

AI Foundation Layer

Headcount tells you scale. Role structure tells you shape. Skill concentration tells you capability. Signals tell you direction. This is the dataset that gives AI models structured memory of how companies actually work.

A dense, signal-rich company intelligence layer designed for direct use in AI systems. Covers growth dynamics, workforce structure, capability composition, and derived signals across 100+ companies in a clean 2018-2020 observation window. Built to be plugged into RAG pipelines, model training workflows, or any system that needs to understand companies without hallucinating.

— WHAT THIS SHOWS

- How companies grow, contract, accelerate, and recover -- tracked year by year with structured signals, not narratives.
- What a company's workforce actually looks like: which roles dominate, how diverse the org structure is, and where capability concentrates.
- Derived classification signals (growth acceleration, early scaling, contraction, recovery) that are directly usable as model features or screening filters.
- How two companies with identical headcount can look completely different at the capability layer, revealing what each can actually execute.

— SIGNALS INCLUDED

- Growth signals: headcount, YoY growth rate, growth bucket, growth percentile
- Workforce structure: primary role, role diversity, role coverage
- Capability layer: top 3 skill buckets per company, skill coverage
- Derived flags: growth_acceleration, early_scaling, contraction, recovery_signal

— FORMAT

CSV or Parquet. One row per company per year. All fields populated, no NULLs in key columns. Clean snake_case naming. Production-ready on first open. Includes all growth, structure, and capability fields in a single flat file.

— WHAT YOU CAN DO WITH IT

- Ground an AI research tool in real company data so it stops guessing about workforce dynamics, growth patterns, and organizational structure.
- Use as a structured knowledge base in RAG pipelines -- query by company, year, industry, or signal and get validated answers.
- Train or fine-tune models on real company evolution patterns: scaling sequences, contraction cycles, capability shifts across industries.
- Screen companies by workforce signal rather than just financials -- detect early scaling, structural pivots, and capability concentration before they reach public data.
- Build competitive intelligence layers based on what companies can actually execute, not just what they announce.

— BUILT FOR

- AI teams building company research, analysis, or intelligence tools that need structured ground truth data.
- Builders and developers integrating company knowledge into AI applications, agents, or copilots.
- Investment research platforms grounding their AI in real workforce signals instead of scraped fragments.
- Data platforms and enrichment providers looking for a unique, high-density company intelligence layer.

Start with a sample. See the data before you commit.

data@vivameda.com