practus
Results. Delivered.



Technology Carve-Outs and Separation

Why Most IT Carve-Outs Quietly Lose Value After Day 1

And How to Avoid Getting Stuck in the TSA



Ameya Waingankar

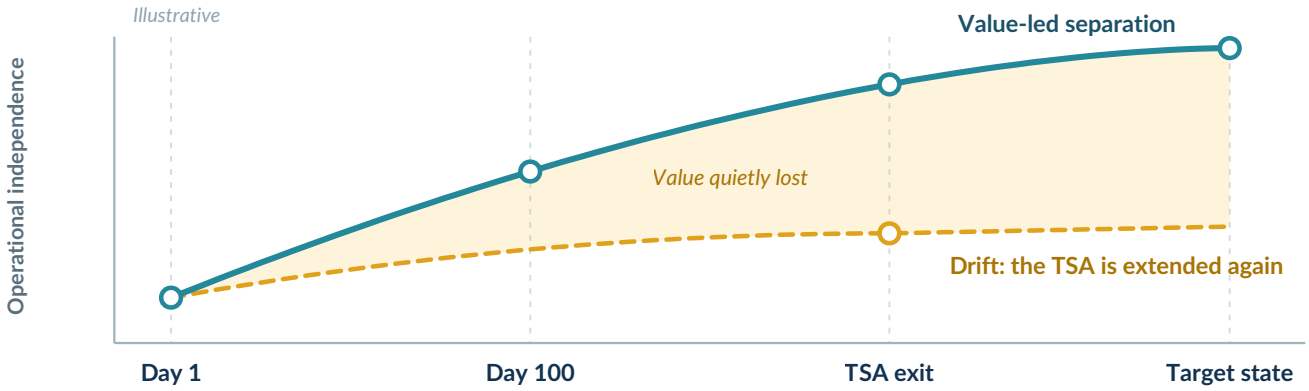
Engagement Partner
Performance Improvement

www.roibypractus.com

Executive Summary

In a technology carve-out, Day 1 is often treated as the finish line. By that point, the entity has been legally separated, staff can log into email, customers can still place orders, finance can still close the books, and networks and applications are running. On the surface, the deal appears to have cleared its biggest operational hurdle. In reality, Day 1 marks the start of a longer period, one in which value is either protected or slowly eroded.

Where value is protected, and where it erodes



A large number of carve-outs achieve a smooth Day 1 but stay tied to the seller's technology environment for months or even years afterward. Applications keep running on shared platforms, data stays lodged inside the parent's systems, licences remain unresolved, and infrastructure continues to be delivered through loosely defined service arrangements. The Transitional Service Agreement, or TSA, is meant to bridge this gap. Left unmanaged, that bridge can turn into a permanent home.

This paper's central argument is straightforward: a technology carve-out should not be treated purely as an IT disentanglement exercise. It should be approached as a strategic reset for the business. The carve-outs that succeed use the separation as an opportunity to build a leaner, faster, more scalable technology model.

What the successful ones do

Begin with the future business in view

Define only the capability truly needed for Day 1

Plan the Day 2 operating model early

Treat every TSA as a time-bound program with a defined exit and clear value targets

Day 1 Is an Operational Milestone, Not a Value Milestone

Getting to a working Day 1 matters. It proves the business can keep functioning once ownership changes hands. But Day 1 usually answers just one narrow question: can the business keep running once ownership changes?

A Day 1 environment is, by design, often imperfect. Continuity is prioritized over optimization, a reasonable trade-off in the moment, but the danger is letting those temporary compromises harden into permanent habits.

Any single item on that list might be tolerable as a stopgap. Taken together, they can leave a business that is legally its own entity but operationally still tethered to its former parent.

> *That gap between legal independence and operational independence is where post-close value quietly disappears.*

It does not answer the questions that matter more

- Can the business run on its own?
- Is the technology cost base right-sized for a standalone company?
- Can leadership actually deliver its value-creation plan?
- Can the business make technology calls without asking the seller first?
- Will the systems scale as the business grows?
- Can the buyer walk away from the TSA without another disruptive event?

At Day 1, a business may still be running on

Still provided by the former parent

- Email and identity services
- Shared ERP platforms
- Data centres or cloud tenants
- Cybersecurity monitoring
- Application support
- Parent-held software licences
- Manual data extracts and reconciliations
- Informal help from the parent's employees

What the deal created

NewCo

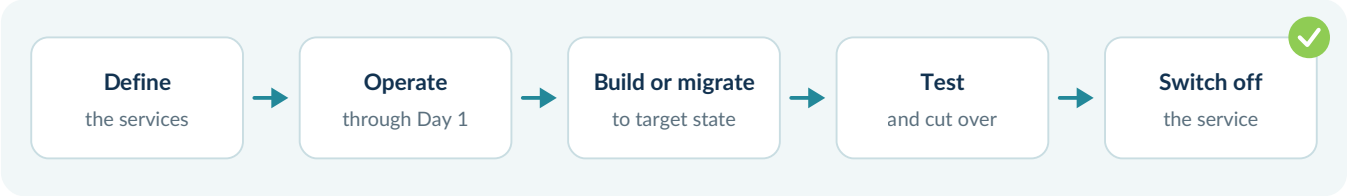
Legally independent

Operationally still tethered

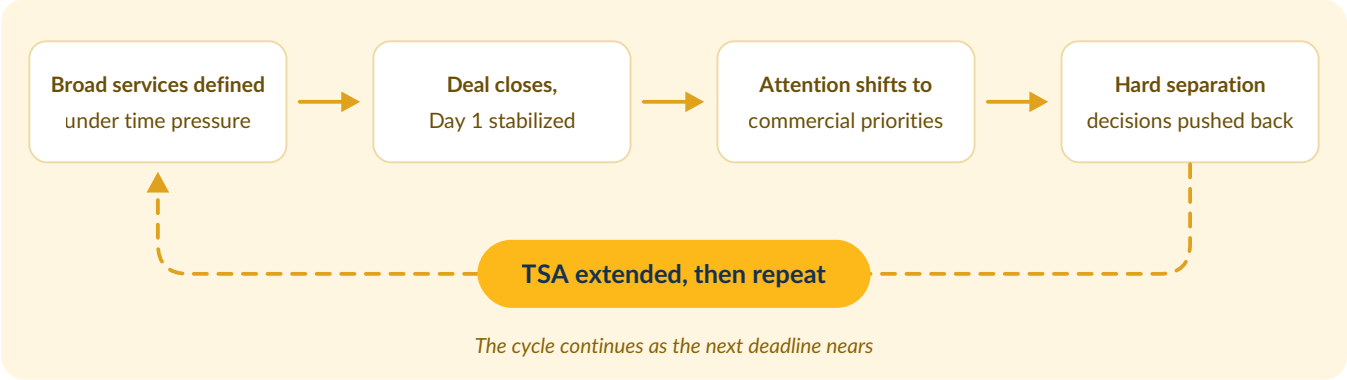
The TSA Trap

A TSA is meant to keep the business running while the buyer builds the capabilities it needs to stand alone. In theory it follows a clean path. In practice, many programs follow a different one.

The intended path



What often happens instead



Extensions are rarely triggered by one dramatic failure

More often they are the sum of several unresolved issues:

- An undefined target-state architecture
- TSA scope that is too broad or vague
- Exit milestones not tied to objective readiness criteria
- Underestimated application and data dependencies
- Vendor consents or licence-transfer requirements surfacing late
- Unclear data ownership and quality
- Diverging buyer and seller priorities
- Weak governance
- Under-resourced separation teams
- Late testing
- Treating TSA exit as an IT project instead of a business transformation

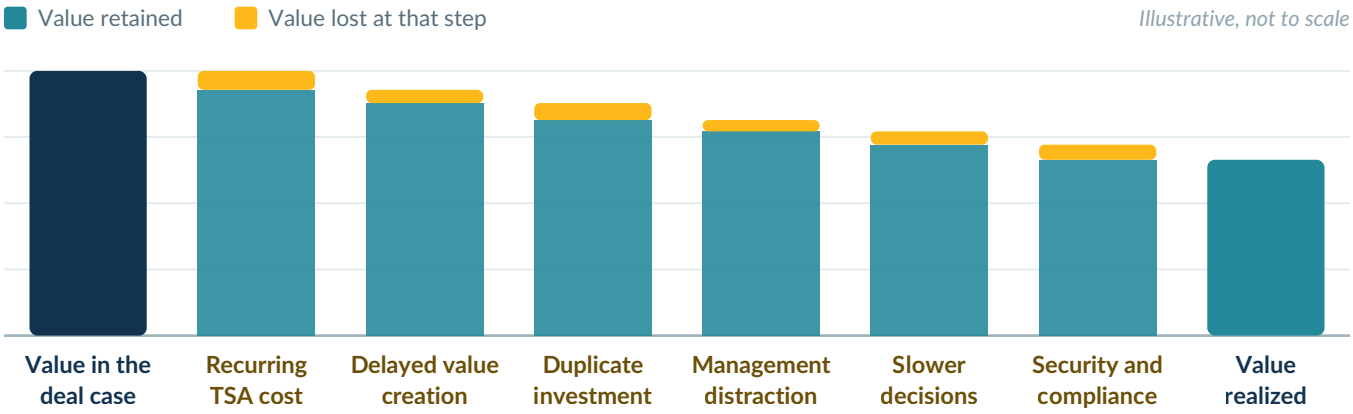
A TSA can hide risk rather than remove it

The business looks stable because the parent is still quietly absorbing complexity on its behalf, but that support comes at a price.

The buyer may be paying for services that lack transparency and don't fit the standalone business, while the seller keeps carrying stranded infrastructure and people costs. Both sides can end up stuck operating a model built for a company that no longer exists.

How Value Leaks After Day 1

The erosion of value after Day 1 is usually gradual and hard to spot. It rarely shows up as a single line item in the deal model; instead it appears as higher running costs, delayed initiatives, duplicated work, and leadership attention pulled in the wrong direction, across six areas.



Recurring TSA cost

Service fees, extension charges, one-time restart costs when an exit slips, parallel-run costs, and duplicated support, on top of the delayed earnings and productivity gains the exit was meant to unlock.

Delayed value creation

Architecture decisions needing extra sign-off, restricted access to data, integration options staying limited, and hesitation to invest in platforms that will soon be replaced anyway.

Duplicate investment

Interim data warehouses, temporary reporting platforms, and parallel support teams that get discarded once the real target state finally goes live, so the business effectively pays twice.

Management distraction

The CEO, CFO, CIO, and COO spending time on service-level disputes, billing questions, access issues, and TSA extension decisions instead of customers and growth.

Slower decisions

The buyer not fully controlling release schedules, security policy, infrastructure standards, vendor selection, or incident prioritization, so the business depends on the seller's process as well as its technology.

Security and compliance exposure

Unclear ownership of the access-control model, vulnerability remediation, incident escalation, logical data separation, and whether the seller can still reach systems once the TSA ends.

Treating Separation as a Strategic Reset

A technology separation is one of the rare moments when an organization can question assumptions it inherited rather than chose, since ownership, governance, reporting requirements, leadership priorities, cost expectations, and the technology perimeter are all already changing. That makes a carve-out a genuine opportunity to redesign, not just copy what already exists.

Replication-led

How do we recreate the seller's environment for the business being carved out?

Value-led

What does this business actually need, technologically, to run its strategy on its own?

The answer is often smaller, simpler, more cloud-based, more automated, and more modular than what the parent runs today, favoring SaaS platforms over heavily customized legacy applications, cloud infrastructure over parent-owned data centres, and a leaner internal IT team backed by managed services and standardized controls.

The lean greenfield principle

Build a fresh environment wherever it adds value, while keeping or recreating only the capabilities that are genuinely essential. For every major system, the organization should make a deliberate, explicit call:

Retain

Strategically important, fit for purpose, and can be separated efficiently.

Replicate

Essential, but the seller's instance cannot be shared long term.

Replace

A simpler or more scalable option better fits the standalone strategy.

Re-platform

Still needed, but should move to a modern architecture.

Retire

Redundant, low-value, or no longer required.

Access via TSA

A short-term dependency is necessary and has a defined exit path.

Designing the Standalone Technology Model

The standalone model should not be dictated by the current architecture alone. It should be shaped by the business strategy, the logic of the deal, and where the business expects to grow.

Before touching system design, be clear on

- **What the business is trying to achieve over the next three to five years**
- **Whether the priority is growth, margin, expansion, resilience, or getting exit-ready**
- **Which capabilities actually differentiate the business**
- **Which are purely transactional and can be standardized**
- **What needs to be scalable from day one**

Day 1

Operate legally, safely, and commercially

- Identity and access
- Email and collaboration
- Network connectivity
- Cybersecurity monitoring
- Core business applications
- Finance and reporting
- Payroll and HR technology
- Business continuity and disaster recovery

Day 1 scope should not be mistaken for the final target state; it is the minimum viable operating platform, not the whole transformation agenda.

Day 2 model

What independence actually looks like in practice

- Target applications
- Data ownership and migration
- Infrastructure and cloud strategy
- The security operating model
- IT organization and skills
- Sourcing strategy
- Vendor and licensing model
- The sequence for exiting the TSA

Without a Day 2 design in place, there is no way to tell whether a given TSA service is genuinely temporary, strategically important, or simply the path of least resistance.

Organize the Roadmap Around Business Capability

The roadmap should be organized around business capability and exit outcomes rather than isolated technical tasks.

Capability	Current dependency	Target state	Exit milestone
Order management	Shared ERP	Standalone ERP or SaaS platform	Orders processed independently, interfaces validated
Employee productivity	Seller's identity and collaboration tools	Independent identity and collaboration environment	All users migrated, access controls validated
Management reporting	Parent's data warehouse	Standalone reporting and analytics platform	Certified reports without parent extracts
Cybersecurity	Seller's SOC and tooling	Buyer-managed or outsourced security operations	Independent monitoring and incident response, tested
Vendor management	Parent's contracts	Contracts re-papered or replaced	Critical suppliers assigned to NewCo

Each exit milestone is an objective test rather than a date: the standalone environment is built, data migration is complete, interfaces are tested, security controls are validated, and the business owner has signed off.

The Separation Blueprint

A disciplined separation program typically moves through five connected phases, each producing a concrete output the next phase depends on.

1

Diagnose the entanglement

Entanglement map

Assess applications, infrastructure, networks, data, identity, cybersecurity controls, licences, contracts, and personnel dependencies. The output is an entanglement map, not just an application list, showing what transfers to NewCo, what stays with RemainCo, and what must be replicated, replaced, retired, or covered temporarily under a TSA.

2

Define the target state

Realistic target state

Set business and technology principles, enterprise architecture, application disposition, the data and integration model, infrastructure, security architecture, IT operating model, sourcing strategy, and governance. The target state must be commercially realistic within the deal's constraints, not just architecturally elegant.

3

Build the roadmap

Phase-gated roadmap

Tie every major activity to a responsible owner, target date, readiness criterion, dependency, cost estimate, risk rating, and cutover decision, phase-gated through scope confirmation, design approval, build, testing, business readiness, cutover approval, hypercare, and TSA termination.

4

Execute an integrated business-and-IT cutover

Tested cutover

Test whole business processes, not just individual systems, including order-to-cash, procure-to-pay, record-to-report, and month-end close, backed by multiple test cycles, data reconciliation, security testing, dress rehearsals, and a time-bound hypercare model.

5

Stabilize and optimize

Formal handover

Confirm actual run-rate costs, close temporary support arrangements, decommission redundant infrastructure, optimize cloud and licence consumption, embed service-management processes, and measure benefits against the original business case, with a formal handover into the permanent NewCo operating model.

Building a TSA That Exits

A TSA should be designed as a mechanism for exiting dependency, not as a general-purpose support agreement. Every service schedule within it should answer four questions:

- What exactly is being provided?
- At what service level?
- For how long?
- What needs to happen for this service to end?

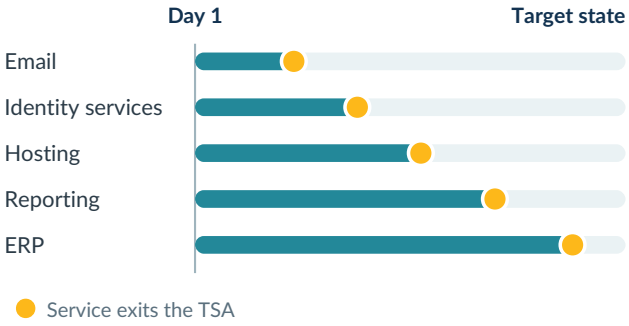
A date by itself does not guarantee readiness

Each service should have objective exit criteria, linked to a detailed exit roadmap from day one:

- The standalone environment has been built
- Data migration is complete
- Interfaces have been tested
- Security controls have been validated
- The business owner has signed off

Not every service has to end on the same day

Staggering exits reduces dependency and frees up management capacity, so a delay in one workstream does not drag out every service.



Commercial discipline pushes toward exit

- **Fees that decline over time**
- **Rights to terminate early**
- **Pre-agreed extension rates**
- **Transparency on cost**
- **Service credits when commitments are missed**

Each service schedule should specify

- Systems and locations covered
- Users and volumes
- Service levels
- Security responsibilities
- Change-control process
- Pricing and invoicing
- Exit activities
- Knowledge-transfer requirements
- Termination rights

Governance That Prevents Drift

Exiting a TSA needs governance that can make decisions without delay.



Two in a box

Each workstream should have an accountable owner from both the buyer and the seller side, a two-in-a-box model that improves visibility and speeds up resolution of dependencies, reducing the risk of each side assuming the other one owns a problem.

The recurring question

It should not simply be whether things are on track. It should be: what has to be true for this service to end, and what is currently standing in the way?

Every week, each TSA service is reviewed against

- Current status
- Exit date
- Readiness criteria
- Risks
- Cost to complete

Measuring What Matters

Standard IT metrics such as uptime, ticket closure, and SLA compliance remain useful but are not sufficient on their own for a carve-out. Progress should be tracked across four dimensions.

Continuity
Is the business still running?

- Availability of critical business processes
- Major incidents
- Transaction success rate
- Disruption experienced by customers and suppliers

Separation progress
How far is the business from exit?

- TSA services exited
- Percentage of users migrated
- Applications separated
- Data migrated and reconciled
- Contracts transferred or replaced

Financial performance
What is the separation costing?

- TSA spend against budget
- Extension costs avoided
- Standalone run-rate cost
- One-time separation cost
- Duplicate or stranded cost eliminated

Strategic value
Is the business becoming more capable?

- Speed of decision-making
- Time to deploy new capabilities
- Automation adoption
- Cybersecurity maturity
- Progress against the value-creation plan

A technology carve-out succeeds when the business becomes genuinely more capable, not merely when the project is declared closed.

Common Failure Modes

Eight patterns appear repeatedly in technology carve-outs. Each has a practical fix.

The failure mode	The fix
Treating Day 1 as the finish line, so stabilization gets mistaken for independence.	Set separate milestones for Day 1, Day 100, TSA exit, and the target state.
Copying the parent's architecture wholesale, so legacy complexity simply transfers into the new organization.	Make an explicit retain, replace, re-platform, retire, or TSA decision for every system.
Signing vague TSAs with undefined scope, cost, or exit responsibilities.	Build measurable service schedules with named owners, clear assumptions, and defined exit criteria.
Planning only after close, so the TSA becomes embedded before any exit plan exists.	Design the target state and TSA exit roadmap before closing, wherever possible.
Separating technology from the business, so technical completion does not guarantee process readiness.	Test end-to-end business processes and run one integrated business-and-IT cutover.
Ignoring contracts and licences, letting vendor consents and change-of-control clauses delay the exit.	Build a licence and contract-transferability matrix early.
Underestimating data, assuming it can be extracted whenever needed.	Define ownership, extraction logic, reconciliation rules, and test cycles from the outset.
Measuring activity instead of value, so a technically busy program still misses its goals.	Link every major workstream to cost, continuity, independence, and value-creation outcomes.

A Practical Playbook for Leaders

What leaders should do, and when, across the four stages of a separation.

Before signing

- Understand the technology perimeter and major dependencies, and estimate standalone operating costs.
- Identify critical vendor and licence constraints and sketch the likely target-state architecture.
- Build separation assumptions, including one-time costs and TSA exposure, into the transaction model.

Signing to closing

- Mobilize the Separation Management Office and confirm the scope of NewCo and RemainCo.
- Define TSA schedules and exit principles, and design the target operating model and cutover plan.
- Secure resources and decision rights, and run integrated testing and readiness reviews.

Day 1 to TSA exit

- Run a dedicated hypercare model and track each TSA service independently against its own exit plan.
- Execute partial exits wherever feasible, and protect the value-creation roadmap from separation delays.
- Decommission redundant services promptly and measure cost and benefit realization along the way.

After TSA exit

- Confirm the target operating model is fully embedded and complete the operational handover.
- Optimize vendors, licences, and cloud consumption, and measure outcomes against the original case.

Who owns it

The CIO owns delivery of the technology; the business owns the outcome.

The Role of Leadership

Technology separation is often delegated entirely to the CIO's organization, which is understandable but not enough. The consequences of a delayed or badly designed separation ripple across the whole enterprise.

The CIO

Owns delivery of the technology.

The business

Owns the outcome.

What a delayed or badly designed separation means for each leader

CFO

Higher costs and delayed synergies

COO

Process disruption and slower execution

CEO

Leadership distracted and value creation pushed back

Board

Higher execution risk

Customers and suppliers

Inconsistent service

Leadership needs to treat the carve-out as an enterprise-wide transformation, not an IT project

That calls for clear, early decisions, communicated consistently across the organization.

Decisions to make early

- ✓ How much continuity is truly necessary
- ✓ How much legacy complexity should be carried forward
- ✓ Which capabilities are genuinely strategic
- ✓ Where the business should buy rather than build
- ✓ How fast independence actually needs to happen

The Carve-Out Advantage

A carve-out creates pressure, but it also creates permission. It gives the organization a reason to revisit technology architecture, the design of the operating model, vendor relationships, security controls, data ownership, organization structure, and investment priorities, changes that in a stable business could take years to get approved.

That urgency needs to be used carefully

The goal is not to turn every carve-out into a sprawling transformation program, nor to add unnecessary complexity to an already risky transition. The goal is to make deliberate choices:

Keep
What is essential.

Simplify
What is not.

Replace
Whatever is holding the strategy back.

Automate
Wherever it buys the business speed.

Outsource
Wherever owning it internally offers no real advantage.

Build
The capabilities the business will need for its next phase of growth.

Done well, a separation does more than protect continuity; it can genuinely improve the business's cost base, decision speed, resilience, and strategic flexibility.

Conclusion

A technology carve-out is not complete when the business survives Day 1. Day 1 proves operations can continue; it does not prove the business is independent, efficient, or ready to create value. That real work happens afterward, when the organization has to separate its data, systems, contracts, people, processes, and decision rights from its former parent.

The TSA is often necessary, but it needs to stay a bridge. If it turns into a permanent operating model, it can quietly drain deal value through ongoing cost, delayed transformation, duplicated investment, weak accountability, and leadership distraction.

The more effective approach starts with

- A clear picture of the standalone business the organization wants to build
- The Day 2 operating model defined early
- The TSA used to protect continuity while creating a genuine path to independence
- Exit managed by service, dependency, and business capability
- Progress measured through cost, continuity, independence, and value created

> *The question leadership should be asking is not whether the business got through Day 1 without an outage. It should be how quickly this separation can be turned into a simpler, faster, more valuable standalone business.*

That difference is what separates simply completing a carve-out from actually creating value through one.