zensar

# DTCC Tokenization Service: Operational Readiness Framework

*How to prepare?*

White paper

An RPG Company

# Editorial note

This paper provides an operational readiness framework for DTC participant firms — custodians, broker-dealers, and asset servicers — to evaluate their preparedness for DTC's announced tokenization service. It draws on publicly available regulatory documents, DTCC publications, industry research, and Zensar's consulting work with capital markets participants.

Qualifications:
(1) The DTC tokenization service is an announced pilot — timelines and specifications are subject to change.
(2) Regulatory obligations are evolving; specific compliance obligations should be confirmed with legal counsel.
(3) This paper is not legal advice and not independent research — Zensar has commercial services in this space.
(4) Statements are distinguished as confirmed facts, communicated roadmaps, analytical inferences, reasonable assumptions, or areas still evolving.

**Series navigation**
Paper 1 (Why it Matters): Executive framing and decision framework.
Paper 2 (this paper): Full six-domain operational playbook.
Paper 3 (How to Validate): Pre-production validation, break injection, and QA playbook.

# 1. Program context: Confirmed facts and key inferences

This section provides a concise summary of the context. For the full structural analysis, including the dual-record consequence analysis and the traditional vs. tokenized reconciliation comparison, see paper 1 of this series.

## 1.1 Confirmed facts

| Category | Statement | Basis |
|---|---|---|
| **CONFIRMED** | DTC has received SEC no-action relief to operate a three-year tokenization pilot. | SEC NAL, December 11, 2025. |
| **CONFIRMED** | Eligible securities: Russell-1000, major ETFs, US Treasuries. | DTC's request letter. |
| **CONFIRMED** | DTC retains the root wallet keys to modify tokenized positions without participants' instructions. | Explicitly stated in the NAL request letter. |
| **CONFIRMED** | Tokenized entitlements carry no collateral or settlement value during the pilot. | Explicitly stated in the NAL structure. |
| **CONFIRMED** | DTCC has convened 50+ firms and communicated limited production trading begins July 2026, full commercial launch October 2026 | DTCC press release, May 4, 2026. |
| **CONFIRMED** | DTC's tokenization service runs on the ComposerX suite; DTCC partnered with Digital Asset to use the Canton Network, with additional approved networks (e.g. Stellar) added during rollout.* | DTCC announcement, 17 Dec 2025; network additions, 2026. |
| **COMMUNICATED ROADMAP** | LedgerScan provides near-real-time event monitoring. Its off-chain record constitutes DTC's official books and records. | DTCC public documentation. API specifications are not yet fully published. |
| **STILL EVOLVING** | ComposerX API schemas, CA processing framework, and participant-level regulatory evidence obligations. | Not yet communicated. See Appendix C for monitoring guidance. |

**This paper currently refers Canton Network only*

## 1.2 The structural consequence

A participant who elects tokenized record-keeping maintains security positions in two systems simultaneously: the traditional DTC book-entry entitlement and the tokenized entitlement in LedgerScan. These two records update at different frequencies, use different security identifiers, are governed by different authorities, and can be modified by DTC through the root-wallet mechanism without a participant instruction. The six-domain framework in section three provides a structured approach to managing this dual-record environment.

## 2. Why operational readiness warrants early attention

Three factors determine the urgency of readiness investment: production timing, program overlap, and implementation lead time. Until recently all three pointed the same way — begin structured assessment now, even if implementation follows later. That sequencing no longer holds for every reader. With DTCC's limited-production phase now live, assessment and implementation run concurrently for firms in or near the first wave, and sequentially only for those confident of a later entry.

## 2.1 Production timeline and overlap

DTCC's tokenization service entered limited production on 15 July 2026, with full commercial launch targeted for October 2026. This phase runs on real assets and real settlement obligations, not a sandbox: on day one, participants including J.P. Morgan and DriveWealth transacted tokenized positions in production. Firms in the first wave therefore carry live settlement obligations before the broader community has accumulated operating experience — and that first wave is now an operating reality, not a future date.

The strategic question has shifted accordingly. Before go-live it was "must we be ready for the initial wave, or can we observe first?" With limited production under way and full launch roughly a quarter out, it becomes "are we already exposed through counterparty activity, and if we intend to elect, how fast must we close the gap before October?" The answer still turns on working-group status and DTCC's onboarding sequencing, so firms should confirm their specific production timeline directly with DTCC. But the margin for a wrong assumption has narrowed: the observation window the earliest movers relied on has largely closed.

The tokenization program also overlaps with the ISO 20022 SCoRE migration for European markets. Under the AMI-SeCo roadmap, the corporate-actions element runs later than general meetings: SCoRE compliance for general meetings falls due around end-2026, but corporate-actions message migration is targeted for end-2028, with full CSD-to-participant CA communication in ISO 20022 by 2030. Whether this intersects with dual-record CA processing architecture is an analytical inference, not a stated requirement; participants should confirm applicability and current dates with their ISO 20022 program teams.

## 2.2 Talent and build constraints

Building or configuring these capabilities requires expertise at the intersection of post-trade operations and distributed ledger technology — a combination not yet widely available in the market. Canton development uses Daml, for which the certified developer community is limited. Firms planning internal builds should treat talent planning as a parallel workstream, rather than a sequential one. Firms using vendor or accelerator routes have significantly more flexibility on this dimension, which is one reason the build vs. buy vs. configure decision should be made early (see section 6).

## 2.3 What participant operations teams are asking

'We don't know who owns the exception.' Operational accountability for tokenized position breaks is unresolved at most firms. The six-domain framework in section 3 is designed to give that conversation a more precise vocabulary.

'We're waiting for DTCC to tell us what they need.' Many operations heads are deferring planning until ComposerX API specs and CA guidance are published. The risk is underestimating implementation lead time.

'We're not sure what the regulator will ask for.' Firms that incrementally extend their existing audit infrastructure are making a reasonable bet. Firms doing nothing are taking a risk of unknown magnitude.

## 3. Capability domains: The operational readiness framework

The six-domain framework identifies the operational capabilities that participants electing tokenized record-keeping will need to address. These domains are an assessment model, not a mandated architecture.

## The six operational capability domains

The diagram below shows the dependency structure of the six-domain framework. Domains one and two are prerequisites for all downstream capabilities. Domain three is the operational hub. Domains four and five can be sequenced based on the exposure profile. Domain six spans all domains.

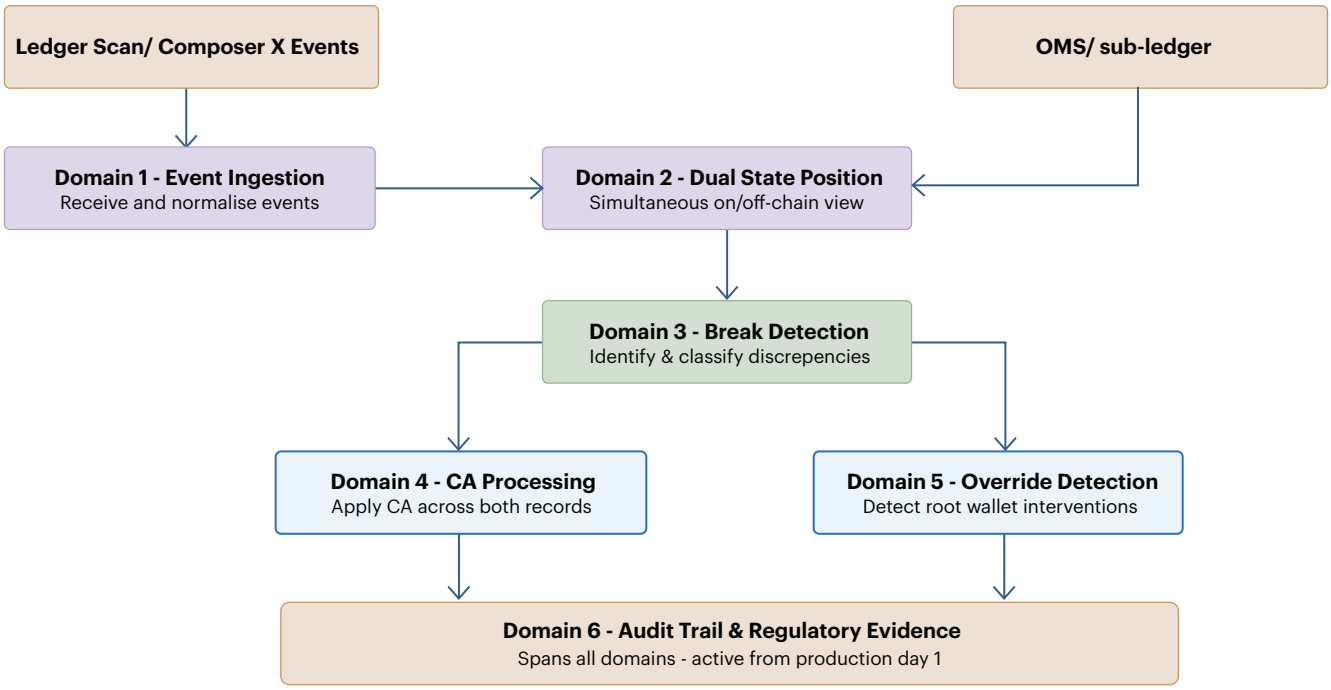


*Six domain dependency structure*

Domains one to three are foundational prerequisites for production entry. Domain three is the operational hub. Without it, position differences cannot be routed. Domains four and five depend on operational exposure. Domain four's urgency is driven by CA complexity. Domain five's urgency is driven by observed override frequency. Domain six spans all domains and should be active from day one of production. This framework is an assessment model, not a mandated architecture, and some participants may find that the existing infrastructure already covers one or more domains.

The framework identifies six operational capability domains that are likely to be relevant to participants who elect tokenized record-keeping. These domains are not prescriptive. They represent an analytical view of the operational functions that the dual-record structure may require. Participants may find that some domains are already addressed by existing infrastructure, that some can be deferred, or that the boundaries differ from those suggested here.

**Note on alternatives.** For each domain, this paper notes potential implementation approaches. These include internal builds, established post-trade technology vendors (Broadridge, Finastra, Murex, Axoni, Vermeg), tools from the Canton/Digital Asset ecosystem, and specialist accelerators. No single approach is universally appropriate. Participants should evaluate options against their own technology stack, timeline, and investment appetite.

## Conditions under which this framework may prove overstated

The six-domain framework rests on assumptions about the operational character of the DTC tokenization service. If the following conditions materialize, the operational complexity described in this section may be significantly lower than this paper anticipates. Participants should monitor for these conditions and adjust their readiness investment accordingly.

| Condition | Implication for readiness investment |
|---|---|
| **DTCC provides a managed participant integration layer for ComposerX** | Per-firm investment in domain one (event ingestion) may be minimal. DTCC has historically prioritized reducing per-participant integration burden; a standardized integration service is a plausible outcome. |
| **Established vendors absorb most operational complexity within existing platform relationships** | This condition is now partly materializing. Broadridge has released integrated tokenized-securities infrastructure covering the CA lifecycle (May 2026), and other platforms are extending. Participants already on such a platform may find domains two, three, and four substantially addressed within existing contracts — but must confirm coverage for their specific platform version and DTC API level rather than assume it. |
| **Root-wallet interventions remain extremely rare** | Domain five (override detection) is operationally significant only if overrides occur frequently enough to warrant systematic management. If the pilot produces a handful of overrides per year, manual detection and documentation may be entirely sufficient. |
| **Participant adoption in the early pilot phase remains limited** | If working-group firms carry only small tokenized position volumes in 2026-2027, the operational pressure on all six domains will be lower than this paper's framing implies. Volume-proportionate readiness is the appropriate response. |

## 3.1 Domain one: Event ingestion from LedgerScan and ComposerX

**Scope:** Reliable, authenticated receipt of token events from DTC's ComposerX subsystems — principally LedgerScan (near-real-time token state changes), factory (wallet permissions), and CMP (lifecycle events). All downstream domains depend on this layer: unreliable ingestion degrades position management, break detection, and audit equally. This is likely the most consequential technical decision in a participant's readiness program.

**Key question:** Participants already using market data feeds or custodian APIs will find the concept familiar. The specific challenge is that ComposerX API specifications — schemas, authentication, rate limits, failover — were, until recently, not fully published. With limited production now live, DTCC has begun exercising these interfaces with working-group participants, so specifications are being released in stages rather than all at once. Confirm current availability via DTCC's developer portal before sizing domain one; detailed build planning should track that release cadence rather than wait on a single complete publication. Participants should also monitor whether DTCC will offer a standardized participant-facing integration layer, which could materially reduce per-firm domain one investment requirements.

**Potential approaches:** Custom API client development; Canton SDK integration using digital asset's tooling; middleware platforms offered by established financial technology vendors; specialist accelerator components from boutique DLT integrators or consulting practices (see appendix A for one example).

## 3.2 Domain two: Dual-state position management

**Scope:** Maintaining a simultaneous, queryable view of both the on-chain tokenized position (from LedgerScan) and the off-chain book-entry position (from the OMS or sub-ledger). Without this, a participant cannot detect discrepancies. The specific approach — a separate store, an OMS extension, or a middleware layer — matters less than the outcome: reliable access to both states at any point in time.

**Key question:** Custodians already maintain multiple position views and may question whether this is genuinely new. The answer depends on the OMS architecture: the near-real-time update frequency, DTC's authoritative status, and the different identifier spaces (Canton token IDs vs. CUSIP/ISIN) may require extension rather than a replacement. Modern cloud-native OMS platforms are better positioned to accommodate a tokenized dimension than legacy batch-oriented systems.

**Potential approaches:** Extension of existing OMS position management; dedicated dual-state position store; capabilities offered by Broadridge, Finastra, or other OMS vendors as those platforms extend to tokenized securities; specialist middleware.

## 3.3 Domain three: Position discrepancy detection and classification

**Scope:** Identifying discrepancies between on-chain and off-chain position states, classifying them (expected timing differences vs. genuine breaks requiring action), and routing actionable breaks with sufficient context for efficient resolution. Without classification, every position difference generates a manual investigation, which is unsustainable at any meaningful volume.

**The role of a break taxonomy:** A break taxonomy that defines categories of discrepancy, expected resolution workflows, and escalation triggers is a useful operational artifact regardless of the technical implementation. One published example is the ZTAOP Break Taxonomy (available as a companion document), which defines 16 break types across four families. Participants may develop their own taxonomies, adapt existing reconciliation frameworks, or draw on industry body standards as they emerge. No single taxonomy approach is universally correct. A well-structured taxonomy also provides the break-type inventory needed for pre-production break injection testing (see paper 3).

**Considerations and counterarguments:** Established reconciliation platforms, such as Broadridge, Finastra Reconciliation Manager, SmartStream TLM, and Duco, already perform break detection and classification for traditional securities. The relevant question is whether these platforms can be extended to handle the tokenized dimension without wholesale replacement. Broadridge's May 2026 release delivered break-detection-relevant capability across tokenized and traditional securities for its existing client base — evidence that at least one major vendor has moved from intent to production.

## 3.4 Domain four: Corporate action processing for dual-record positions

**Scope:** Applying corporate action events, including dividends, coupon payments, rights issuances, mergers, and splits, consistently across both traditional and tokenized records. A CA applied to one record but not the other creates a discrepancy more complex than a timing break: a financial event has been misapplied on one side, requiring coordinated correction across both systems.

**Urgency:** Depends on the CA profile of the tokenized portfolio. A participant tokenizing only US Treasury bills (simple CA profile: maturity and coupon) faces materially lower complexity than one tokenizing a broad equity portfolio with merger activity. Assess CA complexity before setting domain four priority, and defer implementation until DTCC publishes its CA processing framework.

**Key question:** CA processing for tokenized positions is the most operationally complex element of the program. There is no industry consensus on the right architecture. Broadridge's May 2026 release covers the corporate-actions lifecycle — dividends, mandatory and voluntary actions, proxy voting — across tokenized and traditional securities for its platform clients. Digital Asset's Canton platform provides smart contract primitives for on-chain CA automation. The right approach, which may include compensating for a transaction, a smart contract, or enhanced CA notification, depends on the CA type and existing infrastructure.

Illustrative scenario: Dividend event on a dual-record position

| Event stage | Traditional record (OMS) | Tokenized record (LedgerScan) |
|---|---|---|
| **CA event: Dividend declared** | DTC CA notification received; dividend credit queued for batch processing | Canton token event generated; CA processing not yet confirmed on-chain |
| **Processing gap: (minutes to hours)** | OMS sub-ledger updated; client statement reflects dividend | LedgerScan token balance unchanged; no CA event confirmed on-chain yet |
| **Domain three: Break detected** | Position appears complete on the traditional record | LedgerScan does not reflect dividend entitlement. Domain three flags a discrepancy. |
| **Domain four risk: If CA misapplied** | CA credited on traditional record — cannot be uncredited easily | The tokenized record has no CA event. Correction requires coordinated action across both systems — complexity grows with position size and elapsed time. |

## 3.5 Domain five: Root-wallet override detection and documentation

**Scope:** Detecting and documenting DTC root-wallet modifications to participant positions — events described as 'conditions requiring reversal' in the NAL. These are a designed feature of the pilot. A participant who misses them carries an incorrect internal position and cannot provide the corroborating evidence that DTCC's quarterly SEC filings will create, as DTCC's quarterly SEC filings will create an expectation for.

**Key question:** The frequency of root-wallet overrides during the pilot is unknown. If rare (a handful per year), a manual detection and documentation process may be sufficient initially. Automated detection is warranted once transaction volumes and override frequency are better understood. A risk-proportionate approach: monitor LedgerScan for root-wallet event types without full automated evidence generation at first.

Illustrative scenario: Dividend event on a dual-record position

| Time | Participant OMS | DTC LedgerScan |
|---|---|---|
| **09:05** | Position: 10,000 shares XYZ | Token balance: 10,000 shares XYZ |
| **09:08** | No instruction sent. No OMS event triggered. | DTC root-wallet event: balance adjusted to 9,000 shares (Condition Requiring Reversal). |
| **09:09** | Position: 10,000 shares — unchanged | Token balance: 9,000 shares — updated |
| **09:10 Result** | OMS carries 10,000. Any trade or collateral calculation based on this position uses incorrect data. | LedgerScan carries 9,000. The 1,000-share divergence is undetected without Domain 5 monitoring active. |

## 3.6 Domain six: Audit trail and regulatory evidence

**Scope:** Maintaining a complete, queryable record of tokenized position events, discrepancies, and their resolution — for regulatory examination and point-in-time position reconstruction. Participants should plan for domain six to be active from production day one: audit evidence cannot be reconstructed retrospectively, and SEC examiners may request event logs and control evidence for tokenized positions.

**Key question:** Many participants already maintain audit trails for settlement operations. The question is whether those trails extend to on-chain events and LedgerScan state changes. Extending existing audit infrastructure to capture tokenized events is likely more efficient than building a separate system. Participants using an immutable or append-only logging infrastructure — such as Apache Iceberg, Azure SQL Database ledger, Microsoft Purview, or an equivalent — may find this extension straightforward. Actual feasibility depends on platform version and integration architecture.

# 4. Participant readiness maturity model

The following model is a self-assessment tool. Five maturity levels apply across the six capability domains. A firm at level two is not necessarily 'behind.' The appropriate level depends on expected tokenized volumes, production timeline, and available resources.

| Maturity level | Maturity level |
|---|---|
| **Level one: Awareness** | The organization understands that the DTC tokenization service has been announced and may have operational implications. No specific readiness investments have been made. Appropriate for firms not in the initial working group or not expecting to elect tokenized record-keeping in the first production wave. |
| **Level two: Assessment** | The organization has conducted a structured assessment across all six domains. It understands which capabilities exist, which are absent, and which require extension vs. new development. It has evaluated implementation options (internal build, vendor platform, specialist accelerator), engaged legal counsel on regulatory obligations, and has a view on its expected production entry timeline. |
| **Level three: Pilot ready** | Basic capabilities for domains one, two, and three are implemented in a test/UAT environment using the LedgerScan sandbox. The organization can ingest token events, maintain a dual-state position view, and detect and classify discrepancies. It has a documented break-resolution workflow and has conducted synthetic event testing and basic break-injection exercises. Domain four (CA) and six (audit trail) may be at a basic level only. |
| **Level four: Production ready** | Production-grade capabilities across D1-D3 and D5, with documented runbooks and trained staff. Domain four addressed at least four of the CA types in the tokenized portfolio. Domain six provides queryable records of all tokenized events. The organization has completed a structured validation program: synthetic event testing, break injection, root-wallet override simulation, and audit evidence completeness testing. |
| **Level five: Operationally optimized** | The organization has operated in production and optimized based on experience. Break triage is automated (reducing manual intervention rates). Tokenized CA processing is integrated into the standard CA workflow. An operating model exists for onboarding new tokenized position types without bespoke development. DTCC program evolution is actively monitored. |

## 4.1 Domain-level readiness heatmap

Methodology note: Ratings represent Zensar's analytical inference, not primary firm-level research. They draw on known infrastructure characteristics for each participant category, the operational demands of each domain, and analogous readiness patterns from T+2 to T+1 and from CSDR implementation. Within-category variance is significant. Domain 4 (CA atomicity) is L1 across almost all categories because no participant has production experience with dual-record CA processing — this is a market-wide gap, not a firm-specific deficiency. Use as a self-assessment prompt, not a benchmark. Because these are inferred rather than measured, treat the grid as a structured self-scoring template: the value is in prompting each firm to rate its own six domains, not in the pre-filled category ratings, which no firm should cite as a benchmark for its peers.

| Participant type | D1 ingestion | D2 dual-state | D3 break detect | D4 CA atomicity | D5 override | D6 audit trail |
|---|---|---|---|---|---|---|
| **Tier-one global custodian** | L2-L3 | L2-L3 | L2-L3 | L1–L2 | L1 | L2-L3 |
| **Mid-tier custodian/servicer** | L1-L2 | L1-L2 | L1-L2 | L1 | L1 | L1-L2 |
| **Broadridge-platform firm** | L2-L3 | L2-L3 | L2-L3 | L2-L3 | L1-L2 | L2 |
| **Self-clearing broker-dealer** | L1-L2 | L1 | L1-L2 | L1 | L1 | L1-L2 |
| **Fintech / digital-native** | L2-L3 | L2 | L2 | L1-L2 | L2 | L2 |

L1 = Awareness · L2 = Assessment underway · L3 = Pilot-level capability. Domain four (CA atomicity) is L1 across almost all categories — market-wide gap, not a firm-specific deficiency.

## From heatmap level to immediate priority

| Current level | Immediate priority |
|---|---|
| **Level one: Awareness** | Structured assessment. Complete phase one of the action plan (section 8): confirm production wave, identify eligible securities, map existing domain coverage. |
| **Level two: Assessment** | Vendor strategy decision. Complete phase two: obtain written vendor roadmap commitments, select implementation path, and engage legal counsel. |
| **Level three: Pilot ready** | Pre-production validation. Execute phase four before production entry: synthetic event testing, break injection, override simulation, auditor briefing. |
| **Level four: Production ready** | Operational optimization. Begin phase five: automate break triage, integrate D4 CA processing, establish program monitoring cadence. |
| **Level five: Optimized** | Ongoing: Monitor DTCC program evolution (new securities, networks, eligibility) and contribute to industry working groups. |

# 5. Operational risk considerations

> All estimates in this section are illustrative planning anchors, not financial projections. Actual exposures depend on transaction volume, position size, CA complexity, and existing infrastructure.

## 5.1 Settlement integrity risk

If a participant's internal position diverges from DTC's LedgerScan record due to an undetected event — a root-wallet override or an on-chain settlement not reflected in the OMS batch — the participant may transact based on incorrect position data, risking settlement failures, duplicate deliveries, or position over-commitments.

US market settlement failures primarily carry contractual consequences (an obligation to make counterparties whole for buy-in costs). EU markets impose CSDR daily cash penalties for settlement failures; the US has no equivalent mandatory regime. Confirm the specific consequence framework with legal counsel. For firms with low initial tokenized volumes, enhanced manual monitoring during early production may be sufficient.

## 5.2 Regulatory evidence risk

DTCC must file quarterly reports with the SEC covering root-wallet interventions and participating firm activity. Specific participant-level evidence obligations have not been communicated, but it is a reasonable planning assumption that SEC examiners will review participant-level controls. Participants with clear, queryable records of tokenized position history and break resolution activity are better positioned for examination. Confirm specific obligations with legal counsel before committing to an evidence architecture.

## 5.3 Operational cost risk

Participants relying entirely on manual break detection at meaningful tokenized volumes face potentially high operational costs. Break volumes in early production periods are typically higher than steady-state, and a novel dual-record environment amplifies this pattern.

Proportionate approach: identify the maximum manual break investigation capacity within existing operations staff, estimate the transaction volume at which that capacity would be exceeded, and plan automated classification and triage to be in place before that threshold is reached.

Illustrative sizing (participants should substitute their own figures). As a planning device only: if a first-wave book generates B position breaks per day, each requiring m minutes of analyst investigation, daily manual load is B × m minutes. Against a fully-loaded operations cost of roughly C per analyst-hour and an available manual capacity of H analyst-hours/day, the point at which automated triage becomes necessary is where B × m / 60 > H. Early-production break rates typically run above steady state, so B should be stress-tested at a multiple of expected volume. This paper deliberately supplies the structure, not the numbers — a firm's own break-rate, staffing, and cost assumptions are the only defensible inputs, and each should be recorded as an assumption, not an assertion.

## 5.4 Phased and proportionate response

The sequencing below assumes a pre-production starting point. For firms already in the live first wave (from July 2026), "before initial production" items are retrospective checks, not future work — such firms should read the table from the "first 90 days" row onward.

| When | Domains to address |
|---|---|
| **Before initial production** | D1 event ingestion + D2 dual-state position view (prerequisites — cannot safely participate without these). D5 override detection at basic monitoring level. |
| **Within the first 90 days of production** | D3 break classification at a level sufficient to distinguish timing breaks from genuine discrepancies (manual classification is acceptable initially). D6 audit trail at a level sufficient to respond to regulatory inquiries. |
| **Within the first year** | D4 CA Processing — timed to the CA profile of the tokenized portfolio and DTCC's published CA framework — full automation of D3 and D5. D6 is enhanced for proactive regulatory reporting. |

# 6. Implementation approaches

Participants should approach the build vs. buy vs. configure decision as a first-order program question, not an afterthought. The right path depends on internal DLT capability, vendor roadmap status, and production timeline.

## 6.1 Internal build

Maximum control over architecture and integration with proprietary systems. Requires Daml/Canton expertise, post-trade domain knowledge, and significant program management capacity. The implementation duration for core domains (D1-D3) is typically several months from program initiation; D4 adds complexity and is sequenced separately.

Best suited for: Tier-1 participants with established digital assets programs and existing Canton/Daml investment. If your firm has already committed to the Canton ecosystem, internal build is the natural extension.

## 6.2 Established vendor platforms

| Vendor | Capability/status | Best for |
|---|---|---|
| **Broadridge** | Released 12 May 2026: integrated infrastructure for tokenized securities covering the CA lifecycle, proxy voting, and settlement, extending the DLR platform. Now shipping (not roadmap) — confirm eligible-securities coverage and DTC API version support directly with Broadridge. | Participants are already on the Broadridge platform — likely a lower-disruption path for D3 and D4. |
| **Digital Asset** | Canton Network is the primary ledger for DTC's service under DTCC's confirmed partnership with Digital Asset (Dec 2025). Confirm the specific integration scope available to your firm. | Participants engaging directly with the Canton ecosystem; D1 integration. |
| **Finastra · Murex · Vermeg · SmartStream · Axoni** | At varying stages of extending platforms for tokenized securities. Enquire with existing vendors before assuming new infrastructure is required. | Participants already on these platforms — extension may be available within existing relationships. |

This table reflects a fast-moving landscape (current as of August 2026). Vendor positions are changing monthly; re-confirm before any decision.

## 6.3 Specialist accelerators

Pre-built, configurable components specifically designed for the DTC tokenization participant use case. This category is still forming — providers include boutique DLT integrators, digital asset operations specialists, established consulting practices that have productized implementation experience, and smaller capital markets technology firms. Apply the same due diligence as for any emerging vendor: assess reference clients, delivery track record, and upgrade path.

Zensar's ZTAOP (Tokenized Asset Operations Platform) is one example: it includes pre-built components for LedgerScan integration, dual-state position management, break detection and classification, override detection, and an audit trail. Deployed in the participant's environment — not SaaS. See Appendix A for full details.

Best suited for: Mid-tier participants lacking internal DLT expertise for a full build and not exclusively on a single vendor platform with a confirmed tokenized securities roadmap.

## 6.4 Hybrid approaches

Many participants will combine models. For example, CA processing via an established vendor (Broadridge) alongside a specialist or internal ingestion layer (D1). The six-domain structure supports this. The critical architectural decision in any hybrid approach is ensuring reliable event data flow between components and a unified audit trail regardless of where processing occurs.

## 6.5 Readiness validation and QA

Implementation and validation are distinct workstreams. A correctly architected system can still experience material failures during initial exposure to production conditions. Three factors make pre-production validation particularly important: there is no historical break data; API specifications are incomplete; and early-phase failures carry disproportionate reputational consequences.

> Full validation framework in paper 3. Paper 3 (How to Validate) provides the complete pre-production validation framework: synthetic event generation, break injection testing, CA simulation, root-wallet override simulation, audit evidence validation, and regression testing. Paper 3 also defines the maturity gates (levels three and four) that govern production entry from a validation perspective.

## 7. Program monitoring

Five areas are most likely to change materially before a participant needs to make a production commitment: ComposerX API specifications; CA processing framework; regulatory evidence expectations; vendor platform readiness; and root-wallet intervention frequency.

Participants should designate a monitoring owner for these five areas. Given the current pace — live production since July 2026, October full launch, and monthly vendor announcements — a monthly review cadence is appropriate through the pilot ramp, reverting to quarterly once the service stabilizes. A detailed monitoring checklist, covering what to watch for and the planning implications of each development, is provided in Appendix C.

## 8. Five-phase action plan

The entry point depends on your current maturity level (section 4) and the confirmed production timeline. Complete each phase before beginning the next.

| Domain | Objective | Priority actions |
|---|---|---|
| **Phase one: Assess (all participants)** | Understand exposure | 1. Confirm your production wave and current onboarding status directly with DTCC — note limited production is already live (July 2026) and full launch is October 2026, so "which wave" is now a near-term question, not a future one 2. Identify eligible securities in custody/inventory (Russell 1000, ETFs, Treasuries). 3. Map the existing technology stack against the six domains. 4. Engage legal counsel on regulatory evidence obligations. |
| **Phase two: Decide (after phase one)** | Select implementation path | 5. Obtain written roadmap commitments from existing vendors — confirm CA coverage and DTC API version support. 6. Select primary path: internal build · vendor extension · specialist accelerator. 7. Make the build vs. buy vs. configure decision before any implementation begins. |
| **Phase three: Prepare (after phase two)** | Build core capabilities | 8. Implement D1 (event ingestion) and D2 (dual-state position) first — prerequisites for everything else. 9. Add basic D3 (break detection) and D5 (override monitoring), scaled to expected frequency. 10. Establish D6 (audit trail) from production day one. 11. Defer D4 (CA processing) until DTCC publishes the CA framework and your CA profile is clear. |
| **Phase four: Validate (before production)** | Confirm capabilities work | 12. Test using DTCC's sandbox or equivalent synthetic event stream. 13. Conduct a break injection and root-wallet override simulation. 14. Brief internal and external auditors on your control framework. See Paper 3 for the full pre-production validation playbook. |

| Domain | Objective | Priority actions |
|---|---|---|
| **Phase five: Optimize (after production)** | Increase efficiency | 15. Automate break triage based on observed patterns.<br>16. Integrate tokenized CA processing into standard CA workflow once volume and CA type profile are understood.<br>17. Monitor DTCC program evolution (new securities, new networks) and adjust accordingly. |

## 9. Participant readiness profiles

Illustrative participant readiness profiles are provided in Appendix B.
These profiles are intended as self-assessment aids and do not represent assessments of any specific institution.

## 10. Conclusion

The DTC tokenization program, if it proceeds as communicated, will be the first instance in US capital markets where security entitlements are maintained simultaneously in traditional book-entry form and as tokenized records on a distributed ledger. That dual-record structure introduces a qualitatively new level of operational complexity, even for participants with established DLT programs.

The six-domain framework and maturity model in this paper offer a self-assessment tool for calibrating the appropriate response to each firm's specific circumstances. They are not a prescription — not every participant will need to address all six domains with equal urgency. The five-phase action plan in Section 8 provides the sequencing logic.
There is meaningful uncertainty in this analysis. Key technical specifications from DTCC remain unpublished. Regulatory expectations regarding evidence have not been communicated. Vendor approaches are still being defined. Participants should plan for iteration as specifications are clarified and avoid locking into technical architectures before the necessary information is available.

What is not uncertain: the program is live. A working group of 50+ firms is engaged, limited production began in July 2026, and full launch is targeted for October 2026. The question is no longer whether this arrives, but how quickly each participant closes its readiness gap. Participants who begin structured assessment and planning now will be better positioned to make informed decisions, and adjust them as program details become clearer.

This is Paper 2 of a three-paper series — How to Prepare.
Paper 1 (Why it Matters): Executive framing, decision framework, and initial readiness heatmap.
Paper 3 (How to Validate): Pre-production validation framework, break injection guidance, and QA playbook.
To discuss operational readiness, contact info@zensar.com. Participants are also encouraged to engage directly with DTCC's working group program for current implementation specifications.

# Appendix A: Zensar ZTAOP: One implementation approach

This appendix describes Zensar's Tokenized Asset Operations Platform (ZTAOP) as one example of the specialist accelerator approach discussed in Section 6.3. It is presented as an appendix to preserve the analytical independence of the core framework. Evaluate ZTAOP alongside the approaches in Section 6 and conduct independent due diligence.

# Appendix A: Zensar ZTAOP: One implementation approach

What ZTAOP is: An internally pre-built accelerator platform for participant-side tokenized securities operations. Configurable components for LedgerScan event ingestion, dual-state position management, break detection and classification (implementing ZTAOP Break Taxonomy v1 — 16 break types across four families), override detection and evidence generation, and audit trail management. Deployed in the participant's own environment — not SaaS. Integrates with existing OMS, accounting, and risk systems.

| Offering | Scope |
|---|---|
| **DTC/ComposerX integration** | Domains 1-3, 5-6. LedgerScan ingestion, dual-state position management, break detection, override detection, audit trail. |
| **Dual-book corporate actions** | Domain 4 and multi-regime regulatory reporting. CA processing across both traditional and tokenized records. |
| **Multi-rail interoperability** | For participants connecting to Clearstream D7 DLT, SWIFT Shared Ledger, and public chains in addition to Canton. |
| **ZenseAI Context Engineering** | Context store: ISO 20022 token event taxonomy, identifier resolution tables, CA enrichment data. |
| **Agent Mesh** | Domain-specific agent trained on the participant's break history to automate triage of position discrepancies. |

Indicative commercial model available separately. Pricing is provided on request as a distinct rate card.
What ZTAOP is not: a replacement for DTC's ComposerX, a DLT network, or an OMS. Not a regulated financial service. Does not provide legal or compliance advice. Does not provide CA processing for Broadridge-platform firms where Broadridge's own capabilities are likely more appropriate.

## B. Readiness validation and QA components

ZTAOP QA is described in full in Paper 3. The readiness validation and QA components — synthetic event generator, break injection framework, CA simulator, root-wallet override simulator, audit evidence validator, and regression harness — are described in Appendix A of paper 3 (How to Validate). Paper 3 also provides the validation framework, mapping each activity to its domain, objective, and test scenario.

# Appendix B: Illustrative participant readiness profiles

The following profiles are hypothetical and illustrative. They are not assessments of any specific firm. Use them alongside the heatmap in Section 4.1 as a prompt for internal self-assessment.

## Profile A: Large tier-1 global Custodian

| Dimension | Assessment |
|---|---|
| **Likely readiness level** | L2-L3 on D1-D3 and D6; L1-L2 on D4; L1 on D5. Established technology teams and prior DLT investment provide a foundation, but dual-record CA processing and root-wallet override detection are new for all firms. |
| **Primary challenge** | Coordination across multiple internal systems and vendors. A Tier-1 custodian may run an established vendor platform for CA, an internal or accelerator-built ingestion layer for LedgerScan, and a separate audit system — with the integration seams between them as the most operationally fragile point. |
| **Recommended first action** | Map existing systems against the six-domain framework and assign clear ownership for each domain. Confirm production onboarding timeline with DTCC. Engage DTC on the availability of API documentation to unblock D1 planning. |

## Profile B: Mid-tier custodian/asset servicer

| Dimension | Assessment |
|---|---|
| **Likely readiness level** | L1-L2 across most domains. Existing reconciliation infrastructure handles the conceptual task but is typically batch-oriented and unlikely to accommodate near-real-time event processing without extension. |
| **Primary challenge** | The build vs. buy vs. configure decision must be answered before any implementation begins. Limited internal DLT expertise means a full internal build is unlikely to be practical. Whether an established vendor will deliver production-grade ComposerX integration on a useful timeline is the critical dependency. |
| **Recommended first action** | Confirm with existing vendors (Broadridge, Finastra, SmartStream) whether and when tokenized securities capabilities will be available for the firm's specific platform version. Engage legal counsel on regulatory evidence obligations for a second- or third-wave participant. Defer D4 planning until the firm's initial tokenized portfolio scope is defined. |

## Profile C: Self-clearing broker-dealer

| Dimension | Assessment |
|---|---|
| **Likely readiness level** | L1 across most domains. Self-clearing broker-dealers have strong OMS infrastructure for settlement but have not historically managed dual-record positions or external authority modifications to their position records. |
| **Primary challenge** | Understanding actual exposure before making investment decisions. A critical first question is whether the firm's operations could be affected by the tokenization program even if it does not elect the service. Clarify with DTCC whether any operational dependencies arise from counterparty participation. |
| **Recommended first action** | Engage DTCC to understand the firm's specific exposure profile. At minimum, confirm that existing settlement exception monitoring can surface any tokenized-format instructions for escalation rather than processing them through standard workflows. Full six-domain implementation can be deferred until exposure is understood and a production timeline is confirmed. |

## Appendix C: Program monitoring checklist

The five areas below are most likely to change materially before a participant needs to make a production commitment. Designate a monitoring owner for each area and establish a quarterly review cadence.

| Area | What to watch for | Planning implication |
|---|---|---|
| **ComposerX API specifications** | Authentication mechanisms, event schemas, rate limits, and failover behavior for LedgerScan, factory, and CMP. Monitor DTCC's developer portal and engage DTC directly on documentation availability. | Defer detailed D1 planning until complete API documentation is published. All six domains depend on reliable event data from ComposerX. |
| **CA processing framework** | DTCC working group outputs and any published guidance specifically on CA atomicity for dual-record positions. | Treat D4 planning as provisional until this framework is published. Sequence D4 implementation after initial production entry rather than before. |
| **Regulatory evidence expectations** | SEC examination guidance, staff speeches, or findings from the early production phase. Engage legal counsel to monitor actively. | Do not assume the traditional securities examination framework applies unchanged to tokenized positions. Evidence architecture decisions should follow legal counsel's view. |
| **Vendor platform readiness** | Confirmed platform releases with specific eligible securities coverage and DTC API version support. | Require written roadmap commitments from existing vendors. Press announcements are not planned inputs. |
| **Root-wallet intervention frequency** | Early production data on override frequency — available through DTCC's quarterly SEC reporting. | High frequency substantially changes D5 urgency. Do not make large investments in automated override detection before this data is available. |

# References


1. SEC Division of Trading and Markets, No-Action Letter to The Depository Trust Company, December 11, 2025. Available at: SEC.gov.
2. DTCC Press Release, 'DTCC Advances Development of New Tokenization Service, Convenes 50+ Firms to Drive Digital Assets Adoption,' May 4, 2026.
3. DTC ComposerX platform documentation, as described in the SEC NAL request letter and DTCC public communications.
4. ISO 20022 SCoRE Corporate Actions Roadmap. Available at: iso20022.org. Application of Phase 2 timelines to tokenized processing streams is an inference.
5. ESMA, CSDR Settlement Discipline Regime documentation. Applicable to EU market participants; no equivalent mandatory penalty regime exists in the US as of this writing.
6. Broadridge Financial Solutions, 'Broadridge Announces Integrated Infrastructure for Tokenized Securities,' May 2026.
7. World Economic Forum / Accenture, 'Tokenization in Financial Services: Unlocking the Next Wave of Innovation,' 2024–2025.
8. International Monetary Fund, Working Paper on Tokenized Finance, April 2026.
9. Bank for International Settlements, various publications on DLT in financial market infrastructure.
10. Mayer Brown, 'SEC Staff No-Action Relief for DTC Pilot Provides a Pathway for Tokenized Securities,' December 2025.
11. Digital Asset, Canton Network documentation. Available at: Canton.network.
12. Celent, various publications on post-trade technology and digital asset operations, 2024–2026.
13. ZTAOP Break Taxonomy

## Author

**Kamlesh Kosare**
Global Head of Business Consulting,
Banking & Financial Services - Capital Markets,
Industry Solutions Group (ISG)

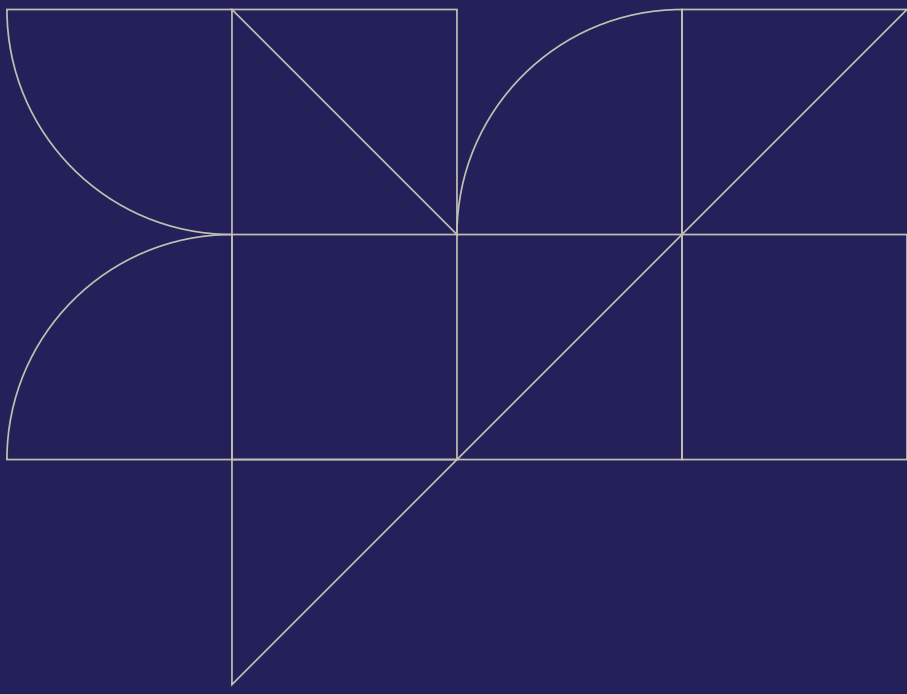

zensar
An RPG Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145+ leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the $4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: **info@zensar.com | www.zensar.com**