ComputeAI

Compute.AI Product Datasheet		
Supported Infrastructure		AWS, Azure, GCP, On-Prem, Hybrid
		Cloud native
Data Model	Memory Model	Default: Hybrid (Columnar/Row)
		Compression: In-memory, on the wire and at rest
	Data Sources	S3, NFS, Azure Blob Storage, GCS, HDFS
	File Formats	Parquet
	Table Formats	Apache Iceberg
		Parquet, Hive Tables
	Catalogs	AWS Glue, Hive Metastore
SQL		ANSI SQL 92
Cloud Architecture		Full Memory and Storage heirarchy, fine-grained paging to different tiers of storage
		Highly elastic, Kubernetes based
Query Processing and Optimization	Cost Based Optimization (CBO)	AI/ML based CBO, scheduling and demand paging
		Optimized scan planning
		Progressive query execution during load
	Join/Windowing Strategies	Automatic
	Predicate Pushdown	Filters, Aggregates, Projects
	Skew Management	Automated
Indexing		Iceberg partitioning
		Real-Time fine grained indexing
SQL clients		Any JDBC SQL client
BI Tools		Tableau, PowerBI, Looker, QuickSight, Superset, etc.
Fine-Grained Entitlements		Table, Column and Row level access using PostgreSQL's Role-Based Access Control API
SLAs		Manage price/performance
		Rolling Upgrades/CICD
SSO		Customizable SSO (OAUTH, LDAP)
Security		OAuth based Authentication (Customizable)
		Data at Rest Encryption on S3
		Admin controlled seamless upgrades
CICD and Rolling Upgrades		24x7, 1 hour response time
Support		Best Practices Guides
Monitoring/Dashboard		Browser based