Beyond the Obvious in Defense Production

How engineering choices, incentives, and more coherent arrangements expand American industrial base capacity.

American Industrial Capability Project

TALBOTWEST

An American Industrial Capability Project field report, produced by Talbot West and drawn from its work with industry as well as its wider research.

Contents

Introduction 4

CHAPTER I **Capacity Is a Property of the Production Route** 4

- Capacity is relational, not additive 4
- Product design is a capacity decision 5
- A supplier's skill carries to a higher rate; its delivery record may not 6
- Operational excellence is capacity 7
- Beyond the factory floor, no function owns output 8
- Fragmented information is a constraint on production capacity 9

CHAPTER II **Developing Manufacturing Capability** 10

- The ability to develop production methods is capital that compounds 10
- Use shared technical evidence to qualify more production sources 11
- Plan workforce development as part of the production system 12
- Flexible capacity pays only if demand shifts and the factory can follow in time 12

CHAPTER III **Demand, Contracts, and Production Readiness** 13

- Price a contract's demands on the factory before accepting it 13
- Reshape the deal before the bid is fixed 14
- Set investment start dates by when the capacity must be working 15
- Shorten customer approvals to give suppliers more time to produce 16
- A customer's ordering practices shape its suppliers' productivity 17

CHAPTER IV **Supplier Development and Industrial Investment** 18

Help suppliers build production capability they can sustain on their own 18

Make supplier expansion worth the supplier's risk 19

Choose what to buy by what is actually scarce 20

Fund interdependent investments as a complete set 21

CHAPTER V **Production Capacity Across Companies** 22

Demand too small for any one program can justify shared production 22

Converting commercial plants depends on the design owner 23

Secure design data and rights with each new purchase 24

Prepare production to move when a supplier fails 24

Build versus buy hides most of a manufacturer's options 25

Common data formats lower the cost of every production transfer 26

CHAPTER VI **Industrial Intelligence and New Production Opportunities** 27

Contract records hide the dependencies that govern production 27

AI lets engineers examine the alternatives they used to skip 27

Industrial AI will favor those who pool process data 28

Procedures written for expensive analysis will cap AI's gains 29

The most valuable production arrangements have no natural owner 29

References 31

PUBLICATION	Beyond the Obvious in Defense Production
EDITION	AICP Field Report No. 1 | October 27, 2026
PRODUCED BY	Talbot West for the American Industrial Capability Project
WEB EDITION	talbotwest.com/aicp/beyond-the-obvious-in-defense-production
CONTACT	jacob@talbotwest.com

INTRODUCTION

A production ramp naturally brings attention to new factories, additional equipment, larger workforces, and more suppliers. These are important sources of capacity, but substantial opportunities also exist in how productive resources are used, developed, and combined.

What appears to be a shortage of manufacturing capacity may in part be a challenge of engineering, timing, investment incentives, or industrial organization. These possibilities often escape conventional planning because the possibility space is distributed across surfaces not normally considered together. To perceive and realize available value often requires innovative perspectives and approaches.

To understand why it's normally difficult to perceive and realize these possibilities, see our Theory of Enterprise Coherence.

Talbot West specializes in recognizing and developing high-value productive arrangements across technical, organizational, and commercial boundaries. This report explores practical interventions and emerging industrial possibilities that become visible when defense production is examined systemically.

CHAPTER I

Capacity Is a Property of the Production Route

INDUSTRIAL OBSERVATION

Capacity is relational, not additive

Industrial capacity is not the sum of machines, workers, and suppliers. It depends on how those resources work together, the technical capabilities that make them productive, and the demands placed on the system. The same assets can deliver substantially different output under different arrangements.

The Army's 155mm ammunition program illustrates the point. In March 2026, its plants could produce 46,000 projectile bodies per month but complete only 36,000 rounds because of a bottleneck in another part of the system. Expanding metal-parts production would not have increased finished output at that point in time. Yet planned downstream expansion to 140,000 rounds per month would make metal-parts capacity the bottleneck.

That metal-parts capacity was itself waiting on a capability the Army did not yet have. It had invested $469 million in a facility in Mesquite, Texas, intended to produce 30,000 projectile bodies per month. As of March 2026, efforts to adapt the facility's equipment to the required projectile had produced no parts meeting specifications. Without Mesquite, the Inspector General found, its three producing metal-parts facilities could support only about 71,000 rounds per month, roughly half the planned rate. The constraint on future output lay in qualifying a process on equipment the Army already owned. [1]

Output depends on relationships among facilities, processes, qualifications, and people that different organizations fund and manage, each judging its own piece by measures suited to its position. Those relationships also change over time; the constraint moves as each piece comes online, and today's commitments become conditions every later decision must work within.

INDUSTRIAL OBSERVATION

Product design is a capacity decision

Design requirements impose capacity constraints, yet the people who write them don’t see the full set of cascading effects of those decisions. Visibility is poor because the cause-and-effect cascade occurs across many companies and over long stretches of time. Visibility can be improved, and greater capacity unlocked at the point of design decision.

A drawing may require precision finishing across a housing even though only its bearing seats need it. That choice consumes specialist machine time and inspection on every unit. Approving a standard finish on the other surfaces releases capacity without adding equipment. Other design choices determine which shops can make a part at all.

The costs often land several suppliers away from the engineer writing the requirement. By the time they appear as high prices or late deliveries, the design may be qualified and expensive to change.

We see a major opportunity in reviewing consequential specifications with the firms that must manufacture them, before the design is fixed. A prime or government

program office could ask those firms which requirements consume scarce processes, lengthen lead times, or unnecessarily narrow the supplier pool. The result may be an approved design change, a different production route, or common blanks and processes shared across programs.

Any change still has to meet performance and qualification requirements. The decision is whether redesign and approval deliver capacity sooner or more economically than expanding the original process.

INDUSTRIAL OBSERVATION

A supplier's skill carries to a higher rate; its delivery record may not

A long, clean delivery record proves a supplier can make the product well, and that know-how and quality discipline will still be there when the buyer needs three times as many. Its on-time performance may not be. The supplier earned that record with spare capacity nobody tracked, such as idle inspection hours, a few technicians who fix problems early, a heat-treat vendor with room in its schedule, and overtime held in reserve. A ramp uses that up first.

Aerojet Rocketdyne had made solid rocket motors in Camden, Arkansas, since 1979 when demand jumped in 2022 and its deliveries fell behind. Raytheon's CEO called Aerojet "the weak link" while praising the quality of its product; Aerojet blamed material shortages and a tight labor market. Both accounts describe a good supplier built for a lower rate.

When planning a higher rate, a buyer should find out:

- how close inspection and test already run to capacity, and whether acceptance can keep up when production speeds up;
- which few people the line depends on, and how much of their time training new hires will take;
- which outside processes and sub-tier suppliers the higher rate needs, and whether other programs have claimed the same capacity;
- how the product mix and qualification workload will change.

A trial run at the higher rate, calls to the sub-tier suppliers and an engineering review of the slowest step answer most of these. Some problems will still appear only when the whole line runs at full rate.

All of this depends on the supplier telling the truth, and the supplier has good reason to hold back. Admitting that inspection runs near capacity, or that two technicians hold the line together, looks like a reason to give the work to someone else. The buyer has

to change that calculation. Where the rules allow, keep what the supplier discloses with the people planning the ramp and away from whoever scores the next competition. Treat each problem found as joint work, with engineering help and money attached. Reward suppliers who report problems early with earlier orders and longer commitments. Some fixes reach beyond the supplier's walls: shared test time, spare heat-treat capacity at another company, a sub-tier commitment the buyer backs. A supplier that tells its customer where it will break gives that customer something to plan around. [2]

INDUSTRIAL OPPORTUNITY

Operational excellence is capacity

Plants in the same industry, using the same measured inputs, produce very different amounts. Across U.S. manufacturing industries, the plant at the 90th percentile of productivity produces almost twice the output of the plant at the 10th, and the gap persists year after year. Much of that gap comes from how a plant's people, organization and technology interact.

For a defense plant, the difference between its current output and what its existing people and equipment could produce is capacity it has already paid for.

Plants usually manage and measure elements separately, while output depends on how they combine, and the losses compound. Fixing one practice alone often moves the loss elsewhere, while practices that do little individually produce much larger gains in combination. Island Components, which makes motors and actuators for aerospace and defense customers, changed four related practices in one effort, within its existing workforce, equipment and space:

- feeder-line replenishment;
- work flow between stations;
- purchasing routines;
- bottleneck management.

It reported daily output rising from 75 to 120 units.

A manufacturer found defects in roughly 6 percent of a 500,000-unit production run and faced scrapping entire lots. Because it had photographed every unit during production, its engineers built an AI inspection tool that identified the defective units individually and saved nearly $4 million of usable inventory.

A ramp strains these interactions before any equipment fails. Recovering capacity this way is faster than buying, installing and qualifying new equipment. [3,4]

INDUSTRIAL OPPORTUNITY

Beyond the factory floor, no function owns output

A plant's output depends on a production system much larger than its factory floor, and some of that system lies outside the company. Each part is run by a different function or firm, measured in its own terms and paid from its own budget, and none of those measures counts units of output.

Boston Consulting Group's 2026 fieldwork with more than 50 defense suppliers attributed roughly 70% of supplier underperformance to conditions outside their direct control. Our own field experience with industrial companies strongly corroborates that finding. Many of the most consequential constraints on production originate beyond the organizational boundaries of those held responsible for output. [5]

Each function can run well by its own measures while the plant produces far less than it could. Engineering batches its reviews to keep its engineers efficient, and production waits for releases. Purchasing holds orders to hit a price target, and a line starves for material.

Worthwhile interventions tend to go unfunded, because the function that would pay for a fix is rarely the one that benefits. For example, a training investment charged to one department raises output measured in another. Supplier development is charged to purchasing, while the deliveries it protects show up in program revenue. A fix that benefits several functions often has no sponsor, while several business cases can claim credit for the same improvement.

The value of any one investment depends on the others. Equipment adds output only if its process is qualified, suppliers can feed it, and people can run it. Judged one at a time, investments are misvalued in both directions. A machine gets approved and then waits for qualified operators. Meanwhile, an engineering change that would release capacity across several operations is never proposed, because no function's business case includes its full effect.

The limit on output also moves between functions. Relieving one dependency shifts the constraint to the next, often into a part of the organization that wasn't watching for it. A capacity plan built from one function's view goes out of date as soon as it succeeds.

None of this requires modeling the whole enterprise. It requires drawing the analysis around the relationships capable of changing a decision, while leaving the rest out. The change with the greatest effect is often small and located at a handoff between functions. Examples include an approval rule, the timing of a supplier commitment, or the order in which engineering takes its reviews. A capacity plan should list every

remedy for the constraint on accepted output, whichever function owns it, and compare them by output added, cost and arrival date, alone and in combination.

INDUSTRIAL OPPORTUNITY

Fragmented information is a constraint on production capacity

When a plant's systems don't share reliable information, the plant loses production capacity as surely as when a machine goes down. Capacity plans count equipment, people and suppliers. Every operation also depends on knowing the current revision, whether material is released, how a nonconformance was dispositioned, and what is actually on the floor. When that knowledge is scattered across systems that don't reconcile, it becomes a constraint like any other.

The drag takes several forms. Lines wait while someone confirms what they are allowed to build. Machines run work that a later change makes obsolete, and parts are reworked or scrapped because one system still held the old configuration. The engineers and planners a ramp depends on spend their hours chasing answers through records that dead-end.

Less visible is the capacity a plant holds back to cover for what it can't see. Extra inventory, padded lead times, schedule buffers, overtime and expediting all compensate for information that arrives late or can't be trusted. Each is a sensible response, and together they tie up equipment hours, people and working capital that could be producing.

Planners often build their own spreadsheets and trackers to hold distinctions the official systems lose. Those tools are part of the plant's capacity. A system replacement that retires them without capturing their distinctions loses output its business case never counted.

These losses rarely register as a capacity problem, because they appear as ordinary delay, waste and buffer. Treating information flow as part of the production route changes the investment question. Connecting the systems whose disconnection idles lines, forces buffers or causes rework adds output much as new equipment does, and sometimes faster and more cheaply.

CHAPTER II

Developing Manufacturing Capability

INDUSTRIAL OBSERVATION

The ability to develop production methods is capital that compounds

A manufacturer's ability to develop, qualify and adapt production methods is a productive asset in its own right. It governs how quickly new equipment yields output, how many alternative routes the company can open when a process falls short, and which future work it can take on. Most companies account for it as the cost of each launch.

During a ramp it is often the binding constraint. Equipment can be bought faster than engineers can establish reliable methods for running it, and the same engineers are usually supporting current production.

The capability compounds when each introduction builds on the last. Developing a process teaches engineers more than the instructions they hand to production. They learn the reasons a method works, its sensitivity to changing conditions, and the points that need investigation before it is adapted. When that reasoning is kept, the next team starts from it and spends its effort on what is actually new. When it is lost, each introduction starts over, and the company pays for the same learning again.

A method that works only while its developers attend it is not yet capacity. Its programming, inspection, maintenance and engineering support fall to a production operation already committed to deliveries, so that operation has to shape fixtures, recovery procedures and diagnostic access while those choices are still open. Transfer is complete when the receiving operation produces accepted work and recovers from defined problems without the original expert present.

Fleet Readiness Center Southwest built a Navy F-16 structural-sustainment repair line this way. The depot adapted equipment previously used for F/A-18 maintenance, obtained training and technical knowledge from Air Force and industry partners, and developed local repair specifications. NAVAIR reported that the line was established in 18 months and delivered its first repaired F-16C in February 2026, eight months ahead of schedule. The line came together quickly because the depot built on equipment and knowledge that it and its partners already had. [6,7]

RECOMMENDATION

Use shared technical evidence to qualify more production sources

Qualification cost keeps capable manufacturers out of defense supply. A shop with the right equipment and skills may be unable to justify an extensive test program against the orders it expects. When several prospective suppliers must each establish the same technical findings, the repeated expense limits the capacity available to customers.

Some qualification evidence belongs to a particular producer or part. Other evidence establishes material properties or process behavior that hold under defined conditions. Once that evidence is accepted, a later applicant shows that its own material and process are equivalent and tests what differs.

The National Center for Advanced Materials Performance (NCAMP) provides an established aerospace model. It develops shared composite-material specifications and property databases, and the Federal Aviation Administration recognizes their use when an applicant demonstrates that the evidence applies to its materials, processes and intended applications. A later manufacturer runs a limited equivalency program in place of the original material investigation. America Makes' CAM-QT project applied similar logic to overlapping customer requirements. ASTRO America worked with GE Aerospace, Honeywell and Pratt & Whitney to consolidate their common installation and operational qualification requirements for additive manufacturing. A supplier assembles evidence against the shared requirements once, and each customer keeps its own further approvals. [8,9,10]

Acceptance determines the value of either arrangement. A shared database saves little if customers still require every applicant to repeat the original tests. Producers and approval authorities have to agree on the evidence they will recognize and the further demonstrations they will require.

Shared evidence gets built where one party profits from many others using it. Composite material suppliers sponsor NCAMP qualification of their own systems because an accepted public database makes their material easier for every fabricator to adopt. Qualification for processes such as machining, joining, coating or heat treatment often has no such sponsor. Each shop pays to prove what others have already proved, and the cost is spread across the industrial base outside any one company's budget. Primes, program offices and industry consortia are positioned to fill that gap by funding shared evidence and settling its acceptance. A purchaser seeking more supply should weigh qualifying existing manufacturers against buying and qualifying new equipment, and count the return across every producer able to use the evidence. It can also pay several capable shops to qualify in parallel under

testing-only contracts and award production separately, which buys evidence and options before a shortage forces the choice. [11]

RECOMMENDATION

Plan workforce development as part of the production system

Most manufacturers run workforce development as a separate function. Human resources recruits, a training group teaches, and success is counted in hires and course completions. Workforce capability belongs in production planning, alongside equipment and methods. Planned there, it relieves capacity constraints that hiring alone cannot.

A hiring plan that only counts heads overstates near-term output. Each new worker consumes expert time before adding to production, and the experts who teach are the same people keeping today's deliveries on schedule. Every hour they spend teaching comes out of current output, and every hour they protect slows the growth of the workforce.

The plant has more levers than hiring. Engineering decides how much skill each job requires. A fixture that fits a part only one way, or tooling that holds tolerance without adjustment, shortens training for every operator who follows. The expert hours it saves can go to the exceptions new workers cannot yet handle, such as an out-of-tolerance reading or an unfamiliar defect. A new shift then runs on its own sooner. One engineering change relieves the training load, the expert bottleneck and the shift problem together.

Retention belongs in the same plan. A worker who leaves takes the expert time invested in them. Firms also tend to train narrowly, because a transferable skill can leave for a competitor. Each of these choices changes how much a plant can produce, and none of them shows up in a hiring count.

INDUSTRIAL OBSERVATION

Flexible capacity pays only if demand shifts and the factory can follow in time

Flexible capacity lets a manufacturer move production between products. Its value depends on two things outside the equipment. Other demand has to arrive while the capacity is free, and the factory has to switch before that demand goes elsewhere. Equipment specifications and single-program business cases show neither.

Demand timing matters most. Products whose orders peak at different times keep shared capacity busy, and products that peak together compete for the same hours. Defense demand tends to surge together, so a line flexible across several munitions or missile components gives little relief in a conflict, when every program needs more at once. Capacity shared with commercial work can be the better reserve, because commercial work can be displaced when defense demand surges.

Switching time decides whether the flexibility is usable. A machine may accept several products while the factory needs months to move between them. Each switch requires process programs, tooling, trained operators and often customer qualification. When those steps take longer than the opportunity lasts, the manufacturer has paid for a capability it cannot use. Keeping methods and qualifications current shortens the switch, and that upkeep has its own cost.

Most factories need a mix. Large, steady, lasting demand favors dedicated equipment, which cuts changeovers and supports rate. Early steps shared by many products can justify dedicated equipment on their combined volume, and flexibility earns its cost downstream, where products diverge. Compare architectures against realistic demand patterns and switching times, and count both the cost of keeping flexibility ready and the cost of dedicated equipment sitting idle.

CHAPTER III

Demand, Contracts, and Production Readiness

RECOMMENDATION

Price a contract's demands on the factory before accepting it

Every contract changes the manufacturer that accepts it. The award ties up engineers, test rigs, special processes, suppliers and floor space, often for years, and it builds methods and qualifications that carry into later work. Most companies judge a contract by its margin, which counts the contract's own revenue and cost and leaves out its effects on the rest of the factory.

Most of those effects are visible at bid time. A company knows which of its resources limit output today, such as a test rig that sets delivery for two programs or the few

engineers who disposition nonconformances on a difficult alloy. The solicitation shows which of them a new award would draw on, and whether the work runs on existing methods, equipment and suppliers. Reading a contract against the factory as it stands requires no forecast.

Two awards with similar margins can have opposite effects. One needs a qualification test that only one rig in the plant runs, and that rig already sets the delivery rate for two existing programs. Winning it slows both programs unless the company adds test capacity, a cost the bid never priced. The other runs on equipment and methods the company already uses and adds load where the plant has room. Margin shows neither effect.

The same reading applies across the portfolio. Related contracts share equipment, skills and suppliers, so each award helps pay for capability the others use. Unrelated awards each bring their own tooling, suppliers and engineering exceptions, and the few people who hold the production mix together stretch thinner with each win. Bidding has a cost as well. The engineers who write proposals are the same engineers who launch products and fix production problems, so the bid list is the company's first allocation of engineering capacity.

Future capability is the uncertain part. A program that develops a new method may open related work, or the follow-on may never come. A company can accept a lower margin for a capability with several plausible uses, up to what it would spend to build that capability on its own. A method that applies across parts the company already bids on justifies some allowance. A method useful to one program earns none.

RECOMMENDATION

Reshape the deal before the bid is fixed

A bid usually prices a fixed specification on fixed terms. The customer decides what to buy, and the supplier then decides how to make it and at what price. That sequence splits one decision into two, and neither party can reach the best combination of design, production method, investment and risk on its own.

Each side holds half the picture. The supplier knows which production methods would cut cost or raise rate, and what they would take to qualify. The customer knows which requirements and terms matter to it and which it inherited from an earlier program or a standard contract. Working together before the bid lets them change what is still cheap to change.

The most valuable changes pair a technical choice with a commercial one. A supplier may know a method with much lower recurring cost that first needs development and qualification. The proposed orders are too few or too uncertain to repay that work, so

the supplier quotes the conventional method. A customer contribution to qualification, shared tooling or a longer commitment could make the better method pay, and both sides would get a lower lifetime cost or a higher rate.

The window closes quickly. Once the customer fixes the design and the supplier prices the work and orders tooling, a better alternative has to overcome the cost of reversing those decisions. Draft solicitations and early industry engagement are the time to raise it. The supplier still has to price the whole deal, including development and continuing support, because a price too low to recover its investment weakens the supply the customer is trying to secure.

RECOMMENDATION

Set investment start dates by when the capacity must be working

Every capability a ramp needs has a date when it must be working and a lead time to get there. The need date minus the lead time is the last date to start. Most companies schedule investment by a different clock. They wait for a firm order, or they fix today's constraint and then the next one.

Lead times can reverse the order of starts. Expanding machining might take three months, and qualifying more test capacity might take ten. If test capacity is needed only a few months after machining, test preparation has to start first, even though machining binds first. A plan that addresses constraints in the order they bind starts the long-lead item too late.

Waiting for a firm order fails the same way. Most last-start dates fall well before the order arrives. Once a date passes, the production route that depended on it is gone for that ramp, and the company delivers late or through a worse route. Information has the same deadline. A forecast or supplier assessment helps only if it arrives before the last start of the action it would change.

The buyer controls part of every lead time it faces. A supplier's quoted lead time starts when someone can commit, so the lead time a program experiences runs from the date it could have committed. Advance procurement, multiyear authority and conditional orders against a demand range make commitment possible earlier and shorten that lead time without changing anything in the factory. Approvals spend the same calendar. Capital thresholds and hiring approvals sized for a slower business consume weeks during a ramp, and the cost of each approval includes any option that expires while it waits. Delegating decisions that are small, reversible and within an approved plan can add more capacity than a machine. Waiting on outside money does the same damage. The hope of a grant or a Title III award can hold an investment that already pays for itself past its last start.

Aerojet Rocketdyne's solid rocket motor deliveries fell behind in 2022 as demand rose. In April 2023 the Defense Department committed $215.6 million to modernize its facilities and consolidate production lines. L3Harris acquired Aerojet that July, raised its internal investment in the business by 40 percent and put more than $25 million into smaller suppliers. By September 2024, the company said it had cut its backlog of late motors roughly in half. The investment helped, but it was funded only after deliveries had already slipped. [12,13,14]

An early start does not have to commit the full investment. Qualification planning shortens the remaining approval time, and preliminary engineering settles technical questions that would otherwise delay the decision. Each step buys time for a fraction of the investment's cost. The full commitment can wait until its expected benefit and the cost of waiting justify it.

RECOMMENDATION

Shorten customer approvals to give suppliers more time to produce

A delivery date fixes the end of a supplier's production window. Every week the customer spends on its own approvals comes out of that window before the supplier can start. Engineering release, funding, purchase orders and material authorization each take time on the customer side. Follow-on orders often repeat approvals of drawings and terms that have not changed.

Time returned to the supplier is worth the most where the supplier has open capacity and material. There, a faster release moves delivery earlier. At a saturated supplier, an earlier release only moves the order ahead of another customer's work. The customer gains a place in the queue, and total output stays the same. Buyers should know which case they face before they count the gain.

Much of the preparation can start before the final release. Planning, tooling design and material sourcing need only stable geometry, materials and quantities. A customer can authorize that work within a set scope and agree in advance who pays if the design changes. The authorization works only if it holds. A customer that revises dates and quantities after an early release teaches the supplier to discount its releases, and the preparation time is lost again.

The same loss occurs after the parts are made. Finished work waits on customer test witnessing, inspection, disposition of nonconformances and receipt. Each wait delays delivery and adds no production. The customer controls these clocks and can measure them like any supplier lead time.

INDUSTRIAL OBSERVATION

A customer's ordering practices shape its suppliers' productivity

A supplier's productivity depends partly on how its customers order. Unstable schedules, shifting batch sizes, unnecessary variants and slow answers to technical questions add setups, queues and scrap at the supplier. Each expedite breaks a planned run and pushes other work back, often work for the same customer, so it pays only when the order it advances matters more than the deliveries it delays. The same plant can run well for one customer and poorly for another. A buyer can raise its suppliers' output by changing these practices, without funding a machine.

Forecasts follow a similar logic. Suppliers learn which customers' forecasts hold and plan capacity to match. When forecasts prove unreliable, suppliers wait for firm orders, deliveries slip, and buyers inflate their next forecasts to secure a place in line. Each response is rational, and together they make the shared forecast worthless. A study of one large chipmaker's forecasts to 78 equipment suppliers from 1999 to 2001 documented this cycle. Suppliers delivered later on orders whose requested dates kept moving, and the buyer inflated forecasts to suppliers that had been late. Nearly half the forecasted orders were eventually cancelled. A buyer breaks the cycle by honoring its forecasts and by sharing quantity ranges and delivery windows that narrow as demand becomes clear. [15]

Some variation is worth its cost. A late change that protects a fielded system, or a variant a mission requires, can justify the disruption. The decision should weigh the customer's gain against the full cost at the supplier, including the work the change displaces. Neither party sees that full cost on its own books, which is why so much avoidable churn persists.

CHAPTER IV

Supplier Development and Industrial Investment

RECOMMENDATION

Help suppliers build production capability they can sustain on their own

Many constrained suppliers have the equipment they need and lack the methods to use it well. Chronic scrap, slow setups, an inspection step that cannot keep pace or weak production scheduling can hold output well below what the machines allow. A customer, specialist firm or peer manufacturer that helps fix these problems adds output without new hardware. Customers have good reason to pay for this help. A shortage at the supplier can cost the customer far more than the supplier earns on the parts, so the customer gains from improvements the supplier would never fund alone.

The help should leave the supplier able to keep improving. A visiting team can lift a line quickly, and the gain can fade once the team goes home. The lasting result is a supplier whose own engineers and managers keep the new methods running and extend them to other lines. That outcome depends on a narrow problem the supplier's staff solve alongside the helpers, and a full-time owner inside the supplier who carries the work forward.

Honda's supplier-development program in the early 1990s produced both outcomes. Small Honda teams worked for weeks or months on single production lines at U.S. suppliers. Honda reported average productivity gains of 50 percent on the lines it reengineered at 53 suppliers by 1994. At Capitol Plastics, a Honda team spent nine months in the plant, and one line reported 45 percent higher productivity and 67 percent less scrap. Problems crept back after the team left, despite more than $1 million of Honda assistance. Progressive Industries, a small stamper, kept learning on its own, developed its own die-sensor improvements and more than doubled its Honda sales from 1992 to 1995. Honda learned to screen for capable management and a full-time internal owner before committing a team, and to stop using line improvement on problems that were really managerial. [16]

Assistance has real costs. It takes up the supplier's scarce engineers and supervisors, often during the ramp that made the help necessary. Intensive outside help can also teach a supplier to wait for the next visit. And it cannot replace an

expansion the required rate truly demands. Assistance pays off where a capable supplier's methods are holding its equipment below the output the customer needs.

RECOMMENDATION

Make supplier expansion worth the supplier's risk

When several manufacturers need the same scarce component or process, each tries to secure enough of the supplier's capacity to protect its own production. Without firm commitments, several buyers may count the same capacity in their plans. Even when the commitments hold, competing for capacity adds none.

A bottleneck supplier often has less reason to expand than its customers have for it to expand. The supplier carries the equipment cost, the qualification expense and the risk of idle machines when the surge ends. Most of the benefit lands downstream, where a missing part can stall a program worth far more than the part. Holding back is rational for the supplier, even when the expansion makes clear economic sense for the supply chain as a whole.

Customers can change that calculation. The useful contribution answers the supplier's actual objection. A supplier worried about idle equipment needs a purchase commitment or an agreed share of that exposure. A supplier with ample demand and capital may need engineering or qualification support. Contributions can also build up the chain. A constraint that costs a Tier 1 heavily can threaten far more production value at the prime, and that difference can justify a contribution the Tier 1 could not afford alone. Each contribution should remove a specific barrier and earn its cost through the added output, with clear rights to that output.

The commitment also has to be one a lender will count. Practitioners writing for Business Executives for National Security point out that a $250 million contract with a $50 million base period and four unexercised options is a $50 million contract for underwriting purposes. The buyer keeps its flexibility, and the supplier loses the dependable revenue it needs to borrow for expansion. A customer that wants new capacity has to give up some of that discretion, and only as much as the investment requires.

In July 2025, the Defense Department assembled a package that made MP Materials' planned 10X rare-earth magnet plant financeable. It set a ten-year price floor of $110 per kilogram for the company's main rare-earth product and agreed to buy all of the plant's magnets for ten years. It also bought $400 million of preferred stock and lent $150 million. JPMorgan and Goldman Sachs then committed at least $1 billion of financing. The government took on real risk in return. CSIS estimated that with market prices below $60 per kilogram, the floor could cost the government more than $50 per

kilogram. As of mid-2026, the plant was under construction, and magnets from the company's existing plant were still in customer qualification. [17,18,19,20]

RECOMMENDATION

Choose what to buy by what is actually scarce

A purchaser protecting against a shortage can buy finished inventory, a claim on a supplier's future output, a supplier's continued readiness to produce, or new production capacity. These purchases do different jobs, and only new capacity adds to the total output available. Inventory delivers units at once and costs money to hold, store and replace when designs change. A reservation moves a buyer up the queue. When capacity is fixed, that buyer's gain is another buyer's delay.

The form of inventory decides how much demand it covers. Stock held as finished variants covers only the variants a buyer guessed right. Stock held as common blanks, before the point where variants diverge, covers every variant, so an uncertain mix strands less of it. Its value depends on what happens after release. The blanks still need finish machining, heat treatment and inspection, and those steps have to deliver within the response window. Blanks waiting upstream of an overloaded furnace preserve material and leave deliveries where they were.

Readiness payments keep trained staff, qualified processes, tooling and test capability in place when ordinary orders would not sustain them. That can be the right purchase, because rebuilding a lost capability often takes longer than the emergency that exposes the loss. GAO found that nearly two-thirds of completed Defense Production Act Title III investments with reported outcomes from 2018 to 2024 sustained existing capacity. Preservation was the main result of the government's own industrial investments. It has real value, and a purchaser should know whether it is paying for preservation or for new output.

A claim on capacity is only as good as the other claims against it. The Defense Logistics Agency's Warstopper program pays suppliers to keep surge capacity available, which the agency reports costs less than buying and storing the stock itself. In 2024, GAO reported agency officials' warning that overlapping domestic and deployed demand could overwhelm those suppliers. Some vendors worried that added stock would expire before it could be sold, and at least two said they would not join an expansion.

A purchaser should define the purchase by what it must deliver under stress. That means the demand the claim must survive, the rule for dividing a shortfall among claimants, who pays for idle capacity and unsold stock, and how the supplier will show it can surge. A reservation that has never been exercised proves little about what the supplier can deliver. When total demand exceeds total capacity, inventory,

reservations and readiness only decide who waits. New qualified capacity is the purchase that shortens the wait. [21,22]

RECOMMENDATION

Fund interdependent investments as a complete set

Many production investments have no value until the other investments they depend on are also in place. A missile without its rocket motor is worth nothing as a missile. Similarly, a new forging press adds no accepted parts until heat treatment, a qualified process, trained operators, and inspectors are also in place.

Until then, the finished pieces cost money and earn nothing. Equipment depreciates, new staff draw wages, and the capital tied up in both cannot fund other work. If the missing piece never arrives, the money is lost.

Separate capital reviews misprice these investments. The first piece reviewed shows almost no return, and the last appears to earn the whole gain. Different sponsors also fund the pieces. The capital budget buys equipment, operations trains people and engineering qualifies processes, so each sponsor pays for a benefit that appears elsewhere. No program office would buy a seeker without a motor, yet companies often buy a press with no plan for heat treatment.

Many investments across the industrial base therefore remain incomplete. Equipment waits for a qualified process, a qualified process for test capacity, and capacity for a demand commitment. Funding the missing piece can return more than a new project of the same cost, because it puts existing capital to work. A modest expansion at a lower-tier processor may make a much larger investment elsewhere productive.

Only analysis at the portfolio level prices these investments correctly. It is also difficult, because the cost, capacity and demand information it needs is spread across budgets, functions and often companies, and no single sponsor has the authority to act on all of it. Define the investment by the output it must deliver. Approve the smallest complete set that reaches the required rate as one decision, priced as a whole and scheduled so each piece is ready when the others need it. Improvements that pay for themselves should proceed without waiting.

CHAPTER V

Production Capacity Across Companies

INDUSTRIAL OPPORTUNITY

Demand too small for any one program can justify shared production

Several programs together can often justify an investment none can justify alone. Their parts may differ in drawing, customer and acceptance standard, yet share machining, heat treatment, inspection or tooling. Program-by-program planning hides that shared demand. Each program office sees only its own orders, and the supplier that sees the combined work cannot commit the programs to it. Machines run short batches with repeated setups, and nobody proposes the investment the combined volume would support.

Manufacturing requirements define the useful part families. Parts that look different may run through the same operations and fixtures, and parts that look alike may need different inspection. A supplier serving several programs may have enough compatible work for a dedicated cell, an automated inspection station or a better process. Merging purchase orders adds nothing unless production changes too.

Ingalls Shipbuilding applied this across its ship programs. Under a Navy ManTech project, it grouped covers, filters, vents, lifting hardware and other repetitive components into fabrication families, built common tooling and fixtures, and installed automated work cells serving several programs. Production use began in fiscal 2021. The parts remained distinct, and their shared fabrication work supported the cells. Ingalls could act because one yard held every program. Across companies, someone has to see the compatible work and secure commitments from customers who do not normally plan together. [23]

The same logic applies to skills. Several manufacturers in a region may each need a few people qualified in the same scarce process, such as specification welding or nondestructive inspection, too few for any one of them to justify a training program. Because a trained worker can leave for a competitor, each firm also has reason to train narrowly and hire away from the others. Pooled, that demand can support shared instruction, practice equipment and credentials that every participating employer accepts. Funding and credential recognition decide whether it adds skilled people or only moves them between firms.

Sharing works only within compatible materials, tolerances and qualifications, and each product keeps its own acceptance requirements. A shared process also becomes a shared point of failure. Programs whose schedules peak together will compete for it, so participants should set priority rules before work begins.

INDUSTRIAL OPPORTUNITY

Converting commercial plants depends on the design owner

A commercial plant needs the design owner's help to produce defense goods quickly. The plant may already have equipment, trained workers, process control and a supplier network. These are often the slowest and costliest parts of capacity to build. The plant usually lacks detailed knowledge of the product. The design owner knows the critical features, the tolerances with margin and the ways the product fails. It can also adapt the design to the plant's processes. The design owner can often transfer this knowledge in less time than new capacity takes to build.

In 2020, General Motors converted a components plant in Kokomo, Indiana to build ventilators designed by Ventec Life Systems. Ventec removed functions from its 700-part ventilator to create a 400-part model for GM to build, and the FDA authorized it under emergency rules. A few dozen Ventec engineers helped run the Kokomo line. GM sent parts lists to about 70 of its suppliers and secured every part within about a week. Some 250 workers started after a week of training. The government ordered 30,000 ventilators for $489.4 million under the Defense Production Act on April 8, and GM delivered the last of them on September 1. The line shut down once the order was filled. [24,25,26]

Defense programs often leave a new producer without this help. The producer may receive only a technical data package and then qualify on its own, rebuilding knowledge the design owner already holds. The design owner is often a prime with little reason to help a potential competitor or to reopen a qualified design. The commercial plant compares small, irregular defense orders with its steadier commercial work. Defense orders also bring cybersecurity, export-control and cost-accounting obligations. The government's rights to the technical data may limit its ability to arrange a transfer. To make conversion worthwhile, the buyer can pay the design owner to teach the product, allow design changes that suit the new plant, and order enough, for long enough, to cover the plant's cost of entry.

RECOMMENDATION

Secure design data and rights with each new purchase

A design owner gives up data and rights most readily while it is still competing for the work. A buyer that writes them into the award keeps the ability to bring in another manufacturer later. A buyer that waits faces a sole source that knows it has no alternative. Many buyers wait, because the data costs money today and pays off years later, when the sole source raises prices or cannot keep up. A Defense Department survey of more than 200 sustainment programs found that about half had data shortfalls blocking competition or government repair.

The Navy went more than a decade without rights to the F/A-18's radio-frequency cables. Reverse engineering and stockpiling spares were judged too costly, so maintainers took cables from grounded aircraft. The program later bought a product-level data package with the aircraft's last production lot and negotiated a three-tier license with the prime. The prime kept first right of refusal on new and critical safety parts. Using the data remains hard. The program rejects about 80 percent of data deliveries for improper markings. [27]

The buyer's chances recur after the original contract, whenever the design owner still wants something. A competition for an upgrade can score offers on the data rights they include. A final production order gives the maker one last sale to win. An engineering change creates new design data to buy. Each purchase should cover the right to manufacture, data detailed enough to build from and a check of the files on arrival, since any of the three can fail alone. A data package becomes a working option once a second producer has built a part from it. Primes buying from subcontractors and manufacturers buying tooling can use the same moments. Design owners can use them too, offering tiered licenses that keep their most valuable data and their claim on new work while releasing the rest.

RECOMMENDATION

Prepare production to move when a supplier fails

A buyer can protect production against a supplier failure by preparing other plants to take over the part quickly. With that preparation in place, the buyer can keep its volume with its strongest supplier. The usual remedy, a qualified second source, divides volume between two suppliers and can weaken each one's reason to invest. It also protects production only if it fails independently; two approved suppliers that depend on one heat treater fail together. A backup that has not run the part in years may have lost its people, its tooling or its approval.

In February 1997, a fire destroyed the Aisin Seiki plant that made brake proportioning valves for nearly every Toyota vehicle. Toyota's network held about two days of stock and had no second source. Within three days, other firms in the network had begun producing the valve, most of them for the first time, on general-purpose machining centers. Aisin's drawings lacked detail that first-time makers needed, and its know-how was tied to its own dedicated machines. Denso engineers rewrote the process instructions for machining centers, and Aisin spread their fixes. The first prototype came from a 320-person supplier whose executives trained at Aisin. Sixty-two firms made the part without any agreement on payment; Aisin later reimbursed their direct costs, and Toyota paid each first-tier supplier 1 percent of its quarter's sales to Toyota. Toyota's plants were back to normal output by February 10. [28]

Toyota's suppliers stepped in because they shared methods, trusted that their costs would be covered and could ship parts Toyota accepted on inspection. Defense production often lacks those conditions. A new producer usually needs qualification, rights to the data and often security or export clearance before its parts count. A supplier that steps in during a shortage also risks losing the work and stranding its inventory when the incumbent returns. A defense buyer has to settle these questions in normal times. It can write process instructions that a capable shop could follow on general-purpose equipment, agree on an emergency acceptance route with the approving authority and secure the right to share the data. It can also settle in advance how emergency work will be paid and what happens when the incumbent recovers, and keep working ties with the shops most likely to take over.

A shortage also disrupts production for longer than the shortage lasts. While parts are missing, operators move to other work and test slots go to other programs, and when the parts return, production waits for those resources. A 2015 Defense Department Inspector General audit found that when the Defense Logistics Agency obtained C-130 wing-flap parts, the Air Force did not order them right away, because the repair personnel had been reassigned. [29]

INDUSTRIAL OPPORTUNITY

Build versus buy hides most of a manufacturer's options

Manufacturers that depend on a critical supplier usually weigh in-house production against continued dependence, and they overlook much of what falls between. The wider range of arrangements includes customer-owned tooling at the supplier's plant and a dedicated line funded in return for first claim on its output. It also includes a specialist's equipment installed in the manufacturer's own plant and run by its own people, and a process licensed to a second producer under the manufacturer's change authority. Each arrangement transfers a particular form of control and leaves the rest of the operation with the company best able to run it.

A company that owns a supplier controls its equipment, workforce, process and output, and a manufacturer seldom needs all four. A manufacturer that buys the supplier anyway takes on its debts and overhead, and the resource that limits deliveries stays the same. Specialist suppliers have the same range of options, from placing equipment in a customer's plant to operating equipment the customer owns.

A ramp also distorts the case for ownership. A merchant supplier spreads its equipment and specialists across many customers, while a captive operation keeps those costs after defense demand recedes. Building and qualifying it pulls engineers off live production during the ramp that prompted the purchase. The test is cost per accepted unit at sustainable volume.

Hadrian, a California company, builds and runs automated factories that machine precision parts for aerospace and defense customers. In October 2026, Lockheed Martin announced that Hadrian would install a 24-machine automated machining cell in Lockheed's munitions plant in Grand Prairie, Texas. Lockheed's workforce will run it starting in early 2027. Lockheed gains dedicated capacity on its own schedule without developing the automation. Hadrian extends its production methods beyond its own factories. The cell has yet to make parts.

When companies divide ownership and operation, responsibilities that a single owner would carry need new homes. An operator that doesn't own the equipment has little reason to upgrade it, and an owner that doesn't run it learns about wear late. Durable arrangements settle maintenance, responsibility for defects, and who keeps the equipment, data and trained staff when the arrangement ends. [30]

INDUSTRIAL OBSERVATION

Common data formats lower the cost of every production transfer

When production moves between companies, someone has to translate the design, process instructions and inspection results from the sender's format into the receiver's. Circuit board fabricators have long re-entered design changes received by email and phone. The receiving shop bears most of that cost in time and errors, so the buyer seldom sees it. A neutral format that every company reads removes that translation from every future transfer. New producers become cheaper to bring in, and buyers can spread work across more of them.

The companies that control formats often have little reason to open them. A design owner or software vendor whose format others must use holds an advantage over every producer downstream. Buyers shift that balance when they require technical data in neutral formats with each purchase. Consortia of buyers and producers fund standards that no single company will. [31]

CHAPTER VI

Industrial Intelligence and New Production Opportunities

INDUSTRIAL OBSERVATION

Contract records hide the dependencies that govern production

A buyer or manufacturer can find production options in records it already holds by reorganizing them around processes, material sources and physical parts. Purchases sorted by process reveal demand large enough to support a shared investment. Suppliers traced to their own sources may prove to share one foundry and fail together. Parts matched by specification expose duplicate part numbers that split orders into smaller lots.

The standard categories miss these relationships because they serve other purposes. Programs, contracts, part numbers and supplier names organize payment, accountability and compliance, and the people who keep those records answer those questions accurately. No one's job covers the relationships that cut across them.

The federal contract database lists the United States as the country of origin on 114 of the 115 active Defense Department contracts for the F-35 from fiscal 2020 through 2024. It records the prime contracts and none of Lockheed Martin's subcontracts. Through Lockheed, the department found that magnets on some F-35s came from China. In April 2025, officials estimated that a dedicated tracing effort had origin data for fewer than 10 percent of the component and raw-material suppliers behind the aircraft's parts. Data on those tiers has to come from companies with little reason to share it. When the Defense Logistics Agency asked 63 companies for supplier data, five returned complete responses. Buyers get that data more reliably when they write it into contracts than when they request it later. [32]

FORWARD-LOOKING JUDGMENT

AI lets engineers examine the alternatives they used to skip

AI lowers the cost of engineering investigation, so companies can examine production problems and alternative routes that were never worth an engineer's time. Engineers now ration their investigations. They take on the defects, substitutions and process

changes large enough to justify weeks of searching drawings, test histories and supplier records, and leave the rest unexamined. Those alternatives never appear in any plan, and no one counts what they would have yielded.

Language models that search technical histories and machine-learning systems that read process data cut the cost of each investigation, so smaller questions become worth asking. An engineer can check whether a defect that looks unique has a precedent in another product's test records. An engineer can also check whether another shop could make a part now tied to one qualified supplier by adapting its process.

In an MxD pilot with GE Aviation and Arconic, measurements of individual turbine-blade castings fed a machine-learning classifier that chose machining parameters for each part. Each casting got its own machining decision instead of a standard setting. MxD reports that all 120 high-risk castings in the trial were machined successfully.

AI makes an alternative cheaper to find. Proving it still costs as much as before. Companies gain the most where qualified routes are scarce and one proven alternative would change a program's supply. [33]

FORWARD-LOOKING JUDGMENT

Industrial AI will favor those who pool process data

AI in defense production will favor the organizations that pool process data across programs and companies. A model that predicts weld defects or recommends machining parameters improves with the number of examples it learns from. Defense plants make few units of each part, and each program keeps its own records, so the data behind any one part number is thin. The same weld procedure, alloy or heat-treat cycle often runs across many part numbers, programs and suppliers. Organized by process, those examples add up to enough for a useful model.

Programs own their records, and contracts limit how data moves between them. Export controls and rules on controlled information keep defense process data out of commercial models. Suppliers guard their process histories as competitive knowledge. Each holder has too few examples to justify building a model and lacks the rights to combine them with anyone else's.

The organizations positioned to combine the data gain an advantage that compounds. A manufacturer can pool data across its own programs. A prime can pool data across its suppliers when its contracts give it rights to their process data and the suppliers receive better models in return. A consortium or government laboratory can hold data that competitors won't give one another. Each new participant adds examples, the

models improve, and better models draw more participants. The first organization to build that pool for a given process is positioned to set the reference models for it.

FORWARD-LOOKING JUDGMENT

Procedures written for expensive analysis will cap AI's gains

In defense production, procedures written when analysis and expert attention were expensive will hold back much of what AI could save. Inspection plans sample a fraction of parts because checking every part cost too much. Qualification programs run large physical test matrices because predicting material behavior was unreliable. Engineering changes pass through serial reviews because no single reviewer could see the whole record. Customers and quality authorities wrote these methods into specifications and contracts. A manufacturer must follow them even after a cheaper method meets the same purpose.

The people who can change these procedures have little reason to. An approval authority that accepts a new method takes on the risk if it fails and gains none of the savings. The supplier that would capture the savings has no authority over the rule. The old method stays in place long after its cost rationale has gone, and the manufacturer pays for both the AI and the procedure it was meant to replace.

A manufacturer releases those gains by bringing the approving authority a revised control with evidence that it meets the original purpose. One example is AI-assisted inspection of every part, validated against parts already inspected by sampling. Authorities that write new requirements around the risk to be controlled, and leave the method open, let later methods qualify without another rule change. Primes and program offices that fund the validation work give the approving organization grounds to accept the change.

INDUSTRIAL OPPORTUNITY

The most valuable production arrangements have no natural owner

Many of the most valuable production arrangements in the defense industrial base go unbuilt because no company is positioned to propose them. Programs and customers each hold part of the demand that would justify an investment. One firm holds the method, another the idle equipment and skilled people, and a third the approvals. Each sees its own piece accurately and would earn only part of the value, so none has reason to assemble the whole.

Assembling it takes work. Someone has to find the combination, test whether it holds up technically and set terms under which each participant does better than its alternative. That work costs time and expertise before any agreement exists, and whoever does it captures nothing unless it takes a stake in the result. Supply-chain illumination tools and supplier scorecards map the firms and routes already in use, and an arrangement nobody has built appears on none of them. Origination is undersupplied across the industrial base, even where the ingredients are already in place.

The natural originators are organizations that can capture part of what they create. A prime can assemble arrangements among its suppliers and recover the value in deliveries and cost. An investor can take equity in the venture it builds. A specialist manufacturer can design an arrangement around its own process and become its operator. A public institution can originate where the benefit spreads too widely for any one firm to claim it.

A larger industrial base will require new factories, equipment, workers and investment. America's capacity to produce also depends on recognizing relationships that existing arrangements conceal, developing manufacturing capabilities that don't yet exist, and putting complementary resources to work together. Some of the most consequential opportunities for expanding production have yet to become anyone's plan.

The American Industrial Capability Project is convened by Talbot West to strengthen America's industrial capability. It brings together original research, industry leaders and systemic understanding to discover and develop productive opportunities.

To discuss this work, contact Jacob Andra at jacob@talbotwest.com.

References

1 Department of War Office of Inspector General, *Evaluation of the DoW's Capability and Capacity to Produce 155-Millimeter Artillery Ammunition*, DOWIG-2026-095, July 2026. ↑

2 Defense One, *Aerojet Rocketdyne Struggling to Deliver Rocket Motors, Raytheon CEO Says*, December 7, 2022; Arkansas Economic Development Commission, *Aerojet Rocketdyne Breaks Ground on Large Solid Rocket Motor Development Facility as Expansion in Arkansas Continues*, April 25, 2019. ↑

3 Chad Syverson, *"What Determines Productivity?"*, Journal of Economic Literature 49, no. 2, June 2011. ↑

4 NIST MEP, *Success Lean Process Implementation*, 2024. ↑

5 Boston Consulting Group, *The Defense Industrial Base Is Not Maxed Out. It Is Misconfigured for Ramp*, April 2026. ↑

6 NAVAIR, *Teams Complete First F-16 Sustainment Repair Ahead of Schedule*, February 10, 2026. ↑

7 NAVAIR, *US Navy's First F16 - Fleet Readiness Center Southwest*, February 11, 2026. ↑

8 FAA, *Material Qualification and Equivalency for Polymer Matrix Composite Material Systems: Updated Procedure*, DOT/FAA/AR-03/19, September 2003. ↑

9 FAA, *Acceptance of Composite Specifications and Design Values Developed Using the NCAMP Process*, Memo No. AIR100-2010-120-003, September 20, 2010. ↑

10 America Makes, *Common Additive Manufacturing Qualification Template (CAM-QT)*. ↑

11 Wichita State University NIAR, *National Center for Advanced Materials Performance*. ↑

12 U.S. Department of Defense, "DoD Strengthens Supply Chain for Solid Rocket Motors", April 14, 2023. ↑

13 L3Harris, "Aerojet Rocketdyne Awarded $215.6M to Supplement Ongoing Modernization Efforts and Increase Solid Rocket Motor Manufacturing Capacity", April 14, 2023. ↑

14 Defense One, "Aerojet digging 'out of this hole' as it clears rocket backlog, president says", September 27, 2024. ↑

15 Christian Terwiesch, Z. Justin Ren, Teck H. Ho and Morris A. Cohen, "An Empirical Analysis of Forecast Sharing in the Semiconductor Equipment Supply Chain", Management Science 51, no. 2, February 2005. ↑

16 John Paul MacDuffie and Susan Helper, "Creating Lean Suppliers: Diffusing Lean Production through the Supply Chain", California Management Review 39, no. 4, Summer 1997. ↑

17 Business Executives for National Security, "Credit as Strategy: Mobilizing Private Capital for America's Defense Industrial Base", 2026. ↑

18 MP Materials, Form 8-K, July 10, 2025. ↑

19 Gracelin Baskaran and Meredith Schwartz, CSIS, "Stabilizing Cobalt Markets: A Price Floor for U.S. Minerals Security", December 10, 2025. ↑

20 MP Materials, "MP Materials Reports Second Quarter 2026 Results", August 6, 2026. ↑

21 GAO, *Defense Production Act: Use and Challenges from Fiscal Years 2018 to 2024*, GAO-25-108497, June 2025. ↑

22 GAO, *Warstopper Program: Comprehensive Cost Estimate Needed to Evaluate Potential Expanded Access to Emergency Medications*, GAO-24-106109, July 2024. ↑

23 Navy ManTech / NSAM, *Ingalls Shipbuilding Develops Improved Fabrication Processes to Increase Quality and Efficiency*, Project S2781, Work Cell Development. ↑

24 James Somers, The New Yorker, "The Engineers Taking on the Ventilator Shortage", May 18, 2020. ↑

25 CNBC, "GM to build 30,000 ventilators for US for $489.4 million", April 8, 2020. ↑

26 Detroit Free Press, "GM and Ford deliver the last ventilators amid coronavirus", September 1, 2020. ↑

27 GAO, *Weapon System Sustainment: DOD Can Improve Planning and Management of Data Rights*, GAO-25-107468, September 2025. ↑

28 Toshihiro Nishiguchi and Alexandre Beaudet, "The Toyota Group and the Aisin Fire", Sloan Management Review 40, no. 1, Fall 1998. ↑

29 DoD Inspector General, *Opportunities Exist to Improve the Inventory Management for Defense Logistics Agency Aviation C-130 Spare Parts*, DODIG-2015-132, June 11, 2015. ↑

30 Lockheed Martin, "Lockheed Martin Expands Advanced Munitions Manufacturing in Texas, Driving Thousands of New American Jobs", October 14, 2026. ↑

31 MxD, "Better Boards Make Better Birds for the Air Force", July 10, 2018. ↑

32 GAO, *Defense Industrial Base: Actions Needed to Address Risks Posed by Dependence on Foreign Suppliers*, GAO-25-107283, July 24, 2025. ↑

33 MxD, "Jet Engine Manufacturing Gets A Boost in MxD Pilot", December 3, 2018. ↑