

**Kove:SDM™**

—  
The New Standard  
of Performance

# Maximizing Performance Beyond Hardware Limitations

Discover the key to unlocking the full potential  
of your hardware and your people



Achieve More™

## In this document, you will learn:

- 01 Why server memory is so important to – and currently a top limiter on – performance
- 02 The benefits and drawbacks of the alternative memory technologies available today
- 03 Understand the transformative capabilities of software-defined memory and Kove:SDM™




---

The information in this publication is provided “as is.” Kove IO, Inc. (“Kove”) makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

©2017-2024 Kove IO, Inc. or its subsidiaries. All Rights Reserved.

Kove provides the following notice in compliance with 35 U.S.C. §287(a), see patents: [www.kove.com/legal](http://www.kove.com/legal).

Kove®, Kove:SDM™, ,  KOVE and  are trademarks of Kove or its subsidiaries. Other trademarks may be the property of their respective owners.

Published in the USA [08/23 v1.6] [White Paper].

Kove believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

## The Promise of Software-Defined Memory Becomes a Reality

We can all agree that server memory is important. After all, overcoming memory limitations is one of the most vexing problems facing technologists today, especially since adding memory is quite expensive, and utilization and capacity needs get larger every day. Many technologists are hesitant to get close to maxing out the memory in any given server because they don't want to run out and cause a crash, or worse, hit swap and take weeks to complete. As a result, they tend to only target using smaller portions of available memory.

To complicate matters, there are many memory technologies to choose from. All of them come with distinct benefits and troubling drawbacks. And none fully address the core issues we described above.

Until now.

That's because software-defined memory truly addresses these issues, providing an extensive array of benefits to its users – without meaningful drawbacks.

But despite the obvious benefits of boosting your processing power while requiring less overall server memory, technical challenges kept software-defined memory unrealized for decades. For instance, a common solution was hardware-based Symmetric Multiprocessing (SMP) and then vSMP (software SMP). The key problem: These solutions didn't effectively scale cache-coherency, and performance degraded as memory moved farther away from the CPUs.

## Introduction

Finally, however, the challenge of software-defined memory has been solved.

Kove resolved the issues with its patented Software-Defined Memory, called Kove:SDM™, which doesn't have cache coherency scaling challenges and enables memory to be located hundreds of feet away from the CPUs without affecting predictable, well-understood, real-memory performance.

This whitepaper looks at the various memory technology options available today, explores the key benefits of software-defined memory in general and Kove:SDM™ specifically, and gives a roadmap for implementing the solution that can transform computing capabilities.

As a result of Kove:SDM™, enterprises and their leaders can finally maximize the performance of their people and infrastructure.

## Software-Defined Technologies

While memory was the most vexing to solve, all four software-defined technologies are now available.

01 Computing

02 Memory

03 Networking

04 Storage



Table of Contents

01 The Competitive Memory Landscape ..... 01

02 One More: Software-Defined Memory ..... 02

03 What About Symmetric Multiprocessing & Software SMP? ..... 06

04 A Look at In-Memory Databases, Caches, & Other Options ..... 07

05 Now A Quick Look at CXL ..... 09

06 A Short Overview of How Kove:SDM™ Works ..... 10

07 Eliminate Stranded Memory ..... 11

08 Enjoy Foundational Memory Resiliency ..... 13

09 A Memory Technology That’s Easy To Install ..... 14

10 Enhance Security ..... 15

11 An On-Prem Alternative to the Cloud ..... 16

12 Find Your Kove:SDM™ Use Case ..... 17

13 Results You Can See ..... 18



## The Competitive Memory Landscape

Numerous memory technologies exist. Each approach has supporting use cases and tradeoffs:

- DRAM and SDRAM provide the highest performing volatile memory.
- Storage Class Memory bridges memory from Enterprise Flash to DRAM and is a candidate technology to replace DRAM in the future.
- PMEM, or persistent memory, includes battery-backed DIMMs, flash-backed DIMMs, and technologies that provide byte-addressable memory residing on the memory bus, trading degrees of performance, cost, and capacity.
- Enterprise Flash can speed-up disk-based data, but cannot replace RAM as memory, either in a box or across a fabric.
- Symmetric Multiprocessing (SMP/vSMP) architectures have aggregated CPUs and memory into a larger machine for years, but face scaling challenges.
- Message Passing Interface (MPI) can scale effectively but requires significant, in advance data modeling.
- Supercomputer-type architectures use big iron, custom silicon, and other custom integration, but users do not benefit from server market efficiencies (cost, scale, logistics, etc.).

The problem with these technologies is that any infrastructure – small, large, or on the edge – can never use server memory 100% of the time. While memory costs ~65-85% of a server, it frequently is only used ~15-30% of the time. What is more, without sufficient memory, CPUs are used 5-10% of the time.

As a result, unused memory is stranded inside servers – bonded to particular CPUs, unsharable and isolated from resources elsewhere in the data center.

## One More: Software-Defined Memory

Market forces continue to drive efficiencies, but full software-defined memory has remained a challenge until Kove:SDM.™


It should be noted that software-defined memory is a subset of software-defined technologies (SDT), which is the management of virtualized resources. Other areas of SDT include storage, computing, networking, and more. Software-defined memory enables organizations to manage memory like these other virtualized data center resources, which have been fully operationalized for years.

Realizing the theoretical promise of software-defined memory, Kove:SDM™ enables individual servers to draw from a common memory pool, so jobs receive exactly the amount of memory they need, including amounts far larger than can be contained within a physical server.

When your current job is completed, memory returns to the pool and becomes available for use by other servers and jobs, increasing your memory utilization. In this way, enterprises require less memory in individual servers and less aggregate memory, because it is using memory more strategically on-demand, where and when it is needed.

## Get More Out of Your Hardware (Before)


Technologists are hitting constraints with their hardware, forcing them to buy additional memory off-schedule so they can complete their tasks.






## Get More Out of Your Hardware (After)

With Kove:SDM™, improve your computing performance by 10x to 100x on existing hardware, so technologists can get the memory they need.



## Get More Out of Your Hardware (After)

With Kove:SDM™, improve your computing performance by 10x to 100x on existing hardware, so technologists can get the memory they need.



BARE METAL  
+ KOVE:SDM™

VIRTUALIZATION  
+ KOVE:SDM™

CONTAINERIZATION  
+ KOVE:SDM™

## What About Symmetric Multiprocessing (SMP) and Software SMP (vSMP)?

Kove:SDM™ is not the first attempt at scaling computing infrastructure. Other attempts at it tried symmetric multiprocessing architecture (SMP) and software SMP (vSMP) for aggregating computation. These attempts could be cumbersome to use because they would require specialized hardware, custom programming techniques, and integration efforts. In addition, they also often required hardware and software modifications and did not scale linearly.

Unlike SMP and vSMP approaches, Kove:SDM™ doesn't require specialized hardware or software code changes. It aggregates arbitrarily sized memory across local compute, uses standard, unmodified hardware, delivers nanosecond determinacy for computation, and scales linearly.

A single rack of Kove:SDM™ currently provides the highest performance density in the world, delivering consistent nanosecond determinacy. For example, using a rack of 1U standard servers, Kove:SDM™ delivers to the aggregated PCIe bus limits, more than 1.7 billion IOPS and 1 terabytes/second bandwidth on-demand, from ~2  $\mu$ s to <8  $\mu$ s depending on the interface, and consistent nanosecond computational determinacy. Our software facilitates convenient right-sizing of the compute-memory relationship across servers, racks, and even the entire data center.

## A Look at In-Memory Databases, Caches, and Other Options

While in-memory databases like SAP Hana, Coherence, MemSQL, or eXtremeDB scale to local system size, and can be well-suited for expensive symmetric multiprocessing (SMP) systems and exotic hardware, some “shard” capacity over scale, reducing performance. Latency targets machine-scale for memory and millisecond-scale for TCP/IP or UDP/IP over Ethernet.

A natural complement to in-memory databases, Kove:SDM™ enables you to manage and scale your memory dynamically across servers. Kove:SDM™ provides nanosecond determinacy for computation, and scales linearly even on unmodified standard equipment. Storage class memory technologies do not, although they may assert some overlapping functionality.

Some in-memory grids may “shard” capacity or compute to improve performance, but the gains do not always scale efficiently. While in-memory grids provide storage of data in memory, with co-located compute, they typically scale horizontally, not arbitrarily in all directions like Kove. In-memory grids do not thrive in dynamic or on demand contexts. They target millisecond-scale latency over TCP/IP or UDP/IP. Kove:SDM,™ on the other hand, scales memory arbitrarily and dynamically across servers in the data center, providing microsecond access latency for storage and transport interfaces and nanosecond determinacy for computation. Kove scales linearly using common, unmodified hardware.



What's more, Enterprise Flash, NVMe, Intel® Optane™, 3D XPoint, and Z-NAND address storage performance, but not limited or stranded memory. With these solutions, the capabilities of the server's memory are determined by multiplying the memory slots by the DIMM size. Increasing the capacity with them increases possible performance but it also greatly increases the potential for memory stranding – unused, isolated, and unsharable server memory.

Finally, although they may assert some overlapping functionality, storage class memory mesh technologies are not real memory.

Kove:SDM™, on the other hand, is real memory, delivering a true software-defined memory solution that effectively decouples memory from the server. It aggregates memory into a pooled resource, provisioning any amount of memory where, when, and how it is needed, reducing memory stranding and at times even eliminating it. Specifically, Kove:SDM™:

- Targets real memory, but can also support storage-class memory
- Pools and provisions memory resources for use beyond server limits
- Reduces memory stranding, allowing provisioning rules to move memory where needed
- Supports byte-addressable media, including DRAM, Intel® Optane™, Z-NAND, and more

## Now a Quick Look at CXL

Compute Express Link (CXL) is an open standards initiative, designed to support cache-coherent memory between processors, accelerators, and external memory devices. CXL targets persistent memory for in-rack workloads with natural memory semantics, especially applicable for multi-GPU scaled performance.

Subsuming in-rack memory design, Kove:SDM™ manages memory across racks, on the edge, or across the entire data center. Our software optimizes flexibility in addition to performance, supporting logical (e.g., 1 MiB granularity) rather than device-size (e.g., per DIMM) memory provisioning. It deploys on standard servers and does not require specialized equipment. Kove does not require upgrading CPUs.

While CXL is not yet available and will require new server hardware, Kove:SDM™ will natively support it as well as other evolving and competitive memory specifications and interconnects as they become available.

## A Short Overview of How Kove:SDM™ Works

Kove:SDM™ decouples memory from servers, pooling memory into an aggregate, provisionable, and distributable resource across servers or even the data center, using unmodified commonly available hardware.

Kove:SDM™ provisions memory using policies, similar to virtualized storage, but Kove delivers real memory. Kove:SDM™ provides a global, on-demand resource, delivering memory exactly where, when, and how it is needed. For example, an organization might:

- Use a Kove:SDM™ policy to allocate up to 2 TiB of need-based memory for any of 200 servers between 5p-8:30p
- Develop a provisioning rule that would burst enable creating thousands of additional containers in an instant of need
- Provision a 40 TiB server for a few hours or a 100 TiB RAM disk with RAID backing store for a temporary burst ingest every morning
- Allow an ML job to burst to a size larger than the physical edge server is

Finally, Kove:SDM™ works by using three transparent software components: 1) a Management Console (MC) that orchestrates memory pool usage; 2) the Kove® Host Software that connects applications to a memory pool; and 3) our software that converts servers into memory targets to form a memory pool. Users and applications do not ever need to know that Kove is present. Everything can be fully transparent to the user and application.

## Eliminate Stranded Memory

Kove:SDM™ solves memory stranding by aggregating memory into a pooled resource, shareable and reusable across servers on the edge, the data center and everything in-between. In other words, it decouples or “disaggregates” memory from standard servers, pooling memory into a common resource, for use by any need.

SDM policies then structure the access to the memory pool.  
For example:

- A policy might enable a collection of applications to auto-grow memory beyond physical server size to meet end-of-day needs
- Upon completion, another policy could repurpose memory to upgrade selected servers from, say, 3TiB to 30 TiB of memory


After use, Kove:SDM™ securely zeros-out and returns memory to the Pool for reuse. As a result, Kove:SDM™ provides strong security against attacks targeting memory, and even isn't subject to attacks such as rowhammer.



Kove:SDM™ Memory Pool

Pooled Memory works with any byte-addressable memory, so customers choose whatever they want, and provide for any need they have.

- CPU
- MEMORY



## Enjoy Foundational Memory Resiliency

Kove:SDM™ provides full foundational memory resiliency, which is not available from other memory technologies. And without foundational memory resiliency, recovering from physical memory failure requires:

- For on-prem servers, an organization makes a service call and faces potentially hours or longer of downtime before server recovery
- In the cloud, a user logs back-in and receives an altogether new server (memory failure removes CPU and memory)
- However, Kove:SDM™ provides memory recovery, independent of the CPU. In fact, it can create a replacement allocation within a few hundred milliseconds and does not require server downtime or new server construction. The server simply requests a new allocation and receives memory from another source in the pool.

Meanwhile, Kove:SDM™ also continues to support hardware resiliency technologies, including rank sparing, memory mirroring, and persistent memory, and emerging resiliency technologies.

## A Memory Technology That’s Easy to Install

Kove:SDM™ is implemented on commonly available hardware. For instance, it utilizes two network fabrics: Ethernet for the Control Plane (i.e., command-and-control), and InfiniBand for the Data Plane (i.e., memory data transfer).

Please note that InfiniBand is used with Kove:SDM™ exactly for what it was designed: memory. No special modifications are required to use Ethernet and InfiniBand.

What’s more, Kove engineers anticipated that a plurality of fast “memory interconnects”, such as PCIe, GenZ, Photonics, RoCE, Omni-Path, Slingshot, CXL, and others would be introduced when designing Kove:SDM™. Kove’s design will support new memory interconnects as they become available. Once Kove:SDM™ is installed on servers, customers can create memory targets. Each target can directly connect to a server, or participate in a Memory Pool.

In other words, Kove:SDM™ provides flexible deployment across diverse contexts. For instance, as you can see below, Kove:SDM™ performance running on standard 1U Servers using PCIe 3.0 buses. Combined with provisioning, customers can programmatically match memory to need on-the-fly across servers on the edge, across the data center or anything in-between. This works transparently and requires no code changes.

### Kove:SDM™ Performance

RACK U	MEMORY (TiB)	IOPS	BANDWIDTH (8 μs)
1	3	37 million	50 gigabytes/sec
40	60	1.7 billion	1 terabytes/sec

## Enhance Security

As mentioned earlier, Kove:SDM™ securely zeros memory after use, and returns memory to the pool for reuse.

Similar to “LUN Masking” – i.e., logical unit number masking in storage – Kove:SDM™ provides secure client masking to secure customer isolation and support multi-tenancy. Fabric partitioning is also enforced, and 64-bit keys secure the host-fabric adapters. As a result, Kove:SDM™ provides strong security against attacks targeting memory penetration.

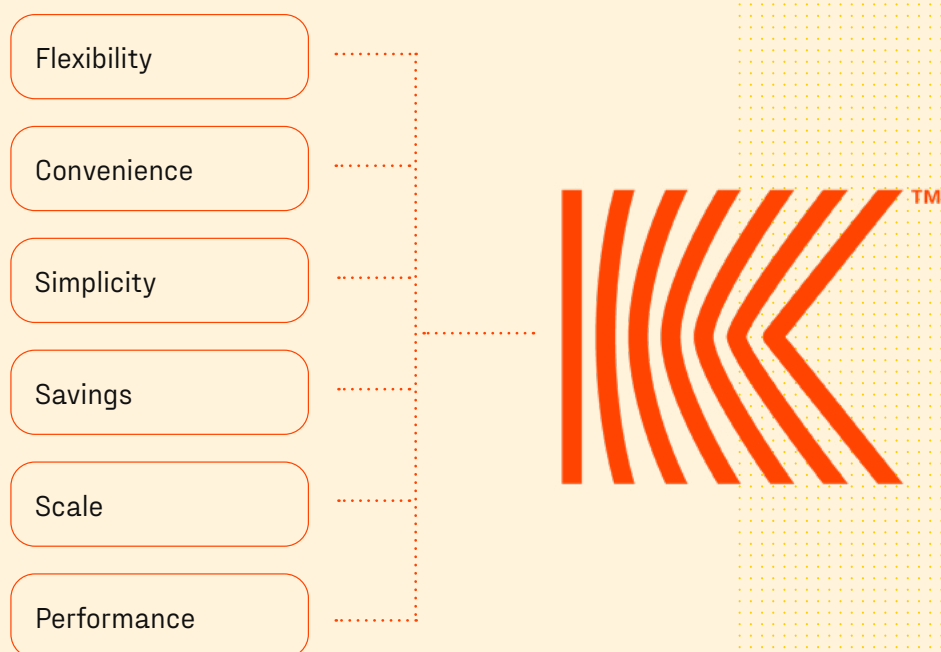


## An On-Prem Alternative to the Cloud

While Kove:SDM™ works with Microsoft® Azure, it can also serve as an alternative to cloud computing for customers who have data that is too large, private or latency sensitive to transport across the Internet, as well as for those who have computation needs at the edge and/or have dynamic memory needs like Artificial Intelligence (AI) and Machine Learning (ML).

### Kove:SDM™ Key Attributes

Kove:SDM™ makes memory a generic, shareable resource, for any need small to large, static or dynamic. Users need not think about DIMMs, out-of-memory (OOM) failures, lousy swap performance, or even memory capacity. Kove:SDM™ abstracts these away from consideration, dynamically providing memory how, where, and when needed.



# Find Your Kove:SDM™ Use Case

Use Kove:SDM™ anywhere you have a memory-favorable application, like analytics, databases, number crunching, VDI, visualization, genomics, Monte Carlo, artificial intelligence, machine learning, enhanced virtual machine infrastructure (increased density, improved performance), and more.

Since memory draws current constantly, power-restricted environments additionally benefit from Kove:SDM.™ By moving memory where and when needed, users can effectively use less capable and fewer servers to achieve comparable results. In so doing, users reduce waste, increase utilization, and receive exactly the required computing resources.

## Kove:SDM™ Use Cases

Kove:SDM™ benefits just about any use case, simple or complex, small or large, inside a server, inside a rack, on the edge, or across the data center.

AI/ML	Containers	Animation
Edge	Analytics	Databases
Cloud	Graphs	Virtualization
Web Services	Life Sciences	Telecommunications
Financial Services	Visualization	Genomics

## Results You Can See

What does all this mean? It means that Kove:SDM™ joins storage, compute, and networking as another foundational technology of enterprise virtualization. It delivers true, flexible, fully-general, fully patented software-defined memory architecture to the marketplace. After all, Kove:SDM™ applies to any context, old or new, simple or complex, performance-enhancing or cost-reducing. You can even make the edge capable like the cloud. Look into it for your enterprise today.

### Key Benefits of Kove:SDM™

Quantitative reports enable customers to measure and validate benefits.

**200%**

or greater ROI

**6 mo**

payback typical

Improved virtual machine and container density

**50%**

power/cooling savings

Processing sensitive workloads on-prem, no tradeoffs

**55%**

effective socket-cost reduction

Virtualize memory, just like any other virtualized resource

**33%**

TCO reduction

**50%**

more compute

Run sensitive workloads on-prem, avoiding risk in the cloud

Reduced solution time with on-demand, any-scale memory

Extend the life of your existing hardware.

## Achieve More™

Accomplish transformative growth with Kove:SDM™, a breakthrough technology that gives your enterprise the memory size and performance it needs when and where it needs it.

Kove:SDM™ empowers individual servers to draw from a common memory pool, including amounts far larger than could be contained within a physical server, so each job receives exactly the memory it needs. It enables 100x more containers, faster processing, smarter AI/ML, data center performance even in edge environments, up to 50% less energy consumption, and unlimited flexibility.

As a result, enterprises and their leaders can do things they could not have done before – including finally maximizing the performance of their people and infrastructure.

