business technology consulting inc.

CLIENT CASE STUDY

Building an AI Portfolio for a Regional CPA Firm

Confidential firm · 105 Employees · Central New Jersey

12 Use Case Discovered | **5 Prioritized in Portfolio** | **~2,400 hrs/yr Recoverable**

BTC
Daniel Furrey
Founder

April 2026

Engagement Snapshot

Client	*Confidential*
Sector	Public accounting — tax, audit & advisory
Size	~105 employees across three New Jersey offices
Engagement	AI Portfolio Development — 6-month program
BTC Lead	AI Portfolio Framework (System, not Tools)
Outcome	Governed AI portfolio with 5 active use cases; overall AI Maturity Score advanced from 1.8 → 3.1

Confidential client is a 105-person regional CPA firm serving mid-market businesses, high-net-worth individuals, and non-profit organizations across Central New Jersey. The firm had accumulated a collection of point-in-time AI experiments — a chatbot pilot here, a document summarization tool there — but lacked any coherent strategy to govern, scale, or connect them to firm-wide value. BTC was engaged to transform this fragmented landscape into a structured, governed AI portfolio using our proprietary AI Portfolio Framework.

Phase 1: Assessment

AI Maturity Matrix — Baseline vs Target

BTC conducted a structured 6-week assessment across all seven domains of the BTC AI Maturity Matrix. Leadership interviews, workflow walkthroughs, and a firm-wide AI fluency survey were conducted across tax, audit, advisory, and operations functions.

Domain	Baseline	Target	Key Constraint
AI Strategy	1.5	3.0	*No enterprise vision*
AI Value	1.0	2.5	*No value tracking*
AI Organization	2.0	3.5	*Informal ownership*
AI People & Culture	1.5	3.0	*Low AI fluency*
AI Governance	1.0	2.5	*No Controls or policy*
AI Engineering	2.5	3.5	*Siloed tooling*
AI Data	2.0	3.0	*Unstructured client data*

Overall weighted maturity score advanced from 1.8 → 3.1 over the 6-month engagement, crossing the Maturity Constraint Line in four of seven domains.

Key Assessment Findings

- **Fragmented ownership**: The firm had no single owner for AI decisions. Purchases and experiments were made independently at the practice level with no central inventory or governance.
- **High appetite, low readiness**: Partners expressed enthusiasm for AI-assisted audit and autonomous tax preparation, but both were rated high-risk relative to the firm's current governance maturity.
- **Strong platform foundation**: Existing tooling (Microsoft 365, CCH Axcess, GoSystem) provided untapped AI surface area activatable without new infrastructure.
- **Data accessibility gap**: Client-facing data (returns, workpapers, engagement letters) was largely unstructured and stored inconsistently across SharePoint, local drives, and client portals.
- **Generational adoption divide**: Staff with fewer than 5 years of tenure expressed high willingness to adopt AI tools; partner-level adoption was the primary cultural risk.

Use Case Discovery - Long List (12 Use Cases)

Twelve use cases were surfaced across four functional areas during discovery workshops. Each was assessed on estimated value, implementation complexity, and risk profile before portfolio structuring.

#	Use Case	Function	Value	Complexity	Risk	Decision
1	Autonomous individual tax return preparation	Tax	**Very High**	High	**Very High**	Deferred
2	AI-assisted audit workpaper drafting	Audit	**High**	Medium	**High**	Deferred
3	Client document ingestion & data extraction	Ops	**High**	Medium	**Medium**	✓ **Included**
4	Tax research summarization (IRC/Regs)	Tax	**High**	Low	**Low**	✓ **Included**
5	Engagement letter & proposal drafting	Advisory	**Medium**	Low	**Low**	✓ **Included**
6	Staff scheduling & utilization forecasting	Ops	**Medium**	Medium	**Low**	✓ **Included**
7	Anomaly detection in client GL data	Audit	**High**	High	**High**	**Deferred**
8	Client onboarding Q&A assistant	Ops	**Medium**	Low	**Low**	✓ **Included**
9	Automated tax notice response drafts	Tax	**Medium**	Medium	**Medium**	Deferred
10	Predictive churn/upsell modeling	Advisory	**High**	High	**Medium**	Deferred
11	AI-generated financial benchmarking reports	Advisory	**Medium**	Medium	**Medium**	Deferred
12	Continuing education content generation	Ops	**Low**	Low	**Low**	Deferred

Why High-Value Use Cases Were Deprioritized

A recurring tension in AI portfolio conversations is the gap between perceived value and actual readiness. Three use cases at *Confidential* client were rated as high or very high value by firm leadership — yet all three were deferred after structured risk assessment. The reasons are instructive and apply broadly across CPA firms.

BTC's position: value potential is real, but unrealized risk is equally real. Launching a high-value use case before the governance, data, and organizational maturity exists to support it does not accelerate value — it accelerates exposure.

1. Autonomous tax return preparation — regulatory & liability exposure

- CPA firms carry direct professional liability for the accuracy of signed returns. Current AI models hallucinate numeric values and misapply multi-state filing rules at a rate that cannot be absorbed into a standard review workflow without material re-work.
- IRC Section 6694 preparer penalties apply to the signing CPA regardless of how the output was generated. AI vendors provide no indemnification for incorrect tax positions taken by their models.
- State CPA licensing boards have issued guidance flagging autonomous AI-prepared returns as a potential competency and independence concern under AICPA professional standards.
- MAG's malpractice carrier had not issued a coverage position on AI-prepared returns at engagement start — proceeding without clarity would create an uninsured liability gap.
- Client data was distributed across incompatible systems. Feeding inconsistent inputs into an autonomous workflow compounds error risk nonlinearly and eliminates the ability to trace output back to source.

2. AI-assisted audit workpaper drafting — independence & documentation risk

- PCAOB and AICPA standards require workpapers to reflect the auditor's own professional judgment. AI-generated narratives may not satisfy sufficiency and appropriateness of evidence requirements under AU-C Section 230.
- If AI generates a memo that a staff associate reviews and signs without genuine independent analysis, the documentation trail is inconsistent with actual audit work performed — a regulatory exposure.
- MAG engagement partners acknowledged that AI-generated prose would lower the bar for staff to skip substantive analysis, increasing audit quality risk systemically over time.
- Auditor independence standards restrict use of third-party tools that create a mutuality of interest with the client. Any AI tool that trains on client-specific data requires formal vendor assessment under ET Section 1.200.
- With AI Governance scored at 1.0, MAG had no model versioning, output logging, or audit trail capability — all of which are prerequisites for defensible documentation in any regulatory review.

3. Anomaly detection in client GL data — data quality & false positive risk

- Anomaly detection models trained on incomplete or inconsistently formatted GL data produce high false positive rates. At MAG's AI Data maturity level (2.0), the overhead of managing false positives eliminates projected time savings.
- Communicating AI-flagged anomalies to audit clients without a robust explainability layer creates relationship and liability risk. Clients will demand justification in professional audit terms that the model cannot provide.

- A false negative — failing to flag a material anomaly during an AI-assisted audit, later discovered in litigation — creates a substantial new exposure vector for the firm.
- The use case required a clean, normalized GL data pipeline across multiple client ERP systems: an infrastructure capability MAG did not have and that would require 6+ months of data engineering work to establish before any model could be deployed.

The Maturity Constraint Line

BTC's framework holds that an organization's AI portfolio cannot sustainably exceed its lowest-scoring domain. MAG's AI Governance and AI Data domains scored 1.0 and 2.0 respectively — placing the Maturity Constraint Line at Level 1.5. Any use case requiring Governance or Data maturity above 2.5 was flagged as exceeding the constraint and required prerequisite domain investment before launch.

This is not a permanent deferral. Each deferred use case has a defined re-evaluation trigger tied to a specific domain maturity milestone. As Confidential firm *advances its Governance and Data scores — targeted at Level 3.5 by Q3 2026 — the high-value use cases will be re-sequenced into the active portfolio.*

Phase 2: Portfolio Structure

Five use cases were selected for the active AI portfolio, organized across four tiers of the BTC AI Portfolio Framework: Foundation (data and operations enablers), Productivity (staff time recovery), Insight (decision support), and Client-Facing (experience enhancement). Each use case was validated against *Confidential* client's current maturity scores before inclusion.

ID / Tier	Use Case, Description & Result
UC-03 Operations **Foundation**	**Client document ingestion & data extraction** Automates ingestion of client-supplied documents (W-2s, 1099s, bank statements) into structured data, reducing manual entry across the tax team. **Result**: ~620 hrs/yr recovered; 94% extraction accuracy on structured documents.
UC-04 Tax **Productivity**	**Tax research summarization** AI research assistant deployed over IRC, Treasury Regulations, and IRS guidance. Staff submit queries; the tool returns summarized positions with cited sources. **Result**: Average research time per complex matter reduced from 3.1 hrs to 0.8 hrs. Partners report higher confidence in cited positions.
UC-05 Advisory **Productivity**	**Engagement letter & proposal drafting** AI drafts engagement letters and service proposals from a structured intake form. Partners review and approve. Standardized template library reduced variability across four offices. **Result**: ~280 hrs/yr recovered. Time from intake to signed engagement reduced from 6.2 days to 2.8 days.
UC-06 Operations **Insight**	**Staff scheduling & utilization forecasting** AI model ingests historical utilization, engagement schedules, and PTO data to generate 6-week forward staffing forecasts and flag imbalances proactively. **Result**: Utilization variance reduced 18%. Weekly scheduling meetings reduced from 90-minute sessions to 30-minute exception reviews.
UC-08 Operations **Client-Facing**	**Client onboarding Q&A assistant** AI assistant deployed on the client portal to handle onboarding FAQs, document upload instructions, and status inquiries — escalating to staff when required. **Result**: Inbound client calls during onboarding reduced 34%. Staff report higher quality of first substantive conversations with new clients.

AI Operating Model — Governance Structure

Alongside the use case portfolio, BTC designed MAG's AI Operating Model: a lightweight governance structure fitting within the firm's existing partner/manager hierarchy without creating a new bureaucratic layer.

- **AI Steering Group**: Designated AI Steering Group (Managing Partner, Tax Director, Operations Director) with a monthly portfolio review cadence.
- **Practice Ownership**: Named AI Practice Owners for each active use case, responsible for adoption metrics, output quality monitoring, and escalation.
- **Risk Register**: Centralized AI Risk Register covering model versioning, output logging, vendor assessment, and client data handling policies.
- **Maturity Re-scoring**: Quarterly AI Maturity Matrix re-scoring to track domain advancement and validate re-sequencing of deferred use cases into the active portfolio.

Phase 3: Execution Results

Portfolio Performance at 6 Months

Metric	Baseline	Target	Change
Total recoverable capacity (hrs/yr)	0	**~2,400**	**+2,400 hrs**
AI Maturity Scope (weighted avg.)	1.8	**3.1**	**+1.3 levels**
Domains above Maturity Constraint LIne	0 of 7	**4 of 7**	**+4 domains**
Active governed use cases in portfolio	0	**5**	**5 activated**
Time to signed engagement (days)	6.2	**2.8**	**-55%**
Tax research time per matter (hrs)	3.1	**0.8**	**-74%**
Onboarding inbound call volume	Baseline	**~34%**	**Reduced**
Utilization variance (scheduling)	Baseline	**~18%**	**Reduced**
Partner-reported AI confidence	Low	**Mod-High**	**Advancing**

Client Perspective

> *"We came in thinking AI meant autonomous tax returns. BTC helped us see that the path to those ambitions runs through fundamentals we hadn't built yet — governance, data quality, and organizational readiness. The five use cases we activated are generating real value today, and we now have a clear, risk-calibrated roadmap to the high-impact cases we actually want."*
>
> **— Managing Partner, Confidential firm**

Looking Ahead — Deferred Use Case Roadmap

The deferred high-value use cases are not abandoned — they are sequenced. As MAG's Governance and Data domains advance toward Level 3.5 (targeted Q3 2026), the following re-evaluation triggers are in place:

- **Autonomous tax preparation**: Re-evaluate when AI Governance reaches Level 3.0, a vendor coverage position from the malpractice carrier is obtained, and a partner-led review workflow with documented sign-off standards is in place.
- **AI-assisted audit workpapers**: Re-evaluate when AICPA issues formal AI guidance for workpaper documentation, MAG implements output logging, and a pilot is scoped for a single low-risk assurance engagement class.
- **GL anomaly detection**: Re-evaluate when the client data normalization initiative completes, a false-positive review workflow is defined, and a test dataset with known anomalies is prepared for model validation.

The BTC AI Portfolio Framework — What Made This Work

Our client's outcome was not the result of selecting the right AI tools. It was the result of applying a structured business system — the BTC AI Portfolio Framework — that treated AI as an enterprise capability to be built, not a collection of point solutions to be purchased.

- **System, not tools**: Used the BTC AI Maturity Matrix to establish an honest baseline across all seven domains before recommending any use case.
- **Constraint-aware sequencing**: Applied the Maturity Constraint Line to prevent the portfolio from being built on domains that would inevitably constrain and break it.
- **Heatmap-driven prioritization**: The Execution Heatmap ensured that high-effort, high-risk use cases didn't crowd out the high-readiness, high-return use cases that could deliver value in the near term.
- **Governance by design**: The AI Operating Model provided governance from day one — not as an afterthought once problems arose.
- **Long-horizon thinking**: Deferred use cases are not rejected — they are on a maturity-gated roadmap with defined re-evaluation triggers.

Ready to build your AI portfolio?

Contact BTC to schedule an AI Maturity Assessment for your firm.

BTCExperiences.com · info@btcexperiences.com