

# Efficient Orchestration of Host and Remote Shared Memory for Memory Intensive Workloads

Juhyun Bae<sup>1</sup>, Gong Su<sup>2</sup>, Arun Iyengar<sup>2</sup>, Yanzhao Wu<sup>1</sup>, Ling Liu<sup>1</sup>

<sup>1</sup>Georgia Institute of Technology, Atlanta, GA30329, USA

<sup>2</sup>IBM Thomas J. Watson Research, New York, USA

## ABSTRACT

Since very few contributions to the development of a unified memory orchestration framework for efficient management of both host and remote idle memory have been made, we present Valet, an efficient approach to orchestration of host and remote shared memory for improving performance of memory intensive workloads. The paper makes three original contributions. First, we redesign the data flow in the critical path by introducing a host-coordinated memory pool that works as a local cache to reduce the latency in the critical path of the host and remote memory orchestration. Second, Valet utilizes unused local memory across containers by managing local memory via Valet host-coordinated memory pool, which allows containers to dynamically expand and shrink their memory allocations according to the workload demands. Third, Valet provides an efficient remote memory reclaiming technique on remote peers, based on two optimizations: (1) an activity-based victim selection scheme to allow the least-active-block of data to be selected for serving the eviction requests and (2) a migration protocol to move the least-active-block of data to less-memory-pressured remote node. As a result, Valet can effectively reduce the performance impact and migration overhead on local nodes. Our extensive experiments on both NoSQL systems and Machine Learning (ML) workloads show that Valet outperforms existing representative remote paging systems by up to 282× throughput improvement and up to 98% latency decrease over conventional OS swap facility for big data and ML workloads, and by up to 5.3× throughput improvement and up to 80% latency decrease over the state-of-the-art remote paging systems. Valet is open sourced at <https://github.com/git-disl/Valet>.

## KEYWORDS

Cloud Computing, Network Memory, Host Memory

### ACM Reference Format:

Juhyun Bae<sup>1</sup>, Gong Su<sup>2</sup>, Arun Iyengar<sup>2</sup>, Yanzhao Wu<sup>1</sup>, Ling Liu<sup>1</sup>. 2020. Efficient Orchestration of Host and Remote Shared Memory for Memory Intensive Workloads. In *The International Symposium on Memory Systems (MEMSYS 2020)*, September 28–October 1, 2020, Washington, DC, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3422575.3422793>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MEMSYS 2020, September 28–October 1, 2020, Washington, DC, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8899-3/20/09...\$15.00

<https://doi.org/10.1145/3422575.3422793>

## 1 INTRODUCTION

Data-intensive and latency-demanding applications [17–19, 24, 25] are typically deployed using the application deployment models, comprised of containers, virtual machines (VMs), and/or executors/JVMs. These applications enjoy high throughput and low latency if they are served entirely from memory. Challenges on these applications increase as workload size becomes larger. When these applications cannot fit their working sets in physical memory of their containers/VMs/executors, they suffer large performance loss in latency, throughput and completion time due to excessive page faults and thrashing.

Most of the existing research studied the above problems and proposed to increase effective memory capacity of VMs/containers by leveraging remote idle memory resources. These proposals promote new architectures and new hardware design for memory disaggregation [26, 28, 30, 32, 60, 62], or new programming models [34, 55]. But they lack of desired transparency at OS, network stack, or application level, hindering their practical applicability. Other efforts [1–7] promotes remote paging with transparency to improve OS paging performance by exploiting the disk-network latency gap via unused remote memory [2, 6, 11, 11, 36–38, 41–44, 79, 81]. However, most existing solutions [2, 6, 11, 42, 62] suffer from high latency limitations due to remote node memory allocation overhead due to receiver-side CPU involvement and the scale-out performance with the large workload. Moreover, existing research efforts have been dedicated either to consolidation of host idle memory across VMs/Containers on the same host or focused on remote memory disaggregation. Very few contributes to the development of a unified memory orchestration framework for efficient management of both host and remote idle memory.

In this paper we presents Valet, an efficient orchestration of host and remote shared memory for big data and machine learning workloads that are memory-intensive in nature. Valet by design aims to address the following three common problems inherent in existing remote memory systems. First, they have latency overhead in the performance critical path due to dynamic connection setup to the remote node(s) and remote memory mapping or disk access scenarios (§2.1). Second, recent effort [6] shows the benefit of remote memory paging with RDMA network and the limitation due to eviction impact when remote node evicts data of local nodes (§2.3). Finally, with the increasing popularity of Container as a Service (CaaS) [9], the container-wide memory imbalance (§2.2) involves managing both node-level memory imbalance and cluster-wide memory imbalance, which pose non-trivial technical challenges [10].

We design and develop Valet to address the above challenges with three original contributions (Figure 1). First, to reduce the hidden latency in the critical path, we redesign the data flow in the




Figure 1: Summary of contributions in Valet.

critical path by introducing a shared memory pool that works as a local cache to remote data. As a result, Valet shortens performance critical path and hides disk access scenarios unlike previous work (§3.3). Second, Valet utilizes idle node level (host) memory across containers via the node-coordinated shared memory pool. This helps to maximize local idle memory utilization and improves application performance on containers (§3.4). Third, Valet provides an efficient remote memory reclaiming technique to minimize the impact of eviction from a remote node on the performance of local containers (local node). Valet achieves the remote memory reclamation by introducing a data migration protocol to move the least active block of data to a remote node of less memory contention. This also helps to maximize remote idle memory utilization across cluster (§3.5).

We evaluated Valet with both memory intensive big-data workloads: Memcached[17], Redis[18], VoltDB[19] on YCSB[20], and memory intensive machine learning workloads: GradientBoosting Classifier, Kmeans clustering, Random Forest Classifier, Logistic regression[22][23][50] and TextRank[45]. Using Valet, throughput improves by up to 282× and latency decrease by up to 98% over conventional OS disk swap. Compared to existing representative remote memory paging system such as nbdX[11] and Infiniswap[6], throughput improves up to 5.3× and latency decrease by up to 80%, demonstrating that Valet is an efficient memory orchestration framework for managing both idle host memory and idle remote memory, and maximizing peek time performance of memory intensive workloads in the presence of transient memory usage variations [49].

In the rest of the paper, we first describe the problems of existing approaches and the challenges to be addressed in Section §2. We present an architectural overview of Valet in Section §3 and §4. We provide discussions in Section §5 and experimental evaluations in Section §6. Section §7 presents the related work and section §8 concludes the paper.

## 2 SOFTWARE CHALLENGES

Before we go into discussions of software challenges in remote paging system, we try to define the term we use in this paper. In

remote paging system, local node(or sender node) handles swapping traffic and remote peer node(or receiver node) allocates memory and registers RDMA Memory Region(MR) blocks as a memory donor for multiple sender nodes. Local node also has multiple peer nodes to distribute paging-out(or write) requests and to read data for paging-in(or read) requests.

### 2.1 Latency Overhead in the critical path.

In in-memory systems utilizing extremely fast DRAM and RDMA, design of critical path in I/O request accounts for the huge portion of overhead in the I/O performance. To understand the burden on latency, we build a prototype of network block device as a baseline. Typical design of RDMA based network block device uses one sided verbs to bypass the kernel at remote side. Before starting I/O operations, connection establishment and mapping to remote MRs are required. We choose dynamic connection and mapping mechanism. We apply power of two choices mechanism for dynamic connection and mapping node selection. Connection and mapping involve querying N remote nodes and selecting the most free node. It also needs address/route resolution, connection establishment and exchanging MR address and keys. Lastly, we add asynchronous disk backup on local side. These design choices are similar to the current state of art remote paging system[6]. We measure latency of each operation to figure out the impact of the latency overhead in general cases. We set our block device as a partition and run FIO microbenchmark on it with the range of 128Kb block I/O size. Write size can be from 4KB up to 128KB and read size is 4KB for both disk and RDMA operations. We run over 10 thousand operations and take an average. Obviously, disk write has the biggest overhead as we expected but we also find out that the latencies for dynamic connection and mapping are not trivial as shown in Table 1.

Operations	Latency (usec)
Disk WR	401336 (58.5%)
Connection	200668 (29.2%)
Mapping	62276 (9%)
Disk RD	20758 (3%)
RDMA WRITE	51.35
COPY	37.57
RDMA READ	36.48

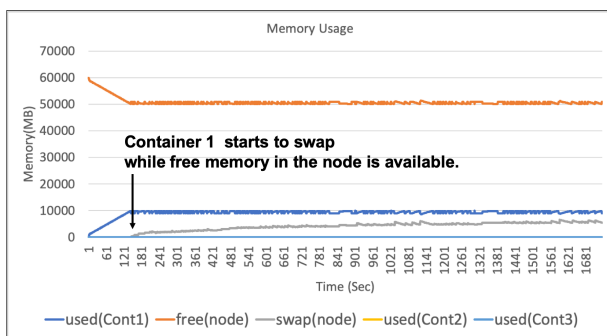
Table 1: Comparison of latency impact on the critical path in typical design of network block device. Connection and mapping have significant overhead on the critical path.

In existing design choices, we find that there are several contributory factors to the inefficiency. First, Performance critical path of I/O is tied with remote sending operation. In one-sided operation, I/O request ends when WC(Work Completion) is polled from CQ(Completion Queue). In two-sided operation, it ends with receiving the response message from receiver node. Second, another latency in the critical path is related to connection establishment and mapping. Connection might not be expensive because it happens only once per receiver node but mapping is. There are two

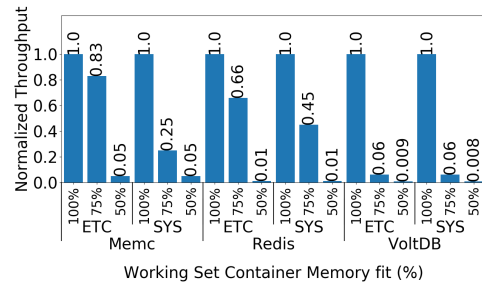
approaches here. One is pre-mapping and the other is dynamic-mapping. Pre-mapping for all possible remote memory in peer nodes removes the mapping latency from the critical path but it is not scalable and also wastes too much resources for internal data structure and buffers that might not be used. Dynamic-mapping is scalable but mapping latency stays in the critical path. As shown in Table 1, connection and mapping cost in the critical path are significant compared to RDMA operations and copy latency. Third, we observe disk access increases during connection and mapping setup because traffic has to be stored in somewhere while remote sending operation is blocked. Those data stored in disk will be accessed by read request later, which causes disk read activities.

### 2.2 Container-wide Memory Imbalance


OS virtualization is a commonly used technology in many cloud servers and datacenters to provide isolated computing environment. There are two ways to set container’s memory constraints. One is to set a limit of memory to each container. Applications on this container can use memory within the limitation. The other is to set unlimited. With unlimited settings, one container can consume all the memory in a node. Then, others running later suffer from performance degradation by swapping to disk. With memory limit, container-wide memory imbalance exists among multiple containers on the same node because cloud systems typically serve heterogeneous guest application workloads and it shows heterogeneous data access patterns during runtime[10]. Figure 2 shows memory imbalance situation where container 1 suffers from swapping while free memory remains on the node. In Figure 3, We run Memcached, Redis, VoltDB with varying the memory limitation of the container. For workload, we use Facebook simulated workload[21] ETC(95%GET and 5%SET) and SYS(75%GET and 25%SET) by using YCSB[20]. Performance severely decreases due to swapping while unused local memory remains in the node. Previous approaches[6, 11] are not free from this container-wide memory imbalance problem.




**Figure 2: Container-wide memory imbalance.** With the container memory limit setting, container cannot use more than its own limitation. We run 3 containers with memory limit in the node and measure memory usage while we run an application in container 1. Container 1 has 10GB memory limit. After 10GB is reached, container 1 suffers from swapping while unused memory remains in the node. Container 2 and 3 are idle at this moment.



**Figure 3: Applications performance with the setting in Figure 2.** Applications suffer from performance degradation while unused memory remains in other containers.



**Figure 4: Experiment setup.** To figure out remote eviction impact on sender node, We run 6 peer nodes for a sender node. Container in the sender node has 5GB limit. When 5GB limit is reached in the sender node, about 18GB workload is evenly distributed into 6 peers in the cluster. We run native applications on M peer(s) at each run to allocate all free memory and cause the remote eviction, where M is 1 to 6. Local memory denotes consumed memory on both sender and remote peer node and remote memory denotes data from sender node.



**Figure 5: Remote eviction impact and imbalanced cluster-wide memory utilization.** Line represents normalized throughput of Redis on sender node and bar represents cluster memory utilization of 6 peers. Remote eviction happens from 1 to 6 peers at each run(Figure 4). Evicted data from peer nodes causes significant performance degradation while unused memory in other peer nodes is not fully utilized(e.g. when only 1 peer evicts all remote memory, it shows 50% decrease in throughput of Redis on sender node).

### 2.3 Remote Eviction Impact

Remote eviction happens due to shortage of free memory when applications in the remote node call for memory. When remote memory eviction happens, performance impact on sender node is inevitable because remote memory is simply deleted from the peer node. Later, all read requests to those deleted data are served from disk in the local node. If the deleted data is highly active one, the impact on sender node is even worse. Another problem is that, finding the most inactive victim is costly. Typical way of handling this is to query write/read activity to multiple sender nodes. This unnecessarily increases communication latency to query sender nodes if the remote memory block is inactive. If the number of queries gets bigger to find the victim well, communication latency increases linearly. In turn, it results in memory pressure on native applications on the peer node due to slow eviction process. Regarding scalability perspective, the impact on sender node due to eviction increases as workload increases. The more pages reside on the peer node, the higher risk of eviction exists and the impact is larger. We measure eviction impact with 23GB workload. We first run Redis with SYS workload to populate 6 peers(See Figure 4). Then, we run native application in the peers until it consumes all free memory. Then, receiver module that manages remote memory evicts remote memory by randomly selecting 1GB sized remote memory block at a time until all blocks are evicted. Figure 5 shows throughput of Redis and cluster-wide unused memory. It shows that eviction causes significant performance loss on the sender node and it becomes worse as the amount of evicted data increases. It also shows that idle memory in the cluster remains unutilized while throughput severely decreases. Addressing remote eviction impact is critical to achieve scalability in distributed in-memory systems.

## 3 DESIGN OVERVIEW

### 3.1 Design considerations

**Maximize CPU utilization** Valet employs asynchronous I/O to maximize CPU utilization. Multi queue block I/O mechanism is working with multiple threads.

**Critical path optimization** Valet achieves shorter latency by optimizing performance critical path. With host-coordinated local mempool, dynamic connection, mapping to remote RDMA MR and local disk access are hidden from the critical path.

**Utilize unused memory** Valet utilizes unused memory both in local and remote memory. Valet tries to utilize the unused memory that is managed by host-coordinated memory pool in a local node first. It exploits container-wide memory imbalance and manages free memory that is not used by other containers. This maximizes idle memory usage in a local node. Valet also utilize unused remote memory in remote nodes by dynamically registering RDMA MR(Memory Region). Local node spreads paging-out data to multiple remote node based on the amount of free memory.

**Reclaim memory efficiently** Reclaiming memory is also crucial for native applications running on both local and remote nodes. Host-coordinated local mempool dynamically expands and shrinks according to the amount of free memory in the local node. Remote RDMA MR also expands and shrinks according to the free memory on the remote node. Valet also provides migration protocol for

remote eviction. It migrates victim data block to other less-memory-pressured nodes. This also maximizes idle memory usage in remote nodes.

**Reliability** Valet uses staging queue and reclaimable queue to maintain the data consistency between local and remote nodes. Unlike parallel reading(paging-in), writing(paging-out) is serialized for data consistency. Valet also provides replication across remote nodes for diskless design. We prefer replication over disk backup. Even though SSD is faster than rotational disk, RDMA is still more than 20 times faster than SSD[83].

**Scalability** Scalability is essential for Valet to process large amount of workload. Valet scales well with multiple remote nodes and distribute workload across remote nodes. Valet also keeps low latency while workload increases. Valet achieves this by removing bottlenecks in the data path.

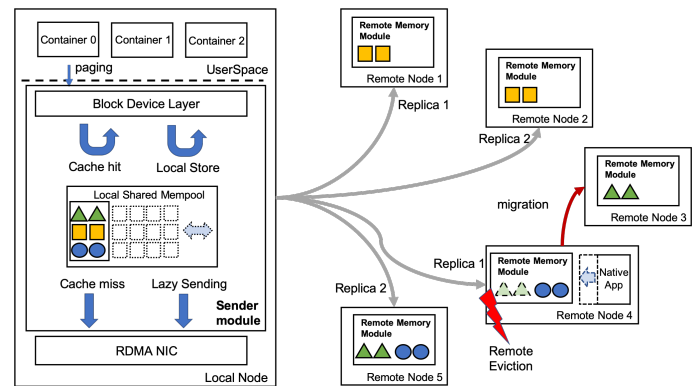


Figure 6: Overall software organization of Valet.

### 3.2 Software organization

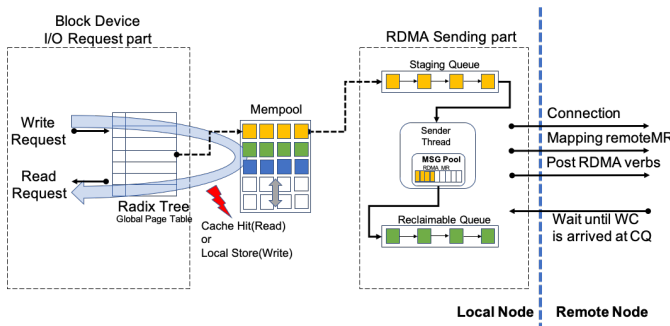
In Figure 6, we show overall software organization in Valet. Valet uses symmetric model. Each node can be a sender and a memory donor(receiver) at the same time although it is not a requirement. Sender module takes swap traffic. Receiver module(Remote Memory module) manages MR blocks as remote memory. Sender node can allocate remote memory across multiple remote nodes. Remote node can serve multiple sender nodes in the cluster.

Valet ends a write request after storing pages in a local shared memory pool(local mempool in short for the rest of the paper). Pages that are stored in the local mempool will be sent out to remote nodes later asynchronously. For read requests, Valet tries to find the page from local mempool first and reads from a remote node if cache is missed. The local mempool extends and shrinks to maximize local idle memory utilization.

Valet tries to spread data evenly across the cluster. If remote eviction happens in a remote node it moves remote memory block to less-memory-pressured node. This maximizes cluster idle memory utilization. Detailed discussion of components in Valet can be found in section 4.

### 3.3 Performance critical path optimization

**Redesign Critical Path.** Valet redesigns performance critical path by having host-coordinated local mempool. For write case, as soon




**Figure 7: Redesign critical path.** With performance path optimization, RDMA Sending part is detached from the performance critical path. Connection, mapping to remote RDMA MR and RDMA verb operations are hiding from performance critical path. For read, Valet shortens read critical path when local cache hit is made.


as it stores pages into local mempool, it can immediately end the I/O request and accomplish shorter write latency. The rest of the remote sending operations are done after the data is written to the local mempool and mempool starts servicing for read request(Figure 7).

Local mempool also functions as a cache for remote data. If data resides in local mempool(cache hit), remote access is not needed. Performance benefit(§3.4) gets larger when local mempool size increases as local hit ratio increases(Figure 8).

**Pipelining the local mempool in the critical path.** Valet also hides connection and mapping of remote MR(Memory Region) latency from the write critical path. This design helps to remove cases that make read latency high too. During connection to a remote node and mapping to a remote memory block, I/O request traffic should be redirected. Valet stores I/O traffic in the local mempool instead of disk. By directly serving read request from the local mempool, it can avoid long read latency due to disk access that is caused by delay of connection and mapping. After connection and mapping are done, local-stored data is sent to remote node to reduce the memory pressure on local mempool.



**Figure 8: Local and Remote hit ratio comparison with various local mempool size.** Local hit ratio increases as local mempool size increases.




**Figure 9:** With performance path optimization, application write latency decreases as the Block I/O size decreases because only I/O request part remains in the critical path.

**Flexible design for input I/O and RDMA buffer size** Unlike previous designs, Valet’s I/O request size is not tied with RDMA MR size. Previous design approaches share the same buffer for RDMA MR and disk writing to avoid extra copies. It is also bounded by max size of hardware disk I/O capacity. **max\_sectors\_kb** determines the number of pages in one Block I/O request. If system has M kb **max\_hw\_sectors\_kb** of hard disk, the size of Block I/O and RDMA MR size for remote paging system are bounded by this hard disk’s physical limitation. Valet can set different value for Block I/O and RDMA MR size regardless the hard disk’s block I/O size limitation if one wishes to add disk backup. The benefit of having different size between Block I/O and RDMA MR is of having a chance to optimize according to various desires. Generally speaking, block I/O size affects the write latency because it adds latency in the critical path while copying pages from Block I/O buffer to RDMA MR. If Block I/O size is set by large number, a Block I/O request has more pages and, in turn, it takes longer time to copy. If the size is small, it takes less time to copy, which leads to shorter latency. See Figure 9. write latency decreases as block I/O size gets smaller. The latency of 32KB is slightly higher than 64KB because of CPU burden due to too many small requests. If RDMA MR size is small, the number of RDMA I/O should increase to send the same amount of data. it may cause WQE(Work Queue Entry) cache miss due to many WQEs injecting to RDMA NIC. It is discovered in previous research[12] that many WQEs cause WQE cache misses in NIC. Valet takes the advantage at this point. Valet can set small size of block I/O to get low latency and use message coalescing and batch sending with large size of RDMA MR to avoid WQE cache miss.

### 3.4 Utilizing unused memory

**Container-wide memory imbalance and Lazy Sending** Local mempool provides a chance to use idle local memory that is not used by other containers by combining them into the local mempool. The local mempool shrinks when the amount of free memory goes below the user defined threshold to guarantee the certain amount of free memory in the node. Then, local pages in the mempool are sent to remote nodes and reclaimed. Before this page replacement happens, this lazy sending scheme best tries to utilize unused local memory and lower the memory pressure on the remote node. Local mempool can grow again when the free memory in the local node goes above the threshold for expansion.

**Impact of the size of mempool** Since the local mempool can dynamically expand and reduce adaptive to the workload dynamics, we first measure the percentage of local hit over remote hit with various size of local mempool to figure out the local mempool’s contribution to local hit. As shown in Figure 8, large size of mempool gets more local hit. If local mempool size decreases, it gets more remote hit. Application latency stays stable with mempool compared to the one without critical path optimization. In Figure 10, we run VoltDB SYS workload with 10 million records and 10 million queries under various ratio of local memory to remote memory by setting container memory limit. 10:0 denotes I/O is served only in local memory and 0:10 denotes only in remote memory.



**Figure 10: Latency comparison with and without critical path optimization.** With performance critical path optimization, application latency stays stable regardless of the various ratio of local to remote memory.

### 3.5 Reclaiming remote memory

**Data migration instead of delete** Valet uses migration protocol when remote data eviction happens in remote node. The major benefit of migration is that it does not hurt the throughput of sender node that maps the data. In order to avoid the I/O blocking during the migration, we allow read requests while migration is in progress. Regarding data consistency concerns between source and destination due to write requests during migration, a local mempool in the sender node can hold the write requests in the local mempool. All the new write requests to the migrating data stay in the staging queue until migration is done. Since these queued write requests are stored in local mempool, read requests to the data are guaranteed to read the latest data by reading from the local mempool. Once migration is done, the sender node can write/read to/from the new destination. Write requests in the staging queue can also be sent out to the new destination (Figure 12). Detailed discussion about consistency is in section §5.2.


**Activity-based Victim Selection on remote node** Unlike read performance, write performance during migration relies on the capacity of local mempool because local mempool is responsible for holding writing requests to the migrating data on MR block during migration. Finding the least-active-MR-block as a victim is crucial factor to lower the memory pressure on the local mempool. To find the least-active-victim, we propose an activity-based victim selection algorithm. We calculate duration since last update for each MR block on the remote node.

$$\text{Non-Activity-Duration} = \text{Time}_{\text{cur}} - \text{Time}_{\text{last\_activity}}$$

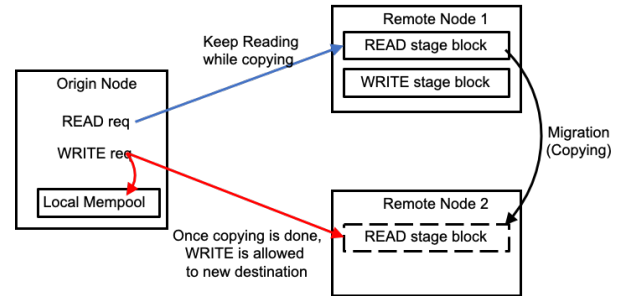
Every MR block on the remote node has small metadata tag and the last write activity is timestamped (See Figure 11). This last active timestamp is updated when this MR block is updated with write requests from its sender node (See Figure 13). Non-Activity-Duration for each MR block will be calculated at the time of eviction process.

Through our observation of write pattern from various workload, we find that the activity cycle of the remote memory block starts with the heavy writes and becomes heavy read state and idle state as time goes by. If a remote MR block starts to receive write requests, it is highly likely followed by read requests. Once heavy read stage is passed, it becomes idle state. This activity cycle is likely repeating by updating with the write operation. The benefits of choosing the


least-active-MR-block are of having low write-request-pressure on the local node while local mempool holds them during migration and reducing communication to query write activity to the sender node. The least-active-MR-block is highly likely to be idle stage. Valet can select this idle block by simply choosing the least-active one without querying to N sender nodes. Then, memory pressure due to holding write operations on the local memory is also limited.



**Figure 11: Format of MR block on remote node.** Tag information is included to calculate Non-Activity-Duration at eviction.



**Figure 12: Read requests are allowed to access remote MR block while copying but write requests stay in the local mempool.** By choosing the least-active MR block as a victim (likely idle or read stage), sender node can lower the memory pressure on the local mempool due to few writes.



**Figure 13: TimeStamp on the MR block is updated by write request.** Then, this block becomes the most-active one in the node. In this figure, the number denotes conceptual last write activity timestamp for each block. The block that has 15 becomes likely the most active block due to recent update. Compared to others, 3 is the most likely read stage due to the longest Non-Activity-Duration among three.

**Sender driven migration protocol** Migration protocol involves many message ping pong and remote procedures. In sender side, it should stop write requests before migration starts and prepare necessary setup with new destination information. In receiver side, source and destination nodes need to communicate each other and



	Linux Mempool	Valet Mempool
Pre-allocation	Guaranteed amount by user define	Guaranteed amount by user define
alloc	Allocation First. If fails, use Mempool	Using Mempool First. If fails, allocates.
free	Freeing memory until It shrinks to the size of pre-allocation(min)	Returning memory to the Pool without freeing
Set Boundary	Only Minimum size can be set.	Grow and Shrink by two user defined threshold(Min, Max)

**Table 2: Comparison between Linux Mempool and Valet Mempool implementation**

**Local Mempool Page Reclaim** Valet uses 24-byte sized `tree_entry` structure to store page references and offset information from one Block I/O request, which represent one transaction in Valet. Staging queue and Reclaimable queue are responsible for tracking these entries that are already sent to remote and that aren't yet. When a write request arrives, the entry for the request is put into Staging queue. Remote Sender Thread takes an entry object from the Staging queue and sends pages to remote nodes. When message coalescing and batch sending are done, those page references are put into the Reclaimable queue. At this moment, pages tracked by Reclaimable queue are safe to be reclaimed because sending is done and a replicated copy is on the remote node. When local Mempool reaches 80% of its size, mempool grows. If mempool cannot grow anymore, it starts to reclaim and provide free pages to new requests directly. For replacement policy, we use LRU in our prototype. Since reclaiming is just moving a page pointer, it takes only a few CPU cycles.

## 4.2 Remote Memory Module

To reduce the CPU overhead on the remote peer node, Valet uses well-known one-sided RDMA verb to bypass kernel on the remote side. Remote Memory module maintains only necessary components and works as passive participant. The main purpose of Remote Memory module is to provide unit sized remote memory registered as MR to multiple sender nodes. Remote Memory module runs in user space and monitors free memory capacity in the remote node. Kernel space MR can utilize physically contiguous memory and reduce PTE cache miss in RNIC but allocation of large physically contiguous memory is challenging. User space MR requires RNIC to cache PTE to access the page because it uses virtually contiguous memory. However, user space allocation is much easier than allocation of large physically contiguous memory in kernel space. We use large MR blocks to reduce the number of MR mapping. Therefore, we choose user space receiver module design for MR block provider. It can dynamically expands and shrinks MR blocks based on the free memory. Remote Memory module also has listener to communicate with other receiver modules when they receive migration protocol messages.

## 4.3 How to track remote pages

Valet provides block device interface. It can be registered as swap space or mounted as a partition with a linear address space. To track the location of remote pages, Valet defines global page address starting from 0 to the end of the user defined space size. This doesn't have to fit the remote memory capacity in the cluster. Then, this virtual address space spans across in the cluster. Mapping partitioned address space to remote peers happens on demand with round-robin or power of two choices. We use power of two choices in our prototype. Each unit sized address space and the same size remote memory blocks on remote nodes is dynamically mapped and internal data structure tracks this mapping information.

## 5 DISCUSSIONS

### 5.1 Fault Tolerance

**Remote node failure.** Valet provides several options for fault tolerance. Either remote replication and local disk backup or mixed approach can be selected as one wishes their fault tolerant level(Table 3). Each combination provides different semantic when remote node failure occurs.

**Local host node failure.** For permanent data store in local host, disk backup option is provided. Then local host writes backup on disk either always or only when writing to remote fails. In paging system example, we provide the same semantic to other paging systems when local node fails.

	w/ Replication	w/o Replication
w/ Disk Backup	1. Access Replica first 2. Access Local Disk if #1 fails	Access Local Disk
w/o Disk Backup	Access Replica	Remote data loss

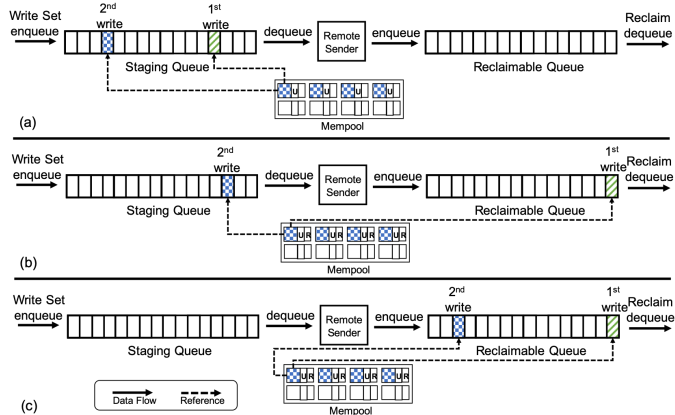
**Table 3: Different level of fault tolerance is provided by combination of replication and disk backup**

### 5.2 Data Consistency

**Between local memory and remote replicas.** An incoming write request's write sets are enqueued into Staging queue as the data is written to the local memory. If an incoming read request finds a page in the local mempool, it is always served from the local mempool directly. The remote pages are accessed only when local mempool does not have the pages due to reclaiming. This guarantees incoming read requests always get the most updated data. Remote Sender Thread takes write sets from Staging queue and sends out to remote nodes in incoming order. Once WC is received, bitmap for the remote page indicates that remote page is ready to read. This guarantees remote node has the same data when it is read. When remote sending is done, the write set is removed from the Staging queue and enqueued into Reclaimable queue. Page slots in the local mempool are reclaimed only through Reclaimable queue. Then only page slots that has replica on the remote node are guaranteed to be reclaimed for the next use.



**Problem with multiple updates on the same page.** There are cases that multiple update write sets are coming on the same page. Then, there are multiple write sets in the Staging queue. The local mempool guarantees the latest data even with the multiple update write sets because the local mempool is always updated immediately and then write sets are enqueued into Staging queue(Figure 17 (a)). The problem may occur between remote sending and reclaiming(Figure 17 (b)). After 1st write set is sent out and enqueued into Reclaimable queue, 1st write set can be reclaimed before 2nd write is sent out. Then reference pointer for 2nd write set is no longer valid. This is solved by having a simple flag. Each page slot in the work entry has 'Update' and 'Reclaimable' flag. Update flag is set on the pages when the multiple write sets are issued on the same page. When 1st write set is reclaimed, Update flag is examined and skipped. When 2nd is sent out, Update flag is removed from the page slots and page slots will be reclaimed when 2nd write set gets reclaimed. The size of Staging queue and Reclaimable queue is the same. The case that the distance between two write sets is longer than or equal to the queue size can be solved by the Update and Reclaimable flags(Figure 17 (b)). Regarding the case that two write sets have shorter distance from each other than the queue size, there is no chance that the 1st write set is reclaimed before 2nd write set is sent out.(Figure 17 (c))



**Figure 17: Data consistency problem in local mempool and remote replicas due to multiple update requests on the same page.** (a) It is solved by a reference counter and an update flag. (b) The case where the distance between two updates are larger than the queue size. (c) The case where the distance is smaller than the queue size.

**Between replicas and disk.** Read is always served directly from the local mempool first. Remote node is only accessed when local mempool doesn't have the page. Likewise, disk is only accessed when remote node fails or pages don't exist in remote nodes including replicas. Pages in local mempool can be deleted only when remote sending or disk backup is done and reclaimable flag is set to the pages. Reclaimable pages are tracked by reclaimable queue. If there is an update write set in local mempool and it is not sent out to remote node or disk yet, a reclaimable flag is removed and an

update flag is set. The latest page is still served from the mempool until an update write set is sent to remote node or written to the disk.

### 5.3 Replication and disk backup

Valet uses replication as default. Compared to disk writing, replication using RDMA is still faster than writing to disk(Table 1). We use replication for all experiments in evaluation.

**Cost of replication and disk backup.** With the local mempool, replication and disk backup do not directly add latency to the critical path because replication and disk backup are behind the local mempool and they are out of the critical path. The cost of having replication and/or disk backup approach is memory pressure on MR pool. Slow releasing of unit MR to the MR pool causes shortage of MR in the MR pool and, in turn, it can make getting MR for incoming requests slow too. Another cost of replication is space cost on the remote node. It requires N time larger remote memory space with N replication.

## 6 EVALUATION

**Setup.** We evaluate Valet with eight popular memory intensive applications listed in Table 4. We run five machine learning applications and three big data applications. We run our experiments on **32 machines with 56Gbps Infiniband cluster** on Cloudlab[51]. Each machine has Xeon E5-2650v2 processor(32 virtual cores 2.6Ghz), 64GB memory(DDR-3 1.86Ghz), 1TB SATA 3.5" rpm hard drives and Mellanox Connect X-3. We run **90 containers on a 32-machine RDMA cluster** and randomly assign one application on each container. We use 4 different memory limitation on each container. We measure the peak memory usage of each application first. The input dataset sizes are from 10GB to 15GB and these create in-memory working sets from 22GB to 35GB. Each machine has 2 to 3 containers and each container fits workload 100%, 75%, 50% and 25% in memory respectively. This makes each container to have memory limitation setting from 5GB up to 24GB and paging-out traffic from 5GB to 27GB according to configuration. Unless stated otherwise, we set 64KB block I/O size, 512KB RDMA message size and 1GB as an unit size of remote MR block. The size of the local mempool will be specified in each experiment. For stable measurement, average is taken from 5 times run for each case. We compare our system with Infiniswap[6] and nbdX[11]. We set Infiniswap as default as their paper mentions and nbdX uses remote ramdisk for storing data.

### 6.1 BigData Workload Performance

In this experiment, we measure Valet performance on Memcached[17], Redis[18] and VoltDB[19]. Memcached and Redis is in-memory distributed caching system through simple key-value interface. VoltDB is ACID-compliant in-memory transactional database. We compare Valet(Figure 19c) to Infiniswap(Figure 19b) and nbdX(Figure 19a). For workload, we use Facebook simulated workload[21] ETC and SYS by using YCSB[20]. We use zipfian distribution for both workload. We first populate the applications with 10 million records in advance and run 10 million queries with ETC and SYS workload. Dataset size is 10GB and working set memory with this dataset




Figure 18: Big Data Workload Latency Comparison of nbdX, Infiniswap and Valet. (VoltDB scale is in right side)




Figure 19: Big Data workload Completion Time Comparison of nbdX, Infiniswap and Valet

Workload	Application/Frameworks	Dataset
Logistic Regression	Scikit learn	87 million samples
Random Forest	Scikit learn	50 million samples
Kmeans	PowerGraph	4 million samples
Gradient Boosting classifier	Caffe	87 million samples
Text Processing	TextRank	1.4 million words
YCSB-Memcached	Memcached	10~50 million records
YCSB-Redis	Redis	10~50 million records
YCSB-VoltDB	VoltDB	10~50 million records

Table 4: Applications and workload used in evaluation.

Valet's improvement over other systems (BigData)			
WorkingSet Fit	Linux	nbdX	Infiniswap
75%	41x(86x)	1.19x(1.5x)	1.17x(1.7x)
50%	59x(101x)	2.5x(3.7x)	2.1x(3.2x)
25%	90x(138x)	3.7x(5.3x)	3.3x(4.5x)

Table 5: Summary of performance improvement comparison of Valet with other systems in Figure 19 and Linux. It show improvement on average and on best case in brackets.

ranges from 15GB to 22GB. Each application takes different amount of working set memory after we populate and run the same 10GB workload. Peak memory for Memcached is 15GB and 22GB for both Redis and VoltDB. Compared to simple key-value structure such as Memcached, its complicated data structure in VoltDB requires more memory. For local mempool setting, we set local mempool dynamically expands and shrinks based on free memory on the host node.

First, Valet shows more stable performance than Infiniswap and nbdX. See Figure 19. nbdX and Infiniswap's completion time increases superlinearly as more pages are sent to remote nodes whereas Valet shows steady performance. Table 5 shows summary of performance improvement comparison in Figure 19. Valet outperforms up to 5.3× over nbdX and up to 4.5× over Infiniswap. Valet also outperforms Linux conventional OS swap by up to 128×. Note that we didn't put the figure of completion time of Linux conventional OS swap.

Second, the performance gap between Valet and other systems increases as more pages are sent to remote nodes. See Table 5. nbdX and Infiniswap's performance is not scalable well compared to Valet as percentage of working set fit decreases.

Third, we also measure average latency of each application on three systems (Figure 18). Compared to 100% working set in-memory fit case, Valet latency increases 1.5×, 2.4× and 3× in 75%, 50% and 25% fit case respectively. nbdX latency increases 2×, 4.9× and 12× in 75%, 50% and 25% fit case respectively. Infiniswap latency increases 2.3×, 5.5× and 12.7× in 75%, 50% and 25% fit case respectively. Conventional OS swap facility latency increases 39.8×, 58.8× and 93.4× in 75%, 50% and 25% fit case respectively.

## 6.2 ML Workload Performance

We use various popular Machine Learning workload (Gradient Boosting classifier, Kmeans, Logistic Regression, Random Forest and TextRank) to measure performance of Valet and other systems in Figure 20. Datasets we use are from 4 million to 87 million samples and they create from 9GB to 34GB workload. For ML, we use click prediction data from Kaggle[47] and NOAA weather dataset[48].




Figure 20: Machine Learning workload Completion Time Comparison of nbdX, Infiniswap and Valet

Valet’s improvement over other systems (ML)			
WorkingSet Fit	Linux	nbdX	Infiniswap
75%	47x(99x)	1.3x(2.3x)	1.4x(2.5x)
50%	69x(187x)	1.5x(2.7x)	1.8x(3x)
25%	109x(282x)	1.8x(2.7x)	2.2x(3.5x)

Table 6: Summary of performance improvement comparison of Valet with other systems in Figure 20 and Linux. It show improvement on average and on best case in brackets.

For TextRank, we use wiki dataset [46], which includes 1.4 million words. We also apply 75%, 50% and 25% working set fit. For local mempool setting, we also set local mempool dynamically expands and shrinks based on free memory on the host node.

Table 6 shows summary of performance improvement comparison in Figure 20. Valet outperforms by up to 282x over Linux, by up to 2.7x over nbdX and by up to 3.5x over Infiniswap. Valet generally shows more stable performance than Infiniswap and nbdX like BigData workload. An interesting observation is that nbdX’s and Infiniswap’s completion time increases superlinearly as workload increases except Kmean’s access pattern is different from others. It intensively accesses certain MR blocks that are mapped in early stage of running rather than access various MR blocks. Since those intensive accessing memory blocks are assigned in early stage of running, it is highly likely in-memory in the local host. This repetitive access pattern also might increase page cache hit in OS. For now, Valet uses LRU on local mempool. Cache replacement policy like MRU that works well on repetitive access pattern might be useful for local mempool replacement policy for this type of workload. We leave this exploration as a future work.

### 6.3 Effectiveness of optimization

**Host/Remote memory distribution** This section compares performance impact of various host/remote memory ratio on application for conventional OS swap(Linux), nbdX, Infiniswap and Valet. We use 25% working set fit configuration for all four systems(Linux, nbdX,Infiniswap and Valet). 75% of working set workload is distributed across remote nodes via paging. For Valet, Valet-75:25, Valet-50:50 and Valet-25:75 denote ratio of local memory to remote memory working set. Valet-LocalOnly and Valet-RemoteOnly denote all working set resides in local node and remote node respectively. Figure 21 shows the comparison.

We highlight several observations below. First, using Valet-LocalOnly, throughput of VoltDB, Redis and Memcached increase by up to 98.5x 226.26x and 15.7x compared to Linux, increase by up to 5.5x 3x and 1.46x compared to Infiniswap, and increase by up to 5.4x 4.7x and 1.17x compared to nbdX.

Second, throughput increases as the size of local mempool increases from Valet-RemoteOnly to Valet-LocalOnly. However, the performance gap between Valet-RemoteOnly and Valet-25:75 is the largest when increasing the size of mempool. Note that Valet-RemoteOnly does not have local mempool component. It shows that critical path optimization with local mempool is the most effective improvement in this experiment.

Third, even with Valet-25:75 that fits only 25% of workload in memory, its performance is comparable to larger percentage cases. By pipelining local mempool in the critical path, it effectively reduces latency(\$3.3). Pages in the mempool are sent to remote and replaced by newly incoming pages. Bigger sized mempool gets more pages in the mempool and it can provide higher local cache hit and, in turn, provide more performance.

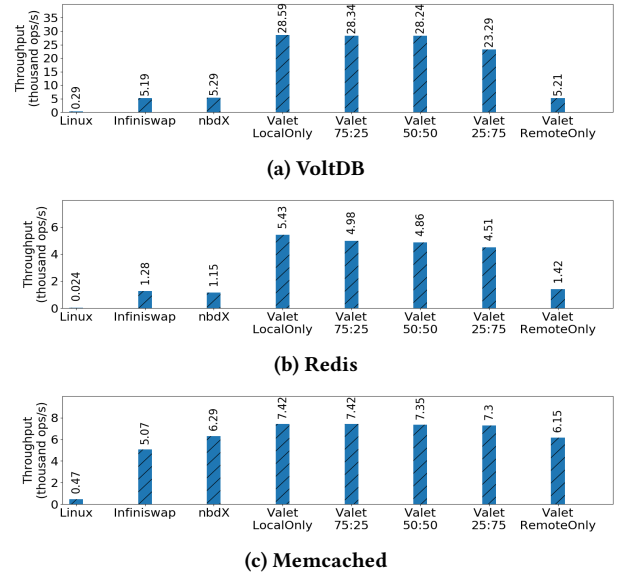


Figure 21: Impact of Host/Remote memory distribution

**Critical path optimization impact on latency** In Table 7, we measure latencies of every events in the critical path with Valet-25:75 setting in Valet and Infiniswap. For workload, we use VoltDB with YCSB SYS workload. ETC and SYS are Facebook simulated workload[21]. ETC is read heavy workload that contains 95% of GET and 5% of SET. SYS is write heavy workload that contains 75% of GET and 25% of SET. In this measurement Valet enables Disk Backup for fair comparison with Infiniswap. Disk access happens when data is not found on remote node(e.g. remote eviction) or there is no connection to node or mapping to MR block. As we expected, Infiniswap’s latency is severely affected by disk access(7b). Infiniswap redirects request traffic to disk while connection and mapping is setup. Valet, on the other hand, avoids disk access due to connection or mapping by having local mempool in the critical path(7a). Request traffic goes to local mempool first and is sent to remote node later. 25% local hit helps to lower the latency further in read request. Write request only spends latency regarding local storing, which is radix tree insertion to track the pages in the local node, data copy from BIO structure to local mempool and enqueueing request to staging queue to track remote sending. Write request doesn’t wait RDMA sending part unlike Infiniswap. Latencies for connection, mapping and disk access are also hidden from critical path. Although connection and mapping are also hidden from write critical path in Infiniswap, the delay causes disk access and, in turn, disk access is not hidden from critical path.

	Read(μsec)		Write (μsec)	
	<b>Average 29.75</b>		<b>Total 35.31</b>	
	Local Hit(25%)	Remote Hit(75%)	Radix	23.9
Valet	Radix 1.39	RDMA 31.39	Copy	9.73
	Copy 2.11	MRpool 0.14	Enqueue	1.68

(a) Valet

	Read(μsec)		Write(μsec)	
	<b>Average 4578.44</b>		<b>Average 19773</b>	
	<b>Total 67480</b>	<b>Total 38.87</b>	<b>Total 1782918</b>	<b>Total 99.45</b>
Infiniswap	Disk (6%)	Remote (94%)	Disk (8%)	Remote (92%)
	Disk 67482	RDMA 36.48	Disk 1782880.4	RDMA 51.35
	Copy 2.2	Copy 2.2	Copy 37.57	Copy 37.57
	MRpool 0.19		MRpool 8.37	


(b) Infiniswap

**Table 7: latency breakdown comparison between Valet and Infiniswap.**

## 6.4 Scalability

In paging system, it is important that sender node handles increasing workload well. In this experiment, we try to figure out Valet’s effectiveness with large workload and scalability(Figure 22). We choose VoltDB because it has the poorest latency among other applications. we measure throughput and 99th percentile tail latency. For Valet, we use 500MB fixed size local mempool to avoid the benefit of the local memory but to include the benefit of critical

path optimization. Throughput decreases as workload increases but Valet still outperforms by up to 7.8× over Infiniswap and by up to 12.65× over nbdX in throughput. 99th percentile tail latency increases by up to 6.45× in Infiniswap and by up to 7.2× in nbdX over Valet. Note that we were not able to measure nbdX with larger workload than 32GB due to unstable running. nbdX uses two sided verb with message pool on both sender and receiver node. We observe sender and receiver side message pool becomes the bottleneck and it severely drops the performance during this experiment.




**Figure 22: Scalability comparison between Valet, nbdX and Infiniswap with increasing workload.**

## 6.5 Eviction Cost

In this experiment, we measure the performance impact on sender node when eviction happens in remote peer nodes(Figure 23). We set the same settings we used in Figure 4. Then, we run Redis with SYS workload because SYS workload has more write operations and this heavy write workload help us observe performance impact when remote eviction happens. After Redis populates peer nodes with about 17GB, we evict certain amount of victim MR blocks selected by Valet with activity-based victim selection. Then, we run Redis with YCSB SYS workload to measure the throughput. We repeat this up to 16GB eviction. Our observation indicates that Valet uses migration instead of eviction when remote eviction occurs and there is no performance impact on local node. However, without migration, one relies on batched-query-based random selection and remote eviction impact is significant on sender node. For example, 2GB eviction(about 8% of workload) results in 50% reduction of throughput on local node.

## 7 RELATED WORK

**Distributed Shared Memory/Disaggregated Memory.** Although Distributed shared memory (DSM) was studied extensively [52–57], DSM suffers poor performance due to high communication overhead. Disaggregated memory has attracted much attention recently and proposed new hardware architecture, and new network protocols to cut down the communication cost [58–62, 64]. Some proposals [65–67, 67–71] show good ways to leverage RDMA technology by exploiting the disk-network latency gap. Remote storage for key-value stores[72–75], distributed objects [76], object replication [82] and swap pages [77–80] show the benefit of RDMA technology in these use cases. Most of these efforts lack of desired transparency and all existing proposals treat and leverage unused host memory as the remote memory, fail to take advantage



**Figure 23: No remote eviction impact in Valet by migration instead of eviction.** We run Redis with 20GB workload. About 16GB is distributed into remote nodes.

of the small performance gap between DRAM and Infiniband compared to disk. Effort to provide transparency at OS, network stack, or application level [1–7] has also been extensively studied. We put summary of comparison of these systems with Valet in Table 8. However, these systems incur CPU overhead at receiver side, fail to handle remote eviction cost, lack of efficient local/remote resource orchestration or optimization in performance critical path.

	Remote MemPager	Nswap	HPBD	Remote Paging for VM	MemX for VM	nbdX	Infiniswap	Valet
Network	TCP/IP	TCP/IP	RDMA	TCP/IP	TCP/IP	RDMA	RDMA	RDMA
Server Side CPU overhead	High	High	High	High	High	High	low	low
Local Resource Orchestration	X	X	X	X	X	X	X	✓
Optimized latency hiding	X	X	X	X	X	X	X	✓
Lazy Sending	X	X	X	X	X	X	X	✓
Remote Eviction Cost	high	high	high	high	high	high	high	low

**Table 8: Comparison with previous approaches.**

## 8 CONCLUSION

Valet addresses three common problems inherent in existing remote memory systems: latency overhead in the performance critical path, remote eviction impact and container-wide memory imbalance. We redesign the data flow in the critical path by introducing a host-coordinated memory pool that works as a local cache to reduce the latency in the critical path of the host and remote memory orchestration. Valet also tries to utilize unused local memory across containers by managing local memory via Valet host-coordinated memory pool, which allows containers to dynamically expand and shrink their memory allocations according to the workload demands. Valet provides an efficient remote memory reclaiming technique on remote nodes by an activity-based victim selection scheme to allow the least-active-block of data to be selected for serving the eviction requests and a migration protocol to move the least-active-block of data to less-memory-pressured remote node. Through extensive experiments on both big data and Machine Learning (ML) workloads,

we show that Valet outperforms conventional OS swap by up to 138× and 282× for big data and ML workloads respectively, and by up to 5.3× completion improvement and by up to 3.5× completion improvement over the state-of-the-art remote paging systems for big data and ML workloads respectively.

## Acknowledgement

The first author thanks the opportunity of the 12-week working experience at IBM T. J. Watson Research Center in Summer 2019 with the group led by Donna N Dillenberger. This work is partially sponsored by the National Science Foundation under Grants NSF 2038029, NSF 2026945 and NSF 1564097, as well as an IBM faculty award.

## REFERENCES

- [1] Evangelos P. Markatos and George Dramitinos (1996) "Implementation of a Reliable Remote Memory Pager" *Proceedings of the USENIX 1996 Annual Technical Conference*
- [2] Tia Newhall, Sean Finney, Kuzman Ganchev, Michael Spiegel "Nswap: A Network Swapping Module for Linux Clusters" *Proceedings of the*
- [3] Shuang Liang, Ranjit Noronha, Dhableswar K. Panda "Swapping to Remote Memory over InfiniBand: An Approach using a High Performance Network Block Device" *Proceedings of the*
- [4] Haogang Chen, Yingwei Luo, Xiaolin Wang, Binbin Zhang, Yifeng Sun, Zhenlin Wang "A Transparent Remote Paging Module for Virtual Machines" *Proceedings of the*
- [5] Umesh Deshpande, Beilan Wang, Shafee Haque, Michael Hines, Katik Gopalan "MemX: Virtualization of Cluster-wide Memory" *Proceedings of the*
- [6] Juncheng Gu, Youngmoon Lee, Yiwen Zhang, Mosharaf Chowdhury, Kang G. Shin "Efficient Memory Disaggregation with INFINISWAP" *Proceedings of the*
- [7] Hikari Oura, Hiroko Midorikawa, Kenji Kitagawa, Munenori Kai "Design and Evaluation of Page-swap Protocols for a Remote Memory Paging System" *Proceedings of the*
- [8] Mel Gorman "Understanding the Linux Virtual Memory Manager"
- [9] MarketIntellica "Global Container as a Service (CaaS) Market Analysis 2013-2018 and Forecast 2019-2024 Understanding the Linux Virtual Memory Manager"
- [10] Ling Liu, Wenqi Cao, Semih Sahin, Qi Zhang, Juhyun Bae, Yanzhao Wu "Memory Disaggregation: Research Problems and Opportunities" *Proceedings of the*
- [11] Mellanox Technology, <https://github.com/accelio/NBDX> "nbdX"
- [12] A. Dragojevic, D. Narayanan, O. Hodson, and M. Castro. "FaRM: Fast remote memory" *Proceedings of the 11th USENIX NSDI, Apr. 2014*
- [13] Mellanox Technology, <http://www.accelio.org> "Accelio"
- [14] C. A. Reiss, Tumanov, G. R. Ganger, R. H. Katz, and M. A. Kozuch. "Heterogeneity and dynamicity of clouds at scale: Google trace analysis" *In SoCC, 2012*
- [15] C.A. Reiss. "Understanding Memory Configuration for In-Memory Analytics." *PhD thesis, UC Berkeley, 2016*
- [16] A. Samih, R. Wang, C. Maciocco, M. Kharbutli, and Y. Solihin. "Collaborate memories in clusters: Opportunities and challenges" *Transactions on Computational Science XXII, Berlin, Germany:Springer, 2014, pp.17-41*
- [17] "Memcached, a distributed memory object caching system" <https://memcached.org>
- [18] "Redis, an in-memory data structure store" <https://redis.io>
- [19] "VoltDB, a translytical in-memory database" <https://github.com/VoltDB/voltdb>
- [20] B.F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. "Benchmarking cloud serving systems with YCSB" *In SoCC, 2010*
- [21] B. Atikoglu, Y. Xu, E. Frachtenberg, S. Jiang, and M Palesczny. "Workload analysis of a large-scale key-value store" *In SIGMETRICS, 2012*
- [22] "Scikit-learn, a free software machine learning library" <https://github.com/scikit-learn/scikit-learn>
- [23] Jia, Yangqing and Shelhamer, Evan and Donahue, Jeff and Karayev, Sergey and Long, Jonathan and Girshick, Ross and Guadarrama, Sergio and Darrell, Trevor "Caffe: Convolutional Architecture for Fast Feature Embedding" *arXiv preprint arXiv:1408.5093, 2014*
- [24] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, et al. "Apache hadoop yarn: Yet another resource negotiator." *In SoCC, 2013.*
- [25] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. Mc-Cauley, M. J. Franklin, S. Shenker, and I. Stoica. "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing." *In NSDI, 2012.*
- [26] "HP: The Machine" <http://www.labs.hp.com/research/themachine>.
- [27] "Intel RSA" <http://www.intel.com/content/www/us/en/architecture-and-technology/rsa-demo-x264.html>.

- [28] "Intel RSA" <http://www.intel.com/content/www/us/en/architecture-and-technology/rsa-demo-x264.html>.
- [29] K. Lim, J. Chang, T. Mudge, P. Ranganathan, S. K. Reinhardt, and T. F. Wenisch. "Disaggregated memory for expansion and sharing in blade servers." In *ISCA*, 2009.
- [30] K. Lim, Y. Turner, J. R. Santos, A. AuYoung, J. Chang, P. Ranganathan, and T. F. Wenisch. "System-level implications of disaggregated memory." In *HPCA*, 2012.
- [31] P. X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy, and S. Shenker. "Network requirements for resource disaggregation." In *OSDI*, 2016.
- [32] M. K. Aguilera, N. Amit, I. Calciu, X. Deguillard, J. Gandhi, P. Subrahmanyam, L. Suresh, K. Tati, R. Venkatasubramanian, and M. Wei. "Remote memory in the age of fast networks." In *SoCC*, 2017.
- [33] J. Nelson, B. Holt, B. Myers, P. Briggs, L. Ceze, S. Kahan, and M. Oskin. "Latency-tolerant software distributed shared memory." In *USENIX ATC*, 2015.
- [34] R. Power and J. Li. "Piccolo: Building fast, distributed programs with partitioned tables." In *OSDI*, 2010. In *OSDI*, 2010.
- [35] P. Zhang, X. Li, R. Chu, and H. Wang. "Hybridswap: A scalable and synthetic framework for guest swapping on virtualization platform." In *INFOCOM*, 2015.
- [36] E. A. Anderson and J. M. Neefe. "An exploration of network ram." *Technical Report, Computer Science Division, University of California, Berkeley*, 1998.
- [37] H. Chen, Y. Luo, X. Wang, B. Zhang, Y. Sun, and Z. Wang. "A transparent remote paging model for virtual machines." In *International Workshop on Virtualization Technology*, 2008.
- [38] S. Dwarkadas, N. Hardavellas, L. Konthanassis, R. Nikhil, and R. Stets. "Cashmere-vm: Remote memory paging for software distributed shared memory." In *IPPS/SPDP*, 1999.
- [39] M. D. Flouris and E. P. Markatos. "The network ramdisk: Using remote memory on heterogeneous nodes." In *Journal of Cluster Computing*, 1999.
- [40] S. Liang, R. Noronha, and D. K. Panda. "Swapping to remote memory over infiniband: An approach using a high performance network block device." In *Cluster Computing*, 2005.
- [41] E. P. Markatos and G. Dramitinos. "Implementation of a reliable remote memory pager." In *USENIX ATC*, 1996.
- [42] Q. Zhang and L. Liu. "Shared memory optimization in virtualized cloud." In *CLOUD*, 2015.
- [43] M. Hao, G. Soundararajan, D. R. Kenchamma-Hosekote, A. A. Chien, and H. S. Gunawi. "The tail at store: A revelation from millions of hours of disk and ssd deployments." In *FAST*, 2016.
- [44] K. Elmelegy, C. Olston, and B. Reed. "Spongefiles: Mitigating data skew in mapreduce using distributed memory." In *SIGMOD*, 2014.
- [45] Rada Mihalcea and Paul Tarau. "TextRank: Bringing Order into Text." *Empirical Methods in Natural Language Processing*, 2004
- [46] "TextRank Dataset : <http://matmahoney.net/dc/textdata.html>"
- [47] "ML Dataset : <https://www.kaggle.com/c/outbrain-click-prediction/data>"
- [48] "ML Dataset : <https://www.kaggle.com/noaa/gsoad>"
- [49] W. Cao, L. Liu. "Hierarchical Orchestration of Disaggregated Memory." *IEEE Transactions on Computers*, 2020.
- [50] Yucheng Low, Joseph Gonzalez, Aapo Kyröla, Danny Bickson, Carlos Guestrin and Joseph M. Hellerstein "GraphLab: A New Parallel Framework for Machine Learning." *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010
- [51] Dmitry Duplyakin and Robert Ricci and Aleksander Mariq and Gary Wong and Jonathon Duerig and Eric Eide and Leigh Stoller and Mike Hibler and David Johnson and Kirk Webb and Aditya Akella and Kuangching Wang and Glenn Ricart and Larry Landweber and Chip Elliott and Michael Zink and Emmanuel Cecchet and Snigdhaswin Kar and Prabodh Mishra. "The Design and Operation of CloudLab." *USENIX Annual Technical Conference (ATC)*, 2019 <https://www.flux.utah.edu/paper/duplyakin-atc19>
- [52] C. Amza, A. L. Cox, S. Dwarkadas, P. Keleher, H. Lu, R. Rajamony, W. Yu, and W. Zwaenepoel. "Treadmarks: Shared memory" *computing on networks of workstations. Journal of Computer*, 1996.
- [53] J. K. Bennett, J. B. Carter, and W. Zwaenepoel. "Munin: Distributed shared memory based on type-specific memory coherence." In *SIGPLAN*, 1990.
- [54] K. Li and P. Hudak. "Memory coherence in shared virtual memory systems" In *TOCS*, 1989.
- [55] J. Nelson, B. Holt, B. Myers, P. Briggs, L. Ceze, S. Kahan, and M. Oskin. "Latency-tolerant software distributed shared memory." In *USENIX ATC*, 2015.
- [56] D. J. Scales, K. Gharachorloo, and C. A. Thekkath. "Shasta: A low overhead, software-only approach for supporting fine-grain shared memory." In *ACM SIGOPS Operating Systems Review*, 1996.
- [57] I. Schomas, B. Falsafi, A. R. Lebeck, S. K. Reinhardt, J. R. Larus, and D. A. Wood. "Fine-grain access control for distributed shared memory." In *SIGPLAN*, 1994.
- [58] K. Asanovic and D. Patterson. "Firebox: A hardware building block for 2020 warehouse-scale computers." In *USENIX FAST*, 2014.
- [59] P. Faraboschi, K. Keeton, T. Marsland, and D. S. Milojevic. "Beyond processor-centric operating systems." In *HOTOS*, 2015.
- [60] P. X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy, and S. Shenker. "Network requirements for resource disaggregation." In *OSDI*, 2016.
- [61] S. Han, N. Egi, A. Panda, S. Ratnasamy, G. Shi, and S. Shenker. "Network support for resource disaggregation in next-generation datacenters." In *HotNets*, 2013.
- [62] K. Lim, J. Chang, T. Mudge, P. Ranganathan, S. K. Reinhardt, and T. F. Wenisch. "Disaggregated memory for expansion and sharing in blade servers." In *ISCA*, 2009.
- [63] K. Lim, J. Chang, T. Mudge, P. Ranganathan, S. K. Reinhardt, and T. F. Wenisch. "Disaggregated memory for expansion and sharing in blade servers." In *ISCA*, 2009.
- [64] P. S. Rao and G. Porter. "Is memory disaggregation feasible?: A case study with spark sql." In *ANCS*, 2016.
- [65] F. Mietke, R. Baumgartl, R. Rex, T. Mehlan, T. Hoefler, and W. Rehm. "Analysis of the memory registration process in the mellanox infiniband software stack. 8 2006." In *Euro-Par*, 2006.
- [66] C. Guo, H. Wu, Z. Deng, G. Soni, J. Ye, J. Padhye, and M. Lipshteyn. "Rdma over commodity ethernet at scale." In *SIGCOMM*, 2016.
- [67] M. S. Birrittella, M. Debbage, R. Huggahalli, J. Kunz, T. Lovett, T. Rimmer, K. D. Underwood, and R. C. Zak. "Intel R omni-path architecture: Enabling scalable, high performance fabrics." In *HOTI*, 2015.
- [68] S.-Y. Tsai and Y. Zhang. "Lite kernel rdma support for datacenter applications." In *SOSP*, 2017.
- [69] Y. Zhu, H. Eran, D. Firestone, C. Guo, M. Lipshteyn, Y. Liron, J. Padhye, S. Raindel, M. H. Yahia, and M. Zhang. "Congestion control for large-scale rdma deployments." In *ACM SIGCOMM Computer Communication Review*, 2015.
- [70] "<http://www.infinibandta.org/>"
- [71] R. Recio, B. Metzler, P. Culley, J. Hilland, and D. Garcia. A "remote direct memory access protocol specification." *Technical Report*, 2007.
- [72] A. Dragojevic, D. Narayanan, O. Hodson, and M. Castro. "Farm: Fast remote memory." In *NSDI*, 2014.
- [73] A. Dragojevic, D. Narayanan, E. B. Nightingale, M. Renzelmann, A. Shamis, A. Badam, and M. Castro. No compromises: "distributed transactions with consistency, availability, and performance." In *SOSP*, 2015.
- [74] A. Kalia, M. Kaminsky, and D. G. Andersen. "Using rdma efficiently for key-value services." In *SIGCOMM*, 2014.
- [75] C. Mitchell, Y. Geng, and J. Li. "Using one-sided rdma reads to build a fast, cpu-efficient key-value store." In *USENIX ATC*, 2013.
- [76] J. Waldo, G. Wyant, A. Wollrath, and S. Kendall. "A note on distributed computing." In *International Workshop on Mobile Object Systems*, 1996.
- [77] D. E. Comer and J. Griffioen. "A new design for distributed systems: The remote memory model." *Technical Report, Department of Computer Science, Purdue University*, 1990.
- [78] M. J. Feeley, W. E. Morgan, E. Pighin, A. R. Karlin, H. M. Levy, and C. A. Thekkath. "Implementing global memory management in a workstation cluster." In *ACM SIGOPS Operating Systems Review*, 1995.
- [79] M. D. Flouris and E. P. Markatos. "The network ramdisk: Using remote memory on heterogeneous nodes." *Journal of Cluster Computing*, 1999.
- [80] S. Koussih, A. Acharya, and S. Setia. "Dodo: A user-level system for exploiting idle memory in workstation clusters." In *HPDC*, 1999.
- [81] S. Liang, R. Noronha, and D. K. Panda. "Swapping to remote memory over infiniband: An approach using a high performance network block device." In *Cluster Computing*, 2005.
- [82] Y. Zhang, J. Yang, A. Memaripour, and S. Swanson. "Mojim: A reliable and highly-available non-volatile memory system." In *ACM SIGPLAN Notices*, 2015.
- [83] Jian Yang Joseph Izraelevitz Steven Swanson "Orion: A Distributed File System for Non-Volatile Main Memories and RDMA-Capable Networks." In *USENIX FAST*, 2019.
- [84] Yibo Zhu, Haggai Eran, Daniel Firestone, Chuanxiong Guo, Marina Lipshteyn, Yehonatan Liron, Jitendra Dattatraya Padhye, Shachar Raindel, Mohamad Haj Yahia, Ming Zhang "Congestion Control for Large-Scale RDMA Deployments." *SIGCOMM*, 2015
- [85] Diego Crupnicoff, Sujal Das, Eitan Zahavi "Deploying Quality of Service and Congestion Control in Infiniband-based Data Center Networks." *White Paper, Mellanox Technologies*
- [86] Radhika Mittal, Vinh The Lam, Nandita Dukkipati, Emily Blem, Hassan Wasel, Monia Ghobadi, Amin Vahdat, Yaogong Wang, David Wetherall, David Zats "TIMELY: RTT-based Congestion Control for the Datacenter" *SIGCOMM*, 2015
- [87] Yuliang Li, Rui Miao, Hongqiang Harry Liu, Yan Zhuang, Fei Feng, Lingbo Tang, Zheng Cao, Ming Zhang, Frank Kelly, Mohammad Alizadeh, Minlan Yu. "HPCC: High Precision Congestion Control" *SIGCOMM*, 2019