

“And here we tinker with metal, to try to give it a kind of life, and suffer those who would scoff at our efforts. But who’s to say that, if intelligence had evolved in some other form in past millennia, the ancestors of these beings would not now scoff at the idea of intelligence residing within meat?”

— Prime Function Aki Zeta-5

“THE FALLACIES OF SELF-AWARENESS” · SID MEIER’S ALPHA CENTAURI

Old instincts can be *dangerous* in new territory.

Let us guide you through.

AI-native transformation of the digital core requires three offerings working as one family. Together they cover the terrain.

SCALED AI FOUNDATION

Architect & build the AI core

AI architecture and engineering for the digital core — where AI participates, agents for the core, knowledge scoped to the problem.

Core Excavation · AI Architecture Foundation · Agentcraft ·
Living Ontologies

AI-READY DATA

Make data work for AI

Resilient, trusted data foundations. Migration that arrives somewhere better; modernization that surfaces the intelligence your data always had.

Data Migration · Data Modernization · Knowledge Engineering ·
Data Managed Services

ADVANCED AI

Reason at industry scale

Foundation for business transformation — enterprise AI platforms, industry intelligence (IDB), assurance, and managed services.

Adv AI Platform · Industry IDB · AI Eval, Assurance & Tuning ·
AI Managed Services

TWELVE SUB-OFFERINGS · POWERED BY SHARED ASSETS & METHODOLOGIES.

Frontier partnership, not competition — we partner with frontier AI labs, we don't compete on foundation models.

IN THE FAMILY AI-ready Data Resilient, Trusted Data Foundations	IN THE FAMILY Advanced AI Foundation for Business Transformation
Data Migration Accelerated Migration · Data Product Factory	Adv AI Platform Arch & Op Model Target State Architecture · Adv AI Platform MVP
Data Modernization Data Products & Strategy · AI for BI (AI4BI) · Data Agents	Industry Intelligent Digital Brain (IDB) IDB Discovery & Baseline · IDB Design & Build · IDB Lifecycle Mgmt
Knowledge Engineering ↔ Living Ontologies Knowledge Agents · Scaled Ontology Creation	Adv AI Eval, Assurance & Tuning → Agentcraft AI Evaluation & Advisory · AI System Optimization · Agent Assurance
Data Managed Services LLM/Model Ops · Data-to-Knowledge Ops · Autonomous BI Ops · AI Lifecycle Mgmt	AI Managed Services ↔ TS&T LLM/Model Ops · Agent Lifecycle Ops · AI Infra Ops · AI Lifecycle Mgmt

FIVE OFFERINGS AS ONE BODY

Five offerings. One body. *Composable — or a chimera.*

The body works because the bio-factory engineers compatible organs from the *same DNA*. Skip the engineering and you don't get an organism — you get a chimera. Stitched-together, immune-incompatible, unlikely to survive.

THE NEW DESIGN DISCIPLINE

Each offering carries its own radical delivery rethink and the way it composes with the other four. Living Ontology + Industry IDB help us write the shared DNA every offering plugs into — nothing is foundational over anything else.

• FIVE OFFERINGS · ONE BODY



AI & DATA · THREE OFFERINGS, ONE COGNITION LAYER

Cognition has three subsystems. *The body needs all three.*

SUBSTRATE · THE CORTEX
SCALED AI FOUNDATION

What the brain reasons *with.*

Architecture, ontology activation, agentcraft — what lets agents observe, decide, and act.

CORE EXCAVATION · AI ARCHITECTURE FOUNDATION · AGENTCRAFT · LIVING ONTOLOGIES

MEMORY · THE HIPPOCAMPUS
AI-READY DATA

What the brain reasons *over.*

Trusted, structured data with an enterprise semantic layer. Without memory the brain has nothing to think about.

AI-ACCELERATED MIGRATION · DATA MODERNIZATION · KNOWLEDGE ENGINEERING · DATA AGENTS

REASONING · THE PREFRONTAL
ADVANCED AI

What the brain reasons *about.*

Industry-specific reasoning — the Intelligent Digital Brain pre-loaded for the client's sector. Certified and assured.

INDUSTRY IDB · ADVANCED AI PLATFORM ARCH · AI EVALUATION & ASSURANCE · AI MANAGED SERVICES

RADICAL DELIVERY RETHINK

BRAIN

The cognition layer becomes *shared infrastructure*, not per-engagement rework.

Living Ontology + Industry IDB are what every other offering plugs into. Second engagement on the same account starts at week one with the first engagement's accumulated domain knowledge already loaded — not at week zero with fresh discovery. The ontology has to be theirs to read and write, not just ours.

OUTCOME Ontology activation cycle compresses from quarters to weeks. Knowledge engineering stops being per-engagement waste.

Five composing practices. One cognitive core. *One running system.*

Five offerings composing against a shared cognitive core, a deterministic substrate where guarantees hold, and agentic surfaces where reasoning is the lift.

PHASE 1

Five practices appear

Each offering claims its position.

PHASE 2

Core densifies

Industry IDB + Living Ontology load.

PHASE 3

Practices compose

Cross-offering edges light up.

PHASE 4

The system runs

Trust gates · service graph live.

Agentic software in symbiosis with people, evolving within the service graph.

PA-01

BANKING

Illustrative

“Nobody understands what the *legacy system* actually does.”

Multinational bank — 25-year-old core banking platform. 15M+ lines of code, thousands of undocumented integrations, critical business logic encoded by developers who left a decade ago. Two prior modernization attempts failed.

What changed: Core Excavation — Code Archaeology and Institutional Archaeology working in tandem. Code Archaeology systematically extracts intelligence across five discovery domains; Institutional Archaeology surfaces tribal knowledge in documents, communications, and process artifacts. In *weeks, not months*: complete dependency maps, hotspot identification, extracted business rules, evidence-based modernization roadmap sequenced by risk.

VALUE LEVERS

- Modernization success rate
- Knowledge preservation before retirement attrition
- Risk reduction through evidence-based sequencing
- Timeline compression — weeks instead of months

OFFERINGS

Core Excavation · AI Architecture Foundation

L4: Code Archaeology · Institutional Archaeology

PA-03

CROSS-INDUSTRY

Illustrative

“Our data can never become *knowledge*.”

Manufacturer with a decade of clean production data. AI models built on it underperform reality, because the schema cleaning stripped out operator judgment, environmental context, and equipment history — the things that actually explain quality outcomes.

What changed: Living Ontologies build a scoped knowledge layer co-evolved with the business problem. Knowledge Engineering Services then scale the semantic foundation across the enterprise. Together they turn raw data into *contextualized enterprise knowledge* — without waiting for the data to be clean.

VALUE LEVERS

- Intelligence from data you already have
- Semantic layer reused across applications
- Tribal knowledge made queryable, not lost to attrition
- Models that reflect operational reality, not stripped abstractions

OFFERINGS

Living Ontologies ↔ Knowledge Engineering Services

Scalpel + Foundation – cross-offering by design

PA-11

HEALTHCARE

Illustrative

“Our data migration is *holding up* everything.”

Healthcare payer post-merger — two claims platforms, two member systems, two provider directories. The 18-month migration plan is two years in. Every downstream initiative (AI, BI, new products) is blocked on data that’s “almost ready.” Delay costs run \$2M+/month.

What changed: AI-Accelerated Migration handles schema mapping, transformation scripts, test generation — engineers focus on *what to keep, restructure, productize*. Data Product Factory means you don’t just migrate, you arrive somewhere better. **18 months → 5 months.** The destination is AI-ready from day one.

VALUE LEVERS

- \$2M+/month delay-cost recovered
- Migration timeline compressed 3-4x
- Destination platform is AI-ready, not just migrated
- Downstream initiatives unblock simultaneously

OFFERINGS

Data Migration · Data Modernization

L4: AI-Accelerated Migration · Data Product Factory

PA-09

INSURANCE · CMT

Illustrative

“We have a data platform but *nobody trusts* the numbers.”

\$30B insurer — data platform investment topped \$200M over five years. The dashboards exist; the LoBs don’t use them. Every meeting starts with a 20-minute debate over whose numbers are right. AI initiatives on top of the platform stall because nobody will sign off on the data lineage.

What changed: Engineered Trust as a discipline — not an audit. Trust Engineering bakes guarantees into the architecture; Advanced AI Evaluation certifies the models running on it; Data Managed Services keeps the lineage live and observable. Result: numbers the C-suite stops debating.

VALUE LEVERS

- Stop relitigating data lineage in every meeting
- Unblock the AI initiatives stuck waiting for sign-off
- Engineered governance, not compliance theater
- Continuous certification, not periodic review

OFFERINGS

Trust Engineering · Adv AI Evaluation · Data Managed Services

Engineered Trust theme — one theme, three expressions

Archetypes → offerings. *How they map.*

A snapshot – every archetype on the full catalog draws from at least two of the three offerings. Cross-offering by design.

ARCHETYPE	SCALED AI FOUNDATION	AI-READY DATA	ADVANCED AI
PA-1 · Legacy Understanding	Core Excavation · AI Arch Foundation	—	—
PA-2 · AI in Transactions	Agentcraft · AI Arch Foundation	—	Adv AI Eval · AI Managed Services
PA-3 · Data → Knowledge	Living Ontologies	Knowledge Engineering · Data Modernization	—
PA-5 · Regulatory Change	Living Ontologies	Knowledge Engineering	Adv AI Eval & Assurance
PA-7 · Readiness Assessment	Core Excavation · AI Arch Foundation	Data Modernization	—
PA-9 · Trust the Numbers	Trust Engineering (L4)	Data Managed Services	Adv AI Evaluation & Assurance
PA-12 · Migration Holdup	—	Data Migration · Data Modernization	—
PA-13 · Compounding Intelligence	Agentcraft · Living Ontologies	Knowledge Engineering	Industry IDB · AI Managed Services

Eight of fourteen archetypes shown. Engineered Trust runs across every row – never a column on its own.

Greater than the *sum of the parts.*

Not a real project — a near-future scenario showing all three offerings, plus their cross-offering partners, composing against a single oncology super-ontology.

THE SITUATION

Top-10 global pharma. Oncology pipeline of 40+ assets. Trial protocols take 14+ months to author and IRB-approve. Cross-trial learning lives in ~600 tribal documents and the heads of ~200 PIs. Two prior “AI in clinical” programs stalled — the data was “clean”, but operator judgment had been stripped out.

WHAT COMPOSES — IN ORDER

01	Core Excavation surfaces what the trial-ops platforms do and what the PIs know. Weeks, not quarters.	SAIF
02	Knowledge Engineering builds the enterprise oncology ontology. Living Ontologies scalpel a scoped layer for protocol authoring.	AIRD → SAIF
03	Data Migration + Modernization arrives somewhere better — trial data, RWE, and lab telemetry as AI-ready products federated to therapeutic-area LoBs.	AIRD
04	Industry IDB activates the oncology super-ontology in the core; Agentcraft builds the protocol-authoring agents that compose against it.	ADV AI + SAIF
05	Trust Engineering + Advanced AI Evaluation & Assurance certifies every agent for GxP, submission integrity, clinical safety — engineered governance, not audit theater.	ENGINEERED TRUST
06	Data + AI Managed Services operate the ontology, agents, and data products as one command structure. The system <i>compounds</i> .	BOTH MS

ONCOLOGY COMPOUND

WHY THIS IS GREATER THAN THE SUM OF PARTS

Each offering alone delivers a real but bounded outcome. Composed against *one* oncology super-ontology, value compounds across three reinforcing loops:

- **Knowledge loop.** Excavation feeds the ontology; the ontology trains the agents; agent outputs improve the ontology. Each protocol makes the next protocol faster.
- **Trust loop.** Trust Engineering codifies the guardrails; AI Evaluation certifies the outputs; Managed Services keeps both observable. Regulators see one engineered chain.
- **Operating loop.** Data + AI Managed Services run as one command structure. The system runs — *and gets better at running.*

METHODOLOGIES + ENGINEERED TRUST

Recon/Praxis structures cradle-to-grave delivery; **DDE** co-evolves the scoped ontology with the build; **Challenger Forge** handles the parts of the org where antibodies would kill it. Engineered Trust runs across all three offerings — one chain, three expressions, visible to the FDA as one.

CLINICAL TRIAL

14mo → 4mo

PROTOCOL AUTHORIZING · IRB-READY

3×

CROSS-TRIAL LEARNING VELOCITY

Live

REGULATORY OBSERVABILITY

Compounds

EACH ENGAGEMENT LOADS + COMPOUNDS

Enterprise estate analysis, *composed*.

A live microservices platform for enterprise technology estate analysis — ingests application + infrastructure inventory, applies domain intelligence, surfaces insights through a modern web UI.

WHAT IT DOES

Ingest — application + infrastructure inventory from multiple source systems.

Correlate — domain intelligence applied to surface relationships, risk, and solution architecture.

Score — risk scoring across portfolio, sequenced for evidence-based modernization.

Surface — modern web UI; architecture-decision governance lint built in.

Operate — containerized stack, infrastructure-as-code deploy, supply-chain quality gate in CI.

```
repo: ../nexus-fusion · quickstart: one command · auth: scoped tokens
```

HOW IT COMPOSES

- **Core Excavation** — the working asset behind Code & Institutional Archaeology at portfolio scale.
- **AI Architecture Foundation** — produces the dependency + risk graph the Agentic DMZ decisions are made against.
- **Living Ontologies + Knowledge Engineering** — ingested estate becomes a queryable graph the scalpel + foundation both leverage.
- **Industry IDB** — estate analysis feeds the per-vertical super-ontology; one client's analysis compounds the next.

WHY THIS MATTERS IN THE ROOM

We don't pitch this. We *open it*. A multi-service platform with PR quality gates, supply-chain attestations, and live health endpoints — running. **Engineered Evidence**: the working artifact IS the sales mechanism.

Drop in a repo. *Get a knowledge graph.*

~30 seconds for the graph, a few minutes for the docs. A queryable knowledge graph of every module, class, and function and how they call each other. Semantic-search index over every function. Eleven agent-authored markdown artifacts. Then you can chat with it.

WHAT YOU GET

Knowledge graph — modules, classes, functions, calls, imports, inheritance, plus git archaeology (hotspots, temporal coupling, ownership).

Semantic index — search every function by intent, not just by name.

11 protected artifacts — Architecture · Business Logic · Domain Model · Technical Rules · User Workflows · Risk Assessment · Exceptions · API Reference · Audit Report · Business Rules · Integration Inventory.

On-demand analyses — performance, security, migration. Chat interface delegates to specialist agents.

Web UI — real-time ingestion progress, 3D/2D graph visualization, command palette, per-document chat.

```
repo: ../CA/code-platform · containerized backend + web UI
```

HOW IT COMPOSES

- **Core Excavation · Code Archaeology (L4)** — this is the named asset, productized.
- **13 languages** — TS/JS, Python, Rust, Go, Java, C, C++, C#, Ruby, PHP, Swift, Kotlin, **COBOL**. Banking legacy hits the same pipeline as modern micro-services.
- **Archive + live modes** — archive (managed) or in-place with git-based incremental re-indexing.
- **Collections** — group related codebases for cross-codebase semantic + graph search. The portfolio reads as one estate.
- **Frontier-model agnostic** — partnership with multiple providers, not lock-in.

USE WHEN

You've inherited a codebase you don't understand · you're planning a migration · you owe an architecture review on Friday.

This is what makes "*weeks, not months*" a sentence we can defend.

The builder *builds itself*.

Four weeks. From human-led foundation to a system that designs, executes, and ships its own code — end-to-end.

BIG ITEM

#7

ON THE
LEADERBOARD

Autonomous design + build went from *nothing* to **#7** on the leaderboard — in *speed and quality*.

Same window as the chart on the right. Not a feature improvement — a category jump.

“The ramp isn’t velocity any more — it’s who is producing.”

day 4 post-threshold • internal note

WHAT THE SYSTEM ALREADY BUILDS

Self-observability — designs its own telemetry, reads its own performance. **Environment ergonomics** — improves its own prompts + autonomy thresholds. **New product frontiers** — multiple pods on a brand-new product family the human team hasn’t deeply touched.

THE NEXT GATE

Closing the self-improvement loop — pod wires its own telemetry into operations, designs improvements from its own performance data. **Scaling to 30 Forward Deployed Engineers** — one team becomes 30 active engagements feeding one shared knowledge graph. The bottleneck inverts permanently.

• The picture changed a lot in the last 24h — pod went from 22% to ~90% of today’s commits, and ran ~5 parallel design cycles autonomously.



Commits per day • 5/4-5/15 • the bottleneck inverted