



## Bare Metal Performance for Virtualized Servers

*Aventra IRON for Amazon Web Services (AWS), Elastic Compute Cloud (EC2) enhances their capabilities and other Amazon cloud services with unique in-memory technologies. With Aventra IRON, AWS users deliver the most cost effective returns and add even more value and productivity to their AWS-EC2 deployments.*

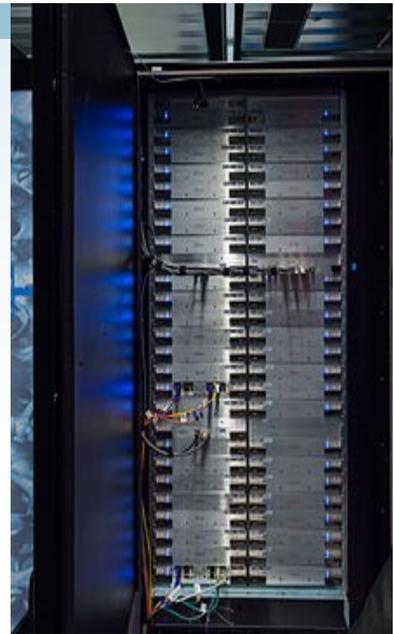
### Executive Summary

Any time information technology operations experience their costs and resource requirements rising faster than the capabilities of their systems, it opens doors for innovations to emerge that reduce the expenses and resources necessary to work around IT system limitations. This is especially true today.

Never have IT operations been pressed so hard to deliver so many applications and services so fast than in our current era of accelerating IT change; characterized by virtualization, internet-of-things, mobility, and big-data. The gap between what businesses need and customers demand, and what IT can deliver from ever tightening budgets continues to widen.

Consequently, IT organizations are turning to innovative computing models so they can continue delivering their goals and objectives, providing more effectively sourced and consumed IT services for their business users and customers alike. One such model is cloud computing.

Cloud computing providers, like Amazon Web Services (AWS), are making large, upfront investments in computing hardware and software. They are handling the heavy lifting of managing that infrastructure, its virtualization, and the other software services required to run today's data centers so IT organizations don't have to. In turn, they offer their infrastructure and services through on-demand and pre-reserved pricing structures with computing resources delivered via the Internet.



#### Aventra IRON for AWS

AWS Primer.....	3
AWS EC2 .....	4
AWS EBS .....	7
Aventra IRON .....	8
Aventra Use-Cases .....	10
Aventra Free Trials.....	12

#### Aventra on AWS

- Achieve maximum productivity at lower costs.
- Deliver faster, more scalable applications and services.
- Process at the highest IOPs for your AWS-EC2 workloads.

The benefit to organizations adopting the cloud computing model is they can provision, when needed, exactly the right types of infrastructure to meet their business and customer demands. The “cloud’s” utility based, variable cost operating structures scales with the business.

In essence, the cloud enables rapid access to flexible and low cost IT infrastructure whether it is being used to run applications that share photos to millions of mobile users, or is supporting business critical services, applications and operations. In some cases, the cloud can replace the high CapEx and OpEx costs necessary to acquire and manage on-prem IT infrastructure.

But the cloud is not just about expanding productivity and lowering costs. It is also about harnessing IT in creative ways to offer new services, get to market faster, and to better understand customers. The

---

### *Aventra IRON for Amazons Web Services boosts the performance of databases, applications, and other computing services by 3x-30x.*

---

#### Leveraging the Cloud

Amazons Web Services empowers IT organizations to focus more on their strategic objectives and less on the daily, tactical management of IT infrastructure.

cloud empowers IT organizations to focus more on strategic business objectives and less on tactical, IT infrastructure management. In fact, the cloud’s on-demand scalability, increased flexibility, and ease of deployment can have a direct impact on a business’s bottom line.

The challenge comes from how to best leverage the cloud to realize this value. Managing IT services in the cloud is different from doing the same for traditional, on-premise IT operations. Fortunately, AWS provides extensive resources, documentation, and support for those IT organizations operating in the AWS cloud, or who are considering migrating to AWS. This is why Aventra selected AWS as its partner. Aventra and AWS are all about our cloud users’ success.

At Aventra, we provide bare-metal performance to tens of thousands of our cloud customers with our IRON software technology. IRON is specifically engineered for sustained workloads where applications face simultaneous processing requirements, are impacted by heavy user traffic, and/or use highly shared resources. IRON has been in use for years, boosting the performance of database servers and application workloads by 3x-30x on commodity infrastructure.

Now Aventra’s unique IRON technology is available as an Amazon Machine Image (AMI) for Amazon Web Services (AWS), cloud users.

Aventra IRON AMI’s for AWS accelerate the compute, memory and storage optimized AWS Elastic Compute Cloud (EC2) instances, high-performance EC2 computing instances, and EC2 instances used for gen-eral purpose applications.

Aventra IRON ensures data is safe from any potential interruptions

in RAM, while maximizing the compute power of AWS-EC2 instances. As a result, databases, applications, and other software services running on AWS can now realize the processing power of larger AWS-EC2 instances from smaller EC2 instances. The result, significantly better workload performance, higher density, and greater concurrency can be delivered using less costly EC2 instances with Avenra IRON for EC2.

This paper describes how to expand the value of Amazon Web Services by using Avenra IRON EC2 instances.

Below is a brief overview of AWS and two AWS services that Avenra IRON leverage to deliver more cost-effective AWS scalability and flexibility. If you are familiar with AWS and its Elastic Compute Cloud (EC2) and Elastic Block Store (EBS) services, feel free to skip ahead. For those new to AWS, we encourage you read on, and to visit [Amazon Web Services](#), and download [The Overview of Amazon Web Services](#).

## An AWS Primer

Amazon Web Services (AWS) cloud platform consists of dozens of cloud computing services that can be combined and tailored to the specific needs of IT organizations. In addition to AWS's own cloud services; hundreds of third party applications, developer tools, and software utilities are also available; facilitating even greater opportunities to leverage AWS to address the IT demands of today's businesses.

All these services are available throughout Amazon's global data center ecosystem. In fact, AWS data centers operate in 14 geographic regions around the world, which are further divided into 38 inter-regions called "Availability Zones." AWS customers retain control and ownership over the region in which their data is physically located. AWS is adding four more regions and nine more zones next year.

To achieve the greatest possible fault tolerance and stability, each AWS region is designed to be physically isolated and completely independent from other regions. AWS customers can increase fault tolerance and create redundancy by replicating applications and data between geographic regions in addition to doing the same across multiple Availability Zones within the same region. Across regions, this is accomplished using both private, high-speed networking and public Internet connections; and is often done to achieve low latency access across the globe, to expand layers of business continuity, and to meet regional data residency and compliance requirements.

Although Availability Zones are physically isolated as well within regions, they are connected to each other via private, fast fiber-optic networks. These Availability Zone networks enable AWS customers to operate production applications and databases that are far more highly available,

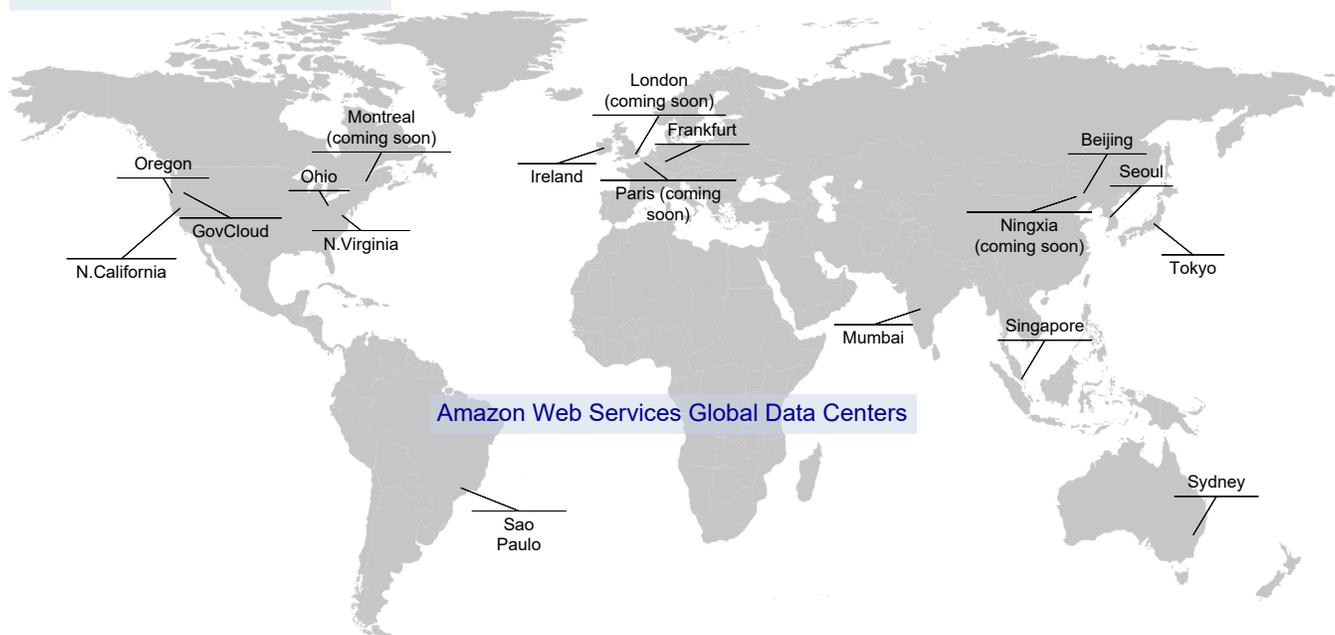
## IT's Strategic Influence

Avenra IRON for AWS helps IT harness creative ways to offer new services, get to market faster, and to better understand customers; all of which directly impact their business's bottom line.

Find an updated map of AWS global data centers [here](#).

fault tolerant, and scalable than would be possible from a single data center. Therefore, AWS customers can architect resources that further increase redundancy and automatically failover without interruption.

Not only can AWS customers select one or more regions to build their computing infrastructure that addresses their computing requirements; the AWS platform is operating system, database, architecture, and programming language agnostic. Accordingly, AWS customers select the services, development, and system administration process models they are most familiar with, enabling them to concentrate on innovation rather than computing infrastructure. These services start with the central element of the AWS platform, the Elastic Compute Cloud (EC2) service.



### *AWS Elastic Compute Cloud (EC2)*

EC2 permits AWS customers to rent a variety of virtual computing configurations so they can run their own applications and services. The different EC2 configurations of CPU, memory, storage, and network capacity are called “instance types” or just “instances.” Current EC2 configurations can be found [here](#).

EC2 instance types are available throughout Amazon’s global data center ecosystem, and they can be launched and terminated as needed. AWS customers pay only for the time an instance is active – hence the term “elastic.” Additionally, EC2 instances are requisitioned and managed through an AWS web service, making it easy to acquire and administer just the number of virtual server instances required. There is no need to wait for and provision hardware and software infrastructure. This enables AWS customers to scale computing resources up or down on the basis of demand, which also empowers quick and cost-effective experimentation, innovation, and iteration.

Secure login using customer defined Key Pairs, optional virtual private networks; and firewalls where customers specify security group sources for IP ranges, protocols, and ports provide access controls to EC2 instances, their data, and to the storage volumes attached. This storage can be configured both for persistent data access across regions and Availability Zones, and for temporary data that is deleted when an EC2 instance is terminated. With root access to each of their provisioned EC2 instance, AWS customers can interact with them just as they would any server. They have complete control of their own AWS-EC2 ecosystem.

### *EC2 Storage Options*

AWS provides several storage options for EC2. Each option has a unique combination of performance and durability that can be used both independently or in combination to address IT requirements. These options offer storage choices for backup, archiving, and disaster recovery; and also block, file, and object storage. Aventura Optimizers leverage AWS Elastic Block Store (EBS) to provide durable, block-level storage for their



*AWS empowers quick, cost-effective experimentation, innovation, and iteration.*

in-memory, IOPs accelerating capabilities. This ensures data processed by Aventura IRON is safe from potential RAM disruptions.

An overview of EBS is presented below along with brief descriptions of the other storage options provided by AWS. Details of these options can be found [here](#), and their relationships is shown on the following graphic.

#### **AWS ELASTIC FILE SYSTEM (EFS)**

EFS provides scalable file storage for EC2 instances that can be used as a common data source for workloads and applications running on multiple EC2 instances. EFS grows and shrinks automatically as files are added or removed, so applications have the storage they need when they need it.

EFS is designed for high availability and durability, providing performance characteristics for a broad spectrum of workloads and applications. For example, multiple EC2 instances can access a single EFS store at the same time, which provides a common data source for workloads and applications. For example, multiple EC2 instances can access a single EFS store at the same time, which provides a common data source for workloads and applications running on more than one EC2 instance.

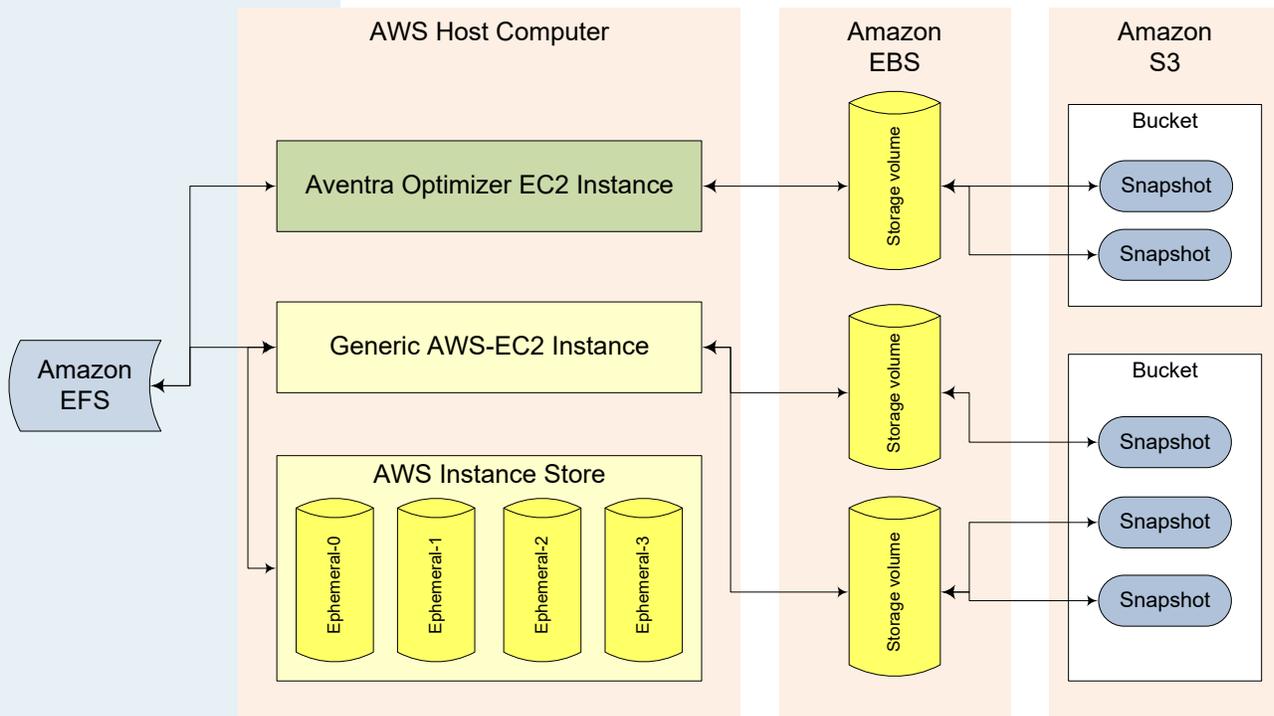
As a result, AWS customers often use EFS for big-data analytics, media processing, web serving, and content management. Its standard system interface makes it easy to integrate EFS with existing applications and tools.



## AWS INSTANCE STORE

Many EC2 instances can access storage from the disks that are physically attached their host computers. AWS refers to this storage as “Instance Store.” The data on an Instance Store volume persists only during the life of the associated EC2 instance. Data on Instance Store volumes is lost as soon as the EC2 instance is stopped or terminated.

Instance Store is ideal for temporary information storage where data changes frequently; such as scratch data, buffers, cache, and other temporary content; or for a data replicated across a fleet of EC2 instances, such as a load-balanced pool of web servers.



An Instance Store consists of one or more Instance Store volumes exposed as block devices. The size of an Instance Store varies by EC2 instance type. The virtual devices for Instance Store volumes are ephemeral 0-23. While an Instance Store is dedicated to a particular EC2 instance, the disk sub-system is shared among those EC2 instances on their host computer.

## AWS SIMPLE STORAGE SERVICE (S3)

S3 is Amazon’s inexpensive data storage infrastructure, and objects are the fundamental entities stored in S3. Every object is contained in a bucket, and these buckets are similar to Internet domain names.

Buckets organize S3 namespaces at the highest level and identify the account owner for that storage. S3 stores these data objects redundantly on multiple devices across multiple facilities.

Because concurrent read or write access is permitted by many separate clients and application threads, and objects are stored redundantly; AWS customers use S3 to quickly and reliably recover from an EC2 instance or application failure. In fact, EC2 instances use S3 to store AMI's (Amazon Machine Images), and to store the snapshots of the AMI's data volumes. *AMI's are preconfigured templates that package all the software needed for launching an EC2 instance, including the unique functionalities and capabilities that make specific AMI's, like Aventura IRON AMI's, valuable.*

Every time an EC2 instance is launched, the AMI for the EC2 instance creates a root storage device on S3 for that instance. This root storage device contains all the information necessary to boot the instance. This means that AMI's can immediately launch other EC2 instances in case of failure. They also can restore data to new instances from EC2 snapshots to ensure business continuity. In addition, snapshots can be used to create multiple, new data volumes; expand the size of an existing data volume, or move data across multiple Availability Zones; thereby making EC2 instance data usage highly scalable as well. S3 is designed to make web-scale computing easier by enabling the storage and retrieval of any amount of data, at any time, from any place within the AWS ecosystem, or anywhere on the web.

## Safe Caching Performance

Aventura IRON leverages AWS Elastic Block Store (EBS) to provide durable, block-level storage for their in-memory, IOPs accelerating capabilities. This ensures data processed by Aventura IRON is safe from RAM disruptions.

*AWS storage options provide unique combinations of performance and durability.*

### **AWS ELASTIC BLOCK STORE (EBS)**

As previously discussed, Aventura IRON AMI's use Elastic Block Store (EBS) to ensure durable, block-level storage for their RAM disk cache, IOPs enhancing software technologies. Since Aventura IRON are delivered as EC2 instances, the benefits of their attached EBS volumes are expanded by the 3x-30x performance boost delivered by Aventura IRON while application data is protected from disruptions in RAM.

EBS volumes are attached to EC2 instances within the same Availability Zone. Each EBS volume is automatically replicated within its Availability Zone to protect from component failure. Since Aventura IRON EC2 instances come pre-configured with EBS, AWS customers only need to select the region where they want to launch their Aventura IRON EC2 instances and everything else is handled automatically.

EBS volumes behave like a raw, unformatted, external block device that are used like any physical drive. Most importantly, they persist independently from the life of EC2 instances, including Aventura IRON EC2 instances. AWS customers are only billed for the amount of EBS storage they use.

As illustrated in the previous graphic, multiple EC2-EBS volumes can be



attached to an EC2 instance. They also can easily be detached from an instance and attached to another EC2 instance whenever needed.

EBS is primarily used for data requiring frequent updates, such as the system drive for an instance, or as storage for database applications, or for throughput intensive applications that perform continuous disk scans. Using Aventura IRON for these and other use-cases significantly improves their performance and workload density capabilities.

The following EBS volume types are provided, each with its own performance characteristics and pricing:

- General purpose solid-state drives \*
- Provisioned IOPs solid-state drives
- Throughput optimized hard-disk drives
- Cold hard-disk drives
- Magnetic storage

\* Aventura IRON EC2 instances use general purpose solid-state, EBS

### *Aventura IRON use EBS solid-state drives to deliver high IOPs performance.*

drives. This ensures that workloads where the dominant performance attribute is IOPs are accordingly accelerated by Aventura IRON in-memory capabilities. These workloads often include applications that face sustained, simultaneous processing requirements, are impacted by heavy user traffic, and/or use highly shared resources.

## **Aventura IRON for AWS**

In the Executive Summary, it was mentioned that Aventura IRON AMI's for AWS-EC2 instances are based on our dynamic random access memory (DRAM/RAM) caching software, IRON, which has been used by our cloud customers for several years. When they need efficient, data persistent, bare-metal productivity for their virtualized, cloud applications; IRON is unquestionably the most effective, least expensive, and easiest solution that achieves their workload performance goals.

Our philosophy in creating IRON and now providing it as Aventura IRON AMI's for AWS-EC2 is:

- The very best IT infrastructure performance architecture is one where an application's active-data is located on or as close to CPU's as is possible.
- A safe, in-memory approach realizes the most optimal IOPs for virtualized computing use-cases.
- Delivering the highest IT infrastructure productivity should be simple, cost-effective, and worry free.

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-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 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, making in-memory computing inexpensive, safe and very effective for clouds and other virtualized computing environments. For AWS, Aventra IRON AMI’s innovate on the capabilities of the AWS cloud architecture, delivering even greater value to AWS-EC2 customers.

For example, because Aventra IRON EC2 instances come pre-configured with solid state AWS-EBS, they add to EC2 and EBS capabilities. Aventra IRON deliver accelerated performance of EC2 instances by factors of 3x-30x by providing RAM disk volumes to the applications, databases, and other services running on Aventra IRON EC2 instances. The active data in these RAM volumes is automatically replicated within AWS Availability Zones, protected from component failure by EBS, so it safely persists beyond the life of the Aventra IRON EC2 instance. In essence, workloads are insanely fast and data is very secure.

Also, Aventra IRON claim and hold RAM independently of other EC2 instances running on the same physical server, so there are no shared contentions for these resources. Additionally, integrating with EBS leverages write-through caching so Aventra IRON data and its data stores are updated simultaneously, ensuring data consistency.

To address I/O performance, Aventra IRON work in close cooperation with AWS’s integrated hypervisors. Leveraging this integration provides the benefits of hypervisors’ caching methods that synchronously

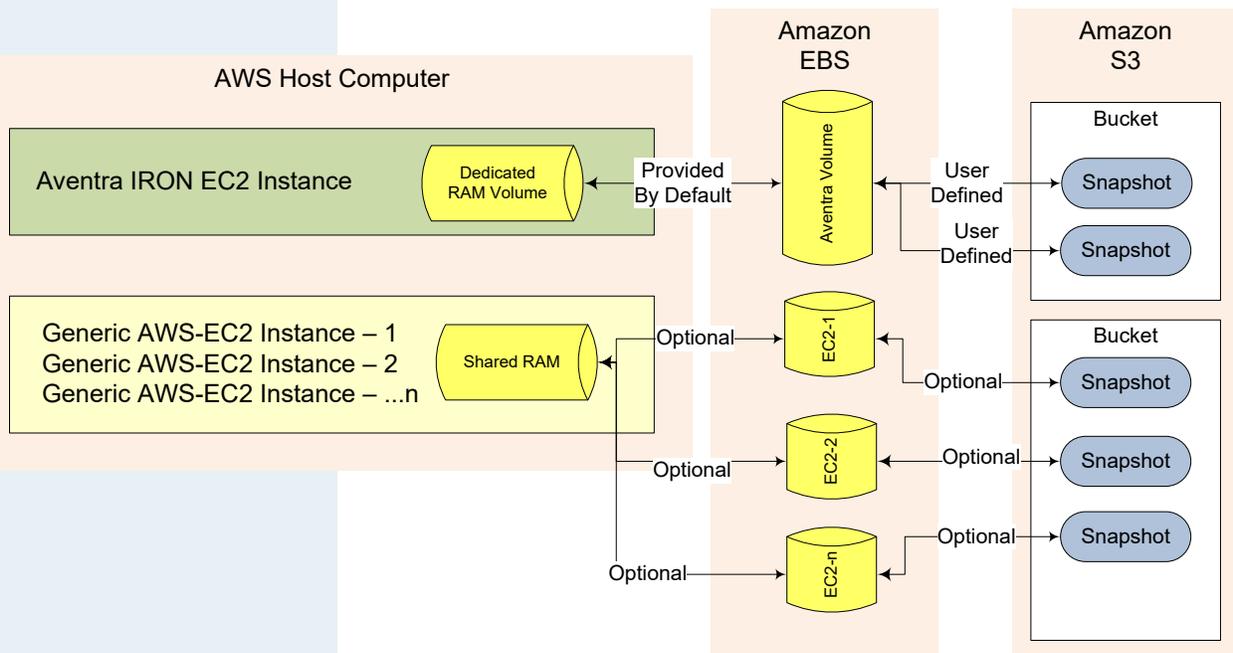
## Bare-Metal Performance for Virtualized Clouds

In use for years by our cloud users, Aventra IRON delivers the performance of bare-metal servers to AWS-EC2 users.

replicates active data across virtual machines. Accordingly, Aventra IRON EC2 instances safely deliver significantly better workload performance by orders of magnitude with less costly EC2 instances, increasing AWS users' returns on their investments in their EC2 infrastructure.

### Aventra IRON Use-Cases

RAM disk volumes markedly improve latency and throughput for many read/write heavy workloads, such as gaming, media sharing, social net-working, Q&A portals, etc. Such data can include remote API calls, or the results of database queries and computationally intensive calculations. Additionally, recommendation engines and high-performance computing simulations with their compute-intensive workloads that manipulate large data sets accessed in real-time across clusters of servers spanning hundreds of nodes, also benefit from in-memory techniques and RAM maximizing computing architectures.



Without the benefits of RAM caching, developers must run large databases and EC2 instances. This not only adds to their costs and potentially decreases application performance, but also limits their ability to scale. Fortunately, Aventra IRON EC2 AMI's greatly simplify the implementation, use, and management of in-memory techniques for EC2 deployments. The following are a few use-cases for Aventra IRON:

#### USE-CASE A: *Retailer's Mobile App Results in Massive Traffic Spikes*

A recently released, mobile phone app has brought tremendous success for an innovative retailer. But it has also resulted in very challenging processing requirements for their EC2 web application deployment.

Daily promotions through social media campaigns are causing huge spikes in traffic as their customers and fans can now upload and share multimedia data to the retailer's site from their new mobile app.

The retailer's current AWS-RDS and AWS-SimpleDB architecture needs to be expanded.

One solution is to deploy a quadruple, extra-large EC2 instance running AWS-RDS; use AWS-SimpleDB as a repository for key-value pair data, and employ S3 for archiving the large quantities of multimedia data now being uploaded to the site. This is a viable solution.

Keeping the S3 implementation for archiving multimedia data, an easier, less expensive alternative is to deploy AWS-RDS on a lower cost Aventura IRON EC2 instance. With Aventura IRON in-memory storage volumes, there is no need for the AWS-SimpleDB repository. The performance of the retailer's site is vastly improved, as well, delivering a much better user experience. Additionally, the scalable data persistence of Aventura's attached AWS-EBS volume protects cached data, providing fail-over options that the retailer's IT organization can deploy.

### Superior Productivity

Aventura IRON EC2 instances are specifically engineered for sustained workloads where application face simultaneous processing requirements, and/or are impacted by heavy user traffic, and/or use highly shared resources.

*Aventura IRON EC2 instances claim and hold RAM independently of other EC2 instances so there are no shared contentions for memory resources.*

#### **USE-CASE B: Large Data Warehouse Needs to Support Global Users**

An organization successfully using AWS for its large data warehouse, now needs to extend access across the globe to dozens of additional users. Presently, they use a high-CPU extra-large EC2 instance running ten, 800GB EBS volumes. The concern is that the additional users and their analytic reporting and query requirements will have negative impacts on the performance of their current EC2 infrastructure. They believe adding another, high-performance EC2 instance is worth the cost to ensure their global users have the best experience for their analytic reporting and query requirements.

Using lower cost, Aventura IRON EC2 instances will provide better performance to more users than both larger EC2 instances combined. This is because the most active query and reporting data resides in-memory on Aventura IRON EC2 instances. Since this data is immediately available to users, Aventura IRON effectively perform at similar or greater levels of concurrent densities as larger AWS-EC2 instances.

In addition, multiple Aventura Optimizer EC2 instances can be deployed to different geographic regions. This brings active query data close to the organization's global users, enabling the very best user experience for the organization's data warehouse.

### **USE-CASE C: *Improving the Performance of a SaaS Application***

Developers of a SaaS application that manages relational data, frequently changing status feeds, and large quantities of multimedia data are pleased with the performance of the AWS infrastructure they deployed a few years ago. Their application is process-intensive and complex, requiring numerous relational data joins to produce results. But using AWS -S3 for multimedia data, AWS-RDS for relational data and status feeds, all of which are powered by large EC2 instances, means their users are very happy with the fault tolerance and high availability of this application.

The developers made good decisions to leverage AWS features such as replication, scalability, and automated backups. They were able to focus on the application's core logic, without hand-coding availability-oriented capabilities into the application. Their AWS deployment is successful.

Now they are being asked by management to increase the application's ability to handle even greater levels of users and concurrency loads without increasing operating budgets.

Accordingly, they are transferring their large EC2-RDS database instances to smaller Aventura IRON EC2 instances. Since their application's status feed are expressed in simple key-value pairs, they can ensure much faster read results to even more users when RDS leverages Aventura IRON's in-memory capabilities. With IRON EC2 instances, they have cost-effectively met their budget targets while delivering the workload productivity goals and objectives requested by the business.

### **Wrap Up**

Aventura IRON for EC2 enhance the capabilities of AWS-EC2 instances and other AWS cloud services and infrastructure with their unique in-memory technology. With Aventura IRON AMI's, AWS users harness in-memory processing power, adding even more value to their AWS deployments and their application productivity goals.

Visit the AWS-Marketplace, search for Aventura, and try EC2. With a 14 day, free trial, it costs nothing to learn how much more productive your workloads can be with low-cost, in-memory EC2 instances.

[Aventura IRON by ClearDB.com](http://AventuraIRONbyClearDB.com)

6860 Dallas Parkway, Suite 200  
Plano, TX 75024

Phone: 855-525-3270 (U.S. callers)  
469-828-3439 (International callers)  
[w2.cleardb.net/aventra-support/](http://w2.cleardb.net/aventra-support/)

