Application Layer

Cognitive Architectures

Built in primitives for Cognitive Architectures								
Compositional AI systems			**Emergent AI** (Req. AIGr.id + AgentGr.id)			**Agentic** (Req. AIGr.id + AgentGr.id)		
Neuro Symbolic	Compound AI	Internet of Intelligence	Collective AI	Societies of Minds	Swarm Intelligence	Agents	Mult-Agent Systems	Reactive, Deliberative AI
Modular AI			Evolutionary AI					
Non Emergent AI			**AI Intelligence levels** (Req. AIGr.id + AgentGr.id)					
Neural Network	Symbolic AI	Classical ML	Fuzzy Logic	Open-ended Intelligence	Super Intelligence	AGI	Narrow AI	

Topologies of ownership, access & governance									
AI Commons	AI Grid	Private AI	Public AI	Federated AI	Sovereign AI	Decentralized AI	Polycentric AI	Cloud AI	Local AI

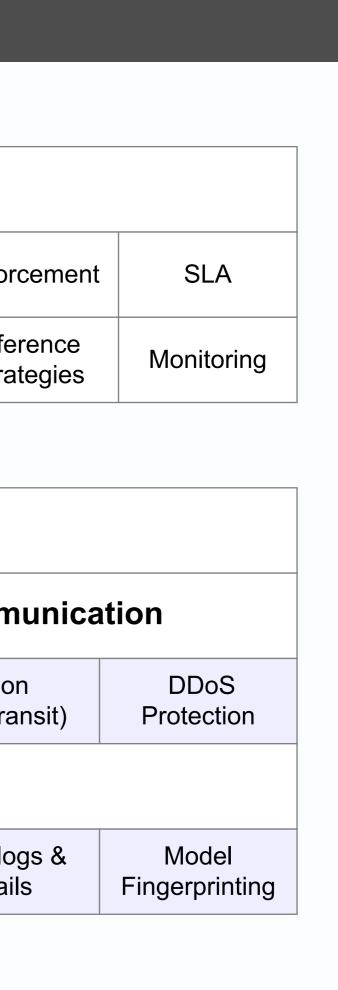| Meta Architectures									
Compositional	Modularity	Pluralism	Heterogeneous	Interoperable	Protocol Standards	Decentralized	Polycentric	Sovereignty	Safe
Alignment	Ethical	Transparency	Open-endedness	Resilience	Explainability	Forkable	Loose Coupling	Accountable	Fractal

Trust, Governance, Safety, Security, Incentive, Reputation

PolicyGrid: Turing Complete, Programmable Policies

Governance	Trust	Guardrails	Security	Incentive	Steerability	Fulfilment Audit	Alignment	Enforcement	SLA
Conflict Resolution	Resource Management	Escalation Handling	Dynamic Deligation	Program Ethics	Program Behaviour	Reputation	Behaviour Audit	Inference Strategies	Monitoring

Security

Identity & Access Control			Secret & Credential Management		Network Security & Communication			
IAM	RBAC	ABAC	Secret Management	Key Management	mTLS	Firewall	Encryption (In rest & Transit)	DDoS Protection

Asset Security			Secure Computing				Others			
Signing & Verification	Asset Encryption	Asset Access Control	TEEs	Sandbox	Confidential VMs	MPC	Rate limiting & Throttling	Abuse Detection	Immutable logs & Audit Trails	Model Fingerprinting

AI Platform Layer

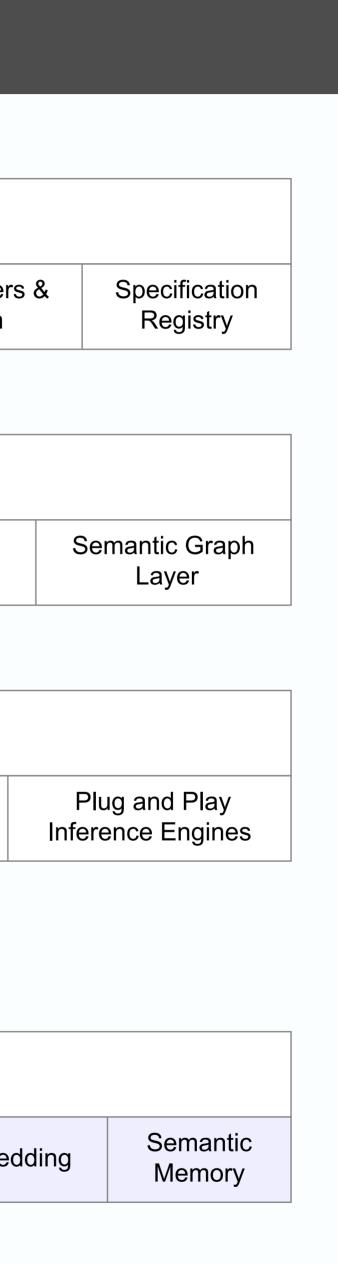## AI Workload Specification

Custom Specification	Specification Validator	Custom Parser	Job Specification	AI Graph Specification	Workflow Specification	Templating Parametrization	Schema Adapters & Composition	Specification Registry

Distributed AI Graph Engine

AI Metagraph	Runtime Metagraph	Compound AI	Auto AI	Static AI Graph	Dynamic AI Graph	Nested AI Graphs	Graph Mutation	Graph Planner	Graph Policies	Semantic Graph Layer

Distributed Elastic Inference

Online Inference	Batch Inference	Adhoc Inference	Stateful Inference	Stateless Inference	Model Mesh	AI Gateway	Serverless Inference	Modelless Inference	Plug and Play Inference Engines
Model Partitioning	Inference Cache	Adaptive Inference	Cold start Optimisation	Resource Optimization	Multi-Tenant Serving	Model Sharding	Inference Isolation		

MemoryGrid

Data Cache	Short Term Memory	Long Term Memory	Local & Global Memory	Working Memory	Episodic Memory	Vector / Embedding	Semantic Memory

Coordination & Orchestration Layer

Coordination	Orchestration	
		Flow Governor.Network (Across Blocks, Nodes & Clusters)
		Network Governor.Network (Multi / Across Cluster)
		Cluster Governor.Network (Multi / Across Node)
		Node Governor.Network
		Block Governor.Network (Multi / Across Nodes & Clusters)

Resource Management

Resource Pooling	Scheduler	Resource Allocation	Resource Selection	Resource Isolation	Quota Management	Auto Scaling	Resource Optimisation	Resource Monitoring	Audit & Log	Priority & Affinity
Resource Negotiation	Metrics	Load Balancing	Fault Tolerance	**Agency**: Network, Cluster, Node, Block Governor Network						

Job Management

Job Scheduler	Job Triggers	Job Queues	Job Executors	Job Resource Manager	Job Isolation	Job Fault Tolerance	Job Status and Tracker	Destructor & Garbage Collection	Job Recovery
Execution Order	Parallelism / Fan out	Concurrency	Conditional Logic	Dependency Resolution	Priorization & Preemation	Result Collection	Job Intervention	Secrets & Config Injection	Audit & Logging
Agency: Flow Governor Network									

AI Graph Management

Decentralized Graph Executor	Decentralized Graph Coordination	Graph Scheduling	Graph Resource Manager	Auto Scaling	Graph Fault Tolerance	Graph Policy Engine	Graph Monitoring	Graph Metrics	Audit & Logging
Graph Load Balancer	Graph Optimisation	Data Router	**Agency**: Flow Governor Network						

Resource, Service, Artifact (RAS) & Data Handling

RAS Asset Registry	RAS Run time Registry	RAS Registration	RAS Discovery	RAS Selection	RAS Gateway	RAS Policy Registry	RAS Container Registry	Data Routing Service	Policy Enforcement
Data Distributor	Data Aggregator	Data Sync							

Managed Platform Services

FaaS	Cache / In memory DB	Persistent DB	Messaging	Queues	Events & Alerting	Pub Sub	Metrics	Logging	AV Streaming	Data Streaming	Data Processing	3rd Party Operators

Distributed Workflow Orchestration

AI as a Services Layer

Kubernetes

Control Layer

Block Management

Agency: Block Governor Network

AI Auto Scaler	AI Load Balancer	Fault Tolerance	Quota Management	Block Monitoring	Block Negotiation	Policy Enforcement	Block Metrics	Audit & Log	Block Executor	Block CI / CD

Pre-Process Sub Block | AI Block | Post-Process Sub Block

Side Car

Reference Block

AI Blocks as Docker

AI Blocks as VM

AI Blocks as MicroVM

AI Blocks as Web Assembly

AI as Discrete Microservices
Scalable across Nodes & Clusters

Compute Aggregation Layer

Node Management

Agency: Node Governor Network

Elastic Compute	Node Registration	Node Monitoring	Node Lifecycle Manager	Configuration Manager	Node Negotiation	Policy Enforcement	Node Metrics	Audit & Log	Topology Awearness	Self Healing

Storage

Distributed File System	Object Storage	Network Attached Storage (NAS)

AI Accelerator

GPU	TPU	Others

Network

SDN	VPC	Overlay Networks

Physical Compute Layer

Public Cloud	Reginal Data Center	Enterprise Private Compute	Research Lab Compute	Crowd Sourced Compute