HarperDB
Feature/Benefit Matrix

Feature	Benefits
Single Model	HarperDB stores structured and/or unstructured data in a single, patent pending data model eliminating the need to transform your data in memory or on disk. This allows for executing queries as soon as data transacts with no waiting for jobs or transforms. The same workload can be handled with less compute, memory and storage.
Dynamic Schema	Gain all the benefits of having a schema with the flexibility of not having one. HarperDB is flexible enough to ingest unstructured and structured data while also providing complex SQL capability for reporting and analytics. Inspect and visualize schemas and tables to more deeply understand your data. HarperDB tracks the metadata around schemas, tables, and attributes allowing for describe table, describe schema, and describe all operations.
Fully Indexed	HarperDB is fully indexed without increasing the storage, memory and compute footprint or adding additional management. This is possible due to our patent pending, exploded data model which natively allows for discrete search and retrieval of individual attributes per row or object.
Exploded Data Model	Our patent-pending data model which stores values and attributes individually, allowing for discreet read/write operations which increases performance, reduces overhead, and enables scale. This is the primary differentiator that enables our fully indexed capability, and ability to transform unstructured data to structured data in real-time.
HTAP (Hybrid Transactional/ Analytical Processing)	HarperDB can handle high volume transactions of data while staying ACID compliant. Structured data can be handled at high volume without row locks. Meanwhile, data can be queried and analyzed in real-time with no transformation or duplication. All of these workloads can execute simultaneously with no performance degradation.
NoSQL/SQL	HarperDB uses standards-based interfaces which developers already know. Whether your projects use only SQL or NoSQL or a blend of both, developers can interact with HarperDB without needing to learn new modalities.
Small Footprint	Because HarperDB is written in Node.js and leverages existing operating system technology, HarperDB has an incredibly small footprint and reduces hardware and infrastructure costs and requirements. HarperDB can be deployed on a Raspberry Pi or other IoT devices while maintaining the ability to scale to a large cloud environment.
Multiple Workloads	HarperDB is designed to function as an edge database, application database, transactional store, document store, time series database, or data warehouse. Most databases are designed to handle only a single workload, requiring companies to purchase multiple licenses, maintain excess knowledge, and incur massive cloud or capital costs.
Scale	HarperDB has both horizontal and vertical scale capability. While Node.js is not inherently multi-threaded, the Node engine allows for finite control of resources, like CPU and threads. HarperDB allows for the maximization of hardware resources. As Node.js was first developed as a web server; HarperDB's horizontal scaling capability is robust and hardened allowing for HarperDB to be easily deployed across multiple resources of varying sizes.
API	Integrate HarperDB into applications without middleware, ORMs, or other complexity. HarperDB provides a native REST API for simple integration with your applications. HarperDB's native REST interface implements standards such as HTTP/HTTPS, WebSockets & JSON and allows for rapid development of solutions in a concise and readable format which natively plugs into any language.
Compatibility	In addition to our standard API, HarperDB Enterprise provides native ODBC/JDBC drivers allowing for the direct integration into third party Business Intelligence and Analytics tools out-of-the-box.
Security	Manage users, roles, and attribute level security. Native OS security features allow isolation and encryption, minimizing potential data breaches and unwanted access.
Studio	HarperDB studio provides a sleek and easy to use management interface with data visualization. Manage security, schema, logs, and node health in a simple graphical interface. Turn SQL or filter searches into real time graphs and charts and share with business end users via live web links.

HarperDB
Simplicity without Sacrifice