Data Computing Laboratory
데이터 컴퓨팅 연구실

The 9th Asian Tier Center Forum

Evolving the CDS: A Proposal for Collaborative Research

Hyeongbin Kang

Ph.D Student,
Chungbuk National University (South Korea),
Hyeongbin.kang@cern.ch
Sep 25, 2025

The 9th Asian Tier Center Forum

CONTENTS

1. Introduction

“ Evolving the CDS: A Proposal for Collaborative Research (Custodial Disk Storage) ”

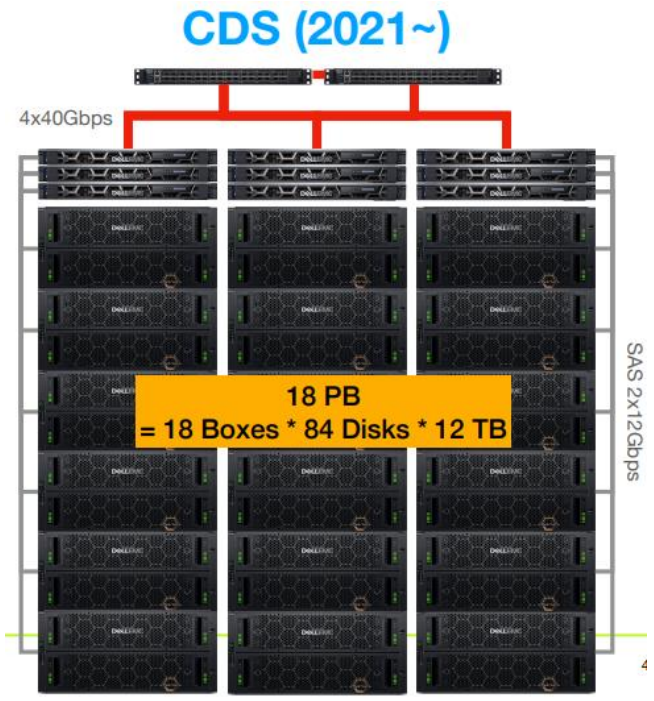

EOS

5.4 Feature

1. Introduction

Overview of Research Experience

- **Main Topic:** Data Center Resource Optimization
- **Participated Research:** Tapeless Storage & Distributed File System

Virtual Machine Placement

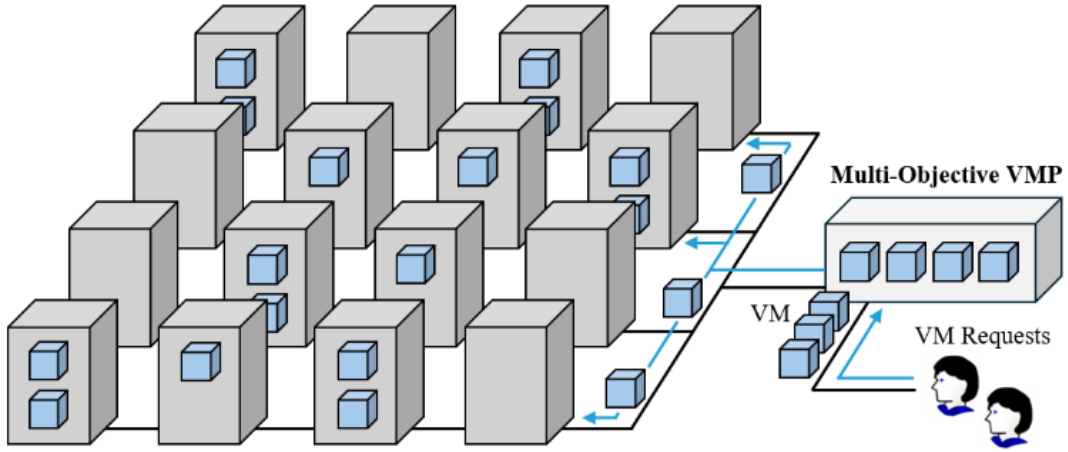

Distributed File System

The Challenge of Data Center Optimization

Virtual Machine Placement (VMP)

Optimizing cloud data centers through resource management techniques

Power Efficiency	Resource Utilization	Quality of Service
1 Minimize Energy Consumption	2 Maximize Allocation of Works	3 Ensure Service Level

Solution

Multi-Objective Optimization-based Virtual Machine Replacement

- Proposed Replacement Framework
- Simulator-based Experimental Design
- Performance Evaluation

The Complexity of the VM Placement Problem

The VM Placement Problem involves mapping virtual machines to physical hosts, subject to multi-dimensional resource constraints.

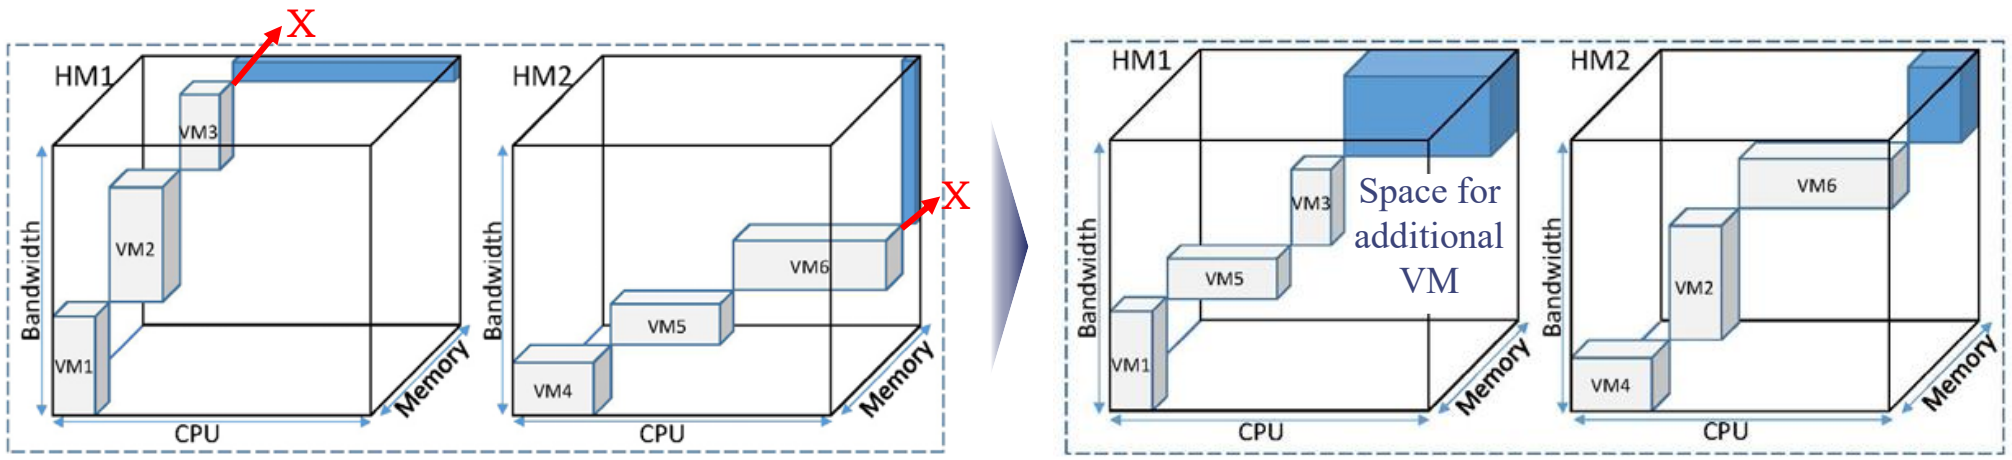

Due to its NP-hard complexity, most real-world systems forgo optimal solutions and instead use rule-based heuristics for VM scheduling.

Multi-Objective Evolutionary Algorithm (MOEA)

A method to find approximate solutions for optimizing multiple, conflicting objective functions.

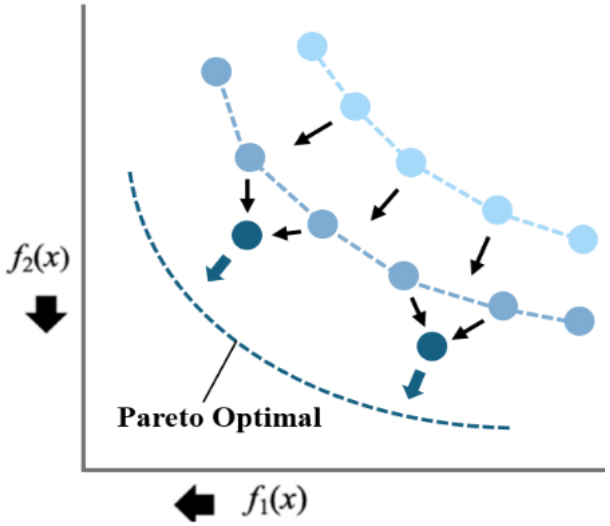

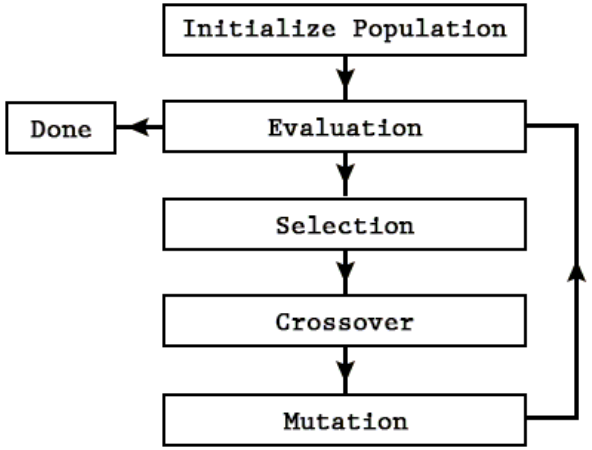

The result is not a single solution, but a Pareto front of optimal trade-off solutions.

Virtual Machine Replacement

This technique optimize objectives by globally replacing VMs within the datacenter using live migration

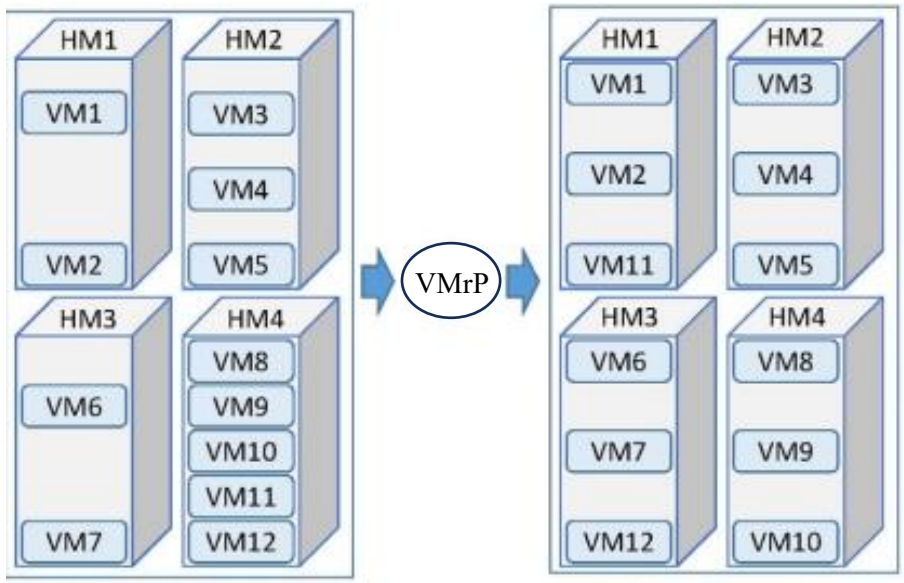

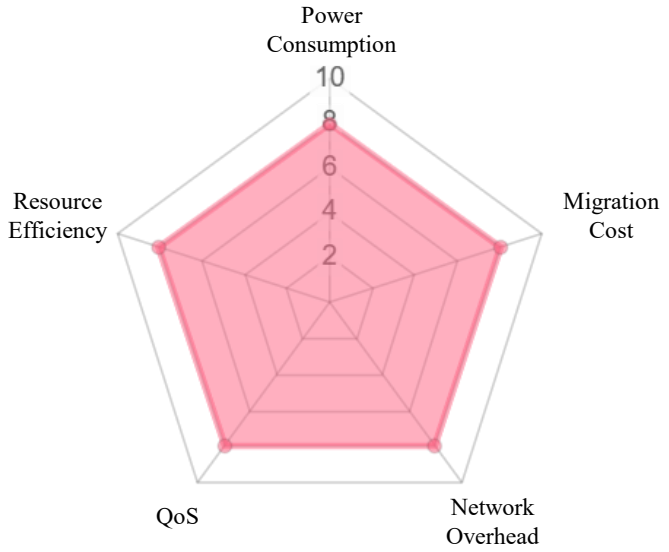

Power Characteristics of Heterogeneous Servers

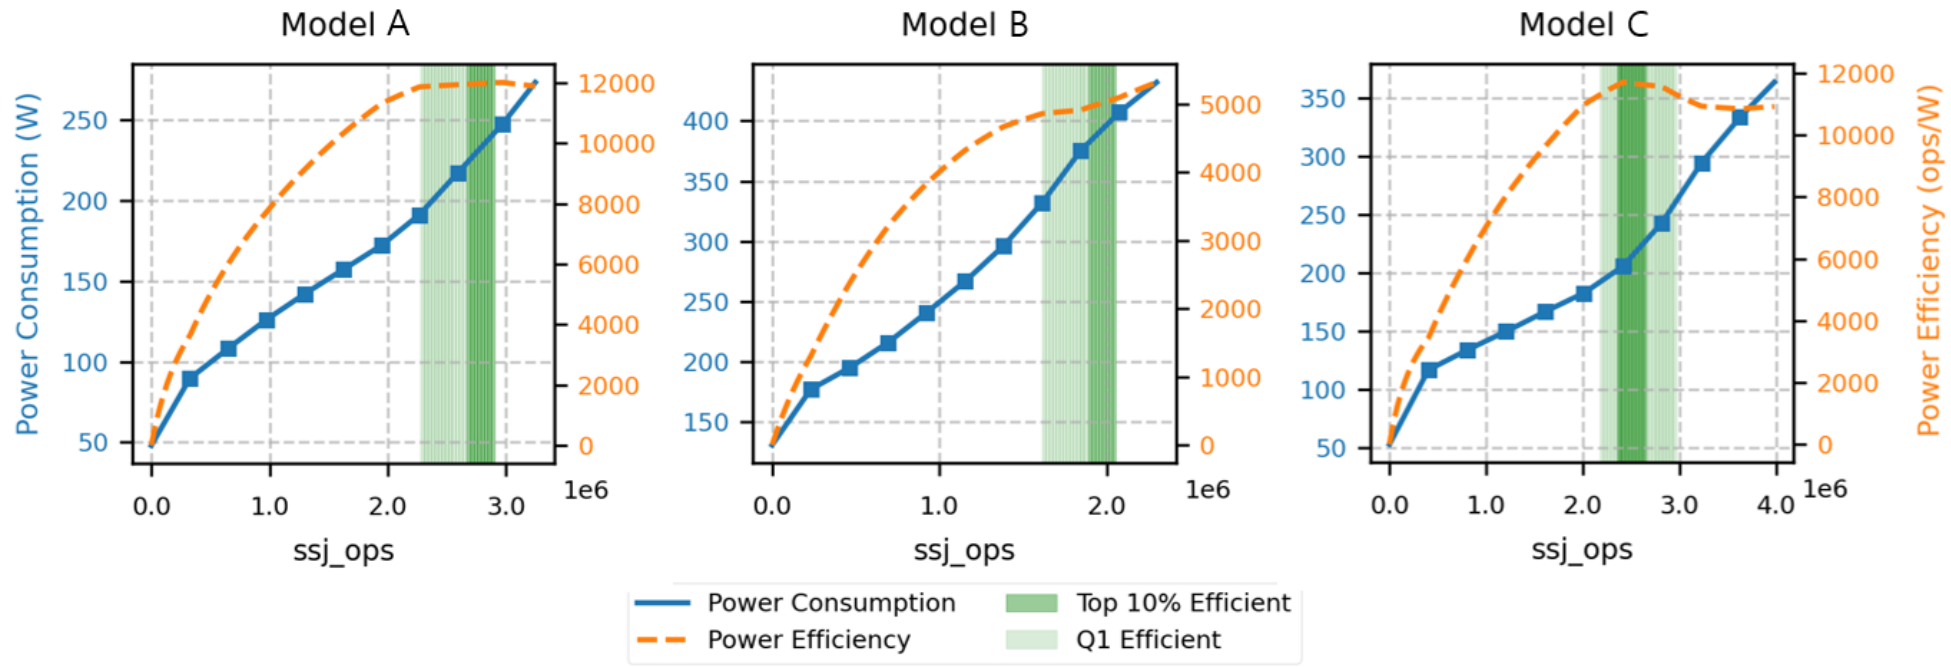

Heterogeneous data centers are composed of various server models with distinct power characteristics. Each model has a unique power efficiency profile and, therefore, a different optimal operating range.

Strategies for Efficient VM Replacement

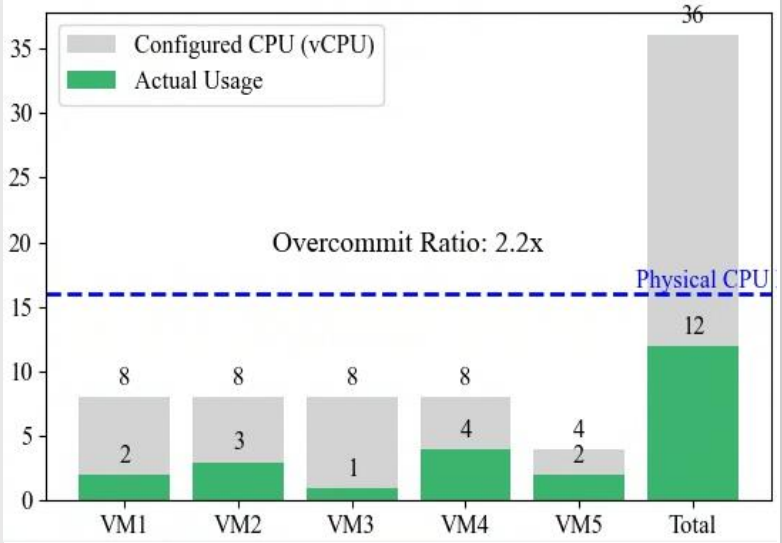

Resource Overcommit

A technique for allocating virtual resources beyond the limits of physical resources.

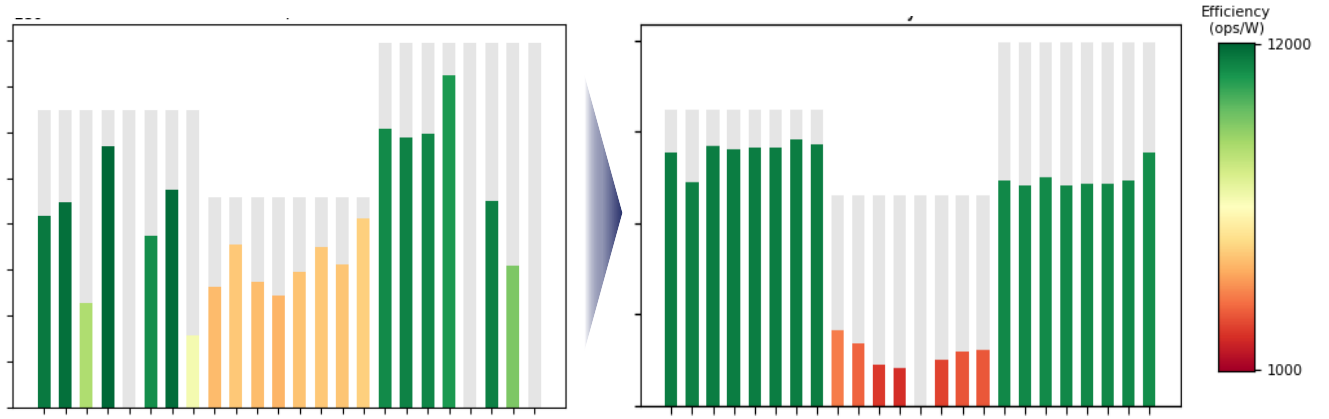

Horizontal & Vertical Optimization

In a heterogeneous environment with varying server efficiencies

- Consolidate workloads onto the most power-efficient hosts
- Maintain the optimal power-efficiency operating range for each active host

VM Replacement Framework

VM replacement technique based on MOEA that ensures QoS
while optimizing objectives in cloud data centers

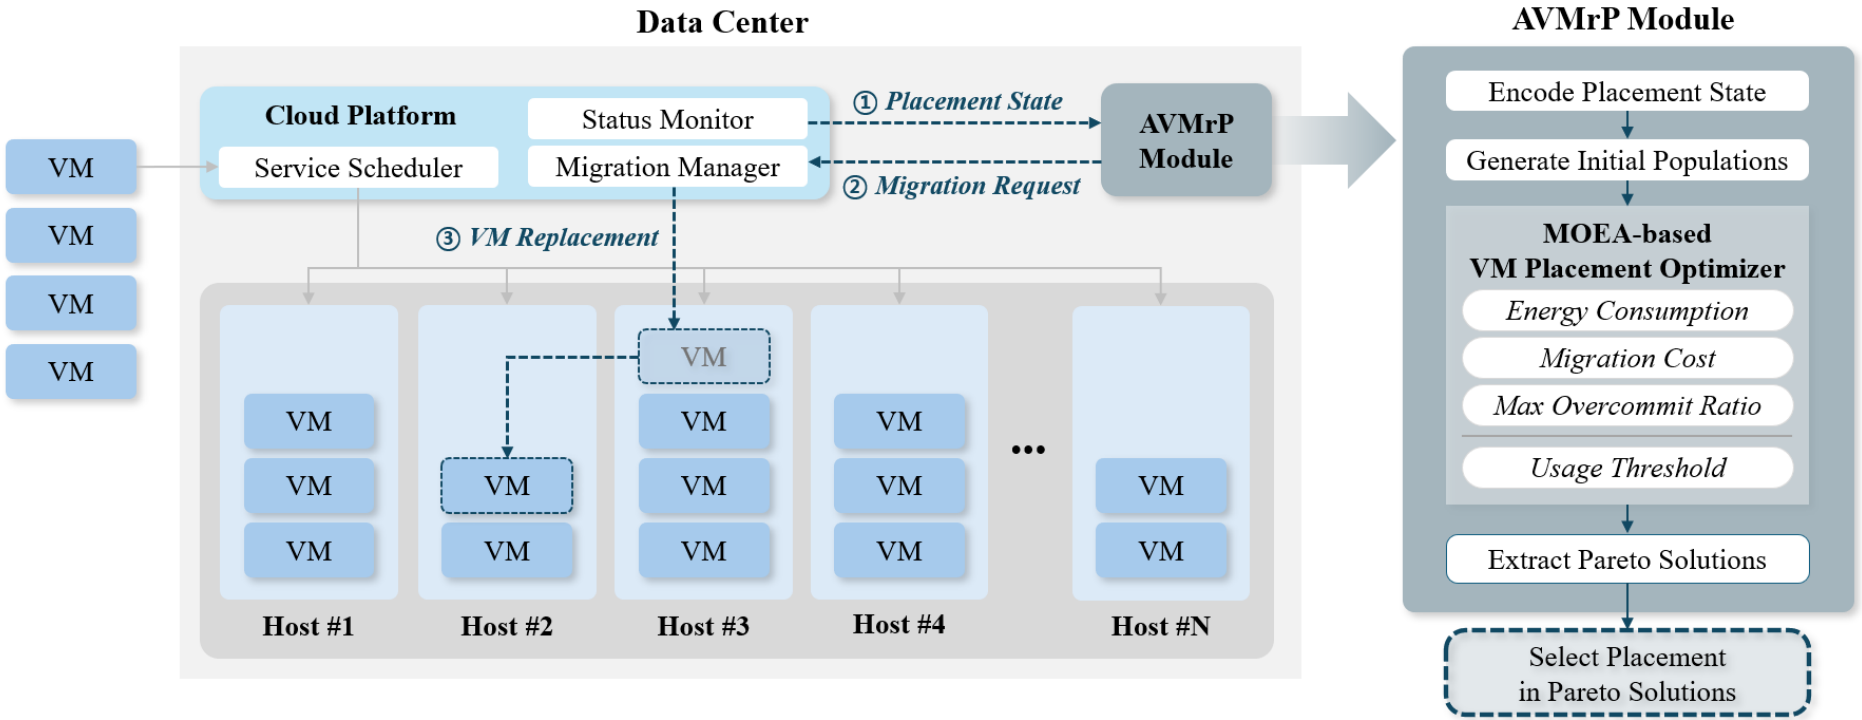

Performance Evaluation (1)

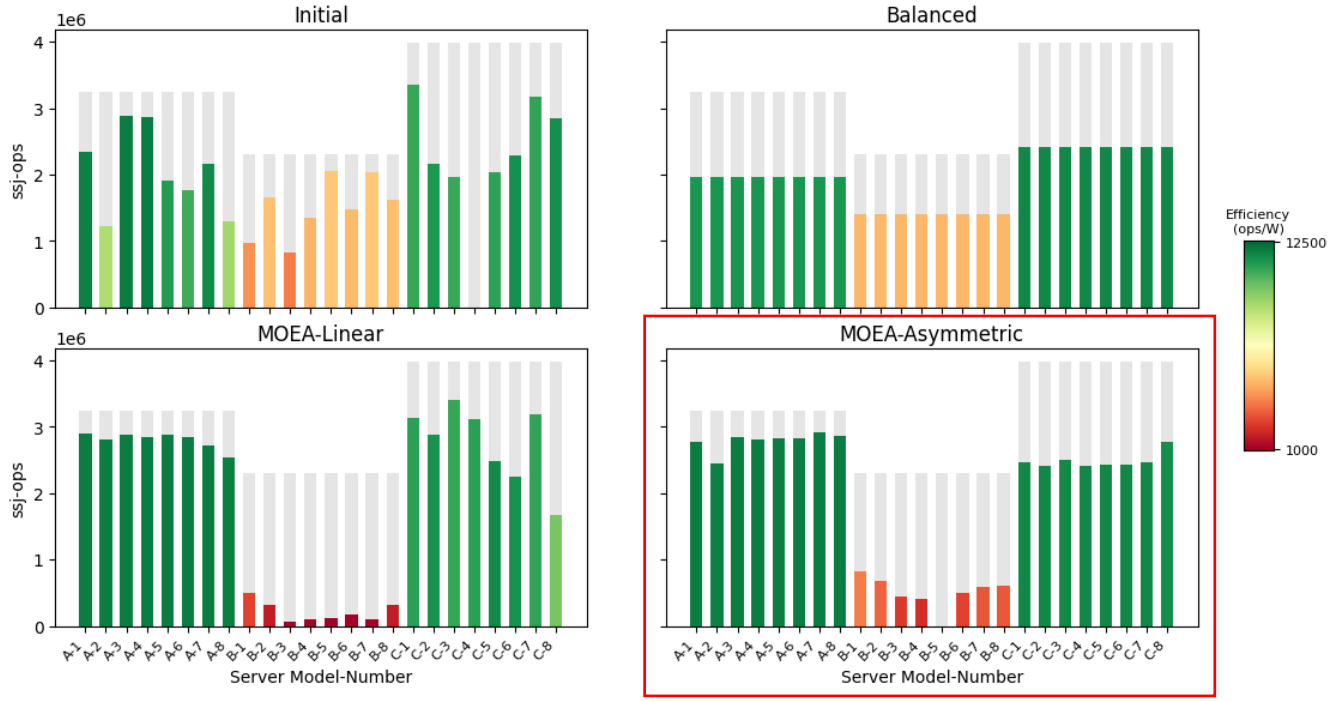

Case	Power Consumption	Ratio
Initial	5683.9W	100.00%
Balanced	5419.9W	95.54%
Linear	5333.7W	93.84%
Asymmetric	5147.9W	90.57%

- Energy-aware placement is more effective than simple load balancing.
- Modeling realistic server characteristics yields superior results.

Performance Evaluation (2)

Simulation of VM replacement was conducted across various initial VM placement algorithms.

The proposed method demonstrates significantly superior results in both power consumption and migration overhead compared to traditional rule-based resource management techniques.

We can customize this methods to fit the unique workload characteristics of scientific environments, such as the disk-intensive tasks.

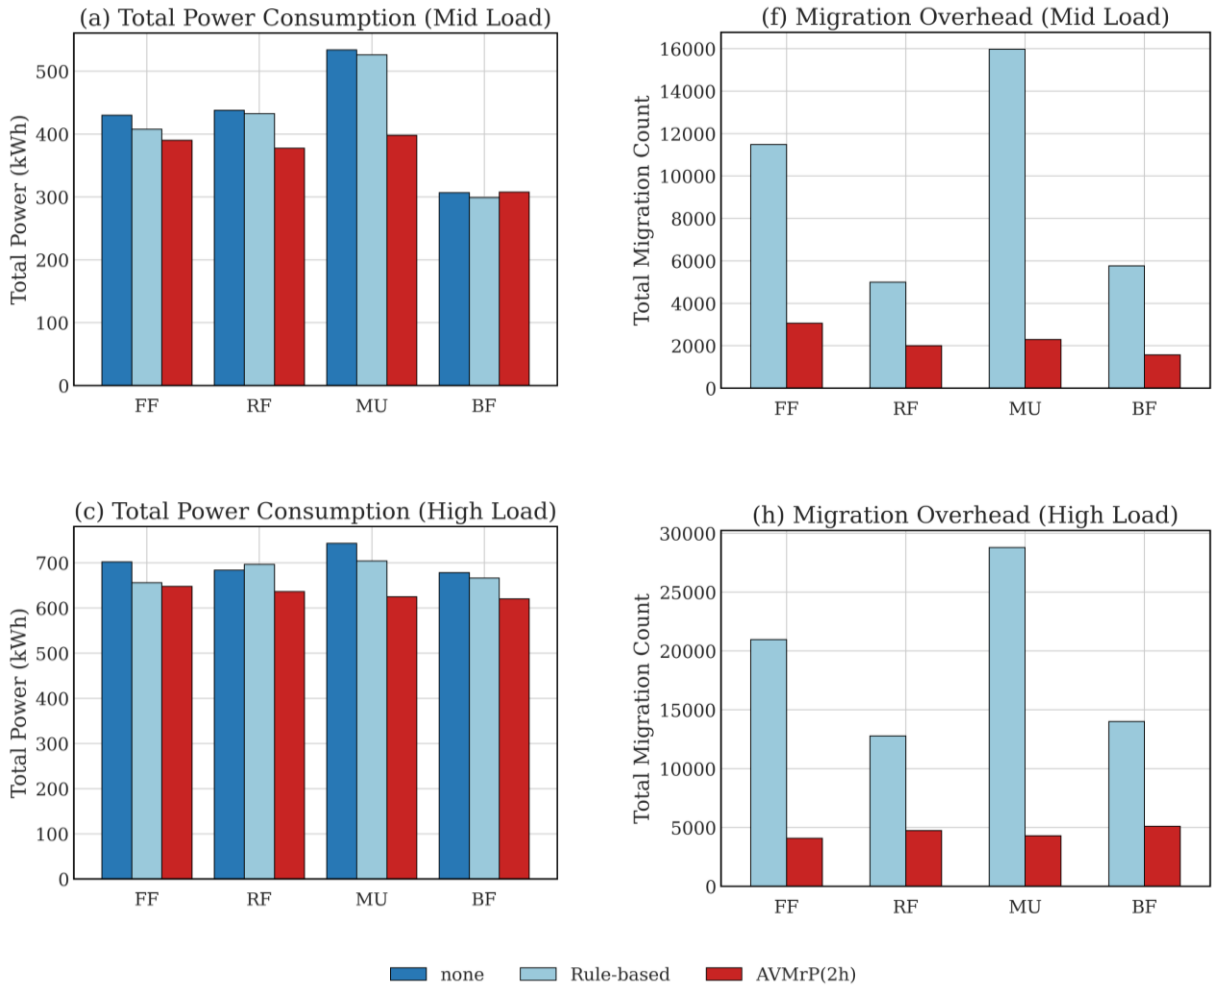

2. Research Background

Competencies

1. **Experiment in Multi-Objective Optimization:**
   A proven ability to balance competing goals such as performance, cost, and energy efficiency in complex systems.

2. **Focus on Resource Efficiency:**
   A track record of developing solutions that maximize efficiency for a given set of resources.

3. **Systematic Methodology for Validation:**
   Proficiency in using a framework of performance analysis, simulation, and benchmarking to test and validate ideas before implementation.

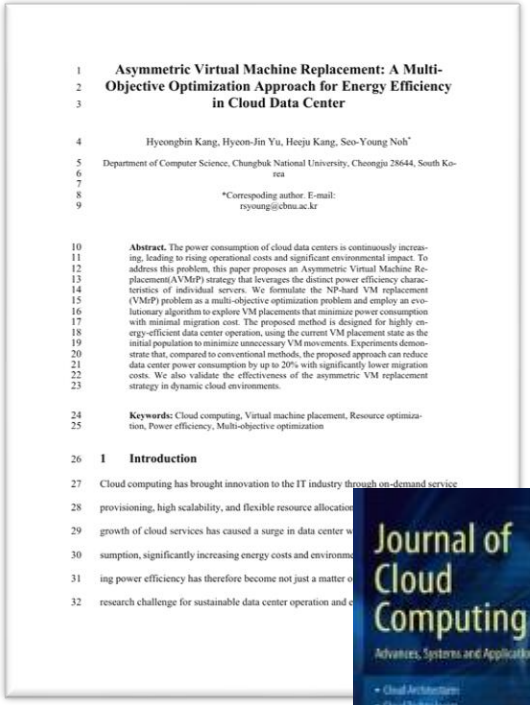

Asymmetric Virtual Machine Replacement: A Multi-Objective Optimization Approach for Energy Efficiency in Cloud Data Center

Hyeongbin Kang, Hyeon-Jin Yu, Heeju Kang, Seo-Young Noh*

Department of Computer Science, Chungbuk National University, Cheongju 28644, South Korea

*Corresponding author. E-mail: rsyoung@cbnu.ac.kr

Abstract. The power consumption of cloud data centers is continuously increasing, leading to rising operational costs and significant environmental impact. To address this problem, this paper proposes an Asymmetric Virtual Machine Replacement(AVMrP) strategy that leverages the distinct power efficiency characteristics of individual servers. We formulate the NP-hard VM replacement (VMrP) problem as a multi-objective optimization problem and employ an evolutionary algorithm to explore VM placements that minimize power consumption with minimal migration cost. The proposed method is designed for highly energy-efficient data center operation, using the current VM placement state as the initial population to minimize unnecessary VM movements. Experiments demonstrate that, compared to conventional methods, the proposed approach can reduce data center power consumption by up to 20% with significantly lower migration costs. We also validate the effectiveness of the asymmetric VM replacement strategy in dynamic cloud environments.

Keywords: Cloud computing, Virtual machine placement, Resource optimization, Power efficiency, Multi-objective optimization

1 Introduction

Cloud computing has brought innovation to the IT industry through on-demand service provisioning, high scalability, and flexible resource allocation

growth of cloud services has caused a surge in data center w

sumption, significantly increasing energy costs and environme

ing power efficiency has therefore become not just a matter o

research challenge for sustainable data center operation and

Under Review

Participated Research: Tapeless Storage

Participated in research on distributed file systems for tapeless storage and has experience with RAID, RAIN, and EOS.

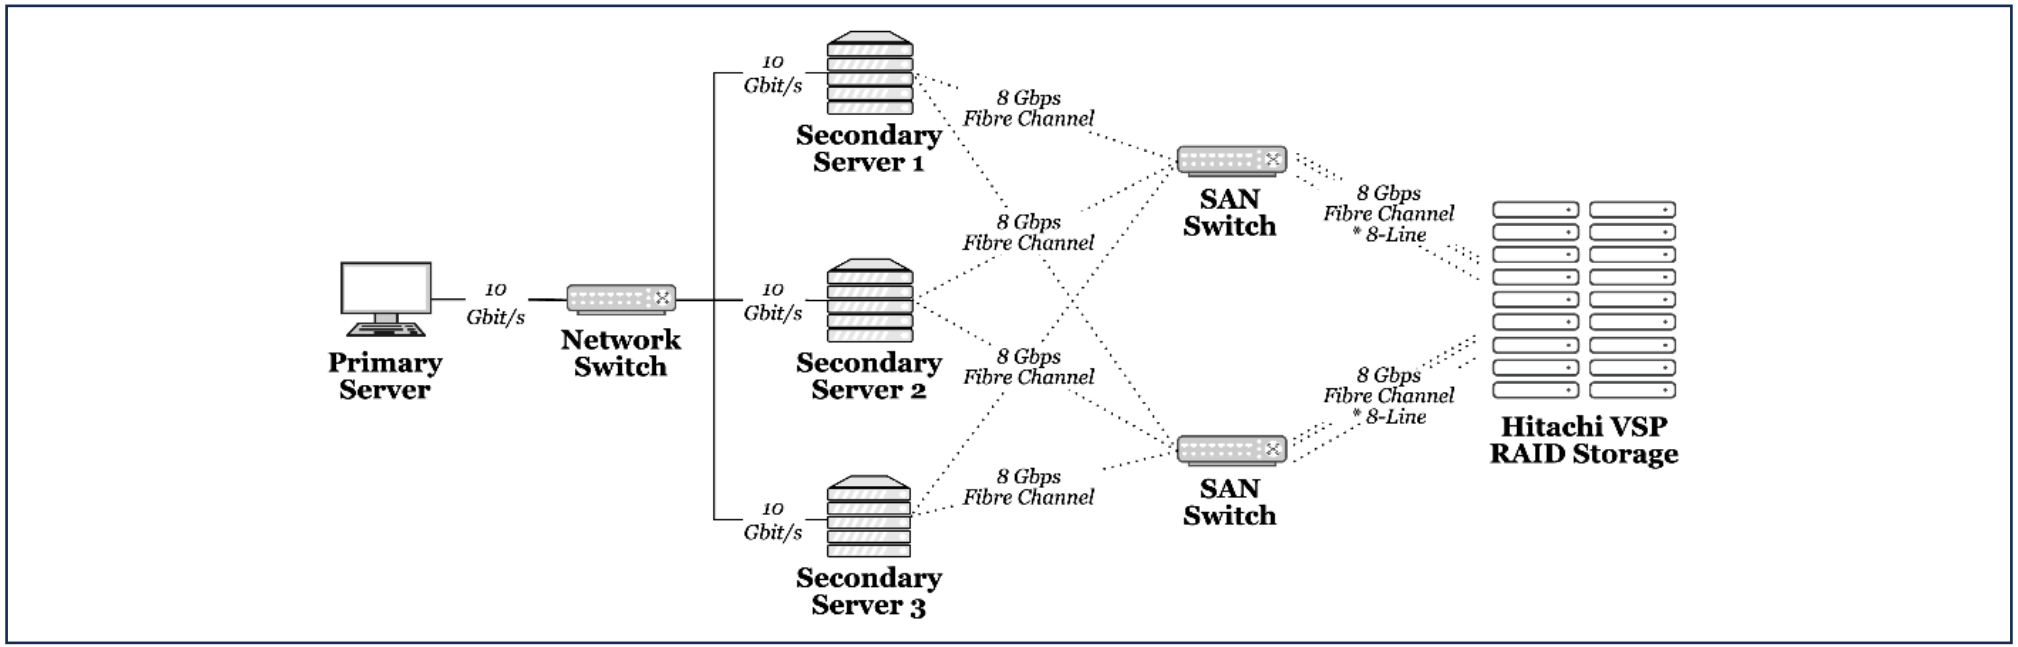

Understanding the Custodial Disk Storage (CDS)

Key Identity of the CDS:

- **Purpose:** Long-Term Data Preservation for the ALICE experiment.
- **Characteristics:** A Disk-Only archiving system proposed as an alternative to Tape-Storage.
- **Core Technology:** Based on EOS, utilizing Erasure Coding for data protection.
- **Primary Data Type:** Composed mainly of Cold Data for archiving purposes.

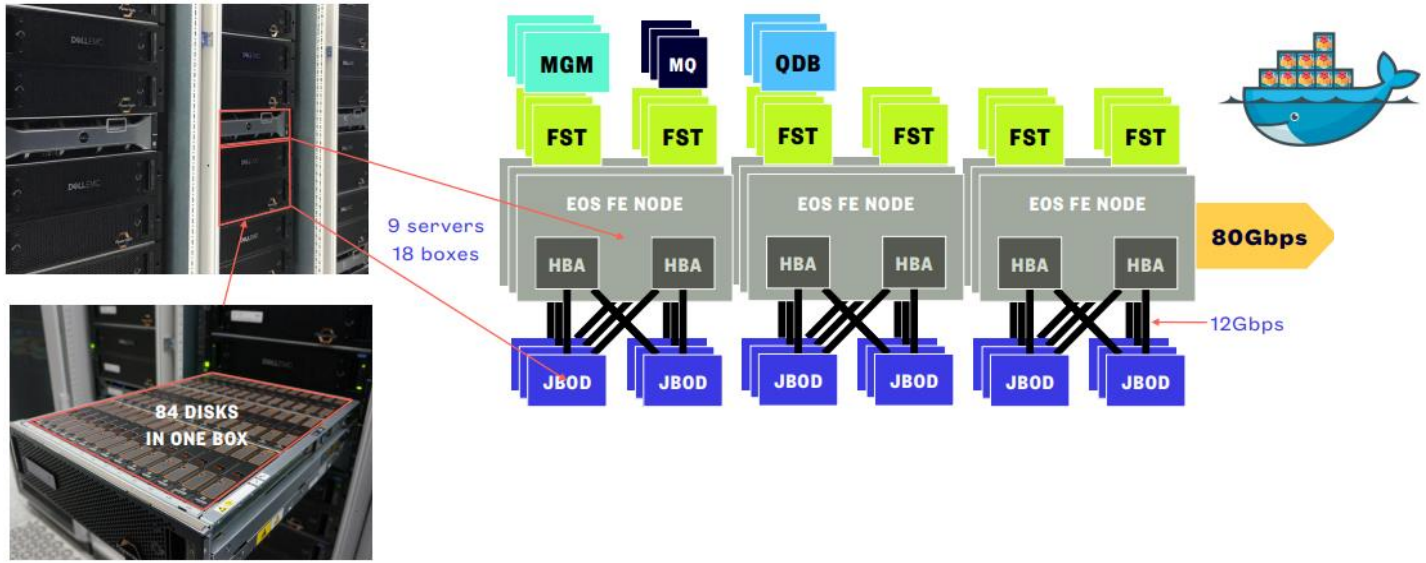

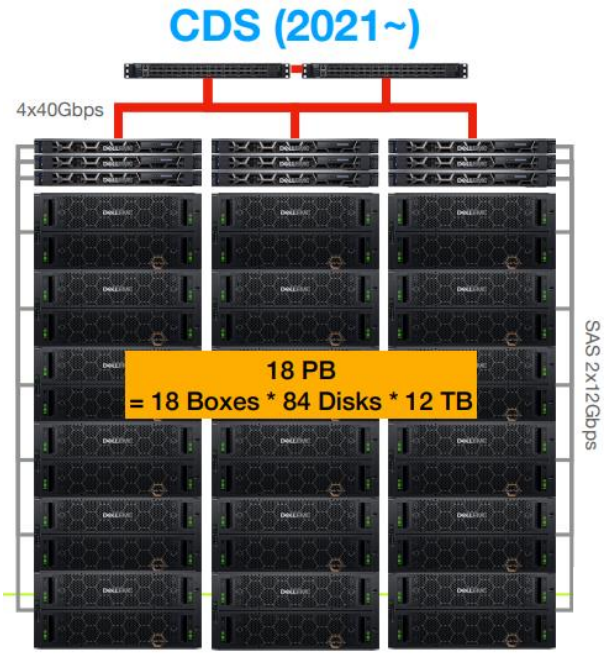

The Evolving Environment

Environmental Changes:

- **New Demands:** Anticipating new data processing requirements
- **EOS Updates:** Advancement of the Storage features

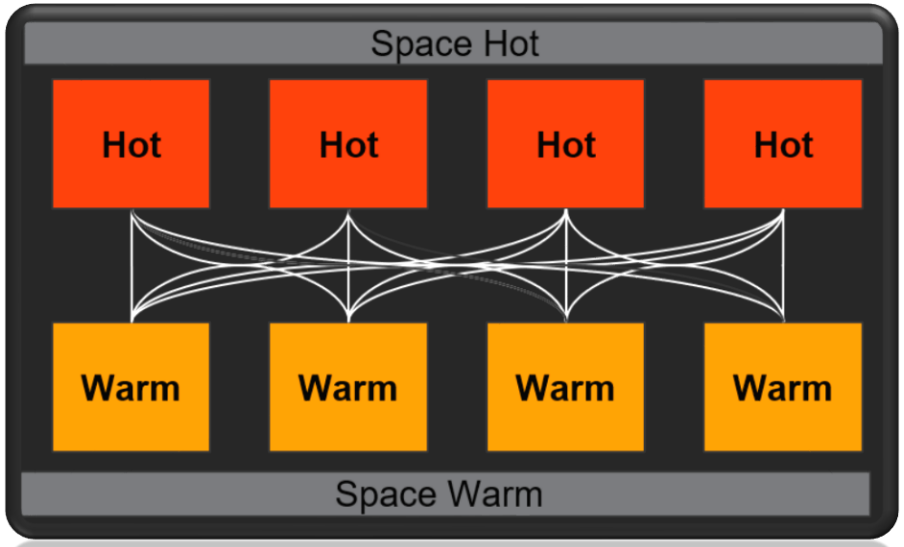

Future Directions: A Study of the Ongoing Discussion

- Selection and prototyping of next-generation CDS technology
- Integration of EOS Disk and CDS into a single system
- Reuse of existing CDS (e.g., Tier 1 cache)
- Introduction of a hybrid SSD/Spinning disks system for highly demanding applications

4. A Proposal for Collaborative Research

Exploring Future Directions for the CDS

Next-Generation

Hybridization

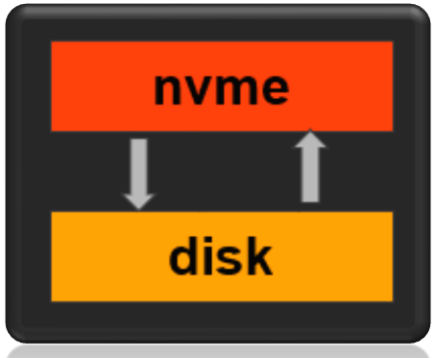

Strategic Reuse

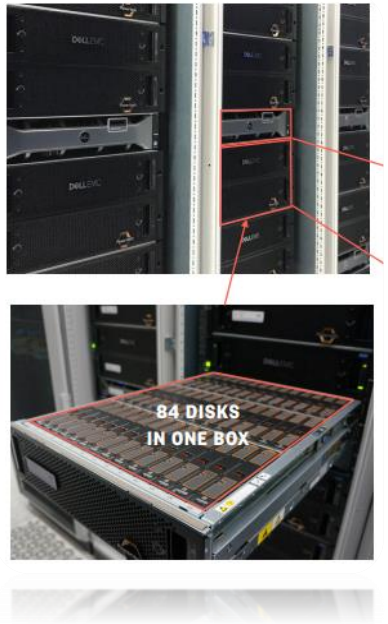

4. A Proposal for Collaborative Research

Architecting the Next-Generation CDS

Technologies to Explore:

- **Next-Gen Hardware:** A comprehensive evaluation of SMR disks and larger JBOD enclosures for their performance, cost-effectiveness, and operational suitability. (supported by EOS updates)
- **RAIN Configurations:** Applying high-availability RAIN architectures based on advanced EOS support.

Contribution:

Support this decision-making process by conducting analysis and providing data through experimentation.

Cost ↓
Capacity ↑

Write Speed ↓
Compatibility ↓
(RAID, RAIN)

Need for Research on Technology Adoption (e.g., Tiering)

4. A Proposal for Collaborative Research

Hybridization: EOS Tiering

Tiering is a powerful technical solution for the evolution of the CDS.

- A Hybrid Configuration of high-performance SSDs and high-capacity HDDs.
- EOS Tiering moves data to the optimal location based on a defined policy.

Contribution:

Research and validate the optimal policy for the workload of the target environment.

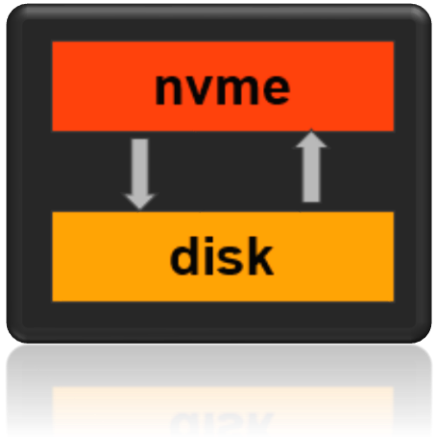

4. A Proposal for Collaborative Research

Strategic Reuse of the Current CDS

The key technologies discussed can also be applied to strategically reuse the current CDS hardware.

Potential Paths for Discussion:

A. High-Performance T1 Cache: Reuse the CDS as a T1 cache. (maximizing access speed with Hybrid Tiering)

B. Integration into New CDS: Incorporate parts of the CDS as the Cold Tier of a next-generation system.

C. Reuse for other RAIN-enabled storage: Repurpose the existing disks to build out other storage solutions.

Evaluating these reuse options involves complex trade-offs,

but also presents an opportunity to apply new technologies to maximize their potential.

5. Conclusion

Challenges

A collaborative research discussion to define the future of the CDS system.

Proposed Idea

- Architecting the Next-Generation
- Leveraging Hybridization with Tiering
- Strategic Reuse of the Current CDS

Today's presentation is the starting point for our collaboration.

The 9th Asian Tier Center Forum

Evolving the CDS:
A Proposal for Collaborative Research

Thank You!