

**SAMSUNG**

# Samsung Cognos Management Console

White Paper | March 26, 2024

Authors: Ron Lee, Mandar Sawant & Mayank Saxena

Data Fabrics Solutions, MSL  
<https://samsungmsl.com/dfs/>



# Table of Contents

---

Abstract	3
Introduction	4
About Cognos	4
Samsung Cognos Management Console – Solution View	5
A Use Case: In-Memory Database	8
Summary	10
References	10

# 1. Abstract

---

AI technologies from companies like Nvidia, Google and OpenAI are moving our consumption of data into new frontiers where technology changes are occurring faster than they have in any time in recent history. AI is moving digital and visualization environments beyond what our current data center capacities are capable of handling. For example, the most recent release of Meta's Quest 3 VR Headset and Apple's Vision Pro is precipitating a digital data explosion and is on trend to accelerate and continue to stress data center capacities. Additionally, the "Magnificent 7" mega-cap companies (Amazon, Microsoft, META, Apple, Alphabet (Google's search engine), Nvidia and Tesla) along with smaller companies with similar goals, are working to analyze AI-generated data sets in real-time to reap financial rewards. These industries mentioned are just a few that are having a profound effect on creating and using very large amounts of AI-generated data.

Advertising on search engines, social media, and online markets is a lucrative industry, worth hundreds of billions of dollars. Mega-cap companies have developed various technologies, such as Hadoop, Spark, and NoSQL, to mine the data they collect in order to get further insights that can be applied to business decisions and improving customer relations. However, they too face significant hardware and software bottlenecks. For example, in cloud data centers, server farms are used to run multiple applications, each with varying memory and storage requirements. It is challenging to allocate the appropriate memory and storage resources for each application while still achieving reasonable total cost of ownership (TCO) goals. In fact, memory is typically utilized less than 50% of the time,<sup>1</sup> with unused memory stranded inside servers, bonded to specific CPUs, and isolated from resources elsewhere in the data center.

Samsung's software-defined memory technology solutions, powered by CXL, aim to provide the scalability, manageability, configurability, and flexibility required as data centers transition to the next-generation of disaggregated system architectures, enabling them to achieve the best TCO.

CXL adoption will come with its own unique challenges. The development of complex software applications will be required in order to write their own CXL-aware software layer to manage caching, prefetching, and tiering between DRAM and CXL Memory. Furthermore, if heterogeneous memory devices are not consistently orchestrated to maintain data coherency, may lead to suboptimal user experience. Samsung's new Cognos software can help overcome these hurdles by providing,

1. A user experience focused pluggable software suite.
2. An easy to use interface for end applications to realize CXL benefits without any significant changes
3. A CXL compliant Full-Stack solution for easy provisioning and management and
4. An application-aware memory orchestration for dispensing memory.

## 2. Introduction

Samsung, being the world leader of DRAM and Flash memory technologies, has recognized early on the potential for a CXL-based solution architecture. Samsung heavily invested into building in-house technologies, know-how, and device solutions to fully enable the capabilities that CXL brings to the next-generation system architecture. These investments include the CXL Memory Module – DRAM (CMM-D) device, CXL Memory Module PMEM/Hybrid devices (CMM-H), and the solution level device CXL Memory Module – Memory Box (CMM-B). Additionally, Samsung has invested in a set of world-class firmware, drivers, APIs and management software to enable customers to quickly adopt these solutions that are aimed at reducing their total cost of ownership.

Samsung Cognos provides the glue that stitches all of the CXL-based technologies and devices into one homogeneous and integrated solution that developers and customers can easily adopt and use in their cloud or private data center solutions and services. Cognos includes a software developer kit that provides the foundational software components that are needed to quickly utilize the new CXL-based devices with customer applications and a host of management software, including a GUI to manage one-to-many server/switch nodes to handle rack-level cluster/system management. Samsung Cognos addresses memory stranding problems by pooling memory into a global resource that is shareable and reusable by resources across the data center. It provides a value-added software layer for Memory Box and CMM-D in the host to enable easier adoption by applications with easy-to-use interfaces. Samsung Cognos achieves this with three main software components:

- **Management Console** — Allows for memory pooling and sharing to be orchestrated
- **Application plugins** — CXL plugins that allow user applications to utilize heterogenous memory devices
- **Agents** — Converts servers with CXL devices into memory targets to form a memory pool

## 3. About Cognos

One of Cognos's objectives is to reduce and eliminate the barriers to customer adoption with a user friendly management software suite including a graphical user interface as well as centralized cluster management functions that can control most aspects of Samsung Cognos provided features. In addition, Samsung Cognos was designed to be easily integrated into customer infrastructure and associated applications. Key features of Cognos Management Framework are:

- Ease-to-use, composable and fine-tuned user interface to manage memory
- Configure and monitor large scale feature rich Samsung memory devices
- Topology view of Devices, Switches and Hosts across rackscale clusters

Cognos tightly couples Samsung's devices and firmware providing features that help simplify diagnosing non-trivial issues arising due to some hidden software or hardware anomalies. Other benefits of Cognos aim to realize better cost efficiencies and improve overall TCO. These key benefits includes:

- Dynamic memory allocation and deallocation as required by application
- Intelligent metrics of resource usage and statistical visualization
- Easily integrate-able into customer infrastructure and associated applications

## 4. Samsung Cognos Management Console – Solution View

Cognos solution architecture is a rackscale solution with support for easy-to-manage features through the Cognos Management Console user interface. It provides flexible fabric management and orchestration and is able to handle one-to-many server and switch nodes to support small-to-medium-to-large scale deployments or configurations. CMM-B provides the fabric interface required to scale both the application hosts and the memory/storage servers.

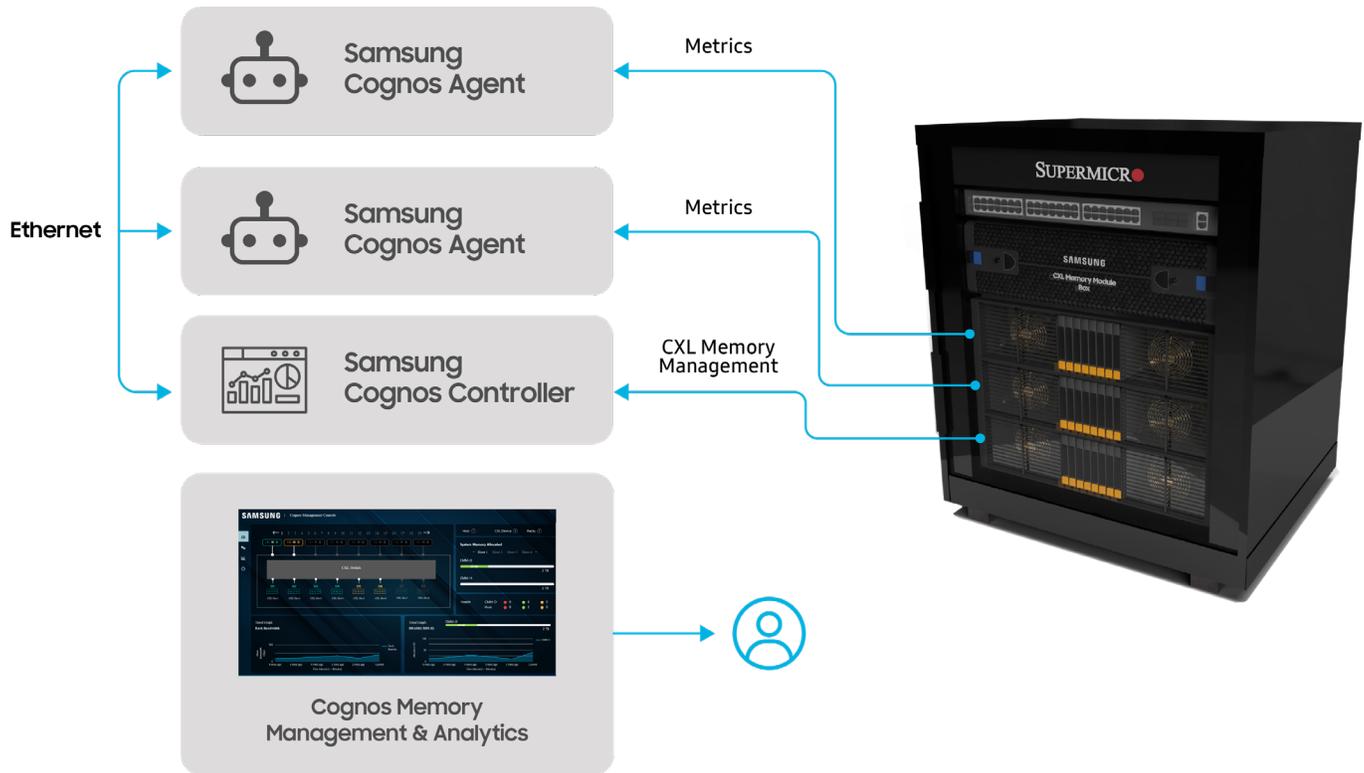


Figure 1

Cognos solution allows customers to quickly integrate their application through management APIs and maximize overall TCO benefit through underlying software defined memory services.

Cognos’s management architecture includes a centralized controller along with a Cognos Agent located in each of the managed servers. The Cognos Agent provides a multitude of node level services required to monitor, control, and diagnose both the server status as well as CXL device statuses. Some of the major services include listing, allocating, attaching, and freeing one or more pools of CXL memories. Available CXL memory can be carved out to desired size memory pools and attached to one or more CXL hosts allowing users to flexibly assign extended memory to applications that need it.

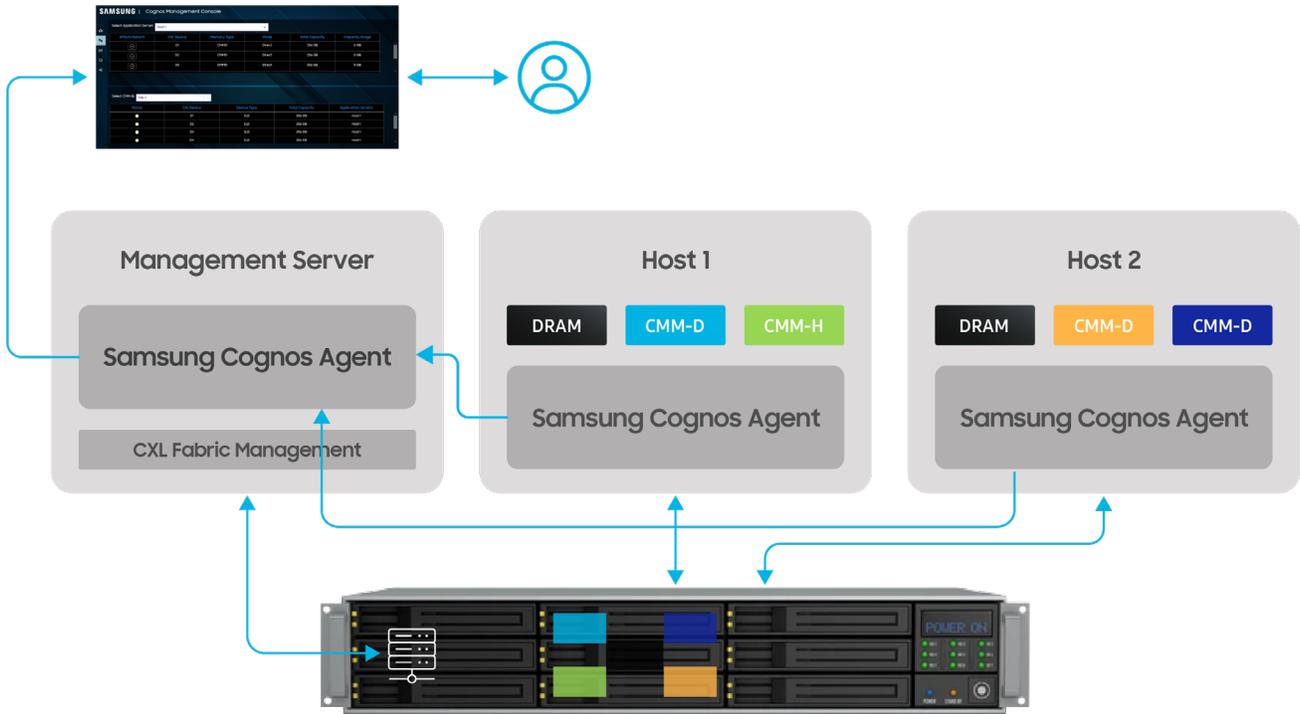


Figure 2

Cognos Fabric Manager aims to simplify the management of CXL based infrastructure by using a user friendly intuitive GUI. The figure below shows how a user can add attach 256GB sized memory to a host.

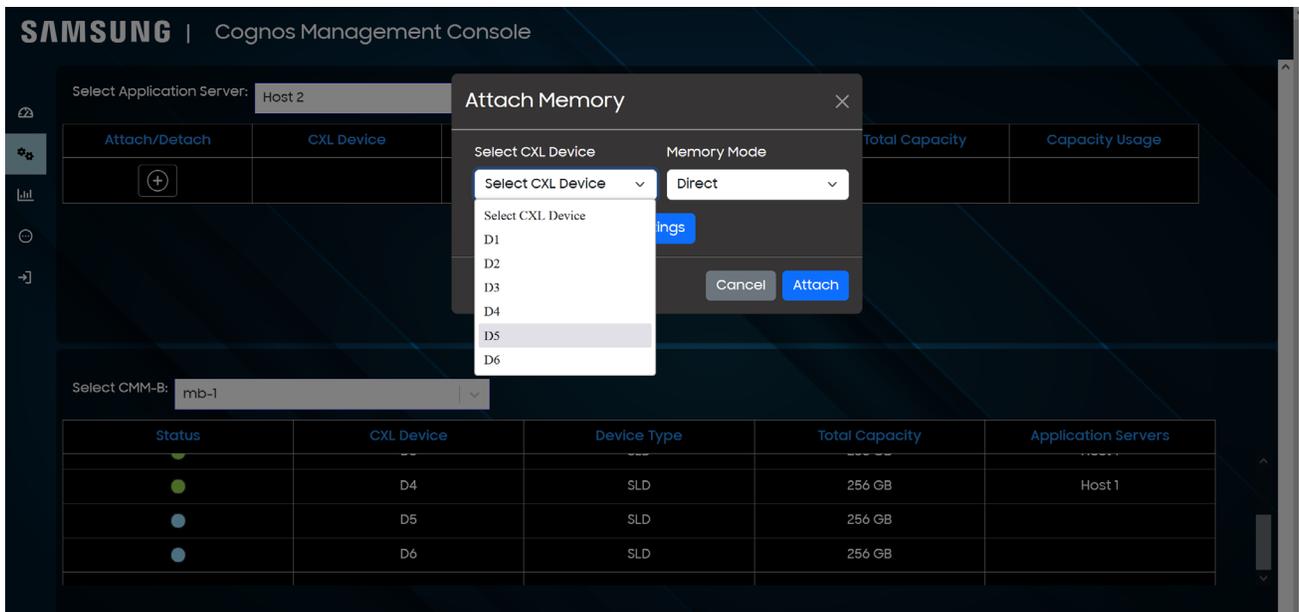


Figure 3

From the GUI, users can obtain detailed status of attached CMM-D, CMM-H and CMM-B devices on the network. Samsung Cognos monitors a vast array of telemetry, performance and event/log data from each of the CMM-D/H/B Devices. These real-time diagnostics provide detailed information that is used by Cognos to pin-point problem devices, predict potential issues, and provide information that will aid in performing root analysis of unforeseen issues.



Figure 4

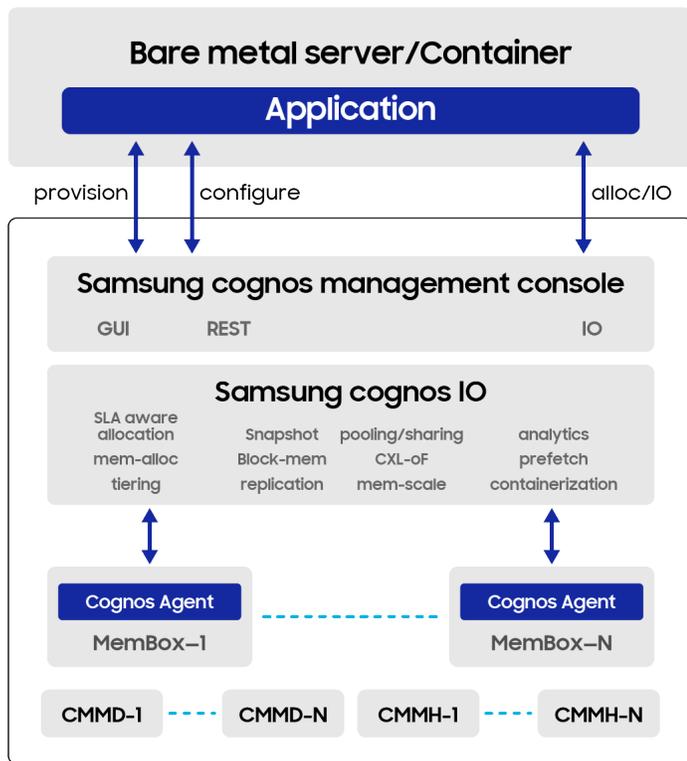


Figure 5: Samsung Distributed Memory-Store

CXL fabric and Samsung CXL servers/devices provide the ultimate ability to scale Memory and Storage seamlessly with best-in-class latency and performance for a variety of memory-hungry applications. Since memory requirements often vary depending on the types of applications and associated data sets. Samsung Cognos aim to provide a solution where fabric attached memory devices and storage pools can be easily and dynamically delegated to applications as needed. Also, Cognos I/O stack abstracts the underlying storage and memory to provide additional intelligent services which we will discuss in the following sections.

Cognos includes a complete set of API services to setup, configure and run I/O operations to/from attached CXL devices. These services create memory or storage volumes by aggregating two or more CMM-D and CMM-H devices. The ability to assign individual volumes to one or more applications on CXL hosts and then servicing the application I/Os with the best-in-class performance. Underlying volumes also manage the CMM-D and CMM-H devices in real-time to improve endurance and enhance performance and have algorithms to predict device problems before they cause potential catastrophic failures. Overall, Cognos's API services and smart algorithms help meet end user's expectations and service level agreements with respect to performance and reliability, from Samsung's solutions.

Samsung is also working with hardware partners like SMC (Supermicro) to integrate Cognos into SCC (Super Cloud Composer)<sup>2</sup> that demonstrates Rackscale management capabilities inherent in Cognos. Although there are some available open-source CXL fabric managers, Cognos is unique in that it can provide a coherent and consistent graphical interface, with REST API interface support, to orchestrate heterogeneous memory devices with complete transparency to end applications.

## 5. A Use Case: In-Memory Database

Samsung Cognos based architecture is an ideal solution for applications such as In-Memory Databases since the software requires large scale high performance memory to perform real-time analytics for business critical applications such as retail, healthcare and cybersecurity.

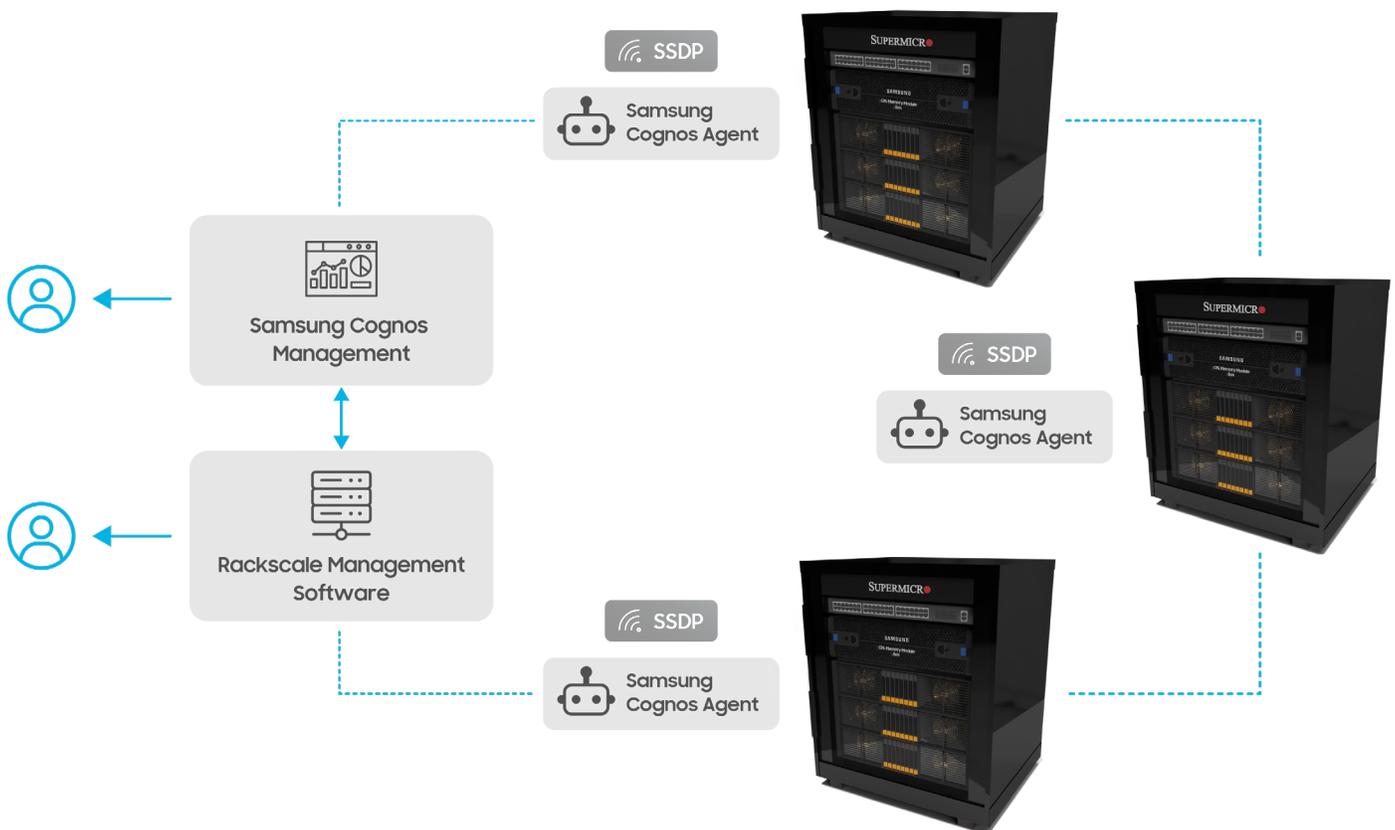
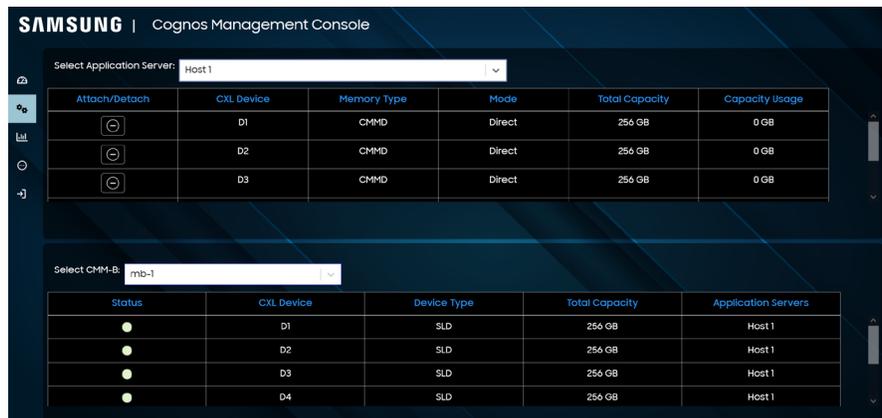


Figure 6

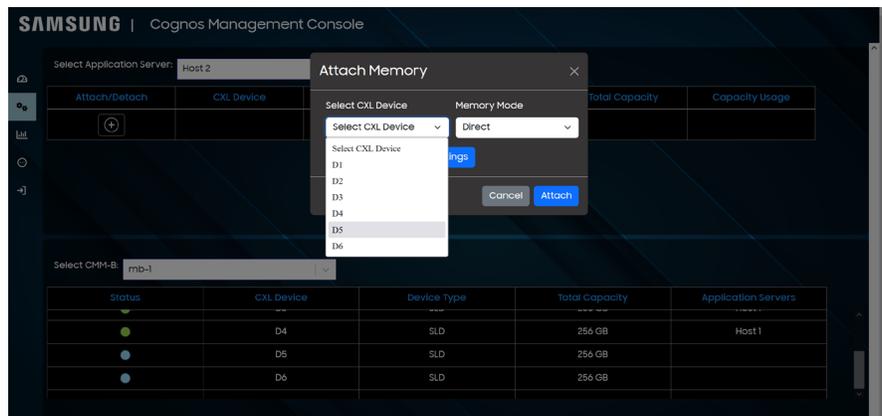
Multiple instances of In-Memory Database applications can run on rackscale Cognos infrastructure where Terabyte Scale CMM-D, CMM-H and SSD storage devices can be interconnected through CMM-B to dynamically provide flexible heterogenous large memory and storage pools to each of the database instances. Also, CMM-D devices with dense DRAM memory sizes can be mapped to CPU address and accessed with nanosecond latencies. With the next generation CXL 3.0 capabilities, these CMM-D devices can be pooled together as a contiguous or interleaved memory to provide the best solution for each of the In-Memory Database instances. Since DRAM based memories are volatile, applications will also need CMM-H and/or SSD devices to provide persistent storage where user data sets can be flushed to in case of failures and permanently backed up for failsafe operation.

Figure 6 shows an example of rackscale solution with Cognos Agent running in each of the servers. Cognos Fabric Manager provides ease-to-use centralized management for the entire infrastructure.

**Step 1**  
Configure Page



**Step 2**  
Attach device D5 to Host 2



**Step 3**  
Device D5 attached to Host 2

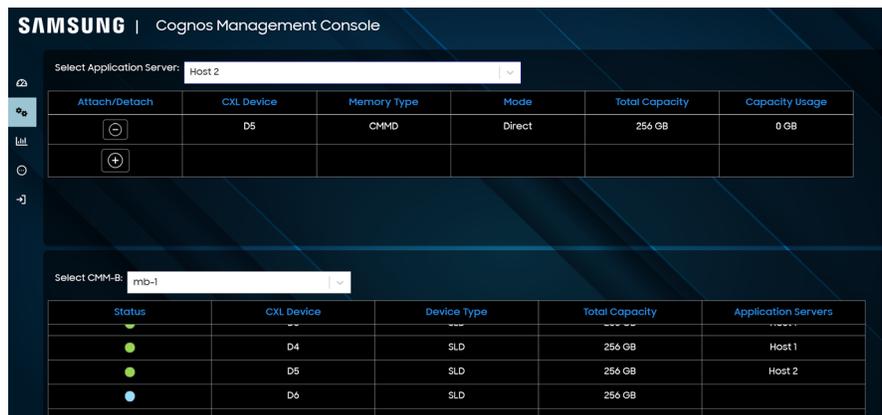


Figure 7

## 6. Summary

Samsung's Cognos solution architecture, in conjunction with Samsung's CMM-H, CMM-D, and CMM-B devices, offers unparalleled flexibility and performance in data centers, allowing for the monitoring and control of next-generation CXL-based rackscale infrastructure. Cloud and enterprise clients can use Cognos to dynamically pool and allocate terabytes of DRAM and petabytes of storage to application instances in real-time, without the need for any hardware reconfiguration or changes to existing infrastructure. This can all be done within a secure environment where Cognos continuously collects telemetry and diagnostic/log data from all devices to monitor performance, health, and other parameters required for 24x7 operations.

Future generations of Samsung Cognos Management Console will continue to enhance and improve the reliability, experience, and usability of this first-generation solution.

## 7. References

[www.samsungmsl.com](http://www.samsungmsl.com)

<https://semiconductor.samsung.com/us/news-events/tech-blog/rethinking-storage-for-the-memory-centric-computing-era/>

### For more information

For more information about the Samsung Semiconductor products, visit [semiconductor.samsung.com](http://semiconductor.samsung.com).

### About Samsung Electronics Co., Ltd.

Samsung Electronics Co. Ltd inspires the world and shapes the future with transformative ideas and technologies. The company is redefining the worlds of TVs, smartphones, wearable devices, tablets, digital appliances, network systems, memory, system LSI and LED solution. For the latest news, please visit the Samsung Newsroom at [news.samsung.com](http://news.samsung.com).

Copyright © 2024 Samsung Electronics Co., Ltd. All rights reserved. Samsung is a registered trademark of Samsung Electronics Co., Ltd. Specifications and designs are subject to change without notice. Nonmetric weights and measurements are approximate. All data were deemed correct at time of creation. Samsung is not liable for errors or omissions. All brand, product, service names and logos are trademarks and/or registered trademarks of their respective owners and are hereby recognized and acknowledged.

Fio is a registered trademark of Fio Corporation. Intel is a trademark of Intel Corporation in the U.S. and/or other countries. Linux is a registered trademark of Linus Torvalds. PCI Express and PCIe are registered trademarks of PCI-SIG. Toggle is a registered trademark of Toggle, Inc.

### Samsung Electronics Co., Ltd.

129 Samsung-ro, Yeongtong-gu, Suwon-si, Gyeonggi-do 16677, Korea [www.samsung.com](http://www.samsung.com) 1995-2021

**SAMSUNG**