



# APPLICATION BRIEF

## ▶▶▶ SCALABILITY

Scalability is the ability of a system to handle larger workloads by enlarging the system in a straightforward manner. In practice, however, it is often impossible to plan for the scenarios that will most benefit from highly scalable systems. Workloads can drastically expand due to business growth, new application features, and usage patterns. How will your application grow eight years down the line?

### Scale from the start

VoltDB was conceived as a high-performance system that grows with application workloads. That means meeting today's application performance challenges with an eye toward preserving tremendous headroom for the challenges of tomorrow.

The way we see it, all apps are ready for scale. Partitioning is a fundamental part of all VoltDB applications, not just a design decision that's deferred until an existing workload hits its performance ceiling. Partitioning at the CPU level not only encourages developers to horizontally partition early on in an application's life cycle, it's also a great strategy for design patterns that can take full advantage of multi-core hardware. As new machines or CPUs are added to a VoltDB cluster, applications will be ready to harness these new resources.

### A database optimized from the bottom up

VoltDB is also optimized for per-machine performance. On a single machine, VoltDB can outperform many data systems by one or two orders of magnitude. Thus, VoltDB clusters can achieve the same level of performance using dramatically fewer nodes.

All nodes are created equal in VoltDB. Many distributed systems designed for scalability assign separate roles to different parts of a cluster. For example, larger Hadoop installations assign some nodes to perform distributed agreement, some nodes to manage file system metadata, and other nodes to do the actual batch processing.

VoltDB has only one kind of node: a cluster member. Like other distributed systems, VoltDB has many different subsystems, but the cluster internally manages what responsibilities are assigned to each node. Any node can fail, and the cluster will quickly reassign the failed node's responsibilities to another node in the system.

Furthermore, VoltDB clusters operate as individual systems. More than simply self-healing after failure, a cluster of many nodes can usually be managed as a single entity. Cluster-wide statistics can be collected from any node, and clients can be connected to one, all, or a subset of nodes.

VoltDB can scale a well-designed application almost without limits, and with incredible cost benefits. VoltDB is designed for — and tested to handle — workloads that run millions of transactions per second. Crucially, as cluster sizes grow, per-machine costs are reduced. The result is a system that isn't much harder to manage at 20 nodes than it is at 2 nodes. second — transactions that require all partitions. This allows scaling highvelocity ingestion workloads, which are single-key by nature, while simultaneously supporting global cross-key transactions for dashboarding and multi-key analytics.



209 Burlington Road, Suite 203  
Bedford, MA 01730  
Phone: +1.978.528.4660  
Fax: +1.978.528.0568  
<http://voltdb.com>

