



Efficiently Unifying Memory and Storage for Better Performance.

Aventra™ IRON increases server or instance workload efficiency up to 30X with a unique memory and storage unification technology focused on optimization. With Aventra IRON, customers harness the full processing power of a server enabling them to consolidate workloads and enjoy the financial benefit of reducing server sprawl or repurposing them while also reducing power and cooling costs from the subsequent reduction in IT load.

Aventra IRON is specifically engineered for sustained workloads where server applications face highly concurrent processing requirements, and/or are impacted by heavy user traffic, and/or use highly shared or virtualized resources. It accelerates the overall performance of general-purpose applications via advanced unification of memory and storage. As a result, workload latency and performance variability decreases due to memory unification, yet are safe from power disruptions.

At ClearDB, we provide bare-metal performance to tens of thousands of our cloud customers with our Aventra IRON technology. This technology has been in use for years meeting our database-as-a-service customer need for efficient, data persistent, highly productive workloads for their virtualized applications.

Aventra IRON delivers the most effective, easiest to implement, and least expensive solution to achieve desired application performance goals. It is now available for any Linux based server, whether on premises or in the public cloud via download for CentOS, Ubuntu, and SuSe distributions. The addition of Aventra IRON will effectively boost server performance while ensuring your data is safe and highly available.

With Aventra IRON, we believe:

- Delivering the highest IT infrastructure productivity should be simple, cost-effective, and worry free.

- The very best IT infrastructure performance architecture is one where an application's active-data is located on, or as close to, the CPU's as possible.
- A safe, unified memory approach realizes the most optimal IOPs for virtualized computing use-cases.

In-memory computing is nothing new. The use of RAM to increase workload performance of databases, applications, and other computing services has been part of computing architectures since the era of mainframes from decades ago.

In-memory computing improves workload performance by storing frequently used data, "active-data," in RAM, so it can be retrieved far faster than accessing slower primary storage.

In fact, today it is possible to cache active data in the operating system, a hypervisor, an application, as well as, within storage arrays. Properly leveraged, caching can result in computing services that not only perform better, but cost less to scale. Unfortunately, obtaining these benefits is not easy.

There are big problems with in-memory computing techniques. RAM is a volatile memory. Its contents are lost when computers are shut down or lose power. In addition, RAM is costly compared to other storage alternatives. This is why caching has traditionally been used selectively, only for specific applications in specialized computing environments.

Aventra IRON addresses the problem of losing data in traditional caching solutions by unifying memory and storage through innovation that leverages advanced software to virtualize Non-Volatile Memory (NVM) technology on server architectures with or without NVM present.

Our goal with Aventra IRON is to help our customers realize even greater returns from their IT infrastructure investments.

Aventra™ IRON

6860 Dallas Parkway, Suite 200
Plano, TX 75024

Phone: 855-525-3270 (U.S. callers)
469-828-3439 (International callers)
E-mail: <http://w2.cleardb.net/aventra-support/>