CUSTOMER WHITE PAPER

Beyond 3D Models: Why Manufacturers Need Decision-Centric Digital Twins

A practical framework for connecting operational state, evidence, authority and verified outcomes around critical manufacturing decisions.

DECISION FRICTION Operational evidence exists, but people repeatedly reconstruct context before they can act.	**TRUSTED STATE** The represented subject stays connected to identity, time-aware state, provenance, freshness and uncertainty.
GOVERNED ACTION Limits, approvals, overrides, fallback and accountability remain explicit around the decision.	**VERIFIED VALUE** The post-action outcome is retained so downtime, quality, throughput or recurrence can prove whether the intervention worked.

BUILT FOR

Manufacturing operations, maintenance and digital-transformation leaders evaluating whether connected operational context can reduce delay, inconsistency, downtime or repeat failure in a recurring decision.

EXECUTIVE SUMMARY

The missing capability is not another visualisation. It is decision-ready operational context.

Most manufacturers already have sensors, dashboards, workflows, 3D models and systems of record. The persistent gap appears when a live case requires several sources to be reconciled, limits to be applied, authority to be established and the result to be verified.

SWATANTRA AI THESIS
A digital twin should not be justified by visualisation or connectivity alone. It earns its place when it keeps a real operational subject connected to trustworthy state, decision context, explicit authority and measurable outcomes.

Four principles for a practical operational twin

Principle	What it means for the customer
Start with one decision	Name the recurring operating judgment, owner, options and outcome before selecting a platform.
Use the smallest capability	Choose a dashboard, workflow, simulation or integration improvement when it is sufficient. Add twin complexity only when connected state materially changes the decision.
Make state inspectable	Identity, source ownership, freshness, uncertainty, conflicts, interlocks and authority must remain visible.
Close the loop with outcomes	Record what was done and what happened afterwards. A restart, recommendation or alert is not proof of value.

VALUE TEST
A twin is justified only when it measurably improves a recurring operational decision or its outcome - not merely when it connects more data.

CHOOSE THE RIGHT CAPABILITY

Not every connected system is a digital twin - and that distinction matters.

The mistake is not using a dashboard, workflow, simulation or 3D model. The mistake is adding twin complexity where it does not improve a named operating decision.

Capability	Use it first when…	Normal limit	A twin becomes justified when…
3D model	Geometry, layout, condition or design must be visualised.	May not maintain time-aware operational state, authority or verified outcomes.	Changing evidence must remain connected to the subject and alter an operational action.
Dashboard	Users need reliable current values, trends or alerts.	Users may still reconcile history, conflicts, ownership and action manually.	Cross-source context changes the preferred decision.
Simulation	The objective is to explore hypothetical behaviour.	A scenario may remain detached from the live subject and current constraints.	The live subject must stay synchronised and the selected action must be verified.
Workflow	Work follows known steps and approvals.	The path is predefined and may not reflect changing state or evidence quality.	The preferred path changes with context, uncertainty or risk.
Integration layer	Systems need to exchange data consistently.	Connectivity alone does not define the subject, authority, decision or outcome.	Connected facts must remain as decision-ready state around a real subject.
Operational twin	State, context, authority and outcome must remain connected.	Introduces modelling, governance and lifecycle overhead.	Decision frequency or consequence justifies the additional complexity.

PRACTICAL DISTINCTION

A 3D model may be the interface. The digital twin is the governed operational capability behind it. Use the smallest capability that fully satisfies the decision requirement.

SEE THE OPERATIONAL PROBLEM

A line stop is rarely just an information problem. It is a context-and-authority problem.

Consider a representative packaging-line interruption. An alarm identifies the immediate symptom, but a restart decision also depends on the current order, recent maintenance, material availability, quality risk, operating limits and who has authority to proceed.

The same operational event, two decision paths

Fragmented response	Bounded operational twin
Historian + MES + CMMS + ERP are checked separately. Evidence is manually reconstructed and may differ by person.	Authoritative facts remain connected to the same line, order and event with visible provenance and freshness.
The alarm is treated as the event, while recent faults, maintenance, quality and material context may remain outside the immediate view.	Current state is interpreted with relevant history, constraints, uncertainty and permitted transitions.
Restart may depend on tacit expert knowledge and informal escalation.	Owner, approval, override, interlock and fallback are explicit before action.
The case closes when the line restarts.	Downtime, recurrence, quality and throughput are checked afterwards to verify whether the intervention worked.

What the current process costs

Decision delay People search several systems and consult experts while the exception remains active.	**Inconsistent action** Different people assemble different evidence, assumptions and interpretations.
Repeat failure Restart is treated as completion even when recurrence or quality later shows the intervention failed.	**Weak auditability** Evidence, rationale, override and post-action outcome are not retained around the operational subject.

BASELINE BEFORE DESIGN
Measure systems consulted, people involved, elapsed decision time, repeat-event rate and post-action outcome. Without an observed baseline, the business case is speculation.

FOLLOW ONE DECISION END TO END

The twin should make the restart judgment explicit, governable and testable.

The line-stop example becomes useful when the twin carries one case from detection through outcome verification. The goal is not to automate every action; it is to preserve the evidence and authority around the judgment.

1 DETECT	2 IDENTIFY	3 ASSEMBLE	4 CONSTRAIN	5 ASSESS	6 ACT	7 VERIFY
Alarm, speed loss or quality event opens a case.	Attach the event to the correct line, equipment, order and material.	Bring current state, history, maintenance, quality and plan context together.	Apply interlocks, operating limits, maintenance status and authority.	Compare restart, hold, inspect, change material, maintenance or escalation paths.	Recommend, approve, override or block inside explicit authority.	Measure downtime, recurrence, quality and throughput after action.

What should be retained for the case

Decision record element	Line-stop example
Subject	Line, equipment, active order, material and event identity.
State and evidence	Alarm, recent condition, maintenance history, quality state, plan, material and source freshness.
Constraints	Safety and quality interlocks, operating limits, approved procedures and authority.
Options considered	Restart, inspect, hold, maintenance intervention, material change, reschedule or escalation.
Action and authority	Recommendation, approver or operator, override reason and execution path.
Verified outcome	Time to recover, recurrence, quality impact, throughput effect and whether the intervention achieved the intended result.

DECISION TRACE
A useful twin does not only show what is happening. It preserves what was known, what was allowed, what was chosen, who had authority and what happened afterwards.

DEFINE THE TWIN AROUND A DECISION

The minimum useful twin connects six things around one named decision.

In this paper, an operational digital twin is a governed, connected and time-aware representation of a real subject used to improve an operating decision and retain evidence of both action and outcome.

1	2	3	4	5	6
OBSERVE Condition, request or exception.	**CONTEXT** Identity, history, plan, policy and ownership.	**STATE** Current, time-aware and uncertainty-aware.	**ASSESS** Options, rules, models, limits and confidence.	**AUTHORITY** Recommend, approve, automate or escalate.	**OUTCOME** Measure the result and retain evidence.

The proposal must answer six questions

Subject What exact asset, line, process, order, site or service is represented?	**State** What is true now, including history, uncertainty, staleness and conflicting evidence?
Connection Which sources own each fact, and how are updates validated and reconciled?	**Decision** Which recurring judgment improves, and what options are compared?
Authority Who may recommend, approve, override, execute, block or escalate?	**Outcome** What evidence proves whether the action worked?

MINIMUM VIABLE TWIN
When these questions cannot be answered, the solution may still be useful - but it is not yet a bounded operational twin for that decision.

HOW THE CAPABILITY FITS TOGETHER

A bounded twin sits between systems of record and governed operational action.

The twin should preserve identity, evidence and time-aware state while existing operational systems continue to own authoritative facts. It should not become an unnecessary replacement platform.

Decision-centric operational architecture

1 SOURCE SYSTEMS	2 CONTEXTUAL STATE	3 DECISION SUPPORT	4 AUTHORITY	5 ACTION + OUTCOME
SCADA / historian	Stable identity	Rules / limits	Owner	Workflow / control
MES / quality	Provenance	Models / simulation	Approver	Audit evidence
ERP / planning	Current + history	Options	Override rights	Downtime / quality
CMMS	Freshness	Confidence	Fallback	Recurrence
Policies	Permitted transitions	Prohibited actions	Escalation	Feedback

What this architecture does not require

- Replacing MES, ERP, SCADA, historian, quality or maintenance systems that already own authoritative facts.
- Making every input real time; freshness should match the operating decision.
- Using a 3D interface unless geometry or spatial interaction changes the decision.
- Representing the entire plant or enterprise when one bounded subject is sufficient.

BOUNDARY RULE
Connect twins through a digital thread only when the selected decision genuinely crosses asset, process, system or site boundaries. “Twin everything” is uncontrolled scope.

AI / REASONING BOUNDARY
A twin maintains decision-ready operational state. A separate reasoning or optimization component may compare alternatives when the use case requires it; the twin itself should not be used as a vague label for every AI capability.

MAKE THE CONNECTION TRUSTWORTHY

As confidence falls, authority must contract.

A twin that hides stale data, conflicting evidence or uncertain recommendations can create more risk than the fragmented process it replaces. Users must be able to inspect state, source, limits and authority without leaving the decision.

Reader question	What must be visible
What is represented?	Stable identity for the asset, line, process, order or service.
Which source owns the fact?	Authoritative source, provenance and conflict-resolution rule.
How reliable is the state?	Freshness, missing data, disagreement, uncertainty and confidence limits.
What may happen next?	Permitted options, interlocks, approvals, prohibited actions and escalation.
Who has authority?	Named owner, approver, operator, override rights and accountability.
When should it step back?	Advisory fallback, human verification, unknown state or blocked action.
How will the organisation learn?	Action, override reason, observed result and controlled rule or model change.

Safe degradation

Normal Required evidence is fresh, consistent and within approved operating limits.	**Advisory** Confidence has fallen or non-critical evidence is missing; the twin may recommend but not execute.
Unknown The state cannot be established reliably; the user must verify through an alternative procedure.	**Blocked** A critical dependency, interlock, policy or safety condition prevents the proposed action.

GOVERNANCE RULE
When evidence becomes stale, confidence falls or a dependency fails, continuing with unchanged authority is not resilience - it is uncontrolled risk.

MEASURE OPERATIONAL VALUE

Value must be proved against a baseline - not inferred from connectivity.

A credible business case connects each twin capability to an operational effect and to evidence that can be measured independently.

Connected capability	Operational effect	Evidence to measure
Stable identity	Evidence stays attached to the correct subject.	Matching errors, duplicate cases and reconciliation effort.
Current state + history	Users see change and recurrence, not isolated readings.	Diagnosis time, repeat-failure rate and intervention effectiveness.
Cross-source context	Operational facts are interpreted together.	Manual evidence-gathering time and decision cycle time.
Decision logic + limits	Options remain consistent with rules, risk and authority.	Avoidable exceptions, overrides, unsafe actions and escalations.
Outcome feedback	The organisation records whether the action worked.	Downtime, throughput, quality, condition and recurrence after action.

Illustrative measurement model - not a promised result

Suppose a site handles 12 recurring incidents per month. Four people spend an average of 45 minutes reconstructing evidence for each incident. A pilot target of reducing that effort by 50% would recover:

ILLUSTRATIVE CALCULATION
12 incidents × 4 people × 0.75 hours × 50% = 18 person-hours recovered per month. The customer must validate incident volume, effort, loaded cost and operational impact before using the result in an investment case.

Evidence discipline

- Use observed values rather than unverified industry percentages.
- Separate labour effort from avoided downtime, quality loss, safety exposure and working-capital effects.
- Define measurement ownership and stop criteria before the pilot begins.
- Stop or redesign when the evidence does not support continued investment.

EVIDENCE STATUS
This paper presents practitioner guidance and an illustrative scenario; it does not claim a customer outcome. A published case study should use verified, customer-approved baseline and result data.

IMPLEMENT A BOUNDED PILOT

Prove one operating decision before scaling the platform.

Begin with a scope small enough to govern, measure and change. A credible pilot must prove that users make a faster, safer or more consistent decision - not merely that technology can connect data.

1 SELECT	2 DEFINE	3 MAP	4 MODEL	5 OPERATE	6 MEASURE
One recurring, consequential and measurable decision.	Subject, owner, options, constraints and non-goals.	Sources, identity, freshness, conflicts and minimum paths.	History, uncertainty, authority, interlocks and fallback.	Run with users; record exceptions, overrides and adoption.	Compare with baseline; scale, simplify, redesign or stop.

Expected pilot outputs

Decision and boundary specification Subject, owner, trigger, options, authority, constraints, non-goals and measurable outcome.	**Evidence, identity and integration map** Systems, authoritative facts, identifiers, freshness thresholds, conflict rules and minimum paths.
Operational state and governance model History, uncertainty, permitted transitions, interlocks, fallback, overrides and audit evidence.	**Measurement and adoption report** Baseline, observations, adoption, exceptions, KPI change and a scale, simplify, redesign or stop decision.

Stop or redesign when...

- The subject cannot be matched reliably across required sources.
- Connected context does not materially change the decision.
- The outcome cannot be measured independently.
- Integration or governance cost overwhelms the plausible benefit.

PILOT RULE
The pilot must prove a better operating decision or outcome - not merely successful data connectivity.

SELECT A USE CASE THAT DESERVES A TWIN

A use case deserves a twin only when context changes a recurring decision and the result can be verified.

Use the readiness test before funding integration, modelling or visualisation work.

Readiness test	Strong signal	Warning sign
Frequency or consequence	The decision occurs often or carries material operational impact.	Rare, low-impact or adequately handled by an existing procedure.
Context fragmentation	People repeatedly rebuild evidence across systems or experts.	One reliable source already supplies the complete decision context.
Stable identity	The subject can be matched across required sources.	Records cannot be associated reliably with the same asset, order or process.
Action ownership	A named owner can approve, reject, override or execute.	Responsibility is distributed or disputed.
Measurable result	The effect of the action can be independently verified.	Success is subjective or measured only through user opinion.
Governable risk	Fallback, interlocks and escalation can be defined.	No safe response exists when evidence is missing or contradictory.

Practical questions

Do we need a 3D model? No. Use one only when geometry or spatial interaction affects the decision.	**Must we replace existing systems?** Normally no. Connect the minimum evidence from systems that already own the facts.
Must every input be real time? No. Freshness should match the decision; policies, plans and history may change less frequently.	**Who owns the twin?** Name business ownership, source ownership, decision authority, rule or model ownership and technical support separately.

GOOD STARTING DOMAINS

Repeated line stops or speed loss; fault recurrence or intervention priority; recurring defects or process drift; supplier or production-plan disruption; abnormal consumption or critical service degradation.

COMBINE CAPABILITIES WITHOUT CONFUSING THEM

Digital twins, reasoning engines, co-pilots and workflows solve different parts of the operating problem.

A strong architecture combines these capabilities when needed but keeps their jobs and authority boundaries clear. Treating every layer as “the twin” makes governance and value measurement harder.

Capability	Primary job	Core question	Typical output
Digital twin	Maintain changing operational state and evidence around a real subject.	What is happening, what changed and what context is trustworthy?	State, context, alert or intervention evidence.
Reasoning / optimization	Compare feasible alternatives using current state, policy and objectives.	Which feasible option is preferable, or should the system abstain?	Recommendation, ranked options or no-decision.
Operational co-pilot	Apply approved knowledge and present guidance in usable language.	What should the person know or do next?	Answer, explanation, next step or draft.
Workflow / control	Coordinate predictable steps, approvals and bounded execution.	What approved action happens next and who performs it?	Task, approval, command or recorded process transition.

A practical combination

1	2	3	4	5	6
DIGITAL TWIN Supplies current operational state, history, provenance and uncertainty.	**REASONING** Evaluates feasible options when trade-offs require judgment.	**CO-PILOT** Presents evidence and recommendation in usable language.	**AUTHORITY** Approves, overrides, delegates or blocks the action.	**WORKFLOW** Executes the approved process and captures the result.	**FEEDBACK** Returns the observed outcome to the case and measurement record.

CAPABILITY RULE

Use a digital twin for changing operational state, a reasoning layer for context-sensitive trade-offs, a co-pilot for approved knowledge and interaction, and a workflow for predictable execution. Combine them only where the decision requires it.

SWATANTRA AI POINT OF VIEW

Start with one recurring operational decision - and keep the evidence around it visible.

The strongest digital-twin projects are not defined by the largest 3D model or the broadest integration footprint. They are defined by a bounded subject, decision-ready state, explicit authority and verified outcomes.

Principle	Practical implication
1. Decision first	Define the recurring judgment, owner, options and measurable outcome before selecting technology.
2. Bounded subject	Model only the asset, line, order, process or service needed to support the selected decision.
3. Existing systems keep ownership	MES, ERP, historian, CMMS and quality systems remain authoritative where they already own the facts.
4. State includes uncertainty	Freshness, disagreement, missing evidence and unknown state must be represented rather than hidden.
5. Authority contracts with confidence	As evidence quality falls, execution should step back to advisory, verification or blocked states.
6. Outcome closes the loop	Capture what was actually done and what happened so value and intervention quality can be tested.

What customers should expect from a decision-first engagement

A credible first engagement should determine whether a twin is justified, what minimum boundary is governable and how value will be measured before major platform investment.

- Opportunity assessment - identify a recurring decision where fragmented evidence creates delay, inconsistency, risk or repeated failure.
- Decision and boundary design - define subject, owner, options, constraints, non-goals, authority and outcome.
- Evidence and identity design - map sources, identifiers, authoritative facts, freshness expectations and conflict rules.
- Governance and safe-degradation design - specify uncertainty, override, escalation, interlocks, fallback, audit and change control.
- Bounded pilot and adoption - connect minimum viable evidence, operate with users and compare the result with the agreed baseline.
- Scale, simplify, redesign or stop - expand only when measured value justifies it.

COMMERCIAL DISCIPLINE
A decision-first engagement should be willing to recommend a dashboard, workflow, simulation or integration improvement when a full operational twin is not justified.

MAKE THE INVESTMENT DECISION PRACTICAL

Invest only when the twin is the smallest capability that measurably improves the decision.

Beyond 3D means keeping a real subject connected to time-aware state, visible evidence, explicit limits, accountable authority and retained proof of the outcome.

Five questions to ask before you invest

Question
1. Which named decision will improve? - Identify its trigger, options, accountable owner and operational consequence.
2. What is the smallest useful subject and boundary? - Limit the first representation to the evidence and authority required for that decision.
3. Who owns the facts and the action? - Define source ownership, decision authority, override rights and escalation.
4. How will the twin degrade safely? - Specify what happens when evidence is stale, missing, uncertain or contradictory.
5. What outcome will prove value? - Agree the baseline, KPI, measurement method, adoption signal and stop criteria before implementation.

RECOMMENDED NEXT STEP
Bring one recurring operational decision, its owner, the systems currently used to support it and an observed baseline to a bounded opportunity review. The output should be a preliminary boundary, evidence gaps, candidate measures and a recommendation to proceed, simplify or stop.

Discuss a bounded digital-twin opportunity review: swatantra.ai

Selected standards and sources

- Digital Twin Consortium - Definition of a Digital Twin.
- NIST - Digital Twins for Advanced Manufacturing.
- ISO 23247 - Digital twin framework for manufacturing.
- NIST AMS 100-61 - Economics of Digital Twins.
- OpenText - The ever-evolving digital twin (2024). The scope taxonomy in the source paper is adapted and clarified from its digital-twin taxonomy.

Practitioner synthesis. Named standards do not certify a specific implementation or replace sector-specific legal, safety, quality or regulatory review. The illustrative scenario and calculation are not guaranteed commercial results; customer outcomes depend on data quality, operating discipline, governance, adoption and the selected use case.