

SOLUTION BRIEF

Lower Development & Infrastructure Costs

Build feature-rich and hyper-performant software at a fraction of the cost.

Save Money, Improve Performance

Unlock Savings for Expensive Workloads

Harper's distributed application saves money and time across the entire software lifecycle. The unified system architecture of Harper delivers database, application, caching, and streaming services in one flexible and distributed technology. This shift away from the multi-system backend standard, offers breakthroughs in ease of development, throughput, latency, resilience, and cost-efficiency. Below are some key ways Harper helps you save money on every workload.

99%
Lower Latency

75%
Lower Cost

Digital Out-of-Home Ads

92%
Lower Latency

90%
Lower Cost

Global Gaming Network


69%
Lower Latency

70%
Lower Cost

Digital Media Streaming

Real World Savings

9 Ways Harper's Unified System Architecture Delivers Breakthrough Cost-Performance



Save During Planning & Development

1 Procurement & Licensing Costs

It takes time to research, test, and select the ideal technology. Additionally, when multiple systems must be used, licensing costs skyrocket since each license must cover the cost of development, marketing, sales, and operations for each organization, making you pay extra overall. A single node of Harper can often replace three to four other existing systems, reducing license and maintenance costs dramatically.

2 Opportunity Cost

Although less tangible of a cost, lost opportunities for innovation are another reason companies switch to Harper. It's all too common for companies to pass on building high-value features because their current architecture and technologies can not effectively support the workload. Since Harper breaks free from traditional capability silos, developers have

significant room to innovate without needing to license, learn, and connect additional technologies. With Harper, your team can innovate at a market-leading pace.

3 Time-Saving Simplicity

Time is money, and spending months building and maintaining connections between systems is expensive. Harper is the only solution that lets you forgo this expense entirely. Imagine how fast your developers could move if their tech stack were ready to go before they even started developing. As an example, one telecommunication customer bogged down by cumbersome maintenance and time consuming development processes has elected to replace their current CockroachDB, Redis, and Kafka stack as they build the next generation of telecommunication systems. The sooner your shift begins to a unified architecture, the greater your long-term savings.

Save During Operations

4 Set Your Workloads Free

Successfully deployed on cloud, multi-cloud, hybrid, 5G, bare-metal, and edge, Harper is truly infrastructure agnostic. This gives you complete flexibility in choosing infrastructure vendors, allowing you to position each workload optimally for your use cases and budget. With true workload portability, your team is free to build and innovate on infrastructure where long-term commitments already exist while simultaneously preparing for a low-lift transition to more affordable infrastructure providers.

5 Reduce Serialization Overhead

Serializing and Deserializing data are two of the most expensive operations in any application. By combining application servers, databases, and middleware in one system, Harper reduces the need to serialize and deserialize data, dramatically reducing the amount of computing required for your workload.

6 Efficient Compute

Most operational databases spend the majority of their resources performing read-and-write operations. The more efficient each operation becomes, the fewer resources the system needs. Being originally designed for lightweight IoT use cases, Harper requires an order of magnitude less memory for most operational workloads compared to systems like MongoDB.

7 Fewer Parts, Fewer Compute Environments

Traditionally you need a compute environment for each layer of your application stack. One for your application logic, one for your database, one for your caching layer, and one for your middleware. With Harper, all of those are combined in a single platform dramatically reducing the amount of compute environments required to run your workload.

8 Reduced Network Costs

Out-of-region egress costs add up for larger enterprises. Calling data from a central data server can be expensive. Since Harper is fully contained and designed for global distribution. Data is retrieved from the nearest node, dramatically reducing cross-region egress costs.

Additionally, by combining HarperDB's capabilities into a single platform, fewer network hops are required. Traditionally, data needs to be sent from your app to your database to your caching layer all through your middleware, combining this in a single platform means the data makes fewer hops, reducing network overhead.

9 Savings Grow with Scale

Data redundancy requires a vertically scaled and cost-intensive central database for most systems. In addition to cost, this legacy paradigm creates a performance bottleneck and slow response times for geographically dispersed users. Alternatively, each node of HarperDB is equipped with everything needed to respond to client requests (including data) and is typically deployed alongside several geo-distributed and synchronized replicas. This ensures that even if an entire region of servers goes down, the request can instantly be routed to the next nearest server to fulfill the request. In addition to deep redundancy and low latency, this system design also offers lower scaling costs. Unlike standard exponential cost models of central data systems, the percentage you save with HarperDB's linear cost model will increase with every user you add.

Start Small, Save Big

Although the most efficient system is built entirely using Harper, a complete migration is not required to start saving. Typically, organizations front-end their most expensive workload to Harper first, allowing them to reap maximum savings and performance gains as quickly as possible. This can look like placing Harper in front of your existing data infrastructure to dramatically reduce the cost of expensive workloads like heavily hit APIs or caching workloads.

Then, as your team develops new features, Harper begins processing more and more workloads, increasing savings and accelerating time to market for new features.

With prebuilt connectors for MongoDB, DynamoDB, and DataStax, the transition to Harper is easier than you might think. Beyond just saving money, Harper will also improve your system's throughput, decrease latency, increase reliability, and get products to market faster.

Start a conversation today with your Harper team to start saving as soon as possible.

Contact Sales at
hello@harperdb.io.

