

TSA Exit by Design: Separating Systems, Data and Operations without Day-Two Failure

A board framework for transition-service scope, migration evidence, operational resilience and controlled cutover

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

A transition services agreement can preserve continuity after a separation, yet continuity at completion does not prove independence. The recipient remains dependent on another company for systems, data, people, controls or operating capacity. The provider carries service risk, retained infrastructure and possible stranded cost. A weak exit programme converts a temporary bridge into an expensive and fragile operating model.

This paper develops a board-level method for designing TSA exit from the target operating model backwards. It connects service scope, application and data ownership, identity, cyber security, privacy, financial control, supplier rights, operational resilience, migration evidence, cutover and economics. It treats exit as a series of service-specific acceptance decisions rather than one contractual expiry date. The framework uses impact tolerances, evidence gates, dual-run tests and an explicit fallback design so that commercial urgency does not outrun operational readiness.

A wholly hypothetical industrial-technology separation begins with forty-two TSA services and an eighteen-month contractual envelope. Management targets substantial exit by month twelve. The assumed programme has USD 74 million of annualised opening TSA charges, USD 128 million of one-time separation expenditure and a severe-but-plausible cutover scenario that can consume USD 31 million of liquidity before remediation. All companies, amounts, service counts, rates, timings and outcomes in the worked case are hypothetical. A live separation requires company-specific legal, accounting, tax, regulatory, technology, cyber-security, privacy, employment and operational advice.

JEL Classification: G34, L22, L86, M15, O33

Keywords: transition services agreement, TSA exit, demerger, carve-out, data migration, operational resilience, systems separation, cutover

1. Define the exit decision

A board does not approve TSA exit merely because a contract reaches its scheduled end. It approves the transfer of operating responsibility from a temporary provider to a recipient-controlled capability. The decision requires evidence that the business can continue serving customers, collecting cash, meeting regulatory obligations, protecting data and producing reliable financial information after the service ends.

Each TSA service therefore needs a defined recipient outcome. Payroll exit means employees are paid accurately and on time from the recipient's authorised system. Finance exit means opening balances, master data, interfaces, controls and reporting are reconciled. Customer-platform exit means orders, entitlements, billing and support work without an unapproved dependency on the provider. A technical deployment without the full operating outcome is incomplete.

The board should govern a portfolio of exit decisions. Services differ in criticality, architecture, data sensitivity, change complexity and fallback options. A low-risk reporting service can exit through a simple transfer. A

tightly integrated identity, manufacturing, treasury or customer service may require a controlled dual run, independent recovery capability and a board-approved cutover window.

The approval unit should be small enough to expose risk and large enough to represent a complete service. Approving an application alone can miss manual work, data feeds and controls. Approving an entire function at once can hide one unsafe dependency among many complete activities. Service outcomes provide a practical middle level for governance.

The exit decision should state the service, owner, impact tolerance, acceptance evidence, residual dependency, contingency, maximum rollback period and financial consequence. This creates a record that can withstand operational, audit and investor scrutiny.

2. Understand what a TSA does and does not solve

A TSA allocates temporary responsibilities after legal completion. It can preserve access to people, systems, facilities, data processing and operational routines while the recipient builds or procures replacements. It can also define service levels, pricing, change control, incident handling, liability and termination.

The agreement does not create the recipient's end-state capability. It may preserve a legacy configuration that was designed for an integrated group. It may exclude projects, enhancements, new markets, security changes or regulatory work. Service levels can reflect reasonable efforts rather than a commercial managed-service standard. Provider personnel may prioritise the retained business when resources are constrained.

The contractual schedule can conceal technical coupling. One named service may rely on several applications, interfaces, databases, licences, accounts and teams. A provider may need access to recipient data after the apparent service exit because another service remains active. A receiving system may technically operate while its data lineage, reconciliation or recovery process remains incomplete.

The operating programme must therefore decompose the TSA into capabilities and dependencies. Contractual expiry remains an important constraint, but readiness is proved through the recipient operating model and tested evidence.

Both parties should maintain one interpretation of the schedule. Disputes often arise when a recipient views an activity as included and the provider views it as project work or an omitted service. A controlled catalogue, decision log and change process reduce ambiguity before it affects continuity. Commercial disagreements should be escalated without delaying urgent operational protection.

3. Build a service and dependency map

The exit baseline should list every service, recipient, provider, service owner, business process, application, interface, data set, identity domain, facility, supplier, control and jurisdiction. The map should include services supplied in both directions and informal support that may not appear in the signed schedule.

Mapping should begin with important business services and customer outcomes. The FCA requires in-scope firms to identify the people, processes, technology, facilities and information needed to deliver important business services, including relevant third parties. Its March 2026 observations emphasise dynamic mapping, governance and quantitative impact measures alongside time-based tolerances. [1] These principles provide a useful design discipline even where the transaction parties are outside the FCA perimeter.

Dependencies should be directional. A billing platform may depend on customer master data from one system, pricing from another, identity services from the provider and a bank interface controlled by the recipient. The exit sequence must respect those directions. Retiring identity before dependent applications are migrated can create immediate failure.

Each dependency should record the source of evidence and its confidence. Architecture documents can be stale. Configuration scans, access logs, interface monitoring, contract records and reconciled data flows provide stronger evidence. Unknown dependencies should be treated as programme risks with discovery actions and owners.

The map should also distinguish hard and soft dependencies. A hard dependency prevents service operation, such as authentication or a required data feed. A soft dependency reduces efficiency or assurance, such as a reporting tool that can be replaced temporarily by a controlled manual process. This distinction supports sequence, contingency and funding decisions.

Table 1. Proposed TSA exit service register

Field	Required record	Evidence	Exit use	Failure signal
Service outcome	Customer or control result delivered	Process map and owner approval	Defines acceptance	Activity listed without outcome
Dependency	System, data, person, supplier or facility	Scan, log, contract or interview	Determines sequence	Undocumented shared component
Impact tolerance	Maximum tolerable disruption and loss	Risk approval and scenario test	Sets cutover limits	Generic severity label only
End-state capability	Recipient-owned or contracted replacement	Design, build record and contract	Proves independence	TSA copied without redesign
Exit evidence	Test, reconciliation and control result	Signed evidence pack	Supports approval	Project status used as proof
Fallback	Rollback, manual work-around or extension	Tested recovery plan	Limits downside	Expiry date is the only response

Original framework. The service register should reconcile to the signed agreement, operating map and technology inventory.

4. Design the target operating model first

Exit design should start with the operating model required after the TSA. The recipient must decide which capabilities it will own, outsource, share under a durable commercial arrangement or discontinue. This decision controls architecture, people, contracts, data and cost.

A copy of the provider's organisation can be excessive or incomplete. The separated business may have different products, jurisdictions, customers and reporting obligations. It may choose a cloud platform instead of a replicated data centre, a managed security service instead of an internal team, or regional operations instead of a group hub. These choices change both the migration path and the control environment.

The target model should identify accountable executives, process owners, system owners, data owners and control owners. Responsibility for a service cannot sit with a project office after exit. The enduring organisation needs budget, competence, access and escalation rights.

The model should include normal operations, peak volumes, incidents, month-end, year-end, regulatory reporting and disaster recovery. A replacement that works during a quiet test period can fail at quarter-end or during a customer incident. Capacity and resilience therefore belong in the design baseline.

Design authority should remain connected to the transaction thesis. A separation intended to create a more focused and agile business can be undermined if the recipient inherits every legacy process. Conversely, aggressive simplification can remove controls or capabilities that investors assumed would exist. Operating choices should reconcile to the financial case and disclosed strategy.

5. Translate the contract into an exit architecture

The TSA should be converted into a service-by-service control sheet. Scope, exclusions, volumes, service levels, charges, duration, extension rights, change rules, incident obligations, audit rights, data terms, intellectual-property rights and termination assistance should be visible beside the operating plan.

Recent public agreements show the range of structures. Kenvue's amended TSA with Johnson & Johnson describes a general service-period deadline of twenty-four months, with a defined extension where regulatory approvals delay transition. [2] The Jacobs and Amentum TSA filed in 2024 includes an administrative fee and formal service schedules. [3] Western Digital disclosed that transition support for Sandisk covered twelve functional areas for periods up to eighteen months, with mechanisms for additions, extensions, termination, governance and dispute resolution. [4] These documents are transaction-specific evidence of contractual design, not universal benchmarks.

The control sheet should identify the last date for notice, the price of extension, the process for omitted services and the consequences of partial exit. A programme that discovers an extension requirement after the notice deadline loses negotiating leverage.

Service-level obligations need measurable definitions. Terms such as materially consistent, reasonable assistance or normal course can be appropriate contractual standards but provide weak programme metrics. The operating plan should translate them into volumes, response times, recovery objectives, evidence retention and escalation thresholds without implying rights that the agreement does not grant.

Contract and build milestones should be linked. Supplier selection, licence transfer, data extraction, testing and cutover must finish before the contractual stop or an approved extension. Legal teams should receive evidence of technical progress early enough to exercise rights.

6. Treat data separation as a controlled transaction

Data separation is more than file movement. The parties must determine what data belongs to the recipient, what the provider may retain, what must be restricted, which records are shared and how historical context will be preserved. The result should support operations, rights, audit, litigation, tax, privacy and regulatory obligations.

The data map should cover source, owner, purpose, lawful basis, jurisdiction, classification, retention, quality, lineage, transformation and destination. It should distinguish structured records, documents, messages, logs, models, backups and derived data. Shared tables and data lakes often require row-level or attribute-level separation rather than a simple database copy.

The UK Information Commissioner's Office states that data sharing following a merger or acquisition should form part of due diligence, that data-protection principles and documentation apply, and that technical advice is needed where different systems create risks of loss, corruption or degradation. [5] Those issues also arise in demergers because control and processing responsibilities change.

Migration evidence should include extraction totals, transformation rules, rejection logs, control totals, sample verification, reconciliation to financial or operational records, security validation and business-owner acceptance. Deletion or retention by the provider should be evidenced separately. A successful import does not prove complete or lawful separation.

Historical data can create a difficult trade-off between operating usefulness and migration burden. The recipient may need detailed history for customer service, warranty, model performance, tax or litigation. Moving every record can increase cost, privacy exposure and testing. A documented archive-access solution can be suitable where ownership, access, retention, retrieval time and eventual disposition are clear.

7. Separate identity and access without creating blind spots

Identity is a critical dependency because it controls users, service accounts, privileged access, applications and data. The recipient needs an independent identity authority, joiner-mover-leaver process, authentication policy, privileged-access control and emergency-access procedure before dependent services exit.

NIST's zero-trust architecture removes implicit trust based on network location or asset ownership and requires authentication and authorisation before access to enterprise resources. [6] In a separation, this means that

inherited network reach or parent credentials should not become the permanent access model. User, device, service and application identities need explicit policies.

Identity migration should distinguish workforce users, customers, suppliers, robots, interfaces, databases, certificates, keys and API clients. Service accounts are frequently overlooked because they do not appear in employee lists. Expired certificates or unrotated keys can cause delayed failure after a seemingly successful cutover.

The parties should reduce standing cross-company access as services exit. Access logs should be monitored during transition, and residual provider access should have a named purpose, expiry and owner. Break-glass access should be tested and independently reviewed.

Privileged access needs separate governance because administrators can change configurations, extract data or disable controls. The recipient should establish its own privileged-access vault, approval workflow, session logging and emergency process. Shared administrator credentials should be retired. Where vendor personnel retain access, contractual authority and technical enforcement should agree.

8. Sequence applications, infrastructure and interfaces

Applications should be grouped by business service and dependency chain rather than migrated as an unrelated list. The programme should identify systems of record, systems of engagement, analytics, integration, infrastructure, monitoring, backup and recovery.


Four broad exit patterns are available. The recipient can clone a segregated instance, migrate to an existing platform, implement a new platform or retain a durable third-party service. Each pattern has different data, licence, control and timing implications. A clone can be fast but preserve technical debt. A new platform can improve the end state but increase implementation risk.

Interfaces require particular discipline. A system can pass standalone tests while failing when real upstream timing, data quality or downstream acknowledgements are introduced. Interface inventory should include direction, frequency, protocol, schema, authentication, error handling, volume and business owner.

Infrastructure decisions should address networks, cloud accounts, domains, devices, monitoring, batch scheduling, storage, backup and recovery. The recipient should own observability before cutover so that it can diagnose failure without relying on the provider.

Decommissioning should be planned beside migration. Duplicate interfaces, dormant accounts, temporary network routes and abandoned environments expand cost and risk. Each exit work package should state what the provider will retire, what the recipient will retain and how both parties will confirm that no required record or service is lost.

Figure 1. Proposed TSA exit dependency architecture



Exit evidence must cover the full operating chain

Original framework. Exit sequence should follow service outcomes and directional dependencies.

9. Build cyber security into the separation perimeter

Separation changes the attack surface. New domains, networks, cloud accounts, remote connections, data transfers and suppliers are introduced while teams are under delivery pressure. Temporary exceptions can become persistent vulnerabilities if they are not recorded and closed.

NIST CSF 2.0 organises cyber-risk outcomes around Govern, Identify, Protect, Detect, Respond and Recover. [7] The framework is useful for assessing both the transition state and the recipient's end state. Asset inventory, access control, data security, platform security, monitoring, incident response and recovery should be tested within the separation programme.

CISA's cross-sector performance goals identify a prioritised baseline of practices for organisations and critical infrastructure, including identity protection, backups and other high-impact controls. [8] A live programme should select controls appropriate to sector, threat and regulatory obligations rather than treating a generic checklist as sufficient.

The recipient needs its own incident command, contact lists, logging, detection, vulnerability management, backup and recovery. The provider and recipient also need a joint incident protocol while the TSA remains active. The protocol should establish decision rights, evidence preservation, regulator and customer communications, cost and post-incident review.

Public-company obligations can shorten the decision window. The SEC's 2023 cyber rules require disclosure of a material incident generally within four business days after materiality is determined and annual information about cyber-risk management, strategy and governance. [12] Separation governance should route facts quickly to the legal and disclosure teams without allowing disclosure considerations to interfere with containment and recovery.

10. Connect exit to privacy, records and legal hold

Personal data, confidential information, legal records and intellectual property require explicit handling. The separation agreement, TSA, data-processing terms and local law should align on controller and processor roles, instructions, sub-processors, locations, incident notice, retention, audit and deletion.

The data team should not assume that every historical record can be copied. Purpose limitation, confidentiality, contractual restrictions, bank secrecy, health information and export controls can constrain transfer. Some records may need redaction, segregation, pseudonymisation or controlled access.

Legal hold and investigation data require continuity. The parties should preserve searchability, chain of custody and responsible ownership. Deleting provider copies before confirming recipient completeness can impair obligations; indefinite retention creates privacy and confidentiality exposure.

Exit evidence should include a data-transfer record, unresolved exceptions, retention schedule, provider deletion certificate where appropriate and business-owner approval. These records should remain accessible after programme closure.

11. Rebuild finance and control capability

Finance separation affects customer and supplier masters, chart of accounts, bank accounts, treasury, tax, payroll, fixed assets, consolidation, planning, reporting and internal control. Technical migration should be reconciled to opening balances and transaction populations.

The recipient needs a close calendar and control matrix for the first independent reporting periods. Interfaces between operational and financial systems should be tested with representative volumes, currencies, taxes, cut-off events, credits and exceptions. Manual work-arounds should have owners, capacity limits and review controls.

IFRS 5 requires separate presentation of certain assets and liabilities classified as held for sale and the results of discontinued operations. [9] Applicable reporting will depend on the transaction and jurisdiction, but the operational separation plan should support the accounting perimeter and traceability of information.

Control testing should cover access, segregation of duties, master-data changes, journal approval, reconciliations, revenue, purchasing, payroll, cash and reporting. A clean technology test without financial reconciliation cannot support finance exit.

12. Protect operational continuity

Operational continuity should be expressed through service outcomes and impact tolerances. A time-based recovery target is useful, but it may not capture transaction backlogs, customer harm, safety, market integrity, financial loss or regulatory deadlines.

The FCA distinguishes impact tolerance from recovery time and encourages additional metrics such as customer categories, transaction values, volumes and estimated losses. [10] A separation programme can apply the same logic to define maximum outage, maximum unreconciled transactions, maximum data loss, maximum customer backlog and maximum manual-processing duration.

Scenario tests should be severe but plausible. Examples include failed data load, identity outage, corrupted interface, unavailable provider expert, supplier delay, cyber incident during cutover, month-end failure and rollback after partial processing. Tests should include decision-makers and communications, not only technical teams.

Continuity evidence should show that the recipient can remain within approved tolerances, recover the service and process the backlog. A system restored after six hours can still create unacceptable harm if three days are required to reconcile missed transactions.

Table 2. Proposed service acceptance evidence

Evidence domain	Minimum proof	Quantitative measure	Accountable owner	Exit blocker
Process	End-to-end scenario completed	Success rate and backlog clearance	Business service owner	Critical step lacks capacity

Evidence domain	Minimum proof	Quantitative measure	Accountable owner	Exit blocker
Data	Population and control totals reconciled	Completeness, accuracy and rejected records	Data owner	Material unexplained variance
Technology	Capacity, monitoring and recovery tested	Availability, latency, recovery and data loss	Technology owner	Recovery exceeds tolerance
Control	Key controls operated with evidence	Exceptions and remediation closure	Control owner	Financial or regulatory control fails
People	Roles staffed and access approved	Coverage, training and escalation response	Functional executive	Single unmitigated dependency
Supplier	Contract, support and termination rights active	Service levels and unresolved obligations	Commercial owner	Required consent or licence absent

Original framework. Evidence should be proportionate to service criticality and jurisdiction.

13. Secure supplier and licence rights

The recipient can build a technically sound platform and still be unable to operate it because contracts, licences or consents remain with the provider. Supplier diligence should identify assignability, change-of-control terms, user metrics, territorial rights, minimum commitments, audit rights, data terms, support and termination.

New contracts should be effective before cutover and should cover implementation as well as steady-state service. A vendor may agree to production support but exclude migration defects. The recipient should understand whether the provider or vendor holds configuration knowledge and whether documentation can be transferred.

Commercial concentration can shift during separation. A provider that previously represented a small part of group spend may become critical to the recipient. Financial diligence, resilience, security, subcontracting and termination plans should be recalibrated for the recipient's dependency. Contract signature should not substitute for operational onboarding and testing.

Licence metrics should be tested against the target model. Named users, processors, transactions, revenue, devices, environments and affiliates can produce different costs. Parallel environments during migration may require temporary licences that do not appear in the end-state budget.

Supplier exit and concentration risk should be assessed. DORA requires financial entities using ICT services for critical or important functions to maintain comprehensive, documented, tested and periodically reviewed exit plans that permit exit without business disruption, regulatory impairment or detriment to service continuity and quality. [11] The principle is directly relevant to replacement providers selected during a demerger.

14. Transfer knowledge and decision authority

Service delivery depends on tacit knowledge, exception handling and decision authority. Documentation alone rarely captures why a process deviates, which customer needs special treatment or how an ageing system is recovered.

The programme should identify critical roles, named experts, decision rights, recurring cycles, known defects, supplier contacts and escalation paths. Knowledge transfer should use observation, paired operation, reverse shadowing and recipient-led execution. Attendance at a training session is weak evidence; successful recipient-led service under realistic conditions is stronger.

Retention arrangements may be needed for provider and recipient personnel. Their purpose, duration, milestone and cost should be explicit. Reliance on one individual should trigger a succession or external-support plan.

Decision authority should transfer before the person who historically exercised it leaves. The recipient needs approved policies, delegated authorities, bank mandates, system roles and regulatory appointments. A capable team without authority cannot operate independently.

15. Define exit tests before build completion

Acceptance tests should be designed early because they shape architecture and evidence. The service owner should define what must be true for exit, the data required, the scenario, the tolerance and the approver.

Testing should progress from components to interfaces, end-to-end processes, performance, security, recovery and operational rehearsal. Representative production-like data should be used lawfully and securely. Test environments should reflect volumes, configuration and dependencies closely enough to support the conclusion.

Defects need severity, owner, target date and retest evidence. A waiver should state the residual risk, duration, compensating control and approval. High-severity defects should not disappear inside an average pass rate.

The final evidence pack should include requirements traceability, results, reconciliations, defects, waivers, capacity, resilience, access approval, operating procedures, training, supplier readiness and business-owner acceptance. Project completion percentage is not acceptance evidence.

16. Engineer cutover and rollback


Cutover converts tested capability into live responsibility. The runbook should specify sequence, entry criteria, data freeze, extraction, migration, validation, interface activation, business checks, communications, decision points, rollback and command structure.

Each step needs a named operator, expected duration, evidence and latest safe completion time. Dependencies should be visible on one integrated plan. Teams should rehearse the runbook and measure actual duration rather than relying on estimates.

Rollback must be technically and operationally possible. If transactions are processed in the new environment, returning to the provider can require data synchronisation and accounting decisions. The rollback point may therefore occur before the full business test is complete. The board should understand when the decision becomes irreversible.

Post-cutover stabilisation should include enhanced monitoring, daily reconciliation, issue triage, vendor presence and senior decision coverage. Exit is complete after the service operates reliably and the provider dependency is removed or formally bounded.

Figure 2. Proposed eighteen-month TSA exit burn-down



Original scenario. Service counts and timing are wholly hypothetical.

17. Govern the exit portfolio

Governance should combine transaction, business, technology, risk and finance perspectives. The board or delegated transaction committee approves risk appetite, funding, material waivers, extensions and final exit for critical services.

An executive steering committee should review the integrated dependency map, milestones, cost, risks and decisions. Service owners should approve requirements and evidence. A separation management office should maintain configuration control, schedule, cross-workstream dependencies and reporting.

Independent challenge is valuable for material services. Internal audit, risk, cyber security, privacy, finance control or external specialists can test whether evidence supports the claimed readiness. Independence should be proportionate and should not remove accountability from management.

Reporting should show outcomes, not activity. Useful measures include services exited, critical dependencies closed, tests passed within tolerance, unresolved severe defects, data reconciliation, supplier readiness, extension exposure, cash spent and residual provider access.

Governance should control baseline changes. Scope, architecture, cutover date and acceptance criteria can change as facts emerge. Each material change should record the reason, cost, risk, dependency and approver. This prevents a late scope reduction from being reported as delivery progress and preserves an auditable explanation of the final outcome.

18. Model TSA economics and incentives

TSA pricing can use cost recovery, cost-plus, fixed fees, unit rates or other negotiated mechanisms. The recipient should compare temporary charges with the cost of replacement and exit. A low TSA charge can reduce urgency even when the provider carries increasing risk and stranded cost.

The provider's cost can become stranded as recipient volumes decline. Shared licences, infrastructure and teams may not scale down in line with charges. The provider therefore needs a resource-removal plan tied to service exits.

Extension pricing should recognise incremental effort and risk without creating a coercive structure. Automatic price increases can motivate exit, but they can also encourage premature cutover if readiness governance is weak. Extension approval should remain an explicit risk and value decision.

The economic model should include TSA fees, build expenditure, dual-run cost, termination cost, stranded provider cost, delay, contingency and operational downside. EBITDA presentation and cash funding should remain separate.

19. Fund the exit and protect liquidity

Separation expenditure is often front-loaded while benefit arrives later. The recipient may pay TSA fees, implementation suppliers, new licences, duplicated infrastructure, retention and working capital at the same time.

The funding plan should map committed and forecast cash by month, currency and legal entity. It should include tax, deposits, prepayments, capital expenditure, operating expense and contingency. Contractual commitments should be distinguished from management estimates.

Operational failure can create liquidity pressure through delayed billing, lost sales, customer compensation, remediation, emergency support and regulatory consequences. The severe-but-plausible scenario should be funded, not merely described.

Liquidity gates should use minimum cash and headroom thresholds. If the downside breaches the approved floor, management should resize scope, add funding, sequence differently or negotiate a bounded extension.

20. Apply the framework to a hypothetical separation

Consider a hypothetical industrial-technology group separating a digital-services business with USD 1.25 billion of annual revenue. At completion, the provider supplies forty-two TSA services across technology, finance, human resources, procurement, facilities, legal, data and operations. Twelve services support critical customer or control outcomes.

The contractual envelope is eighteen months. Management targets exit from thirty-five services by month twelve, leaving seven bounded services for the final period. Opening annualised TSA charges are USD 74 million. The assumed end-state recurring service cost is USD 69 million after redesign and sourcing.

The one-time separation budget is USD 128 million: USD 52 million for applications and data, USD 24 million for infrastructure and cyber security, USD 18 million for operating-model and control work, USD 14 million for people and knowledge transfer, USD 12 million for dual run and cutover, and USD 8 million for contingency.

The programme identifies five high-risk dependency chains: identity, customer billing, product entitlements, finance close and service monitoring. Each chain receives an impact tolerance, end-to-end test, fallback and executive owner. All values, timings and results in this case are assumptions created solely to demonstrate the method.

Table 3. Hypothetical TSA exit economics

Item	Opening or base case	Month 12	End state	Decision use
TSA services remaining	42	7	0	Dependency burn-down
Critical services remaining	12	3	0	Board attention
Annualised TSA charges	74	16	0	Temporary earnings and cash

Item	Opening or base case	Month 12	End state	Decision use
Annualised replacement cost	0	58	69	Sustainable cost base
Cumulative separation cash	0	111	128	Funding requirement
Provider annualised stranded cost	39	17	6	Resource-removal programme
Unresolved severe defects	19	3	0	Readiness gate

Original scenario. All amounts are assumed USD millions and do not represent a forecast or market benchmark.

21. Test the hypothetical downside


The base case assumes a controlled weekend cutover for customer billing and identity at month ten. The service remains inside a four-hour customer-access tolerance and a twelve-hour billing-recovery tolerance. Reconciliations are completed before the next collection file.

The severe-but-plausible case assumes an identity configuration error, delayed rollback and a corrupted outbound billing interface. Customer access is impaired for eighteen hours, billing is delayed for five days and emergency remediation is required. The assumed cash effect is USD 31 million before recovery: USD 17 million of delayed collections, USD 6 million of lost or credited revenue, USD 5 million of remediation and USD 3 million of other working-capital and communication costs.

The scenario does not assign a probability. It tests whether controls, fallback, communications and liquidity can absorb a defined event. Management may choose a lower-risk sequence, an additional rehearsal or a bounded extension if evidence does not support the original cutover.

The decision should compare delay cost with failure exposure. A three-month extension assumed at USD 6 million of additional TSA fees and USD 4 million of dual-run cost can be rational if it closes a credible USD 31 million liquidity exposure and protects customers. The comparison remains company-specific.

Figure 3. Hypothetical cutover liquidity bridge under severe operational stress



Original scenario. All values are assumed USD millions and exclude probability weighting.

22. Use a risk heat map and decision gates


Risk should combine consequence and evidence weakness. A high-consequence service with unproven recovery requires further work even if the implementation date is on schedule. Bubble size can represent cash exposure, customer population or another material measure.

The board should approve entry into final cutover only when critical dependencies are known, severe defects are closed or explicitly accepted, impact tolerances are met, liquidity remains above threshold and rollback is viable. A red service should not be averaged away by many green services.

Gate breaches should lead to defined responses: remediate, resequence, reduce scope, add funding, extend a service or change the operating model. The response should be chosen before commercial pressure peaks.

After exit, the programme should verify provider access removal, data disposition, resource release, actual run-rate cost and stable service performance. This closure gate prevents operational and financial residue from surviving the formal programme.

Figure 4. Proposed TSA exit risk heat map



Original framework. Position and bubble size are hypothetical and should be replaced with transaction evidence.

Table 4. Proposed TSA exit approval gates

Gate	Required evidence	Principal decision	Failure signal	Management response
Architecture freeze	Service and dependency map, target model and contracts	Approve exit route and sequence	Critical dependency remains unknown	Extend discovery and resequence
Build readiness	Configured capability, supplier rights, staffed ownership	Authorise integrated testing	Required licence, role or control absent	Remediate before testing
Cutover readiness	End-to-end tests, reconciliation, recovery and liquidity	Authorise live cutover	Tolerance breached or rollback unproved	Delay, reduce scope or extend TSA

Gate	Required evidence	Principal decision	Failure signal	Management response
Stabilisation	Service performance, issue closure and control operation	End heightened support	Persistent severe incident or backlog	Retain command structure and funding
TSA termination	Recipient independence and provider release evidence	Terminate service and access	Residual operational dependency	Approve bounded support with exit date
Programme closure	Data disposition, cost run rate and stranded-cost evidence	Close programme accountability	Savings or access exist only on paper	Maintain ownership and reporting

Original framework. Thresholds should reflect the business, sector, jurisdictions and approved risk appetite.

23. Execute a phased roadmap

Phase one establishes governance, service inventory, impact tolerances, contractual deadlines and discovery. The programme reconciles the signed schedules with actual support and identifies critical dependency chains.

Phase two defines the target operating model, architecture, data perimeter, supplier strategy, organisation and control environment. It converts each service into a funded work package with acceptance evidence.

Phase three builds and configures the capability. Data is cleansed and rehearsed, interfaces are established, identities are prepared, contracts become effective and operating procedures are written. Component testing begins early.

Phase four performs integrated, performance, security, recovery and operational testing. Recipient teams lead the service, defects are closed and the cutover runbook is rehearsed. The board receives a service-specific readiness record.

Phase five cuts over in controlled waves, stabilises service, reconciles data and closes residual access. The provider releases resources as evidence permits. Actual cost, performance and incidents are compared with the approved case.

The roadmap should remain dynamic. A newly discovered dependency can change sequence without changing the end objective. Governance quality is demonstrated by timely evidence-based changes rather than adherence to an obsolete date.

24. Conclusion

TSA exit is an operating transfer supported by a contract. The recipient must control the people, processes, technology, information, suppliers, controls and funding required to deliver each service. The provider must be able to remove access, infrastructure and resources without harming its retained business.

The strongest programmes design from the end-state service outcome backwards. They map dependencies, define impact tolerances, build measurable acceptance evidence, rehearse severe scenarios and preserve a viable fallback. They treat data, identity, finance, cyber security and supplier rights as operating requirements rather than technical appendices.

They also maintain commercial discipline throughout delivery. Every extension, waiver and scope change is assessed against customer continuity, legal obligations, funding and the transaction case. Actual performance after cutover is measured against the approved design, allowing management to correct cost, capacity or control gaps before they become embedded in the new organisation.

The hypothetical case shows how an eighteen-month contractual envelope can support a twelve-month management target while preserving a controlled final period. It also shows that a cutover problem can

consume material liquidity even when the underlying service is eventually restored. These values are assumptions, not forecasts.

Board confidence depends on evidence at service level. A programme can report high completion while one critical dependency remains unsafe. Exit should occur when the recipient can operate inside approved tolerances, the provider can end its obligations cleanly and both parties understand the residual financial and operational risk.

Frequently asked questions

What is the principal purpose of a transition services agreement?

A TSA preserves defined services for a limited period after a transaction so the recipient can continue operating while it builds or procures replacement capability. Scope, service levels, pricing, data, incidents, changes, duration and termination should be stated explicitly.

Why should TSA exit be designed from the target operating model backwards?

The target model defines the capabilities, owners, systems, controls, suppliers and costs required after transition. Starting with the expiry date can produce activity without proving that the recipient can deliver the service independently.

What evidence supports termination of a critical TSA service?

Evidence can include end-to-end process tests, data reconciliation, access approval, capacity and recovery results, control operation, trained ownership, supplier readiness, cutover rehearsal, fallback and business-owner acceptance.

How should data migration be verified?

The parties should reconcile population and control totals, transformation rules, rejected records, security, business samples and source-to-destination lineage. Provider retention or deletion and recipient ownership should be evidenced separately.

What is an impact tolerance in a separation programme?

It is the maximum disruption or loss the business can tolerate before unacceptable harm occurs. It can include time, transaction backlog, data loss, customer impact, financial loss, safety or regulatory metrics.

When is a TSA extension appropriate?

An extension can be appropriate when remaining risk from cutover exceeds the cost and risk of bounded continued service. Approval should state scope, price, controls, remediation, owner and final exit date.

How should rollback be designed?

Rollback should specify the latest safe decision point, data synchronisation, transaction treatment, authority, communications and provider readiness. It should be rehearsed under realistic conditions before cutover.

When is the TSA exit programme complete?

Completion requires stable recipient operation, terminated or bounded provider obligations, removed residual access, confirmed data disposition, reconciled costs and evidence that provider resources and stranded costs have been addressed.

References

[1] Financial Conduct Authority, Operational resilience: insights and observations one year on, published 27 March 2026, accessed 16 September 2026. <https://www.fca.org.uk/publications/good-and-poor-practice/operational-resilience-insights-observations-one-year>

- [2] US Securities and Exchange Commission, Johnson & Johnson and Kenvue Form of Transition Services Agreement, Exhibit 10.10, filed 2024, accessed 16 September 2026. <https://www.sec.gov/Archives/edgar/data/1944048/000194404824000057/exhibit1010-formoftransiti.htm>
- [3] US Securities and Exchange Commission, Jacobs Solutions and Amentum Transition Services Agreement, Exhibit 10.2, dated 27 September 2024, accessed 16 September 2026. <https://www.sec.gov/Archives/edgar/data/2011286/000095015724001361/ex10-2.htm>
- [4] US Securities and Exchange Commission, Western Digital Form 8-K concerning the Sandisk separation and Transition Services Agreement, filed 21 February 2025, accessed 16 September 2026. <https://www.sec.gov/Archives/edgar/data/106040/000119312525033383/d847507d8k.htm>
- [5] Information Commissioner's Office, Due diligence when sharing data following mergers and acquisitions, accessed 16 September 2026. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/data-sharing-a-code-of-practice/due-diligence/>
- [6] National Institute of Standards and Technology, Special Publication 800-207 Zero Trust Architecture, published August 2020, accessed 16 September 2026. <https://csrc.nist.gov/pubs/sp/800/207/final>
- [7] National Institute of Standards and Technology, Cybersecurity Framework 2.0, published February 2024, accessed 16 September 2026. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>
- [8] Cybersecurity and Infrastructure Security Agency, Cross-Sector Cybersecurity Performance Goals, accessed 16 September 2026. <https://www.cisa.gov/cybersecurity-performance-goals>
- [9] IFRS Foundation, IFRS 5 Non-current Assets Held for Sale and Discontinued Operations, accessed 16 September 2026. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-5-non-current-assets-held-for-sale-and-discontinued-operations/>
- [10] Financial Conduct Authority, Operational resilience: insights and observations for firms, published 28 May 2024, accessed 16 September 2026. <https://www.fca.org.uk/firms/operational-resilience/insights-observations>
- [11] European Union, Regulation EU 2022/2554 on digital operational resilience for the financial sector, Article 28, Official Journal 27 December 2022, accessed 16 September 2026. <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>
- [12] US Securities and Exchange Commission, Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure, Release 33-11216, effective 5 September 2023, accessed 16 September 2026. <https://www.sec.gov/rule-release/33-11216>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and closure.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners, he is Managing Partner and leads the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

He is also an active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.