

MACHINE-READABLE INTELLIGENCE

Differentiated company source material, structured for research.

Synorb continuously monitors over 10,000 source channels across corporate websites, engineering blogs, and business podcasts — and turns each new piece of content into a structured Manifest your research, reasoning systems, and products can use instantly.

STREAMS

SOURCE CHANNELS

CLAIMS / 168H

HOME DOMAINS

Thousands

10,000+

1.05M

12

WHAT SYNORB IS

We deliver Manifests — structured representations of new company source material — via Streams organized by Organization, Person, and Dataset.

01 / MANIFEST

The product we dispatch.

A single piece of new content captured from a source channel and turned into three machine-readable views: a Signal, a Brief, and a Record. One Manifest per piece of content. Always the same shape.

02 / STREAM

The unit of subscription.

A filtered view over the live Manifest graph — pinned to an Organization, Person, or Dataset. Hit a Stream and every new Manifest lands the moment it matches.

03 / DELIVERY

Where it lands and how.

REST API for synchronous reads. MCP server for AI agents. WebSocket firehose for live streaming. S3 bulk export for warehouse loads. Every channel returns the same Manifest shape.

WHY THIS WORKS

Most research teams already maintain a fragile pipeline of RSS readers, news APIs, scraped HTML, and ad-hoc transcription — producing unstructured text an analyst then reads and tags. Synorb replaces that pipeline with a single structured input: a queryable graph of **Manifests**, kept current 24/7, that your models and analysts can both read.

ORGANIZATIONS

Public companies, large privates, asset managers, policy institutes, agencies. Each one a Stream.

PEOPLE

CEOs, founders, allocators, researchers, and policy figures whose primary outputs we track.

DATASETS

Public datasets — FRED, BLS, SEC, sports odds — resolved to the same Manifest shape.

ANATOMY OF A MANIFEST

Three machine-readable views. One Manifest.

From source channel **google-cloud** — routed to **Alphabet** Stream.

MANIFEST
id 1777270819009619611

```

{
  "manifest_id": 1777270819009619611,
  "record_id": 1777270818914478519,
  "stream_ids": [
    17723038993544463
  ],
  "stream_names": [
    "google"
  ],
  "matched_at": "2026-04-27 07:08:44.644821",
  "status": "routed",
  "home_domain": "economics-business-work"
}
        
```

MANIFEST
 The product we dispatch. One per piece of new content — binds the Record, the Signal, and the Brief, plus every Stream the content was routed to.

SIGNAL

```

{
  "story_id": 1777273714134984553,
  "story_type": "signals",
  "headline": "We're launching two specialized TPUs for the agentic era.",
  "claims": [
    {
      "signal": "This indicates a new product offering from Google Cloud, highlighting continuous innovation in AI hardware ...",
      "entities": [{"name": "Google Cloud", "type": "Organization"}]
    }
  ]
}
        
```



SIGNAL
 Structured claims for agents — compact JSON optimized for reasoning.

BRIEF

```

{
  "story_id": 1777273713933287381,
  "story_type": "briefs",
  "headline": "We're launching two specialized TPUs for the agentic era.",
  "summary": "Google Cloud is introducing its eighth generation of Tensor Processing Units (TPUs), featuring two specialized chips. These new TPUs are designed to support the advancement of artificial intelligence, particularly for what Google terms the 'agentic era'."
}
        
```



BRIEF
 Best for RAG, Dashboards, and products requiring humans-in-the-loop.

RECORD

```

{
  "record_id": 1777270818914478519,
  "title": "We're launching two specialized TPUs for the agentic era.",
  "url": "https://blog.google/innovation-and-ai/infrastructure-and-cloud/google-cloud/tpus-8t-8i-clo",
  "source_name": "google-cloud",
  "source_type": "organization",
  "media_format": "text",
  "claim_type": "publication",
  "source_published_date": "2026-04-27"
}
        
```




RECORD
 Canonical data from the source object — enriched with entities and topics.

INVENTORY AT A GLANCE

Thousands of Streams. 10,000+ source channels. 12 home domains. The curve is steepening.

Inventory compounds: every Stream we add resolves more entities, which surfaces more sources, which seeds more Streams. New domains, new asset classes, new geographies — we are still on the flat part of our own curve. If a name shows up in your investment universe today and we don't already cover it, we usually can within days.

12 HOME DOMAINS



DELIVERY SURFACE

- REST API**
JSON Manifests, paginated. Sub-second on cached reads.
- MCP Server**
21 tools for AI agents. Claude, Cursor, GPT-compatible.
- WebSocket Firehose**
Live Manifest stream. Sub-second from capture to delivery.
- S3 Bulk Export**
JSONL files for warehouse loads. Daily and on-demand.

CLAIMS / 168H

1.05M

Atomic assertions extracted in the last 7 days.

INVENTORY GROWTH

Compounding

New Streams every week; cadence accelerating.

UPTIME TRACKING

99.5%

Healed with 20-min Mean Time to Resolution (MTR).