WORKLOADS & USE CASES

Architectural Selection & Production Use Cases

apical labs

To determine whether a production workload belongs on **Apical Arbor** or **Apical Plexus**, evaluate three foundational engineering criteria across feature geometry, execution caller, and lifecycle primitives.

FEATURE GEOMETRY

Separated Embeddings vs. Entangled States

Apical Arbor →

Contrastive text and vision encoders (`BGE` , `MiniLM` , `OpenAI` , `DINOv2` , `CLIP`) already separate classes metrically. Use a single-layer directional readout in **< 0.10 ms** (**93.66%** on MiniLM CLINC150, **81.79%** on DINOv2 Split CIFAR-100).

Apical Plexus →

Raw causal LLM hidden states (`Qwen-2.5` , `Llama-3`), multi-hop agent intents, and raw unencoded sensor arrays are non-linearly entangled. Use multi-layer sparse continual adaptation in **< 0.18–0.22 ms** (**74.66%** on Qwen2.5-0.5B, **+18.68 pp** vs NCM and **+34.40 pp** vs centered-NCM).

EXECUTION CALLER

Backend Pipelines vs. Autonomous Agent Runtimes

Apical Arbor →

Backend microservices, high-throughput ingestion pipelines, and **Apical Edge** ANSI C99 firmware (`apical_edge.h` , **1.2–49.1 KB** SRAM) passing dense floating-point or Q15 fixed-point vectors via the Python SDK or REST API.

Apical Plexus →

Autonomous LLM agents, coding IDEs (Cursor, Claude Code, Antigravity), and multi-agent orchestrators routing natural language prompts and MCP/OpenAI tool schemas via `plexus.Router` , `apical-proxy` , or the MCP server.

LIFECYCLE PRIMITIVES

Backbone Migration vs. Associative Recall

Apical Arbor →

Upgrading equal-dimension foundation backbones (`MiniLM-384D` → `BGE-Small-384D` or `512D` → `512D`) with zero stored customer data via `arbor.Bridge` (**< 20 ms**, **93.02%** relative retention), or **0.01 ms** GDPR Art. 17 class erasure via `arbor.Unlearn` (**0.0000** cross-class leakage).

Apical Plexus →

Reconstructing **50%** occluded sensor frames without a neural backbone via **Apical Edge** / `plexus.edge` associative recall (**+10.75 pp** cosine recovery), instant **< 0.04 ms** route erasure and **< 1 ms** snapshot rollback, or zero-retention style learning via `DeveloperPreferenceAdapter` .

DOMAIN / VERTICAL	APICAL ARBOR Classifier · MemoryRouter · Bridge · Unlearn · Apical Edge	APICAL PLEXUS Router · Network · Preference · apical-proxy · Apical Edge *Note: Plexus plans include all 4 Arbor modules under a unified API key.*
Autonomous Agents, IDEs & Tool Routing	**Fixed-Embedding API Gateway Shortlisting (`arbor.MemoryRouter`)** • **Microservice API Gateways:** Backend services that already compute `MiniLM-384D` embeddings prune 100–250 internal REST endpoints down to Top-3 in `< 0.05 ms` before calling an LLM. • **Voice-Bot Command Dispatch:** Sub-millisecond intent matching from Whisper + MiniLM embeddings to deterministic IVR state-machine transitions.	**Agent Orchestration & Prompt Deflation (`apical-proxy` / `plexus.Router`)** • **IDE Coding Agents (Cursor / Claude Code / Antigravity):** Routes natural language prompts across 500+ MCP tools (`git` , `stripe` , `postgres` , `k8s`) via `apical-proxy` and the MCP server, cutting prompt context from ~15,000 tokens to ~350 tokens (**>92% savings**) and reinforcing winning tools on accepted turns. • **Enterprise 10,000-Tool Mesh:** Multi-layer semantic dispatch across massive internal OpenAPI catalogs where tool descriptions overlap non-linearly. • **Zero-Rot Coding Style Alignment (`DeveloperPreferenceAdapter`):** Continuously learns a developer's 8-dimensional architectural preferences (such as `PARAMETERIZED_SQL` and `STRUCTURED_LOGGING`) from accepted/rejected diffs without bloating static system prompt rules.
E-Commerce, Retail & Catalog Intelligence	**Continual Visual & Text Taxonomy Readout (`arbor.Classifier`)** • **Zero-Downtime Marketplace SKUs:** Classifies newly uploaded merchant product photos downstream of frozen `DINOv2` or `CLIP` backbones into 5,000+ dynamic categories in `< 0.10 ms`. • **Flash-Trend Cold-Start Tagging:** Instantly adds a new seasonal category (e.g., "Tactical Outerwear") from 3–5 human-verified examples online without GPU batch retraining. • **Brand & Counterfeit Logo Detection:** Real-time visual prototype matching at upload time with instant class addition when new brands onboard.	**Multi-Modal & Causal LLM Shopping Assistants (`plexus.Network`)** • **Conversational Commerce State Tracking:** Adapts directly atop frozen `Llama-3` / `Qwen-2.5` hidden layer activations to predict complex, multi-turn buyer intent (distinguishing *"looking for a compatible replacement part"* vs. *"requesting a warranty return"*). • **Cross-Category Bundle Association:** Uses multi-layer associative weights to map non-linear combinations of user session trajectories to high-conversion product bundles in constant $O(1)$ RAM.
Customer Operations, SaaS & B2B Multi-Tenancy	**Per-Tenant Continual Ticket & Document Triage (`arbor.Classifier` / `Unlearn`)** • **Outage-Resilient Support Routing:** Auto-routes high-confidence tickets (≥ 0.82) while escalating novel outage symptoms to human operators whose single-click resolutions train the readout head in `< 0.10 ms`. • **Multi-Tenant RAM Isolation:** Isolates 10,000 B2B SaaS tenant heads at **231–538 KB/tenant** (**2.31–5.38 GB** total vs **~696 GB** on shared RanPAC Gram tables) with `persisted_bytes_delta == 0`. • **Instant Tenant Offboarding (`arbor.Unlearn`):** Executes deterministic **0.01 ms** class erasure (`erase_class()`) with **0.0000 residue** and an `Ed25519` GDPR Art. 17 receipt.	**Entangled Enterprise Workflow & Fleet Coordination (`plexus.Network`)** • **Multi-Hop Incident Escalation:** Routes complex enterprise IT/DevOps incidents where symptoms span logs, stack traces, and natural language descriptions that fail linear separation. • **Shared-Manifold Multi-Agent Fleet Coordination:** A fleet of 10 specialized support subagents (Billing, Technical, Legal, Security) shares a single `Plexus Pro` container so when one agent learns a new error code belongs to Billing, all peer agents route it accurately. • **Live Rule Recovery & Route Erasure:** Automatic `circuit_break` recovers **89.1%–91.8%** when live routing rules change, and **< 0.04 ms** route erasure leaves only **0.47%–1.22%** shared-layer residual (vs **7.2%–13.4%** matrix contamination in RanPAC/SLDA; **0.0000** exact zero residual on Arbor) with an `Ed25519` receipt.
Healthcare, Finance & Regulated Compliance (HIPAA / GDPR)	**Zero-Retention Clinical & Financial Readouts (`Classifier` + `OrthogonalPrivacyProjector`)** • **Blinded Clinical ICD-10 Coding:** Hospital client blinds clinical note embeddings via a client-held secret key (`OrthogonalPrivacyProjector`), adapts an Arbor readout head to local hospital billing codes in constant memory, and retains zero patient text or raw vectors. • **Statutory Right-to-be-Forgotten Webhooks (`arbor.Unlearn`):** Financial institutions or biobanks legally required to prove a specific cohort or identity class has been purged from production weights (`target_class_weight_norm == 0.0`) without spending £40,000+ retraining a foundation model.	**Regulated Multi-Layer Risk & Clinical Triage (`plexus.Network` on Regulated Tier)** • **Real-Time Wire & Payment Fraud Adaptation:** Adapts a multi-layer non-linear head atop transaction graph activations as new fraud rings emerge mid-day, under pre-signed **GDPR Art. 28 DPA** and **Enterprise HIPAA BAA** with an immutable append-only cryptographic receipt ledger. • **Zero-Retention Clinical Trial Cohort Matching:** Matches complex multi-variable patient eligibility criteria atop medical LLM hidden states (`BioMistral` / `Med-PaLM`) where linear classifiers collapse due to non-linear feature interactions (*Drug A + Biomarker B, excluding Comorbidity C*).
MLOps & Foundation Model Lifecycle	**Zero-Downtime Embedding Upgrades (`arbor.Bridge`)** • **Encoder Swapping Without Historical Data:** Upgrades 5,000 tenant classifiers across equal-dimension encoders (`all-MiniLM-L6-v2` 384-D to `BAAI/bge-small-en-v1.5` 384-D, or distinct `512-D → 512-D` spaces) in `< 20 ms` using `M = 1,000` public anchors with **93.02% ± 1.72%** relative retention and zero stored customer samples. • **Cross-Modal Anchor Alignment:** Bridges a classifier trained on text embeddings so it can be queried directly by aligned image embeddings.	**Replacing Nightly LoRA & Experience Replay Pipelines (`plexus.Network`)** • **Eliminating GPU Fine-Tuning Queues:** Replaces nightly LoRA jobs and replay buffers on frozen causal LLMs (`Qwen-2.5` , `Llama-3`) with single-pass CPU updates (`< 0.18 ms/step`), scoring **74.66%** vs **55.98%** NCM and **40.26%** centered-NCM on 15-task `Qwen2.5-0.5B` `CLINC150` . • **Drift-Immune Feature Extraction:** Achieves **84.44% ± 0.26%** across 10 sequential tasks on `DINOv2` Split CIFAR-100 (beating `SLDA 83.10%` & `Offline Ridge 82.62%` with zero replay), and scales **57.01%** → **68.94%** across 100-task streams from Standard to Ultra.
Embedded IoT, Robotics & Bare-Metal Edge MCUs	**Single-Layer C99 Q15 Profile (`Apical Edge` , `1.2–49.1 KB SRAM, 0.52 µJ/step`)** • **Unsupervised Hydraulic & Motor Fault Burn-In:** Runs in **11.30 KB SRAM** on low-cost, ultra-low-power ARM Cortex-M, ESP32, or RISC-V (`0 malloc` , `0 <math.h>` , **0 Flash wear**), catching developing hydraulic faults at **99.64% AUROC** (`UCI #447` , **45×** smaller SRAM than Mahalanobis) using **~2,300×** less battery energy than a `1.2 mJ` BLE packet. • **Wearable Biosignal & Motion Calibration:** On-chip 6-axis activity adaptation reaching **82.29%** in **6.59 KB SRAM** (`UCI #240` , **+13.43 pp** vs TinyMLP-ER-64 in **64×** less SRAM).	**Multi-Layer Sparse Q15 Profile (`Apical Edge` , `69.5–428.5 KB SRAM, 1.31 µJ/step`)** • **Autonomous Drone, Rover & Raw Sensor Streams:** Learns directly on raw sensor windows with zero neural backbone (**83.13%** on raw Split-MNIST in **16.92 KB** Q15; **69.62% / 59.20%** continual 4-fault hydraulic diagnosis in **94.20 KB SRAM**). • **36-Month Sensor Aging & 50%-Occlusion Recovery:** Adapts across 3 years of physical gas-sensor drift (**60.31%** in **69.52 KB SRAM** on `UCI #224` vs **30.19%** EWC) and reconstructs 50% occluded sensor frames via `associative_recall()` (**+10.75 pp** cosine recovery, `0.6836 → 0.7912`).
Cybersecurity, SOC & Threat Intelligence	**Line-Rate Signature & Phishing Triage (`arbor.Classifier` / `Unlearn`)** • **Zero-Day Phishing & URL Classification:** Classifies vectorized email and DOM embeddings at line rate (`< 0.10 ms`) and adapts immediately when SOC analysts flag a new phishing campaign. • **False-Positive Quarantine Erasure (`arbor.Unlearn`):** Deterministically unlearns a benign partner domain accidentally trained into a blocklist class in `< 1 ms` without rolling back subsequent threat intelligence updates.	**Behavioral Telemetry & Prompt-Injection Firewall (`plexus.Network` / `Edge`)** • **Adaptive LLM Jailbreak & Prompt-Injection Gate:** Sits in front of enterprise LLM endpoints and adapts to novel adversarial jailbreak activation patterns in real time without storing sensitive user prompts. • **Encrypted Network Flow Anomaly Detection:** Multi-layer continual classification over raw packet-timing and window tensors to detect C2 (Command & Control) beaconing drift on edge network appliances.
Media, Gaming & Real-Time Personalization	**Zero-History Session Personalization (`arbor.Classifier` / `MemoryRouter`)** • **Zero-Retention Feed Ranking:** Updates a user's interest vector readout in real time as they skip or watch clips while guaranteeing `persisted_bytes_delta == 0` (zero behavioral history database to breach or subpoena). • **Live Moderation Queue Triage:** Fast continual classification of user-generated images and comments downstream of `CLIP` as new community guideline categories emerge.	**Adaptive NPCs & Non-Linear Player Modeling (`plexus.Network` / `Edge`)** • **Continual Combat & Strategy NPCs:** Game AI agents that learn a player's non-linear tactical habits (flanking patterns + cooldown timing) within 3–5 encounters in `< 0.18 ms` on CPU without frame-rate drops. • **Associative World-State Memory:** Recalls complex multi-entity game events from partial contextual cues in fixed memory.

Apical Arbor

Choose **Apical Arbor** when you already have a **contrastive embedding model** (`BGE` , `MiniLM` , `OpenAI` , `DINOv2` , `CLIP`) or a **DSP feature pipeline**, and you need the fastest possible (`< 0.10 ms`) continual readout classifier, deterministic 1-step GDPR Art. 17 class unlearning (`erase_class`), or zero-downtime embedding model migration (`bridge`).

Apical Plexus

Choose **Apical Plexus** when you are building **LLM Agents & MCP workflows** (`apical-proxy` / `plexus.Router`) that require **>92% prompt token deflation**, when you are tapping **raw causal LLM hidden layers** (`Qwen-2.5` , `Llama-3`) where features are non-linearly entangled, or when you are deploying **backbone-free sensor microcircuits** (`Apical Edge`) requiring associative pattern recall. Paid **Apical Plexus** plans and **Enterprise Cloud / Self-Hosted VPC** also include all four **Apical Arbor** modules under a unified API key.