Aventra's family of EC2 optimizers enhance the capabilities of EC2 instances with their unique in-memory technology. With Aventra Optimizer AMI's, EC2 customers harness the processing power of larger instances from smaller instances and deliver targeted workload performance requirements with less costly EC2 instances.

Aventra Optimizers are specifically engineered for sustained workloads where EC2 applications face simultaneous processing requirements, and/or are impacted by heavy user traffic, and/or use highly shared resources. They accelerate compute, memory and storage optimized EC2 instances; high-performance EC2 computing instances; and EC2 instances used for general purpose applications. As a result, workloads are processed close to their CPU's as possible, yet are safe from disruptions in RAM.

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

ClearDB IRON is the most effective, easiest, and least expensive solution that achieves their application performance goals. Now IRON's unique capabilities are available as an Aventra Optimizer AWS-AMI for all kinds of EC2 workloads. They cost effectively boost instance performance while ensuring EC2 data is safe and highly available.

At ClearDB, 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 is possible.
- A safe, in-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. For storage, cache has been used to accelerate I/O performance for the past 25 years.

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. The objective: to better achieve balanced-bound computing that maximizes the productivity of IT infrastructure and increases the performance of application processing.

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 volatile memory. Its contents are erased 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 addresses these problems by innovating on the capabilities of the AWS cloud architecture. This helps deliver even greater value to EC2 customers.

For example, because Aventra Optimizers come pre-configured with solid-state EBS while working closely with AWS hypervisors' caching methods to claim and hold RAM cache independently, Aventra Optimizers deliver 3x-30x workload performance improvements over EC2 instances without Aventra Optimizers.

Our goal with Aventra Optimizers is to help AWS customers realize even greater returns on their investments in their AWS cloud infrastructure usage.

Aventra Optimizers by ClearDB.com

6860 Dallas Parkway, Suite 200
Plano, TX 75024

Phone: 855-525-3270 (U.S. callers)
469-828-3439 (International callers)
http://cleardb.com/aventra-support/