PLAIN OPS CONSULTING, LLC

SAMPLE X-RAY REPORT

IT OPERATING MODEL

Sanitized composite example. Illustrative. Not client specific.

IT Operations X-ray (3-5 weeks, about a 4-week target)

Optional: Execution Leadership (12-16 weeks)

Contact

Nick Vaernhoej, Owner / Principal Consultant

nick@plainopsconsulting.com | plainopsconsulting.com | +1 (605) 595-2135

Plain Ops
CONSULTING

Cut the noise. Fix the work.

CONTENTS

Confidentiality Note (Sample)

This report is a sanitized composite derived from prior operating work.
It is intended to illustrate the structure, evidence discipline, and decision readiness of a Plain Ops diagnostic, not to describe any specific client.

1. Engagement Summary

Duration: ~30 days (3-5 weeks) | Decisions: execution option, owners, targets.

This report is the output of the IT Operations X-ray (3-5 weeks). Execution options below are optional follow-on engagement paths selected at the readout.

Note: Bracketed fields are placeholders shown for structure in this sample.

1.1 Objectives

a) Reduce unplanned work so delivery throughput increases without adding headcount.
b) Reduce repeat incidents and change-driven instability in Tier-1 services.
c) Reduce audit/exam noise by stabilizing evidence capture, ownership, and decision paths.
d) Convert vulnerability/control backlogs into a risk-based queue that produces decisions.

1.2 Measures of Success (Objective-Linked)

Baseline note: Baselines are captured during readout week unless stated otherwise; targets are confirmed at readout.

Objective	Measures	Target timeframe
Objective 1	Unplanned work % (tickets + interrupts)	90 days: reduce unplanned work by 10-20 points
	Decision latency for priority work	60 days: decision latency ≤ 5 business days
Objective 2	Repeat incident rate (Tier-1 services)	6 months: reduce repeats by 30-50%
	Emergency change rate (% of changes)	90 days: reduce emergency changes by 25-40%
Objective 3	Audit/exam prep hours per quarter	12 months: reduce prep hours by 25-40%
	Evidence request response time (business days)	6 months: reduce response time by 30-50%
Objective 4	High/critical vulnerability age % vuln backlog tied	6 months: ≥ 80% within defined SLA
	to known inventory (Tier-1 scope)	90 days: ≥ 95% tied to known inventory

Non-goals: tool selection, RFP support, or full ITSM documentation unless those are confirmed constraints.

1.3 Value (Quantified Where Possible, Conservative Ranges)

Value hypothesis (how value will be quantified)

Time returned to delivery (capacity): measured as the delta in "unplanned work %" multiplied by the team's available capacity hours/month.

- Value formula: (Unplanned% baseline - Unplanned% current) × Team capacity hours/month
- Example (illustrative): If baseline unplanned work is 55% and current is 40%, and team capacity is 1,000 hours/month, then time returned is (0.55 - 0.40) × 1,000 = 150 hours/month.

Repeat incident reduction (stability): measured as the delta in repeat incident rate for Tier-1 services (rolling 30/60/90-day window).

- Value signal: fewer repeats translates to fewer interrupts, less escalation, and lower emergency change rate.

Audit/exam noise reduction (compliance throughput): measured as the delta in audit/exam prep hours per quarter and evidence request response time.

- Value signal: fewer ad hoc evidence chases and fewer cross-team "wait states" during exams.

Quant plan

Baselines captured during readout week; deltas validated at the 30-day metrics review and then monthly (if execution option selected).

Sources

Service desk reports; change calendar; incident postmortems (if any); vulnerability exports; audit request logs; access review artifacts.

Client: ExampleCo
Sponsor: CIO
Period: 4 weeks
Consultant: Plain Ops Consulting

2. Executive Summary

2.1 Environment

ExampleCo is an investor-backed mid-market company operating in a regulated environment. Over the last several years it has:

- Adopted multiple cloud platforms and a new managed service provider.
- Maintained legacy on-prem environments and subsidiary systems.
- Added security responsibilities without a fully defined operating model linking IT Operations, risk, audit, and vendors.

2.2 Engagement Objective

The CIO and leadership team asked Plain Ops to:

- Make the current IT operating model visible at the level needed to explain bottlenecks, risk, and execution drag.
- Identify the 2-3 systemic constraints that drive most instability, audit noise, and unplanned work.
- Separate symptoms (tools, individual findings, incidents) from root causes (operating model, decision authority, foundations).
- Produce a practical 6-12-month roadmap sized to the existing team and vendor reality, designed to run without long-term consultant dependency.

2.3 Top Systemic Constraints

1. Who can decide what is unclear across IT, Security, Risk, and MSP.
    a. Intake/prioritization and “done” criteria are inconsistent; Risk is embedded in operational workflows as an approval gate.
2. Foundations are weak: inventory, configuration knowledge, and access.
    a. Knowledge is distributed and documentation is not authoritative, so triage and evidence work starts with rediscovery.
3. Control work is compliance-driven (activity over stability).
    a. Work optimizes for closing findings/passing exams rather than changing workflows so issues stop recurring.

Leadership and Operating Dynamics (Execution Risk Multipliers)

This engagement also surfaced operating behaviors and cross-functional tensions that increase execution risk (escalation habits, vendor defensiveness, and “urgent-by-default” prioritization). These are captured in Section X: Leadership and Operating Dynamics as mechanisms, not blame.

2.4 High-Level Outcome if the Roadmap is Executed (6-12 Months)

Within 6-12 months, ExampleCo can reasonably expect:

- A basic but durable operating model for intake, change, incident/problem, and control work (with clear owners and decision thresholds).
- Material reduction in unplanned work, duplicated effort, and “reverse-engineering” during incidents and audits.
- Vulnerability and audit backlogs converted into a prioritized, risk-based queue tied to Tier-1 services.
- Security positioned as an operational principle embedded into IT workflows, reducing “compliance theater” and increasing real control effectiveness.

3. Scope, Method, and Evidence Standards

3.1 In Scope

- IT and security practices affecting operational stability and cyber risk.
- Interactions between IT, Enterprise Risk Management, Internal Audit, and the primary managed service provider.
- Use of core security and operations tooling (vulnerability scanning, asset management, remote access, backup/DR, etc.).

3.2 Out of Scope

- Detailed application SDLC and secure coding practices.
- Deep technical penetration testing or control testing for individual systems.

3.3 Method

- Interviews and working sessions with IT leadership, security, risk, internal audit, and key technical staff (including MSP-facing roles) to surface consistent patterns and points of divergence.
- Artifact and evidence review including: org chart (where available), role descriptions, operating cadence documents, ticket samples, change records, vulnerability reports, audit/exam evidence requests, and architecture / concept documents.
- Focused current-state flow mapping for the 2-3 workflows most tied to the primary constraints (for example: change-to-production, incident-to-problem/remediation, and audit finding-to-control remediation). Mapping explicitly targets failure points, rework loops, and “wait states,” and is not a full ITSM documentation exercise unless a specific process is itself the constraint.
- Synthesis into a concise X-ray report with numbered findings tied to evidence, and a 6-12-month roadmap sized to the existing team and vendor reality.

3.4 Evidence Standards

Findings in this X-ray are supported using at least one of the following evidence types:

- Measurable baseline derived from operational data (for example: ticket volumes and aging, cycle time, incident recurrence, change outcomes where available).
- A consistent pattern reported independently across multiple stakeholders, validated against artifacts where possible.
- A repeatable artifact trail (for example: representative tickets, approvals/decision trails, change records, major incident notes/postmortems, and handoff evidence).

This sample report is sanitized and does not include raw client artifacts or identifying references.

In live client deliverables, each numbered finding includes specific evidence references reviewed (for example: export date ranges, representative sample identifiers, and the artifacts used to validate the pattern).

3.5 Accountabilities (Operational, Non-Legal)

These accountabilities describe what materially enabled the diagnostic to produce decision-grade output, and what must remain true if the roadmap is executed.

Plain Ops Accountabilities (What Was Provided)

- Produced the X-ray readout package: current-state operating model view, numbered findings tied to evidence patterns, horizon roadmap, and metric definitions/baseline plan.
- Facilitated interviews and working sessions to surface consistent patterns, resolve points of divergence, and force clarity on who decides what, ownership, and “definition of done.”
- Maintained evidence discipline: findings were anchored to observable patterns across artifacts and stakeholder narratives, with explicit separation of symptoms from constraints.
- Prepared decision prompts and a readout agenda designed to result in named owners, sequencing, and measurable success checks.
- Presented optional execution paths after the readout (advisory vs higher-touch execution), sized to team and vendor reality.

Client Accountabilities (What Enabled Progress)

- Identified an executive sponsor and a clear decision owner for the work system (intake, prioritization, change, incident/problem, and vendor lanes).
- Provided access to SMEs, tooling, and representative artifacts (tickets, change records, incident narratives, delivery tracker impacts, vendor lanes) sufficient to validate patterns.
- Made ownership explicit for the in-scope flows (e.g., incident/problem, change/release, request fulfillment, escalation management, vendor relationship ownership).
- Participated in validation: confirmed what was accurate, where exceptions existed, and what constraints were dominant.
- Made decisions with low latency once recommendations were presented (decision latency is itself a measurable constraint).

Joint Accountabilities (How Work Stayed Decision-Ready)

- Held a short, recurring review focused on decisions and measurable movement (not status theater).
- Agreed on a “stop doing” list to protect capacity for Horizon 1 work and prevent improvement work from becoming “extra.”
- Managed scope by explicit reprioritization when new demands emerged, rather than adding work on top of existing commitments.
- Maintained a single source of truth for decisions, owners, and follow-ups to prevent drift after the readout.

4. Current-State Themes

4.1 Environment and Structure

- Four main technical silos: cloud for the parent entity, cloud for a legacy entity, on-prem/legacy infrastructure, and subsidiary systems managed by partners.

- Multiple parties involved in day-to-day operations: internal IT, a managed service provider, and several vendors.
- A single security specialist acting as de facto CISO, heavily engaged in audits and reactive risk work.

4.2 Observed Pain

- Extensive reverse-engineering of the environment before any meaningful risk decision can be made.
- Overprivileged accounts and unclear reasons for their existence.
- Security and risk activities organized around individual symptoms (a scan result, a finding, a penetration-test issue) rather than around a simple threat-based strategy.
- Tools that consume large amounts of effort (e.g., vulnerability scanners) but produce little operator-grade intelligence.
- Lack of standard procedures and baselines, leading to configuration drift and inconsistent quality.
- Cultural friction between IT, security, risk, and the managed service provider; public criticism and "over-the-wall" behavior are common.

4.3 Current-State Operating Model View

This section makes the work system visible at operational granularity. The point is not to restate symptoms. The point is to show how today's flow mechanics reliably produce those symptoms.

How Work Enters (Observed Channels)

Formal:

- ITSM queues (incidents, requests, problems, changes)
- project/release commitments and vendor deliverables

Informal (side doors):

- executive escalations and direct messages
- "quick favors" and bypass approvals
- vendor pings that skip intake
- shadow queues inside teams

How Work is Prioritized (Observed Pattern)

- Priority is often determined by immediacy and escalation force, not by agreed service impact and decision thresholds.
- Approval-based waiting dominates end-to-end cycle time for non-trivial work.
- The system rewards starting work to demonstrate responsiveness, even when WIP is already overloaded.

Where Flow Breaks (Repeatable Failure Points)

- Intake and routing: inconsistent triage; categories used differently by different teams.
- Ownership: unclear service ownership results in handoffs and reopen loops.
- WIP overload: teams accept more concurrent work than can be completed; aging grows.
- Change enablement: readiness/verification inconsistency increases outages and rework.
- Vendor lanes: acceptance criteria and "definition of done" vary, creating "not us" loops.
- Problem convergence: repeat drivers are visible but do not translate into durable fixes.

Why Better Status Does Not Fix This

Better status improves reporting. Performance improves when:

- intake rules are enforced,
- decision thresholds are explicit,
- WIP limits are real,
- and “definition of done” requires verification and evidence.

5. Findings: Top Systemic Constraints

NOTE (Sample): In live client reports, each finding includes specific references to artifacts reviewed (export dates, sample sizes, representative ticket IDs). This sanitized composite uses representative patterns.

5.1 Finding 1 (Constraint): Who Decides What Is Unclear Across Teams and Vendors

Description

IT and Security operate as loosely connected groups rather than as one system with a shared work model. Risk and audit functions are embedded as operational approvers of technical work. There is no single steering forum that prioritizes work across projects, risk remediation, and run-the-business, and no consistent definition of what “done” means.

Evidence

- Change processes require Risk as a formal approver for technical tasks, causing delay and confusion about who owns outcomes.
- Security is engaged reactively to individual vulnerabilities or audit findings with short timelines, rather than through a repeatable control workflow.
- There is no consistent mechanism for capturing requirements, prioritizing cross-functional initiatives, or declaring success criteria and closure.

Impact

- Conflicting priorities across IT, Security, Risk, and vendors; each group chases its own queue.
- High resource cost to complete basic changes because of unclear roles and repeated rework.
- Initiatives stall or disengage before requirements are fully met, leaving half-implemented solutions in production.

5.2 Finding 2 (Constraint): Inventory, Configuration, and Access Foundations are Weak

Description

The organization lacks an authoritative view of what it owns (systems, IP ranges, domains) and how those systems are configured. Configuration knowledge is tribal. Access is routinely granted beyond what is required, with limited documentation and review.

Evidence

- Asset and vulnerability management depend on each other for accurate data, resulting in neither having complete coverage.

- Documentation that exists is often stale; in practice the production environment is treated as the only source of truth.
- IT and Risk cannot quickly answer basic scoping questions for testing and evidence (for example: external IP ranges, relevant domains).
- Many accounts have broad administrative access with no clear justification or record of purpose.

Impact

- Incident and vulnerability triage starts with “what is this and where is it” instead of assessing impact and containment.
- Significant operational debt: time is repeatedly spent rediscovering known information.
- Over-privileged access increases blast radius and weakens the ability to evidence control effectiveness.

5.3 Finding 3 (Constraint): Control Work is Compliance-Driven (Activity Over Stability)

Description

Security and control work is organized around satisfying external frameworks and passing audits rather than around a coherent internal operating model. Tools were acquired to check boxes but are not integrated into daily workflows; they generate reports and queues without improving cycle time, stability, or readiness.

Evidence

- Vulnerability tooling consumes time to generate static reports that add limited operational value.
- Tool configurations are not tied to asset lifecycle processes; licenses and coverage drift because deprovisioning is not embedded.
- Risk mitigation is framed as “close this finding by the due date” rather than “change the workflow so this stops recurring.”

Impact

- High effort spent satisfying evidence requests without reducing actual exposure or operational drag.
- Security is perceived as noise and friction, not as an enabler of reliable delivery.
- Foundational gaps persist (inventory, access, DR discipline) while resources are diverted to low-leverage tasks.

Note: Leadership and operating dynamics that materially affect execution are captured separately in Section X: Leadership and Operating Dynamics (mechanisms, not blame).

6. Section X: Leadership and Operating Dynamics

6.1 Why This Section Exists

Operating model change requires behavior change: enforcing intake, saying “no,” and holding gates.

This section captures dynamics that materially affect execution and indicates where leadership behavior must change.

6.2 Scope and Method

Evidence Sources

- Interviews: Ten leaders/managers across IT, Security, Risk/Internal Audit, and MSP-facing operations (weeks 1-3).
- Artifacts reviewed (where available): org chart, role expectations, steering/cadence notes, and examples of escalation threads.

Note on Interpretation

- Findings reflect consistent patterns across interviews and observed operating behaviors during the engagement.
- Divergent views are explicitly called out where they materially affect execution.

6.3 Consensus Observations (What Multiple Stakeholders Independently Reported)

Strengths (Consistently Observed)

- Senior leadership recognizes that instability is structural (operating model and foundations), not "a few bad tickets" (evidence: consistent interview themes across IT and Risk).
- Strong individual contributors exist in both IT and Security, but they are trapped in reactive mode (evidence: repeated references to "reverse engineering" and urgent queues).

Friction Points (Consistently Observed)

- Escalation habits and public criticism drive defensiveness (especially with the MSP) and reduce candid problem disclosure (evidence: interview examples + recurring meeting behavior).
- Who can approve what is unclear; Risk is pulled into day-to-day approval, increasing wait states and rework (evidence: change workflow and audit-prep patterns).

Points of Divergence (Where Perceptions Differ)

- MSP view: incoming tickets and requirements are inconsistent, causing churn and delays.
- Internal IT view: MSP operates to minimum viable effort and externalizes ownership; issues become "someone else's problem."

Why it matters: without explicit decision thresholds, intake standards, and escalation rules, execution will default back to blame cycles instead of systemic improvement.

6.4 Operating Tensions (The Tradeoffs Driving Behavior)

Tension 1: "Speed to Delivery" vs "Stop-The-Line for Stability"

- How it shows up: urgent work is accepted by default; recurring issues are patched repeatedly instead of fixed through workflow changes.
- Cost if unchanged: persistent unplanned work, audit noise, and missed roadmap delivery; higher burnout and attrition risk.

Tension 2: "Central Control and Approvals" vs "Delegation with Clear Thresholds"

- How it shows up: routine technical work escalates to senior leaders or Risk for approval; teams wait rather than decide within guardrails.
- Cost if unchanged: slow cycle times, repeated rework, and ambiguity about accountability.

6.5 Constraint Statement (The Mechanism, Not the Blame)

Primary Behavioral Constraint

- Decision thresholds and escalation rules are unclear, causing repeated escalation, rework, and wait states that inflate cycle time and keep teams reactive.

Secondary Behavioral Constraints (If Any)

- Vendor relationship defaults to defensiveness because accountability and success measures are not explicit.

6.6 Implications for Execution (Why This Matters to the Roadmap)

Execution Risk Factors

- Roadmap work will stall if steering, decision thresholds, and intake standards are not established early.
- Improvement work will be treated as "extra" rather than as the operating system, sustaining the unplanned-work trap.

Execution Accelerators Already Present

- Clear sponsor intent to reduce noise and regain predictable delivery.
- Willingness to pilot improvements on a constrained Tier-1 scope.

6.7 Recommendations (Paired: Leader Actions + Org/System Actions)

Recommendation 1: Establish Steering, Decision Thresholds, and Escalation Rules

- Leader behavior shifts:
  - Move "finding discussions" into smaller leadership settings; stop public escalation as the default path.
  - Require success criteria and "definition of done" before work is accepted.
- Operating system changes:
  - Weekly cross-functional steering forum (IT, Security, Risk as policy setter, MSP) with a single prioritized queue.
  - Written decision thresholds (what can be decided locally vs escalated).
- Success indicators (30-60 days):
  - Fewer escalations for routine decisions.
  - Reduced cycle time for common change types.
  - Single queue and consistent prioritization accepted by stakeholders.

Recommendation 2: Reset Vendor Operating Accountability (Without Theatrics)

- Leader behavior shifts:
  - Enforce one owner per outcome; stop "someone else's problem" loops.
  - Hold 1:1 vendor leadership conversations for performance issues; avoid performative group criticism.
- Operating system changes:
  - Standardize intake quality (minimum requirements for tickets/requests).
- Define vendor SLAs/SLOs tied to outcomes (cycle time, repeat incidents, backlog age) not just activity.

- Success indicators (30-60 days):
  - Fewer “reopened” issues and fewer handoffs.
  - Improved ticket quality and reduced churn between teams.

7. Recommended Roadmap (6-12 Months)

7.1 Horizon 1: 0-60 Days (Stabilize and Establish Control)

Recommendation: Create a single steering and prioritization mechanism

- Deliverable: Steering forum charter + single prioritized queue (run work, remediation, projects) + WIP limits + “stop doing” list.
- Decision required: Executive sponsor approves governance and WIP limits by [date].
- Success measure: Unplanned work % trend improves within 60 days; decision latency ≤ 5 business days by day 60.

Recommendation: Clarify approval authority and remove Risk as an operational approver

- Deliverable: Decision and escalation map + change approval thresholds (Risk as policy setter, not day-to-day approver).
- Decision required: Sponsor + Risk leader approve thresholds by [date].
- Success measure: Change cycle time shows measurable improvement within 60 days (baseline and target set at readout); fewer blocked changes due to approval ambiguity.

Recommendation: Stabilize inventory and privileged access for a constrained Tier-1 scope

- Deliverable: Tier-1 scope definition + minimum viable inventory (systems, IPs, accounts, data flows) + privileged access criteria + evidencable reviews.
- Decision required: CIO confirms Tier-1 scope by [date].
- Success measure: ≥ 95% of Tier-1 assets represented in inventory by day 60; incident “reverse-engineering” effort decreases measurably by day 60 (baseline and target set at readout).

7.2 Horizon 2: 2-6 Months (Institutionalize the Operating Model)

Recommendation: Define standards and procedures for repeatable work

- Deliverable: Minimum standards pack (asset lifecycle, change, incident/problem, access, hardening) with owners and “done” criteria.
- Decision required: Process owners approve standards by [date].
- Success measure: Ticket reopen rate and rework loops decrease measurably by month 6 (baseline and target set at readout); audit evidence requests handled within agreed target timeframes.

Recommendation: Implement minimal workflow structure in the service desk tool

- Deliverable: Workflow schema updates (required fields, routing rules, closure criteria) aligned to standards.
- Decision required: Service owner approves workflow changes by [date].

- Success measure: Queue aging improves measurably by month 6 (baseline and target set at readout); cycle time for common change types decreases measurably by month 6 (baseline and target set at readout).

Recommendation: Redesign vulnerability management to produce decisions, not reports

- Deliverable: Inventory-driven scan scope + risk-based triage rules + backlog taxonomy + reporting that supports accept/fix/defer decisions.
- Decision required: Security + IT owner agree on triage rules and SLA targets by [date].
- Success measure: High/critical vulnerability age improves toward target by month 6; backlog is ≥ 95% tied to known inventory.

Recommendation: Implement basic recoverability discipline for Tier-1 services

- Deliverable: Recoverability objectives (RTO/RPO targets) + one executed test with remediation backlog.
- Decision required: Sponsor approves recoverability targets by [date].
- Success measure: One test completed by month 6; remediation backlog prioritized and tracked monthly.

7.3 Horizon 3: 6-12 Months (Optimize and Scale)

Recommendation: Establish durable process ownership and operating cadence

- Deliverable: Named owners for core processes + monthly operating review agenda focused on systemic improvements.
- Decision required: CIO assigns owners with time/authority by [date].
- Success measure: Metrics reviewed monthly; Horizon roadmap updated quarterly with observable movement in core measures.

Recommendation: Lock a small set of operator-grade metrics and enforce them

- Deliverable: Metrics definition sheet + baseline capture plan + monthly reporting package.
- Decision required: Sponsor approves metrics and targets by [date].
- Success measure: Trend movement across unplanned work %, repeat incidents, change cycle time, audit prep hours by month 12.

Recommendation: Convert vendor accountability from activity to outcomes

- Deliverable: Vendor scorecard (outcome measures) + operating cadence (monthly) + escalation thresholds.
- Decision required: CIO + vendor owner approve scorecard and thresholds by [date].
- Success measure: Vendor-related wait states decrease measurably by month 12 (baseline and target set at readout); repeat incident drivers attributable to vendor decrease.

8. Post X-Ray Execution Options (Optional Follow-on; Choose One)

Start point: After the X-Ray readout. Options below assume findings, baselines, owners, and targets have been confirmed.

8.1 Option A: Advisory Control (Light)

- Outcome: Metrics and decisions are brought under control with minimal external lift.
- Inclusions:
  - Baselines and targets finalized at readout; tracked weekly for 6-8 weeks.
  - Weekly 30-minute metrics/decisions review for 6-8 weeks.
  - Decision and escalation map finalized and enforced.
  - Horizon 1 sequencing support (what to do first, what to stop doing).
- Client commitments:
  - Sponsor + decision owner named by readout.
  - Process owners assigned by [date].
  - Decisions within 5 business days of recommendation.
- Duration: 6-8 weeks.
- Investment: Defined in SOW (fixed-fee, outcome-aligned).

8.2 Option B: Operating Model Installation (Standard)

- Outcome: Operating model and cadence installed; Horizon 1 and early Horizon 2 delivered with the existing team and vendors.
- Inclusions:
  - Intake/WIP enforcement + single queue operating cadence.
  - Change and incident/problem discipline installed for Tier-1 scope.
  - Vendor accountability reset to outcome measures.
  - Roadmap delivery support through Horizon 1 and select Horizon 2 items.
- Client commitments:
  - Sponsor + process owners in place by readout.
  - Weekly working sessions.
  - Decisions within 5 business days.
- Duration: 12-16 weeks.
- Investment: Defined in SOW (fixed-fee, outcome-aligned).

8.3 Option C: Stabilize + Execute Horizon 1 (Strong)

- Outcome: Rapid stabilization by executing Horizon 1 with measurable reduction in unplanned work, repeat incidents, and change instability.
- Inclusions:
  - Higher-touch execution leadership for Horizon 1.
  - Backlog triage + stop-doing enforcement to protect capacity.
  - Change/release gating + incident/problem stabilization for Tier-1.
  - Metrics instrumentation + weekly reporting to sponsors.
- Client commitments:
  - Same as Option B plus explicit “stop doing” list agreed at readout.
- Duration: 8-12 weeks (Horizon 1), then reassess.
- Investment: Defined in SOW (fixed-fee, outcome-aligned).

8.4 Next Step (Time/Date-Certain)

- Readout scheduled: [Day, Date, Time, Timezone] (set at kickoff; shown here as a placeholder).

- Attendees required:
  - Executive sponsor (decision owner)
  - IT Ops lead
  - Infrastructure lead
  - Security lead
  - Service Desk lead
  - MSP relationship owner

Decisions required at readout

a) Select execution option (A/B/C).
b) Confirm sponsor, decision owner, and process owners for in-scope flows.
c) Approve metric baselines and targets (or confirm baseline capture plan).

Decision deadline: [Date], no later than 5 business days after readout.

8.5 Working Rules (To Prevent Scope Seep)

- Scope control: new deliverables require written approval and reprioritization of existing work.
- Data handling: define storage location, access, and retention for evidence and artifacts.
- Meetings: cadence and required attendees set at kickoff; cancellations reschedule within the same week.
- Communications: single source of truth for decisions and action items (one doc or tracker).
- Change order triggers: expanding beyond the 2-3 critical flows; adding business units; adding tooling selection work; adding deep technical testing.
- Decision latency: decisions not made within 5 business days are treated as a delivery risk and surfaced explicitly.

Appendix A: Evidence and Artifact Request List (Sample)

A1. ITSM and Work Data (Last 60-180 Days, as Available)

- Incident export (volumes, priorities, MTTR, categories, repeat patterns)
- Request export (volumes, aging, approval counts, cycle time)
- Problem export (open/closed, linkage to repeat incidents)
- Change export (standard/normal/emergency, lead time, failure proxies, verification completion)
- Samples: 10-20 representative tickets per critical queue (to validate patterns)

A2. Reliability and Change

- Last 3-5 major incidents with restoration notes/postmortems (if available)
- Last 30-60 days change calendar and sample change records (including emergency)
- Cab agenda/notes (if applicable) and approval patterns
- Post-change verification evidence examples (if any)

A3. Service Ownership and Foundations

- Service catalog (if any) and ownership map (formal or informal)
- Inventory sources used for routing (CMDB, endpoint inventory, cloud inventories)
- On-call schedules and escalation paths
- Vendor/MSP scope, SLAs, and representative ticket samples (if applicable)

Appendix B: Interview Map and Working Session Plan (Sample)

B1. Leadership Interviews (60 Minutes Each)

- Executive sponsor (CIO/COO/CFO)
- IT operations leader
- Delivery/program leader (if separate)
- Security/risk partner (if control work materially affects ops)

B2. Operator Interviews (45-60 Minutes Each)

- Service desk lead and escalation owner
- Incident/major incident lead (if present)
- Change enablement owner / cab chair
- Platform/service owners for top tier-1 services
- Vendor/MSP relationship owner (if applicable)

B3. Working Sessions (90 Minutes Each)

a) Session 1: Incident and escalation flow map (restoration patterns and recurrence)
b) Session 2: Request intake, approvals, and routing flow map
c) Session 3: Change enablement flow map (tiering, gates, verification)
d) Session 4: Support vs delivery collision map (boundaries and stabilization)
e) Session 5: Future-state operating model workshop + roadmap sequencing

Appendix C: Intake, Routing, and Escalation Thresholds Template (Sample)

C1. Intake Categories

- Incident (unplanned interruption)
- Request (standard fulfillment)
- Problem (repeat driver removal)
- Change (planned modification)
- Project/release (bundled delivery outcome)
- Exception (explicit bypass within defined thresholds)

C2. Minimum Required Fields

- Category
- Affected service (tier-1 or not)
- Owner (service owner) and executor
- Urgency/impact and decision threshold reference
- Definition of done and verification requirement

C3. Escalation Thresholds (Examples)

- Tier-1 outage or safety/regulatory risk: immediate escalation path
- All other escalations: route through queue owner; subject to WIP rules

Appendix D: WIP Limits, Queue Taxonomy, and “Definition of Done” (Sample)

D1. Queue Taxonomy (Example)

- Incidents (by tier)
- Requests (standard vs non-standard)
- Problems (repeat drivers)
- Changes (standard/normal/emergency)
- Delivery/release work (by program)

D2. WIP Limits (Example Patterns)

- Per team lane: max concurrent active items
- Per tier-1 service: max concurrent changes
- Per owner: max concurrent escalations

D3. Definition of Done (Examples)

a) Incident “done” = Service restored + restoration notes captured + repeat driver tagged (if applicable)
b) Problem “done” = Driver removed or controlled + verification evidence + runbook/monitoring updated where relevant
c) Change “done” = Implemented + verified + backout retired + evidence captured
d) Delivery “done” = Readiness gate passed + go-live executed + stabilization complete + exit criteria met

Appendix E: Change Enablement and Go-Live Readiness Gates (Sample)

E1. Change Tiering

- Standard: pre-approved patterns
- Normal: readiness evidence + verification required
- Emergency: explicit criteria + mandatory post-change review

E2. Readiness Criteria (Examples)

- Owner and approver identified
- Risk assessed and backout documented
- Affected services/stakeholders notified
- Verification plan defined

E3. Stabilization and Exit Criteria (Examples)

- Incident rate within threshold for window
- Monitoring coverage confirmed
- Known issues documented with owners and dates

Appendix F: Metrics Pack (Sample Definitions and Baselines)

F1. Reliability

- Major incidents per month
- MTTR P1/P2
- Repeat incident categories count

F2. Flow

- Backlog aging distributions by queue
- WIP by lane vs limit
- Percent work entering via intake vs side door

F3. Change

- Emergency change volume
- Verification completion rate
- Change failure proxies (incidents linked to changes)

F4. Delivery

- Slippage rate and reasons
- Stabilization duration vs planned
- Support-driven delivery interrupts

Appendix G: Readout Agenda and Decision Log Template (Sample)

G1. Readout Agenda (60-90 Minutes)

1) Current-state operating model summary
2) Top constraints and evidence patterns
3) Future-state model (owners, intake, WIP, change gates)
4) Roadmap and metrics
5) Decisions required and deadlines

G2. Decision Log Template

- Decision ID
- Decision statement
- Owner
- Options considered
- Decision
- Date decided
- Rationale (1-2 lines)
- Implications (who/what changes)
- Follow-ups and due dates

END OF SAMPLE REPORT