FORTRESS DIAGNOSE · ROADMAP · NEW VERSION

Where the in-depth diagnostic should focus.

After scoping interviews and a 90-day extract. Themes ranked by impact and readiness — not actions. The deep dive is the product. This pack aims it.

SAMPLE · FICTIONAL · LAKESHORE SPECIALTY PARTNERS · 4-LOCATION PLATFORM

Process, people, or technology — named before any tool. Planning ranges only.

What we read. What this pack is not.

WHAT WENT IN

Four scoping interviews (CFO, VP RCM, two site ADs). 90-day denial, A/R, and auth extract. Public HFMA-type ranges. No PHI.

WHAT IT PRODUCES

Themes, ranked. Nature tag: Process / People / Technology. Planning-range band. Recommended focus for the in-depth diagnostic.

WHAT IT IS NOT

Not a workplan. Not specific actions. Not a forecast. Not a tool mandate. No floor walk. Two of four sites interviewed.

QUALITY GATE

A principal signed every number that remains. Two hypotheses were struck: coding variance as the primary leak, and “the clearinghouse is the problem.” One sizing was cut when the aging file did not support it.

This pack is the baseline after scoping. Specific actions wait for the in-depth diagnostic — virtual or on-site. It is not Prove and not Operate.

A queue. An owner still shared. Process first.

TYPE

Queue — Healthcare RCM

Not structure across functions. Not an empty seat. Do not open TOM or Frac on this finding.

OWNER

VP RCM named. Split in practice.

Sites still run local lists. Metric reviewed monthly, not weekly.

PRIMARY METRIC

Days in A/R · 48

Client target 38. Denial $ as % of gross is the second number.

PROCESS 55%

Four playbooks. Rework at the denial handoff. No single clean-claim definition.

PEOPLE 30%

Ownership split. Capacity on firefighting. Huddle exists; actions do not close.

TECHNOLOGY 15%

Systems can do the step. Interface delay, not a missing model. Last, not first.

HOW THE WORK ACTUALLY RUNS

Charge drop same-day at three sites; one batches twice a week. Clean-claim 86%. Denial land has four inboxes, first-pass review 11 days. Coders and billers both touch. Write-off authority unclear under $2k.
Bottleneck is denial land — not the clearinghouse. Nature: process first, people second, technology last.

Ranked for the deep dive. Not a workplan.

#	THEME	NATURE	IMPACT	READY	DIAGNOSTIC FOCUS
T1	**Documentation gaps on medical-necessity denials**	Process	High	High	Intake checklist and coder handoff
T2	**Four location playbooks, no single standard**	Process	High	Med	Map the four variants; find the core
T3	**Auth queue has no daily owner**	People	Med	High	Role, cadence, what the owner sees
T4	**Rework loop: same denial class returns**	Process	Med	Med	Trace one class end to end
T5	**System flag at entry — unused or missing?**	Tech	Low	Low	Hold until T1–T4 are understood

Technology is last by design. If T1–T4 hold after the deep dive, T5 may never need answering.

Planning range on this queue: $1.4–2.1m / year. Conservative $1.4m. Not a forecast. Not a Fortress performance claim.

T1 — Documentation gaps. Evidence, not a fix.

11.8% of claims denied on medical necessity (90-day extract). Peer range 6–9%. Mean days to overturn: 19. Planning range $0.6–0.9m. Confidence: medium-high.

CLIENT NUMBER	**11.8% med-necessity denial rate**	Extract, March 1–May 31, 2026
SOURCE	**Denial file + EHR order notes**	Named files. No PHI in the pack.
BENCHMARK	**6–9% peer range**	Cited public HFMA-type ranges
PATTERN	**Missing clinical detail at order entry — not coder error**	Interview + sample review

Why the themes will not move without a system.

Ownership of the queue — 1 / 4

VP RCM named. Sites still run local lists.

Standard work / playbook — 1 / 4

Four denial paths. No single reason-code standard.

Cadence and huddle — 2 / 4

Monthly close review. No weekly denial stand-up.

Metric and baseline — 3 / 4

Days in A/R and denial $ exist. Used late.

Variance control — 1 / 4

Dashboard updated. Exceptions not closed.

Systems / native tools — 3 / 4

PM and clearinghouse are adequate. Not the constraint.

Read 1.8 / 4. Strongest: systems and the existence of a metric. Weakest: standard work and variance control. A process-and-ownership gap. Do not buy another denial tool — or stand up AI — from this pack.

Questions the in-depth diagnostic must answer.

1	PROCESS	Where does documentation break between order entry and coder — and which step is the single point of failure?
2	PROCESS	What is the common core across the four location playbooks, and what local variation should stay?
3	PEOPLE	Who should own the auth queue daily, and what does that role need to see each morning?
4	PROCESS	Trace one medical-necessity denial end to end. Where does the same class recur, and why?
5	TECH — HOLD	Only after 1–4. Does the current system lack a control, or is the control unused?

07 · NEXT STEP

Scope the in-depth diagnostic — or stop.

THIS MEETING Both	Accept this as the baseline, or say it is wrong. Name the diagnostic focus (T1–T4 recommended).
IF PROCEED Fortress + client	Scope the in-depth diagnostic — virtual or on-site. Process and people first. Technology held to question 5.
WHAT THAT SESSION DOES Fortress	Walk the work. Prove or kill each theme. Return with specific actions, owners, and a Prove / Design / Frac recommendation.
WHAT IT DOES NOT DO Rule	No tool, no vendor, no Shadow Mode starts from this pack. Prove is a later SOW.

OFF-RAMPS

Seat empty → Fractional COO diagnostic. Decision rights across functions → TOM Assessment. No owner and no metric → name those first; do not spend on a deep dive.

Clarity Session was free. This pack is the scoping baseline. The in-depth diagnostic is the next paid step.