Supply Chain Research · Cornerstone Reference

What a Digital Twin Actually Is

What Is a Digital Twin in Supply Chain, and Do I Need One?

A digital twin is a live, synchronized virtual replica of a physical system, not a model you update by hand or a one-off simulation. The peer-reviewed test is the data flow: two-way and automated. Most systems marketed as twins do not meet it, and many operations do not need one.

CONTENTS

01 The short answer

A digital twin is a virtual representation of a real-world system, synchronized with it at a stated frequency and fidelity, so the twin reflects the current state of the physical thing and can be used to test changes against it. What separates a true twin from the many systems that borrow the name is the data flow: a peer-reviewed classification distinguishes a digital model, with no automated data exchange, from a digital shadow, with an automated one-way flow, from a digital twin, with an automated two-way flow. A static network model, a dashboard, and a one-off simulation are not twins by this test, however marketed, and a large share of what is sold as a supply chain twin is a model or a shadow.

2-way	2002	3
automated data flow is what defines a true twin	the year the digital twin concept was first introduced	distinct things sold under one name: model, shadow, twin

Key takeaways

- **A twin is synchronized, not static.** It reflects the current state of the physical system at a stated frequency and fidelity.
- **The data-flow test settles most disputes.** No automated flow is a model; one-way is a shadow; two-way is a twin.
- **A control tower is not a twin.** It is a decision application that may consume a twin, not the synchronized replica itself.
- **Most marketed twins are models or shadows.** Apply the test to a specific product rather than trusting the label.
- **Many operations do not need one.** Existing planning and visibility often solve the problem a twin is proposed for.

02 What is a digital twin, and who defines it?

The most widely adopted neutral definition comes from the Digital Twin Consortium, which describes a digital twin as a virtual representation of real-world entities and processes, synchronized at a specified frequency and fidelity. Its three defining elements are the virtual representation, the entities or processes it mirrors, and the synchronization between them. The synchronization is the load-bearing part: without an ongoing, defined connection to the physical system, a virtual representation is a model, not a twin.

The concept has a clear lineage worth stating, because marketing often implies it is newer and more settled than it is. It was introduced by Michael Grieves in a 2002 product-lifecycle-management presentation, and the term was later attached by NASA's John Vickers, appearing in a 2014 white paper by Grieves and Vickers. For manufacturing, ISO 23247 provides a reference framework built around an observable physical element and the digital entity representing it. There is no supply-chain-specific twin standard, so a supply chain twin borrows its definitional spine from manufacturing and general-purpose work, which is worth acknowledging rather than obscuring.

Two words in the definition do quiet work and deserve attention: frequency and fidelity. Synchronization at a specified frequency means the twin is updated on a defined cadence, which may be continuous for a fast-moving process or periodic for a slower one, and the right cadence is a design choice rather than a fixed property. Fidelity means how faithfully the model represents the physical system, and a higher-fidelity twin is more expensive to build and to run. A useful twin is scoped to the frequency and fidelity the decision actually needs, not the highest achievable, and a proposal that does not state its intended frequency and fidelity has not yet been specified enough to cost.

03 Model, shadow, or twin: how do I tell them apart?

The most useful neutral test comes from a peer-reviewed classification by Kritzinger and colleagues, which sorts these systems by how data moves between the physical thing and its virtual representation. A digital model has no automated data exchange: a person updates the virtual version by hand, and physical changes do not flow to it automatically. A digital shadow has an automated one-way flow, physical to digital, so the digital side reflects reality but does not act back on it. A digital twin has an automated two-way flow, so the two stay synchronized and a change explored in the twin can inform the physical system.

Figure 1. The model, shadow, and twin distinction after Kritzinger and colleagues. The test is the direction and automation of the data flow, which is what most vendor labeling obscures.

This test matters commercially because the three carry sharply different cost and capability. Building and maintaining automated two-way synchronization with a live supply chain is a substantial undertaking, while a dashboard fed by a one-way flow is far cheaper and frequently labeled a twin. When a product is called a digital twin, the question that cuts through the marketing is simple: is the data flow automated, and does it go both ways. If a person maintains the model, it is a model. If data flows only from the physical system to the screen, it is a shadow. Neither is worthless, but neither is what the word promises.

04 Is a control tower or a network model a digital twin?

No, and conflating them is one of the most common category errors here. A supply chain control tower is a decision application: it aggregates data, provides visibility, raises alerts, and supports or automates decisions. A digital twin is a synchronized virtual replica of the supply chain. A control tower may sit on top of a twin and use it, but the tower is the application layer and the twin is the representation beneath it. The two answer different questions, and a product can be a capable control tower without containing anything that meets the twin definition.

A network design or optimization model is a different thing again. It is typically a static representation of the supply chain used to test structural decisions, such as where to place a facility, and is usually built for a project rather than synchronized continuously with operations. That makes it a model in the classification above, not a twin, however sophisticated the optimization. The leading analyst covering this space draws its own boundary between the control tower, framed as an operational, shorter-horizon capability, and the digital supply chain twin, framed as a near-real-time end-to-end representation supporting decisions across horizons; that framing is useful, though it comes from an interested source that sells research, not a neutral standard.

05 Which supply chain twins are real, and which are claimed?

The honest picture is that operational, continuously synchronized supply chain twins remain relatively rare, and a good deal of what carries the label is a simulator or a dashboard. A recent academic review of digital twins in supply chain management found real deployments uncommon, and simulation models and dashboards frequently described as twins despite lacking the synchronized, two-way character the term implies. This is not a claim that the technology is empty, but a caution that the label runs well ahead of the deployed reality.

The vendor economics deserve a specific flag, because a much-quoted return figure is routinely misapplied. A widely circulated estimate that digital twins can raise revenue by up to ten percent, cut time to market by as much as fifty percent, and improve quality by up to a quarter comes from a consultancy's work on product development, not supply chain, and is presented as an estimate based on company experience rather than independent measurement. The same firm's separate supply chain figures are more modest. Any page that cites the product-development numbers as if they described supply chain outcomes, which many vendor pages do, is conflating two different things.

Where twins are actually deployed, the credible cases tend to be bounded: a synchronized model of a single distribution center, production line, or transportation network, used to test changes before making them. Fully synchronized, end-to-end twins of a whole global supply chain are mostly aspirational today. The fair case for the technology is that the bounded deployments are real and valuable; the unfair case is the implication that an end-to-end twin is an off-the-shelf purchase.

It helps to picture what a credible deployment looks like concretely. A distribution-center twin takes live feeds from the warehouse management system, the labor system, and the automation, maintains a synchronized model of the facility, and lets an operator test a slotting change or a labor reallocation against current state before committing it, with the outcome flowing back into the physical operation. That is bounded, synchronized, two-way, and valuable. The contrast with an end-to-end global twin is not sophistication but scope: the bounded case has a defined data perimeter that can actually be kept in sync, while the global case depends on synchronizing data from dozens of systems and partners that most organizations do not yet have connected.

06 Do I actually need one?

Often not, and the first question is what problem a twin would solve that your existing systems do not. Much of the value attributed to twins, better visibility, scenario testing, and faster response, is delivered in part by planning systems, network design tools, and visibility platforms many organizations already run. If the need is to test a structural change, a network model may suffice; if it is to see current state, a visibility capability may be the real requirement. A twin earns its considerable cost only where continuous, two-way synchronization actually changes what you can do.

There is also a readiness precondition that is easy to skip. A twin is only as good as the live data feeding it, so an organization without accurate, timely, connected operational data cannot build a useful twin regardless of the software: the synchronization the definition requires becomes synchronization to bad data. The honest sequence is to establish the data foundation and be specific about the decision the twin would improve, before treating a twin as the objective. For many operations the more grounded next step is stronger planning, visibility, or data infrastructure, with a bounded twin considered only for a specific, high-value decision that truly requires live two-way synchronization.

There is also a build-versus-buy dimension that the label obscures. A digital twin is not a shrink-wrapped product so much as an integration of a model with live data feeds, and much of the cost and risk sits in building and maintaining those feeds rather than in the modeling software. That maintenance is ongoing: a twin whose data connections drift out of sync quietly becomes a model again, reflecting a state the physical system has moved past. Whoever owns a twin owns the continuous work of keeping it synchronized, and a business case that funds the build without funding the upkeep is funding a demonstration rather than a capability.

07 Frequently asked questions

What is a supply chain digital twin, in plain terms?

A live virtual copy of some part of your supply chain that stays synchronized with the real thing, so you can see current state and test changes against it. The key word is synchronized: it is connected to the physical system by an ongoing, automated data flow, which is what separates it from a model or a dashboard.

What is the difference between a digital twin and a simulation?

A simulation is usually a one-off analysis run on a model at a point in time. A twin is continuously synchronized with the live system. A twin can run simulations, but a simulation on a static model is not a twin, because it lacks the ongoing two-way connection to reality the definition requires.

Is a control tower a digital twin?

No. A control tower is a decision application that provides visibility and supports action; a twin is the synchronized virtual replica. A control tower may use a twin, but they are different layers: the tower decides, the twin represents. A capable control tower can exist with nothing that meets the twin definition inside it.

Are most products marketed as digital twins actually twins?

By the peer-reviewed test, many are not. A good deal of what is labeled a twin is a model, updated by hand, or a shadow, fed by a one-way data flow. The way to check is to ask whether the data flow is automated and two-way. If it is not, the product is something other than what the term promises.

What data do I need before a twin is worth it?

Accurate, timely, connected operational data, because a twin synchronized to poor data simply reflects poor data faster. If the data foundation is not in place, that is the first investment, not the twin. Being specific about which decision the twin would improve is the other precondition.

How much does a digital twin cost to maintain?

More than the software, because the recurring cost sits in keeping the data feeds synchronized rather than in the modeling tool. A twin is an integration of a model with live data, and those connections require ongoing engineering to stay accurate as source systems change. A twin that is built and then left unmaintained drifts out of sync and reverts to being a static model, so the upkeep is not optional.

08 Method, sources, and where to go deeper

Method

- The definition follows the Digital Twin Consortium and, for the reference framework, the manufacturing standard ISO 23247. The model, shadow, and twin classification follows the peer-reviewed work of Kritzinger and colleagues. These are neutral or academic sources rather than vendor material.
- Where a figure or boundary comes from an analyst or a consultancy that sells research, it is flagged in the text as an interested source, and the frequently misapplied return figures are identified as such.
- Supply Chain Research is independent and vendor-neutral. We accept no payment from the vendors or categories covered, and this page names no products.

Caveats

- There is no supply-chain-specific standard for digital twins. The definitional spine is borrowed from manufacturing (ISO 23247) and general-purpose (Digital Twin Consortium) work, which is appropriate but worth stating rather than presenting as a supply chain standard that does not exist.
- The widely quoted returns of up to ten percent revenue, fifty percent faster time to market, and twenty-five percent better quality come from a consultancy's product-development analysis, not supply chain, and are estimates based on company experience. They are cited here only to correct their frequent misapplication, not as supply chain evidence.
- Digital twin market-size figures vary enormously across commercial publishers, with forecasts for the early 2030s differing several times over, so this page publishes no market figure. Figure 1 is a conceptual classification, not measured data.

Where to go deeper

Three SCR resources cover the systems a twin is most often confused with. The SCR control tower and visibility guides cover the decision and monitoring applications that may consume a twin but are not one, and drawing that boundary is the purpose of section 04. The SCR cornerstone on supply chain network design covers the static structural models that are frequently mislabeled as twins. Readers scoping the wider stack should start with the SCR supply chain software category map.

Sources

[1] Digital Twin Consortium. Digital Twin Consortium defines digital twin. Consensus body under the Object Management Group; the widely adopted neutral definition.

[2] Kritzinger and colleagues. Digital twin in manufacturing: a categorical literature review and classification. Peer-reviewed; the model, shadow, and twin data-flow classification.

[3] International Organization for Standardization. ISO 23247, digital twin framework for manufacturing. Standards body; manufacturing-scoped, as there is no supply-chain-specific twin standard.

[4] ScienceDirect, International Journal of Production Economics. Digital twins in supply chain management: scope and methodological issues. Peer-reviewed; finds genuine SCM twin deployments still uncommon.

[5] McKinsey and Company. What is digital-twin technology. Interested source: a consultancy. The oft-quoted return figures here are product-development, not supply chain.

[6] Association for Supply Chain Management. SCOR Digital Standard overview, used for process placement.