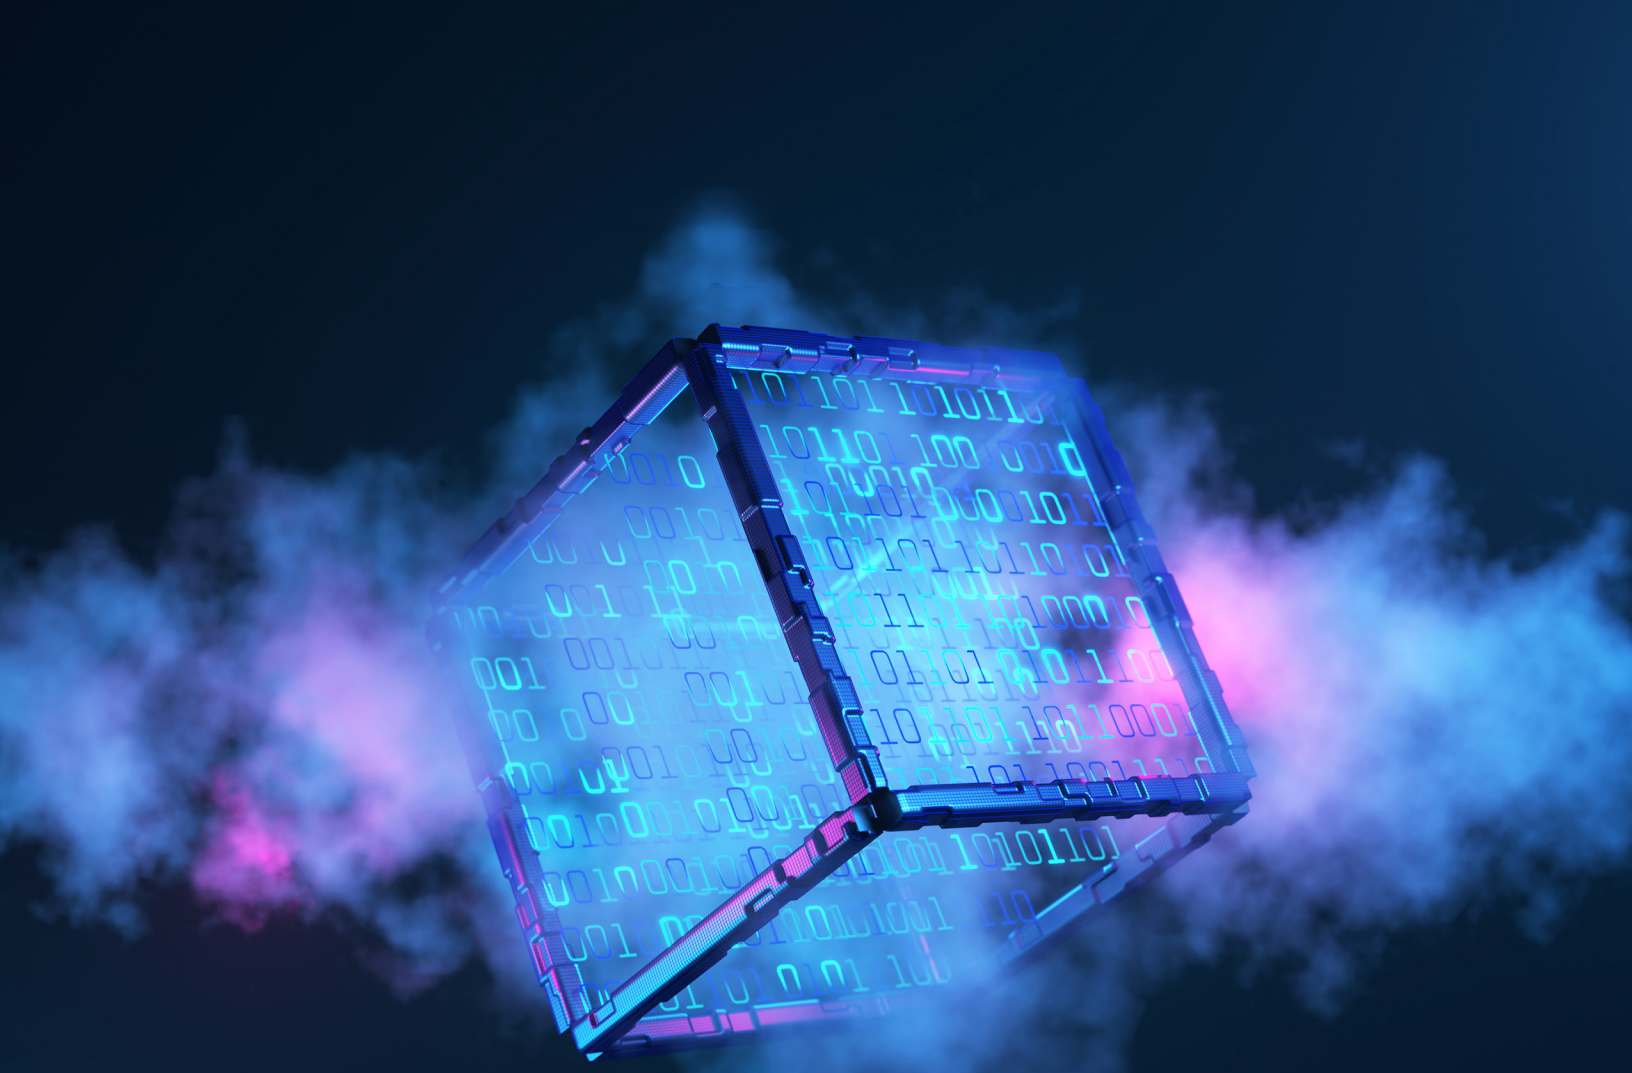


Making T+1 work

Building the operating models for execution.

Settlement failures are failures of information, timing and coordination. Fixing them means more than processing faster; firms need to redesign data quality, operating models and the way their own functions interoperate. With October 2027 approaching, the focus now shifts from understanding where the pressure lies to building the structures, automation and connectivity needed to make T+1 work.

As firms prepare for this transition, Cognizant, together with Microsoft, can help capital markets organisations design and implement resilient, connected and future-ready post-trade operating models, leveraging Microsoft Azure as a scalable cloud, data and AI foundation.

cognizant®

Microsoft



Contents

Introduction:
From diagnostic to design

First principles:
What a trade needs to settle

Zachman framework:
Why the Zachman lens fits EU T+1

From principles to design:
Applying the methods

- STP and exception management
- SSIs and reference data
- Collateral that moves: Turning visibility into settlement-ready liquidity
- Treasury at the settlement edge: Intraday liquidity control

Looking ahead

Introduction: From diagnostic to design

Our previous insight argued that the central risk of the EU T+1 transition is not the compression of the settlement cycle itself, but the cascading effect that follows when interconnected workstreams fail in sequence. A stale standing settlement instruction (SSI) produces a matching failure. A matching failure prevents partial settlement from functioning as designed.

A delayed collateral substitution arrives after the receiving central securities depository (CSD) has closed its window. A securities lending recall, initiated too late, becomes a settlement fail that triggers a margin call in a currency that has already missed its CLS cut-off.

To see how quickly the chain tightens, follow a single trade.



A **dollar-based fund sells a German equity on Monday afternoon (CET)**, for settlement on Tuesday (T+1) in euro.



The stock is out on loan through an agent lender, but the recall needed to cover the sale only goes out an hour after the 17:00 CET deadline because **no real-time link exists between the trading and loan books (operational failure).**



The borrower cannot return the shares in time for the **overnight T2S cycle**, so the first settlement attempt passes with the position still short, and by the Tuesday afternoon DvP cut-off the sale has failed outright, incurring a CSDR penalty.



Meanwhile the fund had already sold the expected euro proceeds into dollars for value Tuesday; with the securities leg failed, that FX trade has no euro to fund it.



Because no one was watching the night cycle, the gap surfaces only the next morning, after the 00:00 CET CLS initial pay-in deadline has passed, leaving the fund to source euro intraday at short notice.



One missed recall deadline has propagated into a securities-lending breach, a settlement fail and its CSDR penalty, an unfunded FX trade and an intraday funding call: four workstreams knocked over by a single instruction.

How One Missed Recall Deadline Cascades Across Four T+1 Workstreams

A single late securities lending recall request can trigger a chain of failures under the compressed T+1 settlement cycle.

Monday		Tuesday (T+1)			
Monday Afternoon	**Monday** > 17:00 CET	**Night** Overnight T2S Cycle	**Tuesday** Afternoon DvP Cut-off	**Tuesday** Value Tuesday	**Tuesday night / Wednesday morning** After 00:00 CET
Trade Executed	**Late Recall**	**Borrower Misses Share Return**	**Settlement Fail**	**Unfunded FX Trade**	**Intraday Funding Need**
Dollar-based fund sells German equity. Settlement on Tuesday (T+1) in EUR.	Stock is out on loan. Recall request sent AFTER the 17:00 CET deadline.	Borrower cannot return the shares in time for the overnight T2S cycle.	First settlement attempt passes. Position still short. Sale fails outright by DvP cut-off. CSDR penalty incurred.	EUR proceeds already sold into USD for value Tuesday. No EUR available to fund the FX trade.	No one monitored the night cycle. Issue discovered next morning after CLS initial pay-in deadline. Fund must source EUR intraday at short notice.

Four workstreams impacted

1 Securities Lending

Recall sent after deadline; borrower cannot return shares in time.

IMPACT: Recall/return breach, agent lender escalation, potential fees.

2 Settlement Operations

Settlement attempt fails by DvP cut-off.

IMPACT: Settlement fail, CSDR penalty, capital & reputational impact.

3 FX Operations

EUR proceeds already sold for USD; no EUR available to fund the FX trade.

IMPACT: Unfunded FX trade, potential market buyback at higher cost.

4 Liquidity Management

Night cycle not monitored; CLS pay-in deadline missed.

IMPACT: Emergency intraday EUR funding at short notice, higher funding cost.

One missed recall deadline has propagated into a securities-lending breach, a settlement fail and its CSDR penalty, an unfunded FX trade, and an intraday funding call: four workstreams knocked over by a single instruction.

These failures are not independent. The solution therefore must be designed to reflect how risk behaves: holistically, with each effect understood not in isolation but across the value stream.

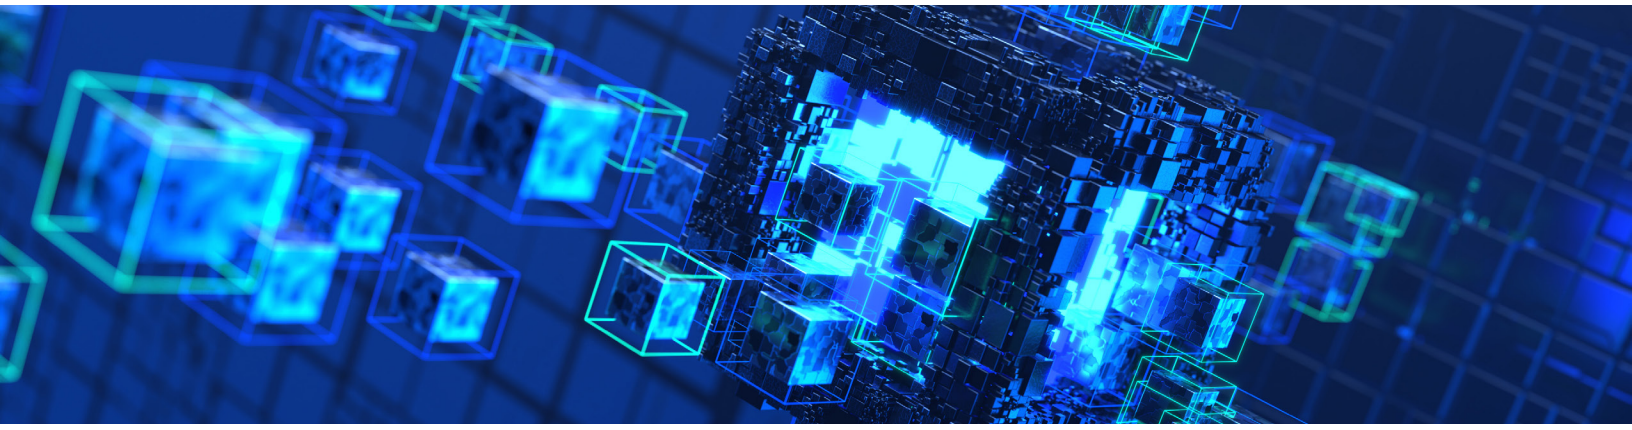
Before prescribing what firms should build, it is worth asking a more fundamental question: why do settlement fails happen at all?

> Reasoning from first principles, rather than from symptoms, means designing an approach that can hold under a one-day cycle.

First principles: What a trade needs to settle

The statement that firms "lose roughly 80% of their operational window" is a quantitative claim and should be supported with an appropriate source, citation, or methodology.

Simply upgrading legacy technology or patching existing workflows (incrementalism) will not work. First-principles thinking allows firms to deconstruct the post-trade lifecycle to its foundational truths and **determine the most effective firm-specific operating model for T+1**. It starts by stripping the problem down to its irreducible truths. Instead of beginning with current systems or organisational silos, firms can begin with the basic logic of T+1: **'A settlement obligation can only be completed if the right security or cash is available, eligible, transferable, funded and instructed before the relevant market cut-off.'**



References

1. AFME | T+1 Settlement in Europe: Potential Benefits and Challenges
   "Post trade activities compressed into shorter time frame: there would be significantly fewer hours between trading and the beginning of the settlement cycle for post-trade operational processes to take place. While it might be assumed that moving from two days to one day would reduce the available post-trade processing time by 50%, AFME actually estimates market participants will be moving from having 12 hours to 2 hours of post-trade operations time, an 83% reduction."
2. New guide sets out operational priorities for Europe's shift to T+1 | PostTrade 360°
   "Moving to T+1 will remove more than 80 percent of the current settlement processing window, according to the report, requiring firms to review workflows across the trade lifecycle."

First-principles thinking in a T+1 context

While a conventional impact assessment may start with questions like 'Which settlement systems need changing?' or 'Which operations teams need additional capacity?', firms must think holistically and apply a problem-solving approach such as first-principles thinking which will strip away all current procedural assumptions ("we do it this way because of our legacy platform"), deconstruct the problem into its most fundamental components or truths and rebuild the solution from those fundamental truths, bypassing existing assumptions.

Using an approach such as first-principles thinking, firms can start with more fundamental questions: What conditions must be true for every in-scope transaction to settle successfully on its intended settlement date? Those conditions are not confined to settlement operations; they extend upstream and across the firm.

For a typical market participant subject to the T+1 mandate, the firm should begin by defining the fundamental truths below:



Truth 1: Settlement is a timed, two-sided resource exchange

This immediately makes T+1 a **treasury, inventory, collateral and connectivity problem**, not just an operations problem.

Truth 2: Information must precede action



The firm cannot allocate, fund, recall, borrow, instruct or escalate a transaction until it has trusted information. So, **data latency sets the minimum possible process latency.** This brings trade capture, client static data, account hierarchy, SSIs, PSET, PSAF, instrument data, eligibility, currency, custody location and legal-entity data into the core solution.



Truth 3: The actual constraint is elapsed recoverable time

T+1 does not just shorten the nominal settlement cycle. It removes the time available for repair, the **recoverable time**, which can be calculated as market cut-off minus detection time, decision time and execution time. The firm needs to reduce detection and decision times. If recoverable time is negative, adding more people to exception management will not solve the underlying design problem.

Truth 4: An asset is valuable only if it is usable in time



The existing collateral formulation is directly transferable to the wider T+1 programme:

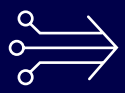
Settlement-supporting inventory = inventory that is visible, available, eligible, legally usable, transferable and mobilisable before cut-off.

An asset in the wrong entity, account, custodian, CSD, currency or collateral arrangement may be economically owned but operationally unusable.

Deconstruct and rebuild: End-to-end solutions

By breaking the lifecycle down to these constituent truths, the firm can reconstruct the high-level approach for the end-to-end solution covering the key features. For example:

Trade allocation and confirmation



- **Deconstruction:** The biggest source of delay is waiting for the buy-side to send allocations, followed by manual matching loops.
- **First-principles solution:** Move from batch processing to **real-time streaming matching**. Implement straight-through processing (STP) triggers so that allocations are automatically processed via institutional matching platforms as soon as the block is filled.

Securities financing (stock lending and repo)



- **Deconstruction:** Sourcing inventory for shorts must happen faster to avoid settlement fails.
- **First-principles solution:** Sourcing liquidity, recalls and collateral optimisation must shift to **early afternoon on T+0**. Leverage the newly proposed EU "Gating Event" (the 11:00 CET net/batch settlement functionality across CSDs) to optimise repo funding and maintain netting efficiencies without draining intraday liquidity.

However, given that T+1 is a fixed-deadline, heavily mandated programme, a pure first-principles approach could be time-consuming and risky for the October 2027 timeline.

So, firms should build their programmes on strong rule-based systems to ensure compliance, speed and interoperability, and deploy first principles with precision: defining the target state around the one irreducible constraint and redesigning the specific subprocesses where legacy assumptions break.

> Europe's T+1 mandate requires a context-driven approach, with problem-solving frameworks tailored to an institution's legacy maturity and operational friction points. First-principles thinking can be used independently, combined with approaches such as **design thinking** and **systems thinking** or replaced by frameworks better suited to the firms' immediate challenges. This helps ensure the strategy addresses firms' actual structural bottlenecks rather than imposing a blanket solution.

Zachman framework: Why the Zachman lens fits EU T+1

While firms should use first-principles thinking to determine what must fundamentally change for T+1, the next step is to use a framework that describes the resulting transformation completely and consistently across business, operating-model, information, technology and implementation perspectives.

A framework such as the **Zachman Framework** can help firms create a **complete architectural view** of the allocation, collateral and funding challenge, ensuring that business, operations, data, technology, control and stakeholder perspectives are not overlooked.

It is useful for T+1 because it forces firms to ask a complete set of operating-model questions: what data must be trusted, how processes must compress, where market dependencies sit, who owns exceptions and when cut-offs become binding.

The Zachman Framework prompts clients to consider the following classic enterprise architecture questions:

What
information does the operating model run on, and is that information clean and machine-readable?

How
is the process sequenced and where do manual handoffs sit in the critical path?

When
must each step complete, and against which hard cut-off?

Where
does **the asset or** the infrastructure sit and how many **boundaries must it cross**?

Who
owns the decision when something deviates?

Using a structured approach such as the **Zachman Framework**, firms can build a complete enterprise architecture view across data, process, location, roles, timing and decision policies. This ensures that treasury, operations, collateral management, repo, securities lending, clearing, custody, risk and technology are working from one integrated operating model.

However, the right way to use Zachman is as a **completeness lens**, not as the primary execution model. It is useful at the diagnostic stage because it helps firms identify blind spots across data, process, location, ownership, timing and purpose.

But once those are identified, firms need to move beyond Zachman into more practical, execution-oriented approaches that translate the diagnosis into design choices, delivery priorities, readiness metrics and accountable change.

Enterprise transformation for EU T+1 requires an architectural framework tailored to an institution's technology maturity and target operating model. Whether the situation demands the structural taxonomy of the **Zachman Framework** for step-by-step execution of standardised service roadmaps or an **agile architecture** approach for accelerated time-to-compliance, firms should adapt the framework to their operational reality. Firms should evaluate their current architectural state to select the model best suited to achieving regulatory readiness without unnecessary friction.

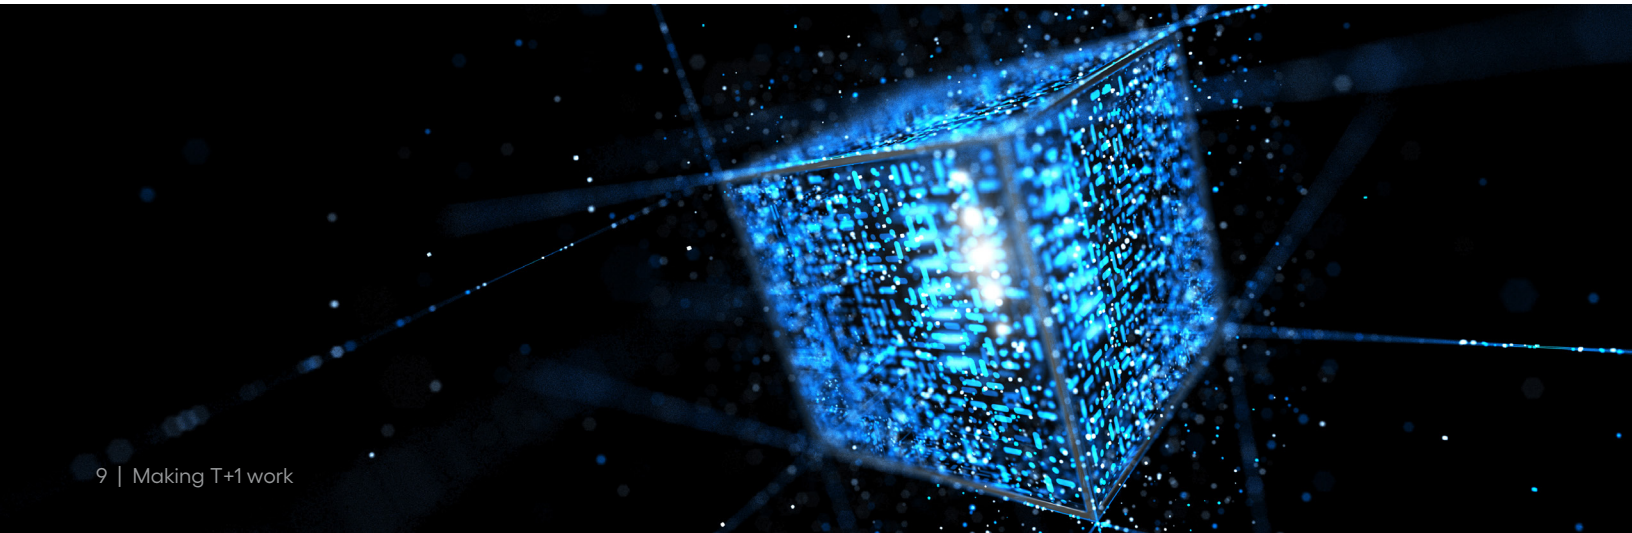
From principles to design: Applying the methods

Moving from diagnosis to design requires more than identifying pain points; it requires a structured way to challenge assumptions and translate requirements into operating reality. By combining approaches such as first-principles thinking and the Zachman Framework, firms can create a practical method for redesigning operating models for a T+1 environment.

Let's look at some of the core processes firms should focus on to meet the T+1 mandate.

STP and exception management

Straight-through processing is no longer an efficiency ambition; it is the operating condition for settlement certainty. The compressed cycle leaves little room for manual repair, re-keying, queue-based investigation or end-of-day clean-up. Every manual touchpoint becomes both a latency source and a potential failure generator. Europe's challenge is particularly acute because the transition must work across multiple currencies, settlement systems, legal frameworks, market practices and CSDs. The EU T+1 Industry Committee's handbook stresses that automation, improved data quality and greater standardisation will be essential to implementation [**EU-T+1 Handbook**]. Firms cannot simply process faster; they must redesign how breaks are prevented, detected, prioritised and resolved inside the settlement window.



Manual touchpoints must be removed from the critical path

Under T+2, firms could often absorb manual allocation repair, confirmation chasing, settlement instruction correction or inventory investigation because there was still time to recover. Under T+1, those same interventions become structurally risky. The target state should be a no-touch processing path from trade capture through allocation, confirmation, matching, settlement instruction generation, inventory validation and settlement confirmation. Manual intervention should be reserved for genuinely judgment-based exceptions, not routine repair. If an issue is predictable, recurring or rules-based, it should be classified, routed and resolved automatically.

Classify exceptions by root cause

A generic settlement-break queue will not be fast enough when allocation, confirmation, funding, collateral and settlement deadlines are compressed into the same operating window. Firms should classify exceptions by root cause. For example, exceptions may include data mismatches, inventory shortfalls, settlement-location errors, funding or FX dependencies and partial-settlement constraints. Each category needs a predefined resolution path. An inventory shortfall may require borrowing, recall, substitution or reprioritisation, whereas a partial-settlement constraint may require threshold logic, hold/release controls or counterparty agreement. The point is not simply to resolve exceptions faster. It is to stop low-value breaks from consuming the same operational attention as settlement-critical risks.

SSIs and reference data

If the settlement instruction is wrong, inconsistent or arrives in a form that a machine cannot read, no amount of downstream automation can rescue it. In this case, data is the precondition for everything.

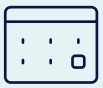
The EU T+1 Industry Committee's SSI Task Force found that inconsistencies in the format and method of SSI exchange are primarily caused by heavily manual underlying processes, with instructions often received by fax or PDF.

SSIs typically live in multiple locations: the settlement engine, the custody/sub-custody system, client onboarding/KYC records, the confirmations platform and a layer of desk-level spreadsheets and email archives. Many firms hold settlement details as free-text narrative fields rather than discrete, validated attributes.

The ValueExchange T+1 Pulse Q1 2026 survey, published in March 2026 [**ValueExchange**], confirms that almost half of EU firms have yet to automate settlement instruction processes.

> The survey also found that only about a **third of firms (35%) have automated PSET population** at the point of allocation or confirmation, meaning that a material share of the market will enter the new cycle with the single most common matching field still unresolved.

To address these challenges, firms should focus primarily on the following core areas:



Pre-match and validate on trade date:

Shift pre-matching of SSIs, PSET and accounts entirely to trade date before custodian cut-offs, rather than treating exception resolution as a T+1 morning task. Same-day validation guarantees trade eligibility for overnight processing such as T2S night-time settlement (NTS) and allows counterparty breaks to be resolved while trading desks are active.



Capture settlement-control attributes:

SSI and reference data should include hold/release status, partial-settlement election, no-partial indicators, instruction linkage, counterparty preferences and CSD/custodian rules. Under T+1, these attributes should be validated on trade date with SSI, PSET and account data, so firms know whether an instruction can be released, partially released, auto-partialled, held back or escalated before the night cycle begins.



One golden source:

One authoritative repository for the critical data elements, including SSI, PSET, PSAF/SAFE, legal entity, account, CSD/sub-custodian and cut-off calendar, held as validated attributes, not free-text narrative.



Enrich and standardise:

Reconcile and enrich against DTCC ALERT and the SWIFT SwiftRef directory and retire email/PDF/fax exchange in favour of standardised, machine-readable allocations and confirmations, in line with the amended RTS on Settlement Discipline, expected to apply from 7 December 2026.



Propagate, don't duplicate:

An enrichment made in one system (settlement engine, custody, KYC, confirmations) must reach all relevant systems before the matching window.

Collateral that moves: Turning visibility into settlement-ready liquidity

Shorter settlement cycles may reduce the period for which cash, securities and collateral are blocked, but that does not automatically make those assets usable for settlement. The real challenge is whether firms can identify, value, optimise, recall, substitute, pledge and move the right collateral before the relevant settlement, margin, repo, triparty, custodian or CSD cut-off becomes binding.

The challenges of managing and optimising collateral start with upstream processes. Today, 70% of firms[1] report daily settlement matching and delivery issues, which affect their ability to optimise collateral because they cannot confidently move the selected asset to the required account, counterparty or settlement venue on time.

More than half of firms[2] are still struggling with internal automation and dependencies on their ecosystem partners. Many EU institutions and asset managers still operate post-trade systems built on overnight batch processing, limiting intraday visibility of asset availability and preventing real-time collateral optimisation. In current operating models, manual collateral substitution can take hours, frequently missing same-day cut-off windows. This compels firms to use defensive buffers: 35% of firms post more than half of their collateral overnight and maintain approximately 6–7%[3] excess collateral to compensate for poor mobility and uncertain delivery.

> Under T+1, collateral management is no longer a product-level support process. It becomes both an internal interoperability problem and a market-infrastructure mobilisation problem. Siloed collateral management functions with fragmented infrastructure limit the ability to view available inventory, requirements and eligibility at the enterprise level.



References

3. https://thevx.io/wp-content/uploads/2025/11/VX-2025-11-The-case-for-collateral-tokenisation-Key-findings.pdf
4. https://posttrade360.com/news/infrastructure/eu-t1-industry-committee-releases-latest-t1-readiness-survey/
5. https://www.nasdaq.com/solutions/fintech/resources/reports/making-the-case-for-tokenized-collateral

To achieve optimisation, **management across business lines must first agree on the programme's objectives and understand the goal of optimising for the enterprise**. Mobilising a cross-functional enterprise programme requires a shift in mindset, but this approach can help FIs achieve significant economic benefits. While leading EU Tier-1 banks have already built enterprise collateral management platforms that ingest real-time positions from internal sub-ledgers, the majority of firms still operate in silos and lack the basic analytics needed to identify overcollateralisation or inefficient collateralisation. Moreover, firms with combined collateral operations across different lines of business can benefit from internal collateral swaps at an internal funds transfer pricing (FTP) rate or use CCP cross-margining facilities across products. To achieve seamless collateral movement across the business, firms should focus on a few major structural changes:



Real-time inventory aggregation and demand consolidation:

As collateral optimisation requires a shift from siloed, manual desk-level allocation to an enterprise-wide, real-time automated workflow, firms should first create a single, real-time inventory engine across all internal systems, including trading books, centralised repo, securities lending, derivatives inventory, prime brokerage and treasury, and external custodians/CSDs. At the same time, the platform should be able to ingest incoming collateral demands across lines of business and create a real-time aggregated view of obligations and collateral demands. An enterprise-wide collateral forecasting system should be established to provide a firm-wide view of funding requirements. Firms should use predictive analytics to forecast settlement fails, margin calls, liquidity shortfalls, securities lending recalls, collateral disputes and FX funding gaps, providing a realistic view of actual demand. This should collectively show where collateral is needed, who controls it, which obligation it can satisfy, whether it is encumbered, how it is valued after the haircut and whether it can be mobilised before cut-off.



Eligibility and constraint normalisation:

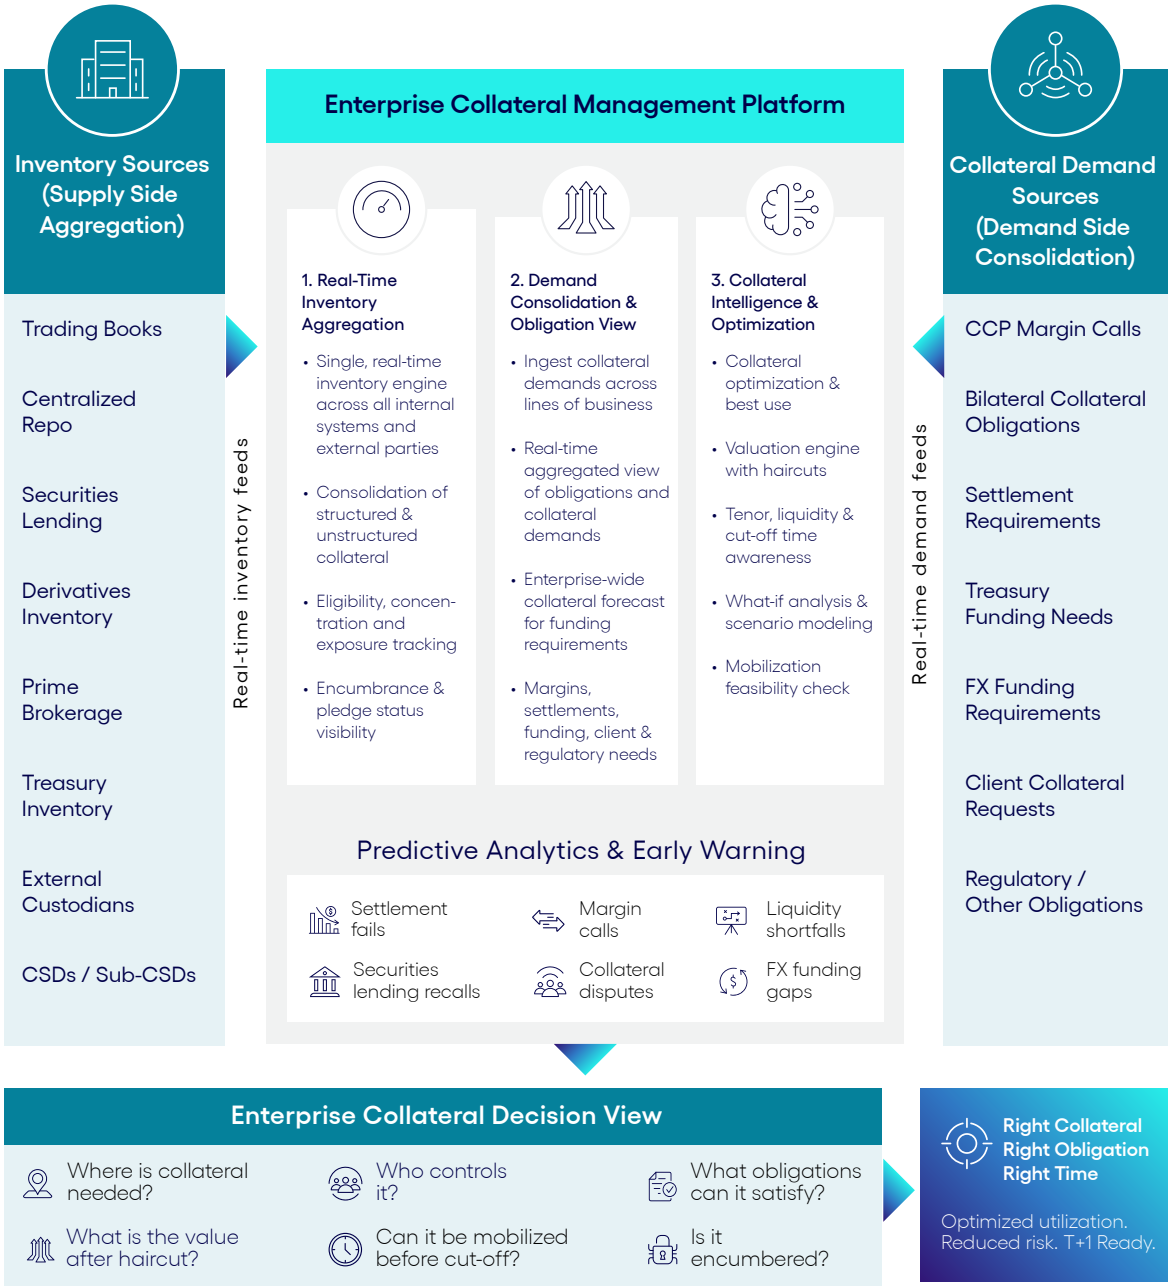
Every collateral obligation is governed by legal agreements (CSAs, GMRAs, OSLA) and regulatory rules. The firm should have a platform that creates a constraint normalisation matrix covering eligibility, valuations, haircuts and concentration rules. CSAs, GMRAs, GMSLAs, CCP rules, triparty schedules and central bank collateral criteria should also be digitised as part of the platform, because manual interpretation of eligibility rules will not scale in a T+1 environment.



Algorithmic optimisation solver:

Collateral optimisation is usually multi-dimensional. It is fundamentally a constrained combinatorial optimisation problem, with the objective of minimising total cost. Firms should not simply post the cheapest asset for each individual obligation. They should optimise at enterprise level using an opportunity-cost model that considers funding cost, liquidity value, HQLA scarcity, haircut cost, balance-sheet usage, settlement cost and alternative use. The collateral optimisation engine should use a mathematical model to determine the globally optimal allocation of assets to obligations and evolve through successive levels of maturity.

Rule-based collateral allocation may be easy to implement but can be suboptimal because it does not solve the global cost-minimisation problem. The cheapest-to-deliver (CTD) method is quite common among firms but also has limitations, because local CTD decisions can create enterprise-level inefficiency. For example, one desk may post a cheap asset that another desk urgently needs for a more valuable obligation. Among other methods, mixed-integer linear programming (MILP) is more realistic because collateral allocation often involves discrete lots, minimum transfer amounts, settlement tickets, concentration limits and binary eligibility decisions. Depending on the situation, firms can use other methods such as linear programming (for divisible asset allocation), scenario-based optimisation (for stress, margin spikes, settlement fails and liquidity shocks) or network-flow optimisation (for cross-custodian, cross-CSD and legal-entity movements).



This target-state architecture could be enabled through Cognizant's collateral transformation expertise and Microsoft cloud and data capabilities, with Microsoft Fabric serving as a unified data and analytics foundation for inventory visibility, forecasting and collateral optimisation.

Treasury at the settlement edge: Intraday liquidity control

Under EU T+1, treasury can no longer operate as a downstream reconciliation function that validates cash after the operational day has closed. It is the place where settlement readiness, FX execution, collateral availability, repo funding and intraday liquidity converge inside a much narrower window. Treasury needs to move from end-of-day cash confirmation to intraday liquidity control.

CLS eligibility is necessary but not sufficient

CLS eligibility is only part of the answer. The real operating-model question is whether each flow can practically meet the full chain of internal FX deadlines, custodian cut-offs, settlement-member requirements and funding timelines. Firms should maintain a **live control matrix** showing CLS eligibility, internal execution deadlines, custodian cut-offs, funding sources, DCA requirements and fallback routes for each material currency and settlement pathway. For non-CLS flows, or those that cannot meet CLS or custodian timing requirements, firms should pre-agree bilateral controls before go-live. These should cover settlement-risk ownership, escalation paths, pre-funding expectations, netting options, payment confirmation standards and fallback procedures.

Repo and SFTs must be pulled into treasury's control layer

Repo and SFTs should not be treated as peripheral to EU T+1 readiness. Even though SFTs are formally outside the T+1 mandate, they are not exempt from its consequences. They provide the cash, collateral and inventory flexibility needed to support settlement completion. A compressed settlement cycle may increase pressure for same-day repo execution, faster collateral mobilisation, reduced netting opportunities and more active intraday funding decisions. Repo desks, collateral teams, securities lending and settlement operations therefore need to operate from a **shared intraday liquidity view**, rather than managing funding, collateral and settlement exceptions in separate workflows.

Follow-the-sun coverage needs decision rights

The follow-the-sun model also needs to move beyond monitoring. APAC and other out-of-region desks may face the heaviest timing challenge because the European settlement window falls within their overnight or early-morning operating cycle. These teams need delegated authority to execute FX, fund accounts, mobilise collateral, arrange repo funding and escalate exceptions within pre-agreed thresholds. A global coverage model that provides visibility but not authority will still fail if every material decision requires European sign-off during a compressed settlement window.

Treasury ownership must be explicit

The target operating model should assign clear ownership of intraday liquidity forecasting, multi-currency FX execution, T2S DCA funding, repo funding, collateral mobilisation and exception escalation. Under EU T+1, treasury's success will not be measured by whether cash reconciles tomorrow. It will be measured by whether the right cash, collateral and funding reach the right account before today's settlement window closes.

Looking ahead

This framework rests on a foundational assumption: that exception-by-rule processes and governance structures can be executed reliably at the speed and volume T+1 demands. However, the overnight window, the recall chain, cross-CSD collateral flow and intraday inventory reprioritisation each generate exceptions that may be manageable individually but, under T+1's compressed timelines, can arrive faster and in greater volume than rule-based systems and staffed desks can absorb without degradation.

That is where the next frontier lies: autonomous AI agents. They reason across live data streams, adapt to novel exception patterns that rule sets have not anticipated, and act within predefined authority boundaries without waiting for human triage. AI solutions can help turn real-world challenges into proofs of concept (PoCs) and support a phased shift from batch-based to event-driven architecture.

Two industry leaders. One post-trade frontier.

The transition to T+1 is more than a regulatory milestone. It is an opportunity to modernise post-trade operations for the future. Cognizant brings deep post-trade consulting and transformation expertise, while Microsoft Azure provides scalable cloud, data and AI capabilities. Together, Cognizant and Microsoft can help firms modernise post-trade operations and build more resilient, efficient and future-ready operating models.

> In our fourth paper, we explore how autonomous AI agents can help firms manage the operational exceptions T+1 will inevitably generate at scale, and what firms must have in place before deploying them responsibly in a regulated settlement environment.



Authors



Anshuman Choudhary

Senior Partner, Consulting
Anshuman.Choudhary@cognizant.com

Anshuman has diverse experience in financial services sector over the last 25+ years. He has strong subject matter expertise in risk management, capital markets and wealth management.



Pramit Basu

Consulting Principal, BFS Consulting
Basu.Pramit@cognizant.com

Pramit brings over two decades of deep domain experience in capital market and commodity trading, bridging the gap between traditional business and modern digital agility, including AI solutioning at enterprise level.



Aniruddha Ranade

Consulting Manager, BFS Consulting
Aniruddha.Ranade@cognizant.com

Aniruddha brings 15 years of consulting expertise spanning capital markets, commodities trading, BFSI, and data standards, connecting financial domain depth with modern digital and enterprise solution.

cognizant®

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or **@cognizant.**

World Headquarters

300 Frank W Burr Blvd
Suite 36, 6th Floor
Teaneck, NJ 07666
USA
Tel: (201) 801-2333

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (0)20 7297 7600

India Operations Headquarters

Siruseri-Software Technology
Park of India (STPI)
SDB Block – Ground floor, north wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North,
North Tower
Singapore 138542
Tel: +65 6812 4000

© Copyright 2025–2027, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.