SYM-AM-24-051

Excerpt from the Proceedings of the Twenty-First Annual Acquisition Research Symposium

Acquisition Research: Creating Synergy for Informed Change

May 8–9, 2024

Published: April 30, 2024

The research presented in this report was supported by the Acquisition Research Program at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School

Leading Practices
Iterative Cycles Enable
Rapid Delivery of Complex, Innovative Products

Brian Smith—is a Senior Analyst at U.S. Government Accountability Office (GAO) [SmithBT@gao.gov]

Abstract

The Department of Defense is increasingly acquiring complex products, such as combined networks of hardware and software, which require innovative processes to design, produce, and deliver. GAO has found that to consistently deliver new warfighting capabilities with speed, acquisition programs for cyber-physical systems, such as aircraft and uncrewed vehicles—must adopt new approaches to its acquisition structure. Solutions, though, are unlikely to originate exclusively within government. Rather, identifying the practices that leading companies rely on to deliver new cyber-physical products can provide crucial, cutting-edge information to acquisition leaders and, in turn, ultimately help DOD frame changes to its acquisition processes. This presentation will focus on GAO's recent work on innovative practices that can inform DOD's ongoing efforts to improve acquisition performance. The presentation will also draw on GAO's annual weapon systems assessments to discuss how DOD is beginning to implement some of these principles and practices.

Keywords: Innovation, iterative development, digital twins

Background

In March 2022, we found that leading companies consistently deliver innovative products to market with speed by relying on four key principles throughout product development. Implementing these four principles positions leading companies to satisfy their customers' needs and correspondingly grow their market share (see fig. 1). Appendix II further details these principles and their associated sub-principles.

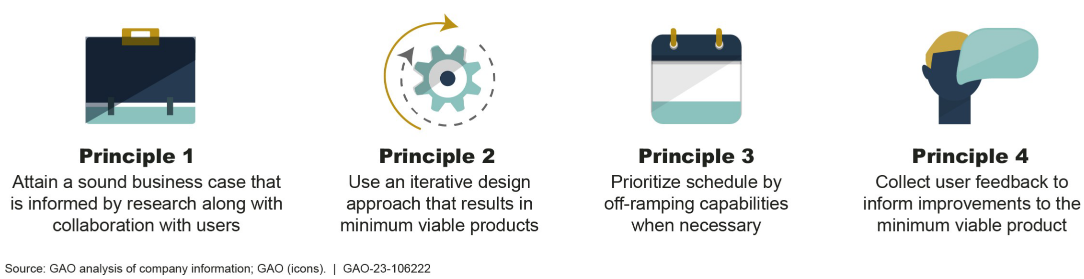

Figure 1. Leading Companies Rely on Four Principles to Deliver Innovative Products to Market with Speed

These four principles, along with several of their sub-principles, provide important context for understanding the analyses included in this report. We describe below how we continue to leverage these leading principles in our work.

Attain a sound business case. Sub-principles address how leading companies conduct market research and obtain and use customer feedback to establish and then continually maintain a sound business case throughout development. This report discusses how the initial business case—one that underpins the very start of a product development—can evolve over the course of the product development effort. Our future work, however, will discuss in more detail how leading companies establish key sub- principles underlying this business case, which include:

- investing time to research the marketable product;
- soliciting early feedback from customers;
- developing high-level cost and schedule parameters; and
- drawing from institutional memory and corporate knowledge to develop initial estimates, avoid earlier mistakes, and build off previous success.

Use an iterative design approach that results in minimum viable products (MVP). Leading companies use modern design tools, such as digital engineering and additive manufacturing, throughout development for both hardware and software. Key concepts within this sub-principle related to this report include:

- The use of iterative design and testing allows leading companies to identify an MVP—a product with the minimum capabilities needed for customers to recognize value and that can be followed by successive updates.
- Digital engineering includes digital twins—virtual representations of physical products. Digital twins incorporate dynamic data of a physical object or a system—meaning the model changes and updates in real-time as new information becomes available. Unlike a digital twin, a 3D model is a static visualization of a physical aspect—meaning it cannot be updated without someone manually inputting new data, and is similar to paper design drawings in digital form. Digital threads are a common source of information that connect stakeholders with real-time data across the product life cycle.
- 3D printing is a type of additive manufacturing, a computer-controlled process that creates physical objects, such as aircraft components, by depositing materials, usually in layers.

Prioritize schedule by off-ramping capabilities when necessary. To achieve speed to market, leading companies use periodic reviews to monitor performance and work to maintain a realistic assessment of development activities. Leading companies will off- ramp capabilities—an industry term for removing them from the planned release—if needed should those capabilities pose a risk to delivering the product on schedule. The off-ramped capabilities can be deferred to a later release or terminated.

Collect user feedback to inform improvements to the minimum viable product. Leading companies establish a process to facilitate ongoing engagement with users and customers after delivery of the first iteration. They use this feedback to identify new features to include in subsequent iterations or new products.

Iterative Development Approaches for Cyber-Physical Systems

The rise of cyber-physical systems in product development has led to new iterative development approaches within industry. These approaches integrate modern software practices with hardware development processes to achieve speed in innovation and capability delivery to users.

Figure 2. Examples of DOD Science and Technology Stakeholders in the Planning, Programming, Budgeting, and Execution (PPBE) Process

Differences between Linear and Iterative Development

Historically, both hardware and software product development progressed through a linear process with sequential milestones. Companies solidified requirements prior to development and delivered capability in a single completed program at the end of the development cycle. However, over the last several decades, software developers have utilized Agile practices, which provide iterations of capability that are continuously evaluated on functionality, quality, and customer satisfaction to increase innovation and speed in delivery. Now, as software increasingly dictates hardware functionality, companies are finding ways to incorporate the same iterative, Agile practices into products beyond software. Some of these methods include Modified Agile for Hardware Development Framework and hybrid models, such as a model that combines aspects of Agile and Stage-Gate®. Table 2 describes some of the differences between traditional, linear development and modern, iterative development.

Table 2: Comparison of Linear Development and Iterative Development

	Linear development	Iterative development
Requirements	Requirements are fully defined and fixed up front	Requirements evolve and are defined in concert with demonstrated achievement
Development	Development is focused on compliance with original requirements	Development is focused on users and mission effect
Performance	Performance is measured against an acquisition cost, schedule and performance baseline	Performance is measured through multiple value assessments—a determination of whether the outcomes are worth continued investment

Source: GAO analysis. | GAO-23-106222

Cyber-Physical Systems

Cyber-physical systems—sometimes called hybrid systems—are co-engineered networks of hardware and software that combine computation, communication, sensing, and actuation with physical systems. Within a cyber-physical system, software does not simply process data; it also interacts with the physical world. The software receives information about the environment through sensors, such as temperature, tire pressure, camera, or radar sensor data. The software then uses these data to instruct physical hardware, such as motors, pumps, or valves. The system's functionality is controlled by software algorithms.

Major government acquisitions at DOD, DHS, and NASA increasingly reflect this close interaction between digital and physical environments. For example, satellites, uncrewed vehicles, aircraft, planetary rovers, and cooperating robots in a manufacturing line are instances of cyber-physical systems. Table 3 defines common elements of cyber-physical systems.

Table 3: Common Elements of Cyber-Physical Systems

Cyber-physical system	Description
Physical layer	Real object in the physical world
Digital layer	Algorithms for managing real objects
Interface	Interaction between physical and digital layers—such as control mechanisms and detectors; and interaction between physical and digital layers with a person
Domains	Different application areas for which stakeholders have subject-matter expertise
Cyber-physical system	**Description**
Modularity	Allows common elements to be combined and reused while retaining security and reliability
Cybersecurity	Helps to guard against malicious attacks.

Source: GAO summary of information from the National Institute of Standards and Technology and Institute of Electrical and Electronics Engineers. | GAO-23-106222

Figure 2 depicts the integration of digital and physical inputs in cyber-physical systems.

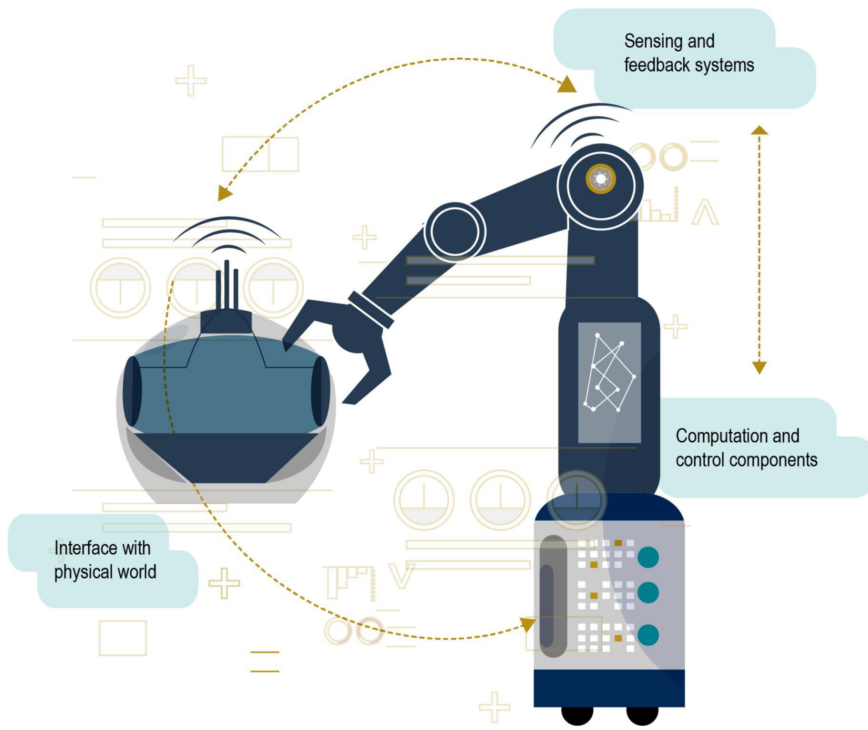

Source: GAO summary of information from the National Institute of Standards and Technology; GAO (illustration). | GAO-23-106222

Figure 2: Cyber-Physical Systems Integrate Continuous Physical and Digital Information

Leading companies develop cyber-physical systems as products for consumer use. As a result, we refer to cyber-physical products and product development throughout this report. Table 4 describes key concepts related to cyber-physical product development that are relevant to this report.

Table 4: Key Terms in Cyber-Physical Product Development

Key Term	Description
Backlog	The backlog is a list of features, user stories, and tasks to be addressed by the team, and is ordered from the highest ranked to the lowest ranked. If the team discovers new requirements or defects during development, these are added to the backlog. A backlog can occur at varying levels; for example, a product backlog is a high-level backlog that contains all the requirements for the entire program. An iteration backlog includes a list of user stories intended for that iteration. See description of user stories below.
Iteration	An iteration is a predefined, time-boxed, and recurring period of time in which product teams develop a working solution.
Stakeholder	A stakeholder is anyone who has an interest in the product. Specifically, stakeholders are parties that may be affected by a decision made by or about the product, or that could influence the implementation of the product team's decisions. A group or individual with a relationship to a product change, a product need, or the solution can be considered a stakeholder.
Sprint	A sprint is a short, time-boxed iteration that is intended to provide distinct, consistent, and incremental progress of prioritized features.
User	Users are the operators of the product. The user is an integral part of development and has specific responsibilities depending on the Agile methods used. The user is often synonymous with the customer, but at times the customer and the user might differ. This definition is organizationally and contextually dependent. For consistency, GAO refers to users throughout the report unless otherwise noted.
User story	A user story defines a high-level requirement by using everyday or business language. User stories are not vehicles to capture complex system requirements on their own. Rather, full system requirements consist of a body of user stories. User stories are used in all levels of Agile planning and execution.
Velocity	Velocity measures the amount of work a team can deliver in each iteration. Commonly, this is measured as user story points accomplished per iteration. For example, if a team completed 100 story points—a unit of measure for expressing the size of a user story—during an iteration, the velocity for the team would be 100. Velocity is a team-specific abstract metric and is generally not compared across teams as a measure of relative productivity.
Verification and validation testing	Verification and validation testing is a set of independent procedures that are used together for checking whether the program meets the requirements and specifications, that is, that it fulfills its intended purpose. To simplify, GAO refers to these procedures as validation in the report.

Source: GAO analysis. | GAO-23-106222

Leading Companies Progress through Iterative Cycles of Design, Validation, and Production

Leading companies employ an iterative structure when developing complex, cyber-physical products. The iterative process involves continuous cycles, which, similar to Agile software development, revolve around companies rapidly developing and deploying products. Key practices are common throughout the iterative cycles. For example:

- Leading companies seek and obtain continuous user feedback—feedback from the actual operators of the product—throughout the iterative cycles.

- Leading companies capture this feedback to determine if the design is meeting user needs and reflects an MVP—a product with the minimum capabilities needed for customers to recognize value.
- Leading companies continually feed this product design information into a real- time digital thread—a common source of information connecting stakeholders with real-time data across the product life cycle to inform product decisions.

Other development activities—such as modeling, validating, and refining specifications—overlap between cycles as product teams design and test sub-components and integrated systems. Figure 3 depicts key elements of this structure.

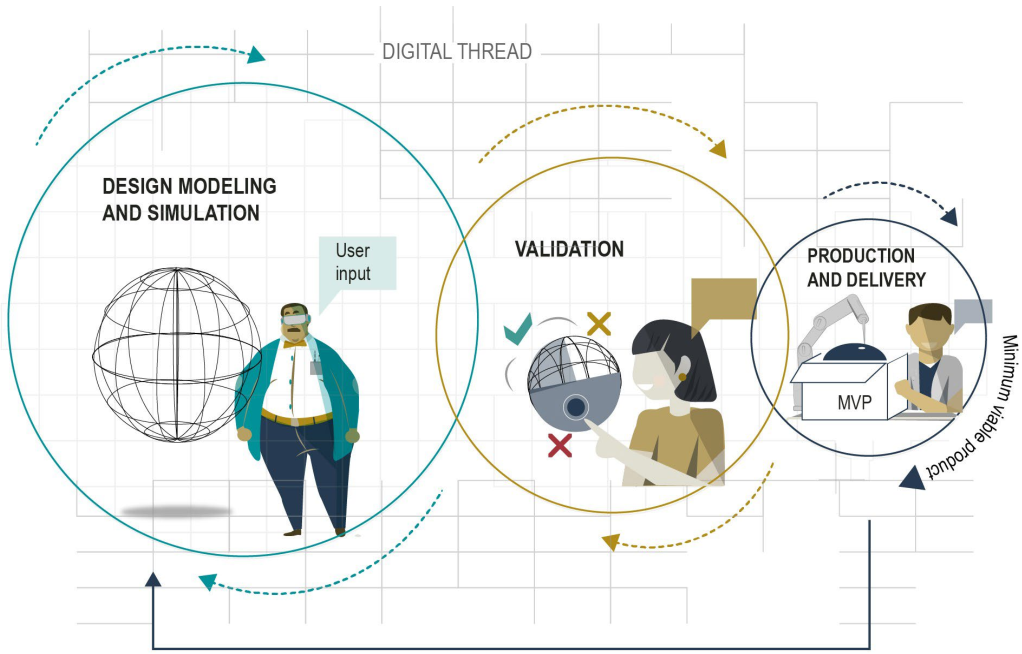

Source: GAO analysis of leading company information; GAO (illustration). | GAO-23-106222

Figure 3: Leading Companies Progress through Iterative Design, Validation, and Production Cycles to Develop a Minimum Viable Product

Knowledge Gained through Iterative Cycles

We found that leading companies increase knowledge about a system's design through each iterative cycle of design, validation, and manufacturing. Leading companies do not attempt to start development with a business case that includes a detailed specification of requirements. This approach differs from traditional linear development, which fixes operational requirements needed to deliver a capability to meet predetermined performance criteria. Instead, development begins with a high-level need statement or idea. Throughout development, this high-level need is progressively refined into distinct requirements.

Leading companies enable the initial business case to evolve over the course of product development. For example, Siemens ensures that the business case connects to research and development and technology management, so that research and development efforts focus on providing key technologies to be utilized in future new products. This means that research and development for a specific product does not end with the product—it continues so that future iterations of the product will have new, innovative, and mature technologies available.

The outcome of the business case development is the high-level framework of an MVP that the company will develop. This framework validates that the planned iteration of a product is responsive to a market need, underpinned by realistic expectations about technology development and achievable within defined cost and schedule parameters for that iteration. Knowledge acquired during design modeling and simulation and validation further refines the business case. Leading companies capture data from these iterative cycles in a digital thread. They then use information in this digital thread to inform decision-making, such as how to refine requirements or whether to make certain changes to the product's design. Table 5 outlines knowledge acquired during iterations in development.

Table 5: Product Development Cycles Characterized by User Feedback and Refined Knowledge Captured within a Digital Thread

	Product Development Cycles		
	Design modeling and simulation cycle	**Validation cycle**	**Production and delivery cycle**
User feedback	Users provide input to define design specifications for the minimum viable product, using multiple iterations as needed	Users agree design meets needs for minimum viable product, or design returns to modeling and simulation	Users provide feedback on desired product improvements to inform subsequent iterations
Knowledge captured in digital thread	Specifications that ensure the design meets most essential user needs	Integrated prototype that is tested in multiple environments to verify performance and can be manufactured as the minimum viable product	Optimized manufacturing tools and processes and insight into efficiencies for future iterations

Source: GAO analysis of company information. | GAO-23-106222

The number of iterative cycles that a product requires varies according to product type and team. For example, for products that are entirely new to develop, NVIDIA anticipates several phases of iteration across the design modeling and simulation, and validation cycles. The product team uses these multiple iterations to ensure all hardware, software, and infrastructure needs are validated through testing and user feedback. When NVIDIA develops improvements and updates to existing products, the product team starts with the existing design and makes updates that continuously optimize the product. Personal computers, for example, are largely in this category. The technology is mostly known, so NVIDIA can leverage this more advanced state and optimize existing designs.

Key Metrics for Delivering Minimum Viable Products

Leading companies structure product development around MVPs to ensure that they deliver essential product capabilities to users with speed. Under the iterative construct, schedule is a key driver, and companies make adjustments on performance, as needed. Accordingly, key metrics and measures track speed to market—generally the time measured from establishment of an initial business case to delivery of the MVP to users. For example, Danfoss measures time to market in its product development model, and seeks to reduce that time through iterative development. The metric begins at project start—which occurs after the initial business case is developed—and ends when the company delivers the product. Companies deliver new products on a schedule needed to meet customer needs and satisfy market demand. We previously reported that this speed to market calculation is relative to different product types and industries.

Leading companies often use metrics for cyber-physical products that reflect those within Agile software development, including velocity, sprints, and addressing user story points. For example:

- NEC uses velocity of development teams to identify the speed of each sprint, and then measures how many sprints are required to build and deliver the MVP. This allows the product team to better estimate the required schedule to build the product and communicate progress to the customer in a more transparent way.
- SAP is developing a metric that measures the time it takes to address customer feedback. The measure begins when the product team receives feedback and ends after the team places the feedback in the backlog and ranks, addresses, and delivers the product.
- For new physical products, Danfoss also measures progress of short, time- boxed sprints, which might be 2 to 3 weeks long, with a cadence that it can readjust depending upon customer need and type of program. This allows Danfoss to focus more on the project's progress and value added based on feedback, rather than simply checking whether it completed tasks and deliverables. This approach has shortened development cycles. For example, Danfoss representatives said that the company shortened its average time to market from more than 35 months in 2017 to less than 20 months in 2021.

Other key metrics used by companies revolve around establishing and verifying key performance specifications that define the MVP during design modeling and simulation and validation, which we discuss later in this report.

Leading Companies Increase Product Development Investments as MVP Design Matures

As the MVP design matures with each iteration, leading companies commit to increasing levels of resource investment for the product. They identify potential problems early through digital modeling and simulation and collaboration with stakeholders. As leading companies decrease risk, they proportionately increase funding.

Leading companies apply feedback about the design from cross-functional teams throughout iterative development—including design engineers, domain experts, cybersecurity teams, manufacturers and suppliers, marketing and sales teams, and customers and users.

For example, cybersecurity stakeholders include cybersecurity controls as specifications early in design and re-evaluate them as development progresses. According to the National Institute of Standards and Technology, cybersecurity is a necessary feature of the cyber-physical system's architecture to help ensure that capabilities are not compromised by malicious agents. Arista's Network Detection and Response (NDR) team builds cybersecurity into its products from the beginning of design through delivery and customer support. The team establishes security measures, such as firewall rules, to ensure there are no external actors affecting daily operations, and ensures its own devices are protected before writing the first line of code. To help ensure continued product security, Arista's NDR team also protects against vulnerabilities from outside sources, such as original equipment manufacturers or subscriptions to third-party code libraries. Arista's NDR team representatives said the codes in these libraries frequently have bugs and vulnerabilities that could be exploited, so Arista's NDR team builds in security features, and also puts a team in place to constantly look for risks from external sources.

Throughout all development cycles, stakeholders have access to real-time information through a digital thread. For example, Siemens' digital threads capture digital records of all states of the product throughout development, manufacturing, and service so that product teams and stakeholders can predict performance and optimize the product.

Users also rely on this information to identify areas where the product's design can provide the most value. The end result is that, rather than having a "relay" with handoffs of the product components to different stakeholders in succession, the digital thread enables parallel collaboration. We discuss the application of digital twins and digital threads later in this report within the context of specific development cycles.

The cross-functional structure also provides real-time knowledge that enables decision- making at the lowest appropriate level. For example, at Alphabet, Inc. (Google), the Product Manager acts as the "Chief Executive Officer" of the product and is responsible for defining the product, working with the technical team, and negotiating on product requirements that are achievable with each MVP.

Investment Decision- Making

Leading companies regularly evaluate the product's value with users to determine whether it continues to meet the initial business case and warrants continued investment. Leading companies provide funding commensurate with the product's design and development progress, rather than give a product development team a substantial amount of funding upfront at development start. For example, Danfoss initially provides a small portion of the product funding. It then scales product funding as the development team develops the design, tests the prototype to refine requirements, and ensures the business case remains valid for the MVP.

In addition, leading companies acknowledge that detailed estimates will change as development progresses, and correspondingly scale funding to ensure the investment provides value. This approach differs considerably from traditional linear development, which generally relies on fully resourcing a project to meet predefined performance requirements at development start. For instance, Volvo Group used to set full budget commitments early at fixed milestone gates. This made ending product development, if needed, "painful" and slow, even if the product was no longer relevant. Now, with the adoption of iterative, Agile practices, Volvo Group scales funding to keep pace with development. As the design progresses and is validated with stakeholders through integrated testing, the product team meets with senior leadership to determine whether the company wants to continue to invest in the product or specific technology.

Through collaboration with stakeholders and early discovery of design risks and vulnerabilities, leading companies are able to increase investment as they minimize changes to the design (see fig. 4).

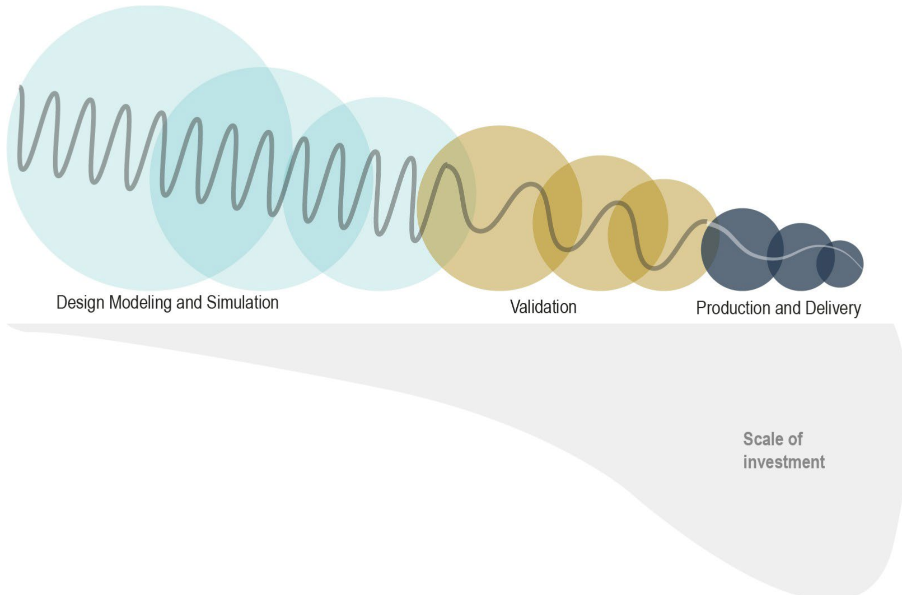

Source: GAO analysis of company information; GAO (illustration). | GAO-23-106222

Figure 4: Leading Companies Scale Investment to Increase as Frequency of Design Changes Decreases

Leading Companies Refine, Validate, and Produce a Minimum Viable Product by Employing Modern Tools and Engaging with Users

Leading Companies Model and Simulate Design Concepts with Users

Using digital models and user feedback, leading companies engage in a design modeling and simulation cycle to develop and refine the initial business case. We found that leading companies work with product users to co-develop requirements and indicators that can change. For example, using digital twins, Volvo Group can identify significant differences between the expected and actual performance of a truck's design and go back to the design team to iterate on the product design to meet the most important needs.

Through the design modeling and simulation cycle, leading companies repeatedly obtain feedback from selected users to inform design specifications. For example, during early design modeling and simulation, Arista's NDR team releases multiple, early iterations of the product to early adopters—the first users of that product—to solicit their input and feedback on product features. This user-centered design means that information gathered from users leads to building, testing, and redesigning through rapid iterations and innovation until the product specifications meet user needs.

Modeling and Simulation Input into Digital Threads

Leading companies leverage this collaboration with users to ensure the early design both provides performance and still has a valid business case. For example, when designing Pixel mobile phones, Google's product development team evaluates the right balance of features that optimizes performance at the target price. To meet the stated needs of North American customers, Volkswagen Group of America, Inc. (VW) made design changes to the interior of its ID.4 electric vehicle and to the exterior of its Atlas SUV. For ID.4 design changes, customer feedback from previous VW models, such as the Jetta, provided VW with the knowledge to change the interior of the ID.4 during design.

Figure 5: Digital Thread Captures Information throughout Product Life Cycle

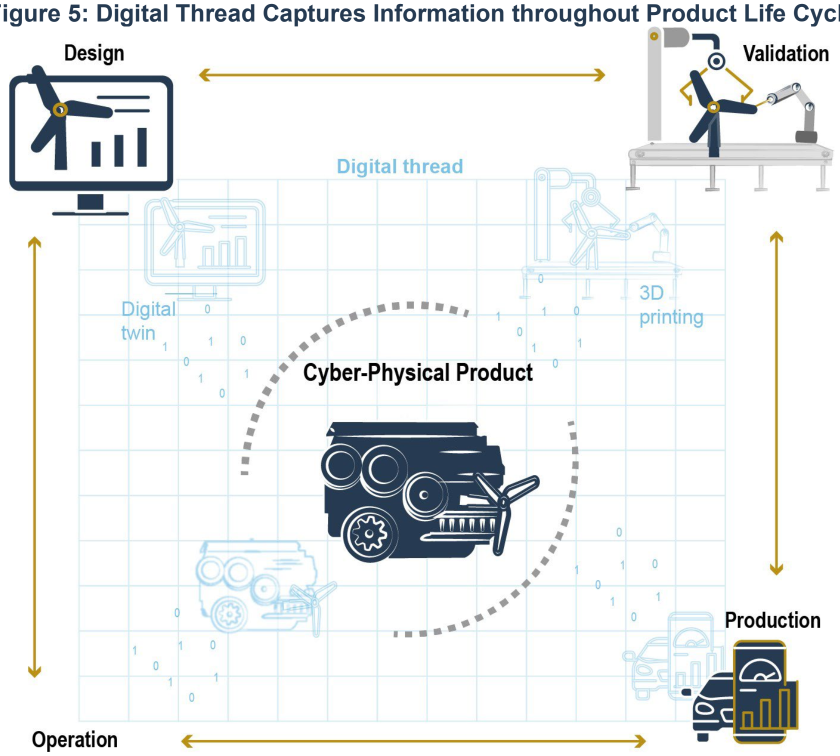

Source: GAO (analysis and illustration); bsd studio/stock.adobe.com (icons). | GAO-23-106222

Companies develop digital engineering models during design modeling and simulation based on specific needs. In particular, leading companies use digital twins and 3D printed models to quickly determine the most optimal design of a product that meets users' specifications. Digital twins, as previously noted, are virtual representations of their physical products, and incorporate dynamic data of a physical object or a system. 3D printing is a type of additive manufacturing, a computer-controlled process that creates physical objects by depositing materials, usually in layers. Creating a new design is easier in a digital environment because it enables faster design iterations, using digital twins and 3D printing. During design modeling and simulation, product development teams refine specifications with user feedback. Doing so can even result in starting over with new design solutions. Table 6 describes how the digital modeling and simulation inputs to the digital thread help inform decisions about the product's design.

Table 6: Design Modeling and Simulation Inputs to a Product's Digital Thread

Model	Description
Digital twin	• During design modeling and simulation, the product development team collects data to build the digital twin, either by connecting the digital twin and the physical world through sensors or by collecting domain data to build the twin digitally. • The digital twin simulates the behavior of different designs and feeds those data into the digital thread. • Stakeholders and users access this information to further define requirements and identify preferred design options.
3D printing	• During design modeling and simulation, the product development team uses 3D printing to prototype early designs of a product, which provides initial validation of the digital model.

Source: GAO analysis of company information. | GAO-23-106222

At Danfoss, digital twinning allows for faster design iterations. For example, Danfoss representatives told us that a product development team can test 500 digital designs using a digital twin in the same time that it could test five designs using traditional design approaches. Through the rapid digital design and test cycle, the product team is able to model and simulate many more possibilities than with physical prototypes alone. HP uses digital simulation early in design as the first step in coding and developing initial use cases that HP engineers can show to users. Microsoft found that digital twinning consistently results in more efficient design. It allows product development teams to design each component of a smartphone to the appropriate thickness and weight.

Through the use of AI, leading companies can create real-time synchronized simulations that are physically accurate and obey the laws of physics. These simulations can aid the implementation of system-level digital twins. At NVIDIA, AI may augment a digital twin, standing in as a good representation of the physical model, such as representing employees in a factory or representing a driver in a digital twin of an autonomous vehicle.

Leading companies' use of specific domain knowledge—particular expertise or skills relevant to the product—and user input into digital twins provide confidence that capabilities can be developed to meet schedule and cost parameters identified in the project's business case. For example, at HP, the most critical aspect of the digital twin is that it reflects the right domain knowledge to understand how the system works. This domain knowledge includes internal factors such as heat as well as the physics of the external environment, which will affect a product's performance. These data, together with rapid digital design and testing, predict expected performance of the product.

Leading companies develop trust in digital twins by inputting high-quality data that capture information about the relevant domains. At Siemens, this requires input from users and understanding the manufacturing capabilities and other domains needed to create the product. Digital twins take fundamentally good information from physical engineering to build a foundation. Then, data from people, processes, and tools feed into those models.

A digital twin becomes more robust and reliable through continuous testing and correlation to the physical model in a real-world environment. NVIDIA trains its engineers not to trust the simulator immediately. Over time, engineers build trust in the model through correlation with a real-world version of the model—each instance of correlation proves that the model is correct in the specific area. Because NVIDIA has run simulations and correlated to the physical model and environment, it can document, demonstrate, and quantify reliability, establishing greater confidence in the model.

One challenge, however, is knowing when a model—and a design—is good enough. At Siemens, knowing the digital twin is good enough revolves around the data. Obtaining the correct product data during early design—such as by ensuring the data used to create the digital twin are accurate and similar to the real-world model—is what makes the digital twin an actual digital asset, and ultimately, what reduces dependence on physical prototyping.

In addition to the use of digital engineering, leading companies also use 3D printing, along with augmented and virtual realities to aid in product design. Augmented reality overlays digital content onto representations of the real world using smartphones, tablets, or glasses. Virtual reality completely obscures the real world, immersing users in digital environments using head-mounted displays. 3D printing allows product development teams at Danfoss, for example, to build early prototypes during design modeling and simulation cycles to obtain early user feedback on design and make early changes to the design based on that feedback. 3D printing is unique in that it enables this early, quick prototyping, resulting in cost and schedule efficiencies. Product development teams also use augmented reality and virtual reality to virtually see a product in its space—for example, the interior of a virtual vehicle——before building a physical prototype, enabling the product team to visualize an integrated design.

Deferring Technologies and Prioritizing Capabilities

We found that leading companies only embark on product development once they assess and establish confidence that the underlying technologies in the product are sufficiently mature to meet user needs and support the product development schedule. Leading companies vigilantly monitor product technologies during design modeling and simulation and will not hesitate to defer any to future design iterations if they prove incompatible with schedule and cost parameters defined in the initial business case. For example, Google has engineering processes in place that balance the development of new product features while prioritizing meeting the target release dates for its Pixel phone launches. Volvo Group employs a common architecture design system that enables product teams to defer technologies from one product and insert that technology into a later product without disruption.

Further, the use of a backlog allows leading companies to organize, rank, and track capabilities for the product. This backlog includes both software and hardware functions. Ranked work is driven primarily by what the majority of users need. For example, Siemens employs risk-based analysis with users to transform input into prioritized development activities based on user needs within initial business case parameters. However, the backlog does not stand alone—it reflects a broader plan to achieve the overall goal of the product. At the start of development, Danfoss uses its backlog to help product development teams identify and rank the features and capabilities that are a part of that development cycle and map back to the overall product development plan.

Developing Design Specifications for Integration and Testing

We found that leading companies sufficiently develop design specifications to enable system integration and prototype testing. The outcome of design modeling and simulation cycles is a solution—in the form of an MVP—that companies can validate through testing. These design cycles give companies more confidence that they have made major changes by the time they are ready to validate the product. Danfoss, for example, starts design modeling and simulation with potential solutions. When the product development team is ready for validation, those solutions have become the product they intend to sell to their customers. By the time Google's Pixel device, for example, is ready for validation, design for that iteration is nearly complete.

Modular design supports prioritizing capabilities for optimal design. For example, Volvo Group's use of modular design allows it to develop different vehicle ranges from a single architecture. This approach enables customized solutions to a single vehicle to meet different user needs. The modular design means that Volvo Group can integrate different hardware components into a new design iteration and still easily produce vehicles at scale.

Leading Companies Validate Product Design with Users

Following design modeling and simulation, leading companies build fully-integrated prototypes—incorporating data from both physical models and digital twins—to test with users in the expected operating environment. As a part of this process, leading companies revisit the business case, assessing whether the MVP remains within cost and schedule parameters and still meets user needs. Leading companies use the results of these tests and user feedback to update the product design, as needed, and prepare the MVP for production.

Testing Fully- Integrated Prototypes

Leading companies build system-level integrated prototypes—either physical, digital, or a combination—to test the MVP's design established during design modeling and simulation. This prototyping incorporates all hardware and software components to test the product's integrated functionality. As a result, testing of the fully-integrated system can uncover problems that were not apparent when subsystems were tested, both physically and digitally, earlier. Prototyping may also be used to test more than one design variation of a product to determine which best meets user needs.

While system-level integrated testing is a long-standing practice, leading companies now combine digital with physical prototypes to test the complete cyber-physical product with users in the operating environment. Digital twins inform the physical prototypes—which are built from digital designs—and also incorporate testing results from the physical prototypes to better simulate the product's functionality. For example, HP creates 3D-printed parts from digital designs to test and ensure structural integrity. Similarly, Danfoss' 3D printing lab prints physical parts from digital designs to observe how they fit together. Danfoss also provides the physical prototypes to its customers so customers can test the prototypes in their own products and ensure they will work together.

Data from the physical prototypes then feed back into the digital twin to continue testing and validating the product's design. HP, for example, incorporates physical data into the digital twin to replicate how the product will behave in different operating environments.

Similarly, NVIDIA captures real-world data from sensors placed on test vehicles, then uses the data to reconstruct the operating environment in digital twins and run simulations for autonomous vehicles. As data are incorporated into the digital twin, they are also incorporated into the product's digital thread and used to validate the design's performance as an MVP (see table 7).

Table 7: Validation Inputs to a Product's Digital Thread

Model	Description
Digital twin	• During validation, the product team conducts systems-integrated tests—such as thermal or drop tests—on a physical prototype connected to the digital twin, or through a fully digital model. • Test data inputs and design updates to the digital twin become part of the digital thread. • Validation data are available to outside stakeholders—those with an interest in the product—to collaborate on design strategies and decisions and determine the minimum viable product.
3D printing	• During validation, the product team uses 3D printing of certain parts or of integrated products to test their performance and collect physical data.

Source: GAO analysis of company information. | GAO-23-106222

By adding physical data inputs into digital twins, leading companies use modeling to simulate potential operating scenarios that have yet to be realized, leading to more robust testing. As a result, leading companies can run scenarios repeatedly with unlimited variations, building confidence that the products they designed will work once produced. For example, NVIDIA can apply data from car accident reports and insurance claims to a digital twin for an autonomous vehicle, and use modeling to create rare and difficult scenarios for a vehicle's operation.

In some cases, leading companies use digital twins to gain insight into a system's design that cannot be obtained physically. For example, in developing Earth-2—an AI supercomputer intended to predict climate change—NVIDIA used a digital twin to simulate the inside of a nuclear reactor, which is physically inaccessible. Danfoss used a digital twin of an industrial motor drive to simulate its overload to the point of explosion. Compared with a physical test, which would have destroyed the prototype, Danfoss could identify the specific point of explosion, locate defects, and fix them in the digital twin.

Assessing Prototype Performance

In design validation, leading companies focus more heavily on how prototypes perform against goals for quality. For example, Arista's NDR team seeks to balance product completeness—the extent to which all planned features are included in the release— with product quality. Similarly, SAP tracks metrics related to defects found once users begin to interact with the product.

Leading companies use prototyping results to help assess whether the product will remain within the cost and schedule parameters established in the business case, and whether the product will still meet user needs. Leading companies may make adjustments to cost and schedule parameters in rare instances, such as delaying product delivery when the company needs more time to develop a key feature that is critical for a majority of customers.

Refining MVP Capabilities

After confirming the maturity of underlying technologies within the MVP, and with schedule as a key driver, leading companies evaluate the most critical functions and off-ramp product capabilities that are not essential and could delay the current release. As they work through validation, leading companies collaborate with customers and users to ensure the capabilities they are testing and the related product requirements are still the right priorities. For example, NEC ensures that all customer "must haves"—the capabilities that customers definitely need for the MVP—are satisfied first, before it adds less-critical capabilities. By maintaining flexibility on product specifications into design validation, leading companies can adapt the MVP to meet cost, schedule, and performance parameters.

Leading companies make off-ramping decisions for a given MVP largely based on customer and user needs, with the knowledge that they can add some of the capabilities in subsequent iterative product deliveries. Because the iterative process provides such opportunities, leading companies delay capabilities that are not ready until the next release or decide not to provide them if they are no longer needed. For example:

- To meet schedule, Microsoft may de-scope a product and deliver a subset of the full set of planned capabilities in the current iteration, then deliver the remaining capabilities in the next.
- HP identifies and off-ramps the capabilities that it does not need to meet the core functionality of the product.
- Siemens uses digital twins to support off-ramping decisions by examining the multiple designs in the digital thread and delivering the one that provides only the specifications that users need immediately.

With the various design options captured in the digital thread, leading companies can use them as a basis for the design of the next iteration and facilitate quick delivery of the next MVP.

Updating Design to Ready MVP for Production

Leading companies incorporate user feedback and results from the integrated prototype testing—including decisions about the minimum set of capabilities—into the product's hardware and software design, modifying it as needed.

For cyber-physical products, hardware design is ready for production when the company and the customer agree that the MVP design has been sufficiently proven in different conditions and still meets user needs. The iterative process leading up to this point directly informs the decision, as leading companies have tested and adapted the design multiple times and incorporated feedback on the user experience. For example, after testing multiple versions with different designs for a keyboard in one of its laptops, Microsoft determined it had reached the final design for the iteration when the material adhered well to enclosures and looked "crisp." For HP, the design must be scalable— that is, verified that it will work at scale in the field—which includes the ability to configure automatically and work without intervention.

Leading companies are willing to accept an MVP that does not include 100 percent of the features they envisioned initially, provided the MVP still meets user needs. This approach helps to ensure the MVP can be delivered on time, and that the user will have critical capabilities in hand. It also sets leading companies up to improve upon products in the future. For example, NVIDIA determines when the design of the optical lens in a camera is "good enough" based on the extent to which simulated temperature changes degrade the image. HP considers whether the design has sufficient high-quality features to provide value. It aims to meet the vast majority—though not necessarily all—of the proposed requirements with the product, including basic requirements and the ability to improve in subsequent iterations.

Once leading companies are satisfied with the MVP design, they stop designing hardware for the given iteration and prepare parts for production, recognizing that they can add functionality through software updates later. For example:

- Microsoft completes the design of the MVP's hardware—such as the display of a touchscreen tablet and the wire mesh on top of it—first, and then tunes software algorithms to enable the device to adjust to its surrounding environment or work with a stylus pen.
- Google and HP intentionally design their consumer electronics to enable software updates once they are in users' hands. Google enables software updates across its products to ensure that products are able to receive improvements throughout their lifecycle.
- VW anticipates providing additional functions and features, such as improved functions for infotainment systems, to vehicles through software updates in the future.

Figure 6 provides a notional example of how leading companies provide these updates to MVPs once they are delivered. By the end of product validation, leading companies have tested multiple design iterations, addressed gaps found in testing with users, and validated the MVP design to ensure hardware is ready for production.

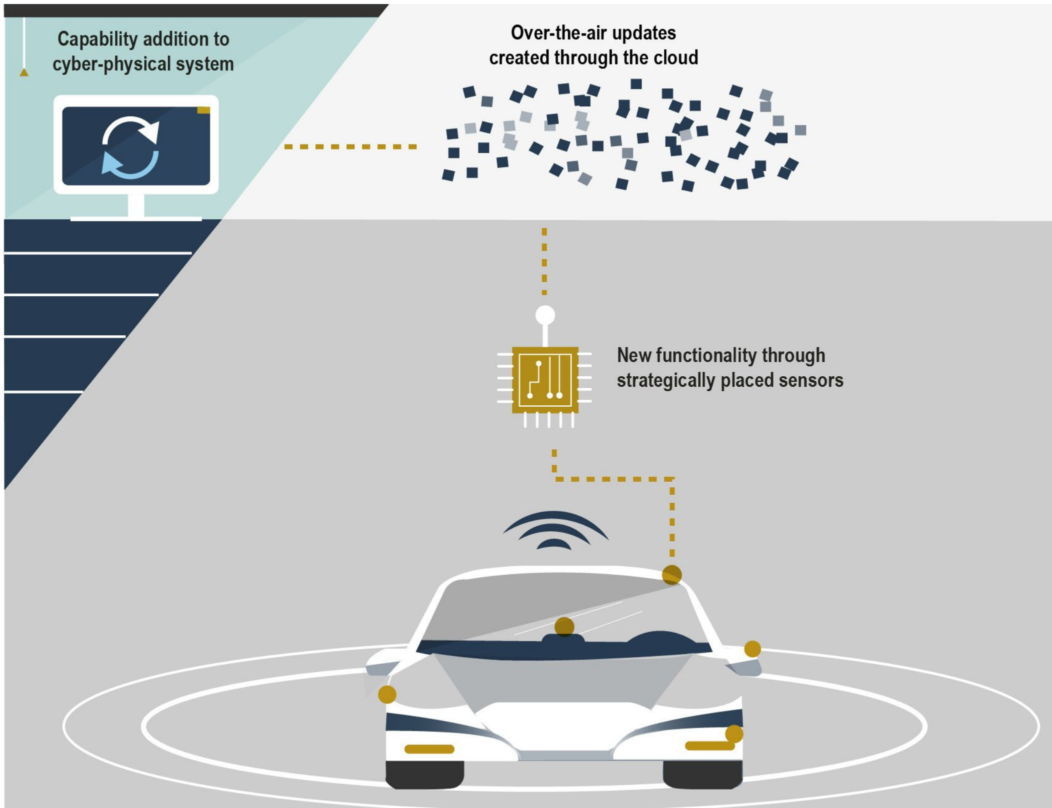

Source: HP and GAO (analysis); GAO (illustration). | GAO-23-106222

Figure 6: Leading Companies Add Functionality to the Delivered Minimum Viable Product through Over-the-Air Software Updates

Leading Companies Optimize Manufacturing of the Minimum Viable Product and Future Iterations

Once leading companies have validated the MVP design, they begin manufacturing products for delivery. The manufacturing planning process begins much earlier in product development, however. Leading companies start this planning while they are still designing the MVP itself. Through this early planning and the use of digital models, leading companies reduce the risk that manufacturing issues will delay product delivery. Leading companies gain further efficiencies and flexibilities through modularity in both design and manufacturing and collect customer and user feedback to continue improving products in subsequent iterations.

Planning for Manufacturing Using Digital Models

We found that leading companies begin manufacturing preparations early, while they are designing the cyber-physical product. As previously noted, leading companies' product design teams are comprised of those designing the product as well as stakeholders who will be producing it after testing and validation. Production stakeholders are involved throughout product design to ensure the manufacturing process can accommodate the design of the product. As a part of planning for manufacturing, product teams use digital twins to design efficiencies into the physical manufacturing complex and the production line that is housed there—which leading companies consider equally important to the design of the product itself.

Digital models optimize factory layout. Digital twins of factories allow for consideration both for workers and machinery before the factory is built. Equipment can be placed and tested digitally to simulate different production processes, changing a worker's position relative to a robot, for example, or the number of steps required to complete an operation. Leading companies have found that this drives greater cost and schedule efficiency. For example:

- NVIDIA is using its products to build a digital twin of a new electric vehicle factory years ahead of breaking ground.
- HP models its manufacturing processes using physics data to simulate an optimal mix of 3D printers and traditional manufacturing technologies. This provides data that HP can use to both confirm that a manufacturing process can successfully be completed and inform adjustments to a manufacturing process in response to irregularities that occur on the factory floor.
- Volvo Group uses digital twinning and virtual reality to test and optimize production flows.

Digital twins reduce risk in planning for production. Digital twins allow production teams to determine ranges of equipment stress and production capacity before production begins, including digitally testing robots to their maximum limits before using them. This knowledge reduces risk to the robots, because the manufacturing process can be adjusted to reflect those limits. Knowing this capacity also reduces the possibility of an expensive equipment replacement. For example, at VW, a robot that can lift up to a maximum amount of weight might exist on the manufacturing line, but a new part could be higher in weight. Process engineers consider these restrictions and possible alternatives during planning; they simulate the robots used in manufacturing to ensure safe and efficient production processes.

Leading companies utilize digital twins for manufacturing to reduce risks involved with advanced manufacturing processes required to produce complex designs. For example:

- Siemens builds electric components, but the company must first build a machine that makes the components. The product team has a digital twin of the machine on the factory floor that they can debug virtually for optimization of the real equipment to manufacture.
- Microsoft uses digital twins to simulate the injection molding production process of hardware components that have very tight variances to the appropriate thickness and weight.
- NVIDIA used a digital twin of the working environment to train robots to operate on the factory floor. It found that the robots complete such training more quickly in a digital twin than in real life.

Digital Monitoring of Production Progress

Once manufacturing begins, leading companies use digital twins to monitor production progress. A Kanban board—a tool developed for Agile project management to observe the flow of work and alleviate bottlenecks—enables teams to keep track of their work, which can be either physical or virtual. Activities are "parked" until the activities ahead of them are cleared, which helps ensure the production team executes key steps. Leading companies monitor the Kanban board and can make adjustments in real time, as needed. Danfoss uses Kanban for product maintenance and improvement, because it tracks process flow. Identifying bottlenecks in that flow supports materials management for production. SAP uses Kanban with smaller teams for high-frequency delivery development projects. Such visibility into operational performance also provides transparency for management and senior leadership, who can track production progress based on real-time data.

The digital twin of the factory accesses the signals of the physical plant and enables production teams to detect anomalies or differences between the virtual and actual factory in real time. For example, if there is divergence between the two factories, the digital twin can identify it and signal the production team, which can then determine whether potential issues, such as a cyber-attack, may lead to breakdowns in operations. Such real-time data analytics contribute to production efficiencies through automation, as well. For example, Siemens' factory design includes automated deviation management, which saves the quality team from manually reviewing paper documents.

By simulating real-time factory operations using a digital twin, leading companies are able to troubleshoot manufacturing challenges and measure output to monitor schedule performance. The result is not only a physical product, but a digital record of the process as well. Volvo Group, for example, records a digital copy of every unique heavy-duty electric truck it produces and places it in a digital "garage," where it stores the digital design so it can provide the building blocks for future digital twins.

3D Printing for Manufacturing

Leading companies use advanced manufacturing processes, such as 3D printing, to solve specialty production challenges by printing parts directly from digital designs. 3D printing is particularly useful for producing low-volume parts that would otherwise be impossible to manufacture because of the precision required, such as Danfoss' manufacturing of equipment joysticks that conform to the grip of a specific operator.

Since a critical element of designing cyber-physical systems is being able to scale the design for production, product teams must identify when a 3D printed part is appropriate for a specific product.

For example, Volvo Group uses 3D printing for low-volume production of spare parts for already-fielded vehicles.

Leading companies also apply 3D printing for hybrid manufacturing, in which product development teams create a 3D component of a part, such as a pump, that is customized and highly efficient, and make millions of that single component to contribute to a larger system.

Table 8 describes manufacturing inputs to the digital thread used to inform current and future manufacturing processes.

Table 8: Manufacturing Inputs to a Product's Digital Thread

Model	Description
Digital twin	• During production, in addition to the completed product, the company also has a data set that describes how the product was manufactured, contained in the digital thread. • The digital thread documents all the steps in the manufacturing process, from the design of the machinery and toolset to the processes for assuring the product meets the company's quality standards.
3D printing	• During production, 3D printers create specialized parts on a limited scale.

Source: GAO analysis of company information. | GAO-23-106222

Modular Manufacturing for Efficiency

Leading companies are transforming their production processes to become more flexible through modular manufacturing—producing individual sections that can be assembled into different finished products. Specifically, modularity relies on basic designs that can be added, removed, or replaced to build different products, effectively speeding up the production process while also providing flexibility to customize products.

To support modular manufacturing, leading companies establish common standards that build on top of each other, which allows them to rapidly replicate production and reuse components already proven to work. For example:

- SAP develops standard software and then customizes the product to specific customer needs.
- Volvo Group uses modular interfaces similar to a building block set, and manufactures modules that the company can readily integrate to respond to similar customer needs with a set of scalable solutions. The application of interchangeable modules with modular interfaces helps the product team provide users with a unique product while at the same time reducing parts in production (see fig. 7). As a result, Volvo Group can mix and match modules in multiple ways to meet unique customer needs.

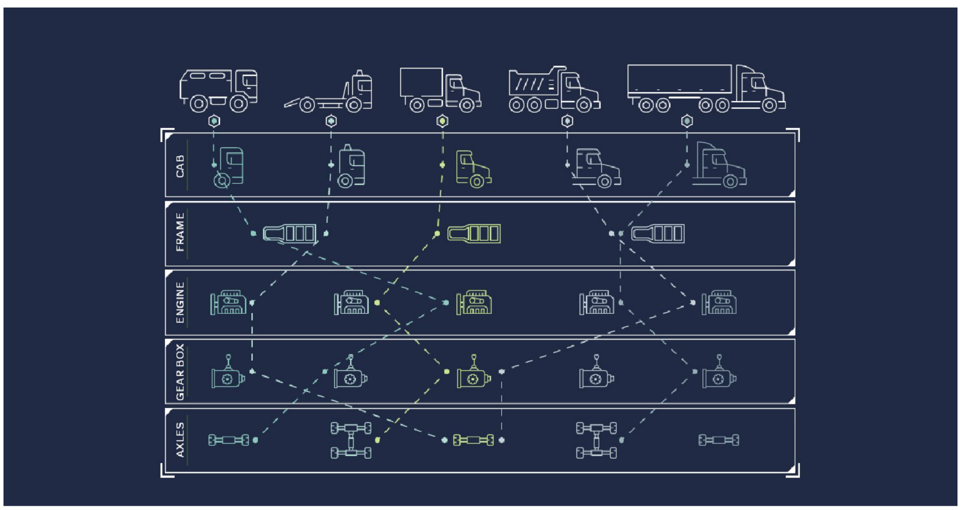

Source: Volvo Modular Interface Graphic (CAST). | GAO-23-106222

Figure 7: Volvo Group Develops Modular Interfaces to Manufacture Vehicles Efficiently

Collecting Feedback to Inform Next Products

After product delivery, leading companies collect user feedback to inform the next iteration of the product or the design of a new product. Leading companies obtain feedback from a variety of sources, including surveys, customer clinics, showcases, and social media. For example, Arista's NDR team asks users open-ended questions about their intended use of the product and its actual performance. At Cisco, product teams solicit feedback from users about the integration and performance of the MVP. For products sold through partners and the reseller community, Cisco collects user feedback indirectly through the seller about how well the application is working. Cisco uses the feedback collected through each of these means to inform improvements to subsequent iterations and products.

Soliciting user feedback about components within a larger system can require several steps. For example, since Danfoss makes the components inside an excavator rather than the excavator itself, customers may not always see the value in their products, so the company showcases how Danfoss products can work in an end product, such as an excavator. This allows Danfoss to talk to two distinct customer groups—the end user as well as the end-product manufacturer.

Real-time data collected through hardware sensors or automated software also provide statistically significant information on system performance, such as how long it takes for the system to perform a certain task. This type of information provides actionable data in conjunction with qualitative responses on user satisfaction.

- Google products are designed and manufactured so that Google knows when certain buttons are pressed and the actions users take. This information can identify less optimal elements of the user interface, popular features that should become more prominent, or whether functions can be streamlined—for example, if it takes multiple “clicks” to accomplish a task.
- Arista’s NDR team also monitors user data to get insight into how well products are working. The NDR team may see that it is taking longer than expected for a user to move through several pages or steps, suggesting that the product could be more intuitive. The team can determine trends, such as whether users seem to be experiencing the same problems, and match that up with feedback to better understand problems.

Ultimately, leading companies do not view delivery as the finish line in product development. Rather, delivery provides a springboard for establishing a new business case for the next iteration of the product. Leading companies will structure this business case around improvements to the already-delivered MVP. Some of these improvements could be software-only in nature. Others could necessitate changes to both the hardware and software of the existing product. Leading companies will make these determinations on the basis of user feedback provided on the existing product, coupled with technical information and new knowledge captured in that product’s digital thread. This knowledge positions leading companies to identify a new MVP for the iteration and quickly progress it through the same design modeling and simulation, validation, and production and delivery cycles described above.

Appendix III details how leading principles of product development underpin iterative cycles to refine knowledge about the product, information that remains critical to both companies and agency programs alike. Accordingly, we expect these iterative cycles and the practices that propel them will provide acquisition leaders in government with an increased understanding of cutting-edge product development practices, which these leaders can, in turn, use to frame changes to their agencies’ acquisition processes.

We are sending copies of this report to the appropriate congressional committees and other interested parties, including the Secretary of Defense, the Secretary of Homeland Security, and the NASA Administrator. In addition, the report is available at no charge on the GAO website at https://www.gao.gov.

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net