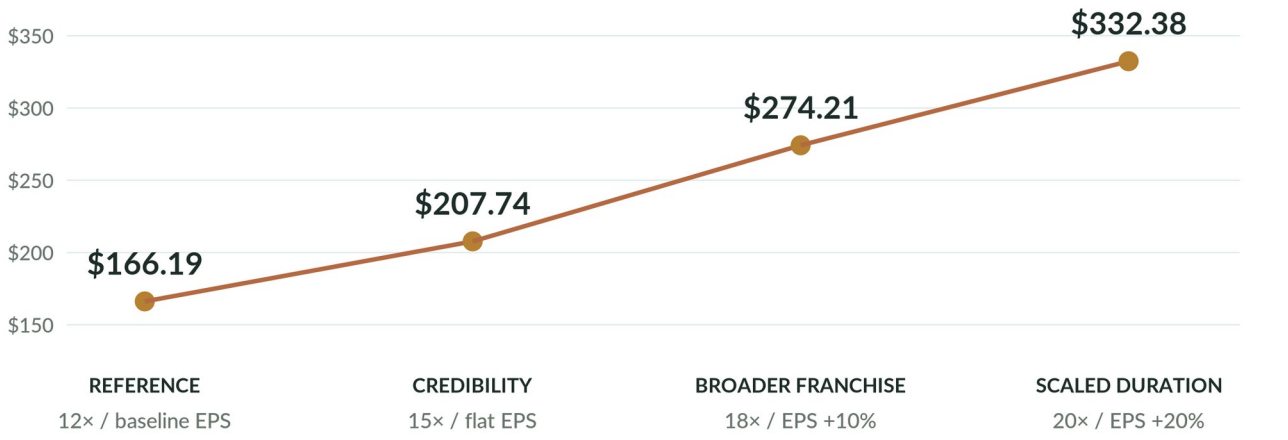
Earnings, duration and the share price

The table shows how a 24–36 month reassessment could combine improved expected earnings with a higher forward P/E. Each case has an operating interpretation rather than relying on a favorable multiple in isolation.

Scenario	Expected EPS versus reference	Forward P/E	Indicated share value	Change from reference
Dated reference	Unchanged	12×	$166.19	—
Independent demand gains credibility	Unchanged	15×	$207.74	+25%
Broader recurring application franchise	+10%	18×	$274.21	+65%
Scaled demand and stronger earnings duration	+20%	20×	$332.38	+100%

ANALYSTS AND SHAREHOLDERS / 24–36-MONTH SCENARIOS

NVIDIA: a conditional earnings-and-multiple framework



Reference: 12 Nov 2026 close $166.19 and ≈12× forward P/E. Original scenarios, not consensus targets. [39–40]

Exhibit 35. Share-value sensitivities from the dated reference price and explicit earnings and multiple assumptions.

At a fixed illustrative 24 billion shares, the reference equity value is approximately $3.99 trillion. The subsequent cases imply about $4.99 trillion, $6.58 trillion and $7.98 trillion. The arithmetic explains why changes in perceived durability can matter more than the initial revenue of an ecosystem program.

The credibility case would be supported by repeated independent purchases, several qualified providers and effective use of established capacity relationships. The broader-franchise case adds growth across customer cohorts and application categories, leading analysts to raise expected earnings. The scaled case assumes that the market supports a significant, sustained execution economy and supplier investment. These are operating milestones against which the thesis can be evaluated.

A $150 billion application market at 20% compute intensity produces $30 billion of execution-provider revenue. One transparent direct-revenue bridge might assume half is incremental activity, equipment expenditure is 40% of that incremental amount, NVIDIA supplies 70% of the equipment value, and incremental operating contribution is 55%. That yields $4.2 billion of NVIDIA equipment revenue and $2.31 billion of operating contribution before tax. The direct equipment channel is $2.31 billion of operating contribution; the broader revaluation case concerns the duration and breadth of the entire franchise.

What shareholders should ask management to do

Treat independent application demand as a strategic use of the ecosystem NVIDIA has already built. Support common deployment profiles, customer-relevant performance work, operator qualification and a bounded credit-to-demand program. Measure the return through collected customer spending, developer proceeds, supplier contribution and confirmed contractual effects.

The most useful disclosure is not the number of available GPUs or registered developers. It is the number of independent businesses earning, the recurrence of their customers, the diversity of their work and the investment justified by those earnings. Those measures make ecosystem strength visible in financial terms.

A successful program would improve the description of NVIDIA's credit network at the same time. It would become a commercial channel connecting prepared supply to new demand. Analysts could model it as an asset in building the next market, not merely a collection of contingent exposures.

The investment case is therefore a stronger, longer-lived franchise. Computraq can contribute by making CUDA useful to a larger population of independently profitable participants. The operating objective is a broader, more durable franchise. The share-price mechanism is the capitalization of that stronger earnings stream.

A disciplined bridge to market capitalization

The central share-price relationship is multiplicative: expected earnings per share times the earnings multiple. In the reference framework, a shift from 12 times to 15 times at unchanged earnings increases value by 25%. An increase of 10% in expected earnings combined with an 18-times multiple increases value by 65%. A 20% earnings increase combined with 20 times produces twice the reference value.

The same arithmetic applies across the equity base. Using the model's fixed 24 billion shares, the $166.19 reference represents $3.989 trillion. The $207.74 case represents $4.986 trillion, the $274.21 case $6.581 trillion and the $332.38 case $7.977 trillion. These figures are a transparent valuation framework. The earnings, share-count and multiple assumptions are all editable in the accompanying workbook.

The direct equipment example supplies an important scale anchor. After an illustrative 30% tax charge, $2.31 billion of operating contribution is approximately $1.617 billion, or $0.067 per share at 24 billion shares. The larger earnings-growth cases therefore describe a wider franchise reassessment, including future applications, replacement demand and other operating developments—not a mechanical extrapolation from one pool of marketplace purchases.

The main long-term asset is the ability to earn across multiple cycles. A chip with superior performance today creates current orders. A platform around which many independent businesses repeatedly form can sustain replacements, upgrades, developer investment and customers beyond the roadmap of any one laboratory. Analysts can express that change through a longer period of attractive growth, stronger cash conversion and a more durable terminal earnings base.

What would change an analyst's model

The leading commercial measures are independent customer purchases, recurrence after onboarding support, service categories added, developer proceeds and the earnings of qualified operators. The financing measures are collected receipts at supported capacity, changes in residual purchase expectations and documented contractual treatment. The strategic measures are the population of firms with profitable reasons to build and operate CUDA services.

Management can report these measures in an application-demand program alongside hardware launches and large customer agreements. Shareholders can then distinguish technology leadership from commercial concentration, and infrastructure commitments from the customer revenue they help create. That is how a qualitative ecosystem story becomes evidence relevant to valuation.

14. Build an economy in which outsiders can succeed

A practical agenda for employees, developers, customers, suppliers and public authorities.

NVIDIA leadership: sponsor the market's productive functions

The decision is to treat independent application demand as an ecosystem business, not merely as another developer-relations activity. Assign an accountable commercial sponsor and an engineering counterpart to a defined portfolio of customer work. Connect the program to qualified operators and channel partners. Let the market determine which Agents and configurations earn repeat purchases.

The first deliverable is a portfolio with identifiable buyers, requirements, qualified services and available execution. The second is a repeatable technical profile that more than one operator can supply. The third is an economic account showing what the customer bought, what the developer earned and what contribution reached the operator.

Commercial cooperation with Computraq can be organized around those outputs. Support can take the form of technical collaboration, introductions, training, qualified operator programs and vendor access. The relationship's value should be measured through the independent businesses it enables and the continuing CUDA demand they create.

Decision to make: establish a dedicated independent-demand program. **Counterpart:** Computraq's marketplace and operator teams. **Measure:** recurring customer spending, sustainable developer proceeds and profitable execution across unrelated buyers.

Employees: extend the reach of the engineering

NVIDIA employees built the compilers, libraries, drivers, networking and support practices that make accelerated computing usable. A broader application economy gives that work another route to impact. A small developer can turn an optimization into a better offer. A regional operator can make a previously awkward service dependable. A customer can purchase work it could not economically obtain before.

The professional incentive is to extend the platform Jensen Huang and NVIDIA's employees built together. A larger population of independent businesses gives that engineering a continuing commercial life. Employees can connect their technical contribution to useful applications, successful developers and customers who keep returning.

Many employees also participate in NVIDIA's future through equity. Stronger earnings and a more durable franchise can increase the value of existing fixed-share holdings and awards as they vest. That financial alignment gives ecosystem work personal significance. A broader population of successful CUDA businesses gives employees a concrete reason to remain engaged with their work and community.

Engineering teams can select customer-relevant workloads and make configurations reproducible. Their measure is accepted work per unit of cost, time and energy, including failures and review. Improvements in memory use, deployment reliability or isolation can be as commercially important as raw model throughput. The objective is to make the capability usable by another team, not to leave every successful deployment dependent on its original specialist.

Developer-relations teams can connect independent builders to defined customer needs. Office hours, optimization challenges and technical training become more valuable when participants can qualify for actual purchasing. Success is a developer operating a useful service and retaining customers, not attendance or downloads alone.

Commercial and cloud teams can bring multiple providers into a common supported environment. Finance teams can identify eligible capacity and map third-party purchases to existing relationships. Together they make an engineering program into something the supply chain can plan around. Approved collaborations can turn those contributions into reusable public deployment knowledge and dependable commercial relationships.

Developers: specialize in work worth buying

An independent developer should begin with a specific job and a complete service proposition. State the inputs, supported tasks, model dependencies, tool permissions, expected capacity and price of the 24-hour reservation. Show how exceptions are handled. Price computing, maintenance and customer support rather than relying on the apparent cost of one successful call.

The benchmark supplies a direction for improvement. Reducing review may create more customer value than increasing a general score. A reliable exception path may qualify the service for a buyer who could not previously use it. A smaller model with a better workflow may earn more than an expensive universal model applied indiscriminately.

The distribution opportunity is equally important. Reuse the relevant evidence, deployment package and service explanations across compatible customers. Measure acquisition cost and retention, not only gross reservations. The shared market is valuable when it makes the next customer less expensive to serve and leaves the developer with sustainable earnings after fees and computing.

Model maintainers, kernel specialists and other component developers can participate through maintained dependencies, qualified configurations and specialist services. The application economy creates more ways to earn from technical contribution than a market organized solely around a few frontier laboratories.

Benchmark creators: become the customer’s replaceable buying expert

A strong benchmark creator does more than publish a test. It maintains representative cases, useful cost measures and a review process that tracks the buyer’s actual needs. It compares services after relevant fees and operating effort. It identifies when a result has stopped being representative and when a new capability deserves consideration.

The 10% fee creates an incentive to do this work professionally. Continued selection must be earned. A customer can choose another creator, negotiate a different arrangement within the available programs, or maintain its own benchmark. Creators should therefore measure customer outcomes and the durability of their recommendations, rather than treating a benchmark as a perpetual toll on spending.

This market is open to specialist consultancies, industry experts, evaluators and integrators. The platform makes their purchasing intelligence usable at a scale beyond one bespoke advisory engagement. It also subjects that intelligence to the same discipline applied to Agents: evidence of value and an available alternative.

Customers: keep the requirement, not the dependency

The buyer’s first commitment is a meaningful recurring portfolio with an accountable owner. Record current AI spending, token consumption, review, maintenance and the cost of changing the service. Define equivalent accepted output and the permissions needed to produce it. Then invite competing offers against that definition.

A second-source program should become large enough to matter to the budget owner. The goal is not to collect another subscription; it is to make a substantial share of spending responsive to measured value. Keep useful incumbent capabilities and direct the rest toward better qualified services as the evidence develops.

Retain the requirements, evaluation history and operating knowledge needed to change suppliers. Use the benchmark creator as an expert, not as a new gatekeeper. Require an explicit all-in price and a clear account of support. The customer’s strongest asset is the ability to make the next purchase without starting from ignorance.

Operators and the industrial supply chain

An execution provider's product is a dependable service, not a theoretical inventory of GPUs. A useful offer identifies the hardware and software configuration, available capacity, region, permissions, isolation, meter, price and recovery responsibility. It lets an independent developer know what it can promise a buyer. Suppliers should begin with configurations they can reproduce and maintain, then expand that set through measured qualification.

The commercial reward is access to customers the provider could not economically acquire one by one. A provider serving many unrelated Agent developers can diversify demand across workflows and industries. Its relevant measure is incremental operating contribution after electricity, facilities, maintenance, financing and idle capacity. The market's shared evaluation and purchasing processes can reduce the sales and support burden per deployment as well as improve utilization.

System builders, distributors, memory and networking suppliers, data centers, utilities and service specialists all have a role in making that proposition repeatable. The path from a personal prototype to a reliable commercial service requires compatible equipment, repairs, understandable software support and available facilities. A broad application market gives those practical services more customers.

Infrastructure finance: connect obligations to productive work

A lender needs to know who owes the money, when it becomes due, which service must be delivered and what remedies exist. Computraq's purchasing records can make that account more legible. Repeat transactions can show revenue quality, concentration, utilization and service performance. A separate capacity agreement can identify the enforceable rights relevant to a particular financed asset.

The public purpose is not to give every equipment owner a guarantee. It is to let independent operators become financeable through useful business rather than permanent dependence on one manufacturer's discretion or one hyperscaler's purchasing decision. An institution capable of recording and organizing that business is a valuable part of the supply chain.

The public-sector decision

Agencies and legislators should assess the independent route before restricting it. A qualified outsider needs a usable path to equipment, models, evaluation, customers and payment. A rule that makes those steps affordable only inside an incumbent's commercial system has allocated a franchise. Corrective policy should reopen the route rather than create an exception available mainly to the well-connected.

Public procurement can reinforce that correction. Agencies can retain their own requirements, invite multiple qualified implementations and record the complete cost of accepted work. They can acquire an operating capability while retaining knowledge that supports future competition. Shared evaluation and appropriately open supplier programs can build a broader national capability base than a permanent attachment to one laboratory's brand.

A common program, distinct responsibilities

Constituent	Immediate contribution	The result to measure
Customer and procurement team	A funded recurring workload, acceptance rules and current AI baseline	Complete cost per accepted result and a usable alternative
Agent developer	A versioned capability, realistic reservation price and support plan	Repeat purchases and proceeds after delivery costs
Benchmark creator	A maintained purchasing standard and useful comparisons	Customer value after the creator's fee
Integrator or specialist engineer	Approved interfaces, tests and transferable operating knowledge	Successful handover and lower repeat deployment effort
NVIDIA engineering and ecosystem teams	Reproducible configurations, optimization and supplier access	Independent applications and operator earnings enabled
Computing operator and vendors	Qualified capacity, transparent terms and dependable service	Diversified profitable utilization
Infrastructure financier	Capital against defined obligations and operating evidence	Collections, service performance and coverage
Public authority	Remove entry barriers that channel independent firms into incumbents	More unrelated firms able to build and compete

A PRACTICAL DIVISION OF RESPONSIBILITY

Different participants. One complete customer transaction.

BUYERS
Define work and approve spending

DEVELOPERS
Prove a supported capability

BUYING EXPERTS
Maintain useful comparisons

ENGINEERING TEAMS
Make configurations reproducible

OPERATORS + VENDORS
Deliver dependable computing

PUBLIC AUTHORITIES
Preserve independent entry

The shared outcome: accepted work, sustainable earnings and another purchasing decision the customer can control.

Exhibit 36. Each constituent has a distinct contribution; no participant must own the entire chain.

The program should be judged through a complete transaction and its repetition. A customer defines work; competing developers demonstrate capability; an operator supplies qualified execution; the buyer authorizes a

reservation; service and money reconcile; a later improvement is considered against the same retained requirement. When the next transaction requires less selling, integration and uncertainty than the first, the platform is doing its job.

Integrators and consultancies: turn expertise into a reusable business

Start with a workflow the firm already understands. Convert its discovery questions, acceptance tests and implementation knowledge into a maintained purchasing standard. Publish an Agent where the firm can deliver the capability itself, or help the buyer select and connect a qualified external service.

The deliverable should outlive the initial project: a benchmark, an approved integration, operating instructions, a handover record and a clear support offer. That makes the firm's contribution more valuable each time it is reused. A new client benefits from previous work; the consultancy can earn from evaluation, implementation and continuing maintenance rather than reproducing the entire discovery process.

A firm can retain direct client relationships while using Computraq as the purchasing and execution market behind them. This gives smaller consultancies access to a wider field of AI capabilities and dependable infrastructure without having to own every component of the service.

Decision to make: select a repeatable practice area. **Counterpart:** customers, benchmark creators and Agent developers. **Measure:** lower deployment effort per client, repeat use of the method and contribution after support.

Equipment vendors: sell a path to productive operation

Offer more than a bill of materials. Begin with a small number of reproducible CUDA configurations matched to qualified Agent workloads. Package installation, software versions, support, repair and an understandable upgrade path. Connect the buyer to the operator program through which that equipment can earn.

A useful first deliverable is a configuration sheet that answers the questions a new operator actually has: which services it can run, how much concurrent capacity it supports, what it costs to maintain and whom to contact when a component fails. It should be usable by a developer selecting execution as well as by a business acquiring equipment.

A vendor can then work with a capital provider on financing for operators with demonstrated demand. Customer revenue provides a stronger basis for expansion than a general expectation that every new machine will find profitable work.

Decision to make: qualify an equipment-and-support package. **Counterpart:** independent operators and financing partners. **Measure:** time to productive deployment, service contribution and repeat equipment demand.

Analysts and shareholders: follow the independent earnings base

Treat the number of independently profitable CUDA businesses as an important measure of the platform's durability. Track customer recurrence, application categories, developer earnings and operator contribution alongside large hardware orders. A broadening market changes the bargaining structure around future equipment purchases.

Keep the direct equipment channel, the effect on supported capacity and the duration of the wider franchise visible in the model. A dollar paid to an Agent developer is not a dollar of NVIDIA revenue. The analyst's work is to identify how increased useful activity affects equipment purchases, cash obligations, margins and the length of the earnings stream.

The first deliverable is a market-to-manufacturer bridge with explicit inputs. The second is an operating scorecard capable of updating those inputs from actual purchases. The share-price sensitivities in the financial chapter then become a framework for interpreting evidence rather than a substitute for it.

Decision to make: add an independent-application-demand module to the NVIDIA model. **Counterpart:** company disclosures, operator results and marketplace commercial data. **Measure:** incremental earnings, cash conversion and franchise duration.

The first operating cycle

A productive launch joins work that already has an owner. Customers contribute a small set of recurring requirements. Developers package services that meet them. Benchmark creators maintain useful comparisons. Operators qualify available execution. Computraq connects the account, offer, reservation and settlement. NVIDIA and its industrial partners can improve the common technical and supply conditions around those transactions.

The first cycle is complete when a buyer has paid for delivered work and the relevant participants have received their agreed proceeds. The second cycle is more important: the buyer renews or selects a better qualified offer using the knowledge retained from the first purchase. Repetition establishes the institution.

AN OPERATING PROGRAM

Ninety days organized around work, purchases and evidence.



Every audience has a concrete contribution: requirements, services, execution, distribution, capital or market access.

Exhibit 37. The operating program advances through completed commercial cycles, each leaving reusable knowledge and a clearer account of value.

A practical first ninety-day program can be organized around three successive achievements. The first month establishes requirements, offers and the first delivered reservations. The second month repeats purchases, improves the common configurations and measures the full economics. The third month expands the successful workflow into additional customers or operators and identifies the equipment investment that the new demand justifies.

The calendar is a management plan. The commercial milestones determine progress: qualified work, completed purchases, repeat demand and sustainable participant earnings. Each organization has a distinct job, and none needs to own the whole chain before it can contribute.

15. The independent institution at work

A company beside NVIDIA, a market for its ecosystem, and a business of its own.

Imagine the market after an ordinary month of operation. A small manufacturer has reduced the cost of processing a recurring body of documents. The developer that supplied the Agent has acquired another customer without repeating a long enterprise sales cycle. The benchmark creator has earned a fee because its method helped the manufacturer choose well. The operator has earned execution revenue from a workload unrelated to its largest tenant. Computraq has earned its marketplace fee. The customer's requirement and results are ready for the next purchasing decision.

No one needed to own every part of the transaction. The manufacturer did not buy a data center. The developer did not create a hardware supply chain. The operator did not become an expert in the manufacturer's business. NVIDIA did not need to administer the customer's procurement. Each participant supplied the part it could perform productively.

Now repeat that pattern across software maintenance, document work, industrial analysis, scientific computation, translation and thousands of specialist applications. The commercial institution becomes valuable because each additional requirement, qualified offer and successful delivery makes further business easier. Its assets are not only software. They are the accumulated relationships and purchasing knowledge through which people continue earning.

This is the institutional opportunity in the C.V. Starr comparison. The private organization can remain distinct from the larger public enterprise while extending its economic reach. It cultivates markets that deserve dedicated attention. It develops its own commercial methods. It earns from business that would otherwise be more expensive, less organized or harder to reach.

Jensen Huang's achievement supplied the technical and industrial foundation for this possibility. Cornelius Vander Starr's example explains why international commercial institutions can become much more consequential than the products with which they began. Computraq brings those ideas into the application economy: privately controlled market-building, a shared technical foundation and separate businesses that prosper by making one another more useful.

The company's success should therefore be visible in the independence of the people around it. Customers retain the ability to choose. Developers retain a route to market. Experts retain a business in improving purchasing. Operators retain a reason to invest. NVIDIA gains a broader field of customers and applications with continuing reasons to use CUDA.

Computraq's role is to build the commercial institution through which the ecosystem can keep earning—independently, productively and at increasing scale.

Economic definitions and essential terms

The definitions behind the economics and the evidence.

Commercial definitions

Computraq's commercial program joins customer-owned requirements, qualified Agent offers and CUDA execution. The following definitions are used consistently in the worked cases and accompanying economics workbook.

Qualifying service spending is the pre-tax price of Agent reservations. It excludes a separately accepted benchmark-creator surcharge and the wider value of transactions assisted by an Agent. The working allocation is 20% to Computraq and 80% to the developer before agreed compute. With compute equal to 20% of the service base, the developer remainder is 60% before other business expenses.

The benchmark creator receives 10% of an expressly attributed service purchase. The fee is additional to the service base and belongs to the creator. The model applies that fee to half the service base where a weighted 5% surcharge is stated. All-in customer comparisons include it once.

Economic illustrations

The AI-to-AI comparison holds output at 10,000 accepted cases. The token baseline uses $2 per million uncached input tokens, $0.20 per million cached-read input tokens and $10 per million output tokens. Token quantities include retries. Labor, tools, adaptation and maintenance are explicit scenario inputs. Review and maintenance quantities are editable scenario inputs; accepted output is held constant for the comparison.

The developer illustration is a first-year acquisition comparison. The reusable $20,000 package is allocated across 100 buyers; computing and support costs are held constant to isolate distribution effects. The advertising-agency analogy concerns paid purchasing expertise. Computraq's 10% is its chosen commercial rate.

The frontier-lab objective uses the issuer's approximately $81 billion October annualized pace: 50% of that work, purchased at 30% lower prices, creates $28.35 billion of service spending. Run-rate is distinct from recognized annual revenue. Net displaced revenue deducts any continuing upstream payments to the incumbent. Application-market scale cases organize the commercial opportunity at specified levels of annual service spending.

Provider receipts are operating revenue, not identical to hardware capital expenditure or NVIDIA sales. The credit model separates eligible usage, conversion to the value recognized under an agreement and the period's contractual cap. Contractual eligibility determines how provider receipts affect an existing capacity arrangement.

NVIDIA valuation framework

The analyst cases use the November 12, 2026 close of $166.19 and the contemporaneously reported approximate 12-times forward P/E. Their quotient is a normalized EPS input, not an independently sourced consensus estimate. Share count is fixed at an illustrative 24 billion. Higher-price cases are 24–36-month sensitivities combining the stated earnings and multiple inputs. The direct equipment bridge is calculated separately from the broader franchise cases.

Essential terms

Term	Meaning
Agent	Models, tools, software and permissions organized to perform a defined job.
Benchmark	Repeatable evaluation of work under specified acceptance and operating conditions.
Contestability	A practical ability for a credible alternative to win the next purchase.
Token / API	A unit processed by a model / an interface through which software uses a service.

Term	Meaning
CUDA / GPU	NVIDIA's accelerated-computing software platform / a processor suited to parallel calculation.
ASIC	An application-specific chip; practical flexibility and independent availability vary by system.
Frontier laboratory	A developer of leading general-purpose AI models.
Hyperscaler / neocloud	A very large cloud provider / a specialist accelerated-computing provider.
Merchant hardware	Equipment made available for others to own and operate.
RFQ / RFP	Request for quotation / request for proposal.
FDE	Forward-deployed engineer: an expert working closely with a customer on deployment.
Run / reservation	A bounded 24-hour Agent service under an accepted offer.
EPS / P/E	Earnings per share / share price divided by earnings per share.
Station / Train	The participant’s operating workspace / the environment running one Agent.
Passport / Router / Connector	Account identity and funding / reservation-to-execution coordination / access to a running service.
Managing general agency	A specialist insurance intermediary performing delegated commercial or underwriting functions.
Underwriting capacity	An insurer’s ability to accept a risk and meet the resulting insurance obligations.
SPV	A separate vehicle holding specified assets and obligations.
Offtake	A contractual commitment to purchase defined future output or capacity.

Public sources and reading notes

[02] Anthropic. Claude Sonnet: availability and pricing.

Tariff reference accessed 16 November 2026: Sonnet 5.5, $2 per million input tokens, $10 per million output tokens; cached-read assumption at 10% of the input rate.

www.anthropic.com / public source

[03] Anthropic. Registration statement, Amendment No. 2.

Amendment No. 2, 10 November 2026. Approximately $81 billion October revenue run-rate and enterprise Applied AI activities. Run-rate annualizes recent consumption and subscriptions.

www.sec.gov / public source

[04] American Association of Advertising Agencies. 2024 Compensation Methodologies.

2024. Agency compensation methodologies and the relationship between client spending and professional purchasing services.

www.aaaa.org / public source

[05] Association of National Advertisers. Fees Persist as Dominant Method of Agency Compensation.

2010. Long-established agency fee-based compensation. The report uses the agency relationship as a commercial analogy; Computraq selects its own 10% rate.

www.ana.net / public source

[06] Analysis Group / Apple-supported study. Apple's Global App Store Ecosystem and Its Growth, 2025.

June 2026; Table 1. Estimated 2025 app ecosystem activity: $149B digital goods/services, $1,137B physical goods/services, $151B advertising. Apple-supported Analysis Group study. Digital activity includes some purchases outside Apple billing.

www.apple.com / public source

[07] Cambridge Centre for Alternative Finance. Cambridge Digital Mining Industry Report.

April 2025; Figure 23(a), p. 53. Forty-six hardware respondents; sampled installed equipment around mid-2024. Bitmain 82%, MicroBT 15%, Canaan 2.1%.

www.jbs.cam.ac.uk / public source

[08] NVIDIA. GeForce equipment guide and CUDA Compatibility.

Equipment distribution through retailers, system builders and manufacturers; CUDA software and driver compatibility. Accessed 16 November 2026.

www.nvidia.com / public source

[09] SpaceX. SmallSat Rideshare Program.

Public rideshare services and external launch access. The ownership/service comparison is this publication's analysis.

www.spacex.com / public source

[10] U.S. Department of State, Office of the Historian. Foreign Relations of the United States, 1969–1976, vol. XXXI, document 245.

August 1973. Contemporary diplomatic account of the soybean export restrictions and relations with trading partners.

history.state.gov / public source

[11] Japan International Cooperation Agency. Across the seas: Japanese migration and international cooperation.

20 August 2026. Japan's diversification after the 1973 soybean embargo and the 1974 Japan–Brazil initiative.

www.jica.go.jp / public source

[12] Japan International Cooperation Agency. Japan–Brazil agricultural development of the Cerrado: PRODECER.

PRODECER project record: September 1979–March 2001; three phases; ¥68.4B project cost, including ¥27.9B official development assistance.

www.jica.go.jp / public source

[13] American International Group. Our history: AIG timeline.

Corporate history: Cornelius Vander Starr established American Asiatic Underwriters in Shanghai in 1919; American International Group incorporated in 1967.

www.aig.com / public source

[14] U.S. National Archives. Records of the Office of Strategic Services, Entry 210, box 388.

Office of Strategic Services records, Entry 210, box 388, WN14971. Records concerning Starr and his companies gathering intelligence in the Far East, December 1941–September 1945.

www.archives.gov / public source

[15] International Energy Agency. Global EV Outlook 2026: executive summary.

2026 executive summary: Chinese automakers accounted for 60% of global electric-car sales in 2025; European and North American automakers approximately 15% each.

www.iea.org / public source

[16] International Energy Agency. Global EV Outlook 2026: manufacturing and trade.

2026 manufacturing and trade chapter: China nearly 75% of electric-car production in 2025; Chinese imports 55% of EV sales outside Europe and the United States.

www.iea.org / public source

[17] U.S. Treasury. Clean vehicle consumer savings announcement.

1 October 2024. More than $2B of upfront clean-vehicle consumer tax-credit savings across eligible manufacturers in the first nine months of 2024.

home.treasury.gov / public source

[18] U.S. Department of Energy. Federal Financing Programs for Clean Energy.

September 2014 guide. Tesla's $465M federal manufacturing loan issued in 2010. Loan principal is identified as financing, not a grant.

www.energy.gov / public source

[19] OECD. Survey on Drivers of Trust in Public Institutions: 2026 results.

2026 results: associations between political agency, financial insecurity and trust in public institutions.

www.oecd.org / public source

[20] Federal Trade Commission staff. Partnerships Between Cloud Service Providers and AI Developers.

January 2025 FTC staff study of Microsoft–OpenAI, Amazon–Anthropic and Google–Anthropic partnerships; investment, spending commitments, rights, switching and information access.

www.ftc.gov / public source

[21] U.S. Department of Justice. Department of Justice Prevails in Landmark Antitrust Case Against Google.

17 April 2025. Civil findings of unlawful monopolization in publisher ad-server and ad-exchange markets and unlawful tying.

www.justice.gov / public source

[23] Anthropic. Another $20 million to Public First Action.

21 July 2026. Additional $20M brings Public First Action support to $40M. Public education and policy advocacy; statement restricts candidate-election use and advocates stronger chip and equipment controls.

www.anthropic.com / public source

[24] Coefficient Giving. We're Urgently Scaling Our Work on AI and Biosecurity.

9 September 2026. Forecast of more than $1B in 2026 commitments across technical AI safety, security and fieldbuilding.

coefficientgiving.org / public source

[25] Survival and Flourishing Fund. 2026 grant recommendations and funder allocations.

2026 grant recommendations. Moskovitz allocation: $785,000 for Foundation for American Innovation; $465,000 for Institute for AI Policy and Strategy. Figures are recommendations.

survivalandflourishing.fund / public source

[26] Coefficient Giving. Abundance & Growth grant table.

July 2026 grants table. $10.5M general support for Institute for Progress from a fund with multiple named and private donors.

coefficientgiving.org / public source

[28] Special Competitive Studies Project. About SCSP.

Institutional disclosure: SCSP is a subsidiary of the Eric & Wendy Schmidt Operating Foundation. Accessed 16 November 2026.

www.scsp.ai / public source

[29] CNN / David Wright. How allies of AI are ramping up political donations for the midterms.

11 February 2026. Reporting based on filings for the 2026 midterm cycle: Brockmans $25M and a16z $25M to Leading the Future in late 2025; Thiel's 2025 congressional giving in the hundreds of thousands.

kioncentralcoast.com / public source

[30] Anthropic. Series G funding announcement.

12 February 2026. Founders Fund named as a co-lead of Anthropic's Series G.

www.anthropic.com / public source

[31] Anthropic. Series A financing announcement.

28 May 2021. Anthropic names Dustin Moskovitz and Eric Schmidt among Series A investors.

www.anthropic.com / public source

[32] Dario Amodei. We Must Pace the Frontier.

September 2026. Frontier-development pacing, embedded evaluation and a proposed narrow antitrust waiver for specified safety discussions.

darioamodei.com / public source

[33] Demis Hassabis / DeepMind Institute. A Framework for Frontier AI and the Dawning of a New Age.

14 July 2026. Proposed industry-funded, federally overseen frontier standards body, possible mandatory review, and non-frontier exemption.

institute.deepmind.com / public source

[34] U.S. Bureau of Industry and Security. AI Diffusion Rule rescission announcement.

13 May 2025. Announcement rescinding and declining to enforce the preceding AI Diffusion Rule, alongside other export-control measures.

www.bis.gov / public source

[35] U.S. Bureau of Industry and Security. Revised license review for certain advanced semiconductors.

15 January 2026. Conditional case-by-case review for qualifying H200/MI325X-class exports.

www.govinfo.gov / public source

[37] NVIDIA. Form 10-Q, quarter ended July 26, 2026.

Quarter ended 26 July 2026. $36B AI-cloud commitments, generally six-year agreements; provisions concerning third-party capacity use. Purchase obligations, leases and guarantees are distinct instruments.

www.sec.gov / public source

[38] NVIDIA / Apollo. AI infrastructure financing platforms.

10 August 2026. Preliminary financing-platform initiative intended to mobilize more than $500B of third-party capital over time.

ir.apollo.com / public source

[39] Reuters. Nvidia PE falls to 14-year low, but CEO exit and AI worries cloud the bargain.

12 November 2026. Reported NVIDIA forward P/E approximately 12 times. Used as a dated reference for the analyst sensitivity framework.

www.reuters.com / public source

[40] Yahoo Finance. NVIDIA historical prices.

Historical close, 12 November 2026: $166.19. Cross-checked against Morningstar historical data. Source date matches the P/E reference.

finance.yahoo.com / source 1 · www.morningstar.com / source 2

[41] North Carolina Historic Sites. Alamance Battleground: History.

Eighteenth-century Regulator movement, grievances against dishonest officials and abusive fees; used for the institutional analogy.

historicsites.nc.gov / public source

[42] Nevada Governor's Office of Economic Development. Annual Report 2023.

Tesla's approved $330M tax-abatement package associated with announced investment.

goed.nv.gov / public source

[43] Tesla. Form 10-K for the year ended 31 December 2025.

Automotive regulatory-credit revenue: $2.763B in 2024 and $1.993B in 2025.

www.sec.gov / public source

[44] Founders Fund. Portfolio.

Published portfolio includes OpenAI and SpaceX; accessed 16 November 2026.

founders.com / public source

[45] OpenAI. Introducing OpenAI.

11 December 2015. Peter Thiel named among the nonprofit's founding financial supporters.

openai.com / public source

[46] The Information. Some OpenAI Investors Sit Out Share Sale, Reflecting AI Divide.

6 February 2024. Reports Founders Fund's 2023 participation.

www.theinformation.com / public source

[47] Bloomberg. OpenAI Joins Groundswell of Support for AI Chip-Control Measures.

12 November 2026. Reports OpenAI's support for bipartisan proposals and deliberation over preceding months.

news.bloomberglaw.com / public source

[48] American International Group. 2004 proxy statement.

Historical relationship with C.V. Starr & Co.; private holding-company status and $173 million paid for the production of insurance business in 2003.

www.sec.gov / public source

[49] American International Group. Correspondence to the SEC, 12 June 2006, Annex A.

Historical account of C.V. Starr & Co.'s establishment in 1950 and the specialist agencies it retained as the AIG structure developed.

www.sec.gov / public source

[50] NVIDIA. Corporate timeline.

Founding in 1993 and introduction of CUDA in 2006.

www.nvidia.com / public source

[51] NVIDIA. CUDA Programming Guide: Introduction.

GPU programmability and CUDA's extension of accelerated computing beyond graphics.

docs.nvidia.com / public source