



TABLE OF CONTENTS

- 2 INTRODUCTION**
- 2 SCOPE**
- 3 SOLUTION OVERVIEW**
- 4 KEY SOFTWARE-DEFINED FEATURES**
- 4 MICROSOFT SDDC CERTIFICATION AND PERFORMANCE TEST**
- 5 HARDWARE COMPONENTS**
 - Servers
 - Processor
 - Memory
 - Storage
 - Networking
- 8 SECURITY**
- 8 APPENDIX: SYS-1028U-S2D BILL OF MATERIALS**

WHITE PAPER

SUPERMICRO SYS-1028U-S2D REFERENCE ARCHITECTURE FOR MICROSOFT STORAGE SPACES DIRECT (S2D) HYPER-CONVERGED INFRASTRUCTURE (HCI)

EXECUTIVE SUMMARY

Supermicro reference architecture *SYS-1028U-S2D* based on 4x 1U Ultra 10 NVMe servers has received the *Microsoft Windows Server Software-Defined (WSSD) Storage Spaces Direct (S2D) Certification*, and is a part of the first wave offerings in the WSSD Solutions Catalog.

Microsoft S2D solution for Windows Server 2016 targets at the cost-effective, cloud-scale software-defined data center (SDDC) market segment, allowing highly available and scalable software-defined storage (SDS) deployments on commodity x86 hardware with direct-attached storage.

Microsoft S2D offers two distinct deployment options: *hyper-converged* and *converged*. Both promote a streamlined IT infrastructure with lowered CAPEX and OPEX, while delivering high scalability and flexibility, and uncompromised performance and efficiency. It is considered a direct replacement of the traditional expensive SAN infrastructure.

Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131 USA
www.supermicro.com

INTRODUCTION

Microsoft has revamped its Windows Server-based software-defined storage (SDS) offerings to address the need of lower-cost, cloud-scale software-defined data center scenarios. Storage Spaces Direct (S2D), a new feature in Windows Server 2016, allows rapid provision of highly available and scalable SDS deployments with industry-standard servers using direct attached storage. S2D offers two deployment scenarios, namely *hyper-converged* and *converged*. Both aim to dramatically simplify IT infrastructure deployment and lower acquisition costs, while providing high scalability and flexibility, and uncompromised performance and efficiency. Therefore, S2D is considered by Microsoft a direct replacement of traditional expensive SAN infrastructure.

Microsoft Storage Spaces Direct is designed from the ground up to take advantage of cutting-edge storage technologies including NVMe NAND flash, and able to dynamically adapt for mixed operations with SSD/HDD with legacy SAS and SATA interfaces. Supermicro's latest server platforms are a perfect match to complement the flexibility and scalability offered by S2D with hot-swappable NVMe, SAS3 and SATA3 hybrid interfaces, so that the most cost and performance-optimized solutions can be customized and deployed at a scale.

The Supermicro reference architecture is developed based on Microsoft design principles and best practices, and is validated and certified for all requirements, providing guaranteed quality and performance.

SCOPE

This white paper provides a high level overview of the new Supermicro reference architecture SYS-1028U-S2D HCI solution, which as an example includes 4x 1U Ultra SYS-1028U-TN10RT+ servers, its main features and hardware specifications. This document is focused on the building blocks and does not discuss other infrastructure considerations, such as the detailed deployment consideration for different customer infrastructure environments, as well as various management services considerations.

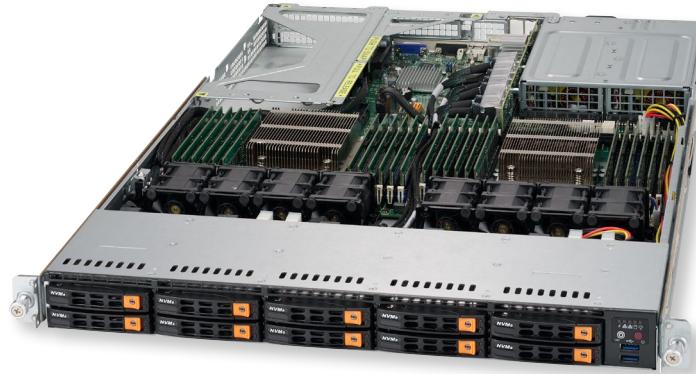


Figure 2. Supermicro 1U Ultra SYS-1028U-TN10RT+

SOLUTION OVERVIEW

Supermicro 1U Ultra HCI solution **SYS-1028U-S2D** is a high-density, all-flash NVMe, Windows Server-based software-defined infrastructure, designed for high-performance computing workloads. It consists of a cluster of 1U Ultra SYS-1028U-TN10RT+ server with NVMe drives, RDMA-enabled NICs, and a low-latency, high-bandwidth TOR switch, along with the software-defined features that are built into Windows Server 2016 Datacenter edition.

Supermicro 1U Ultra HCI solution is designed for hyper-converged deployment scenarios, where multiple server nodes are combined to form a high-availability cluster so that both compute (Hyper-V) and storage (S2D) components can scale-out together. Additional server nodes can be added to the cluster as required to increase overall performance and storage capacity simultaneously. Supermicro 1U Ultra HCI solution supports scaling from a minimum of 2 server nodes up to 16 nodes.

In the HCI configuration, S2D integrates well with other Windows Server built-in features, including Clustered Shared Volume File System (CSVFS), Storage Spaces, Failover Clustering, Hyper-V, and Server Message Block 3 (SMB) direct, to compose a Windows Server-based software-defined storage stack. Figure 3 below demonstrates a 1U Ultra HCI 4-node architecture.

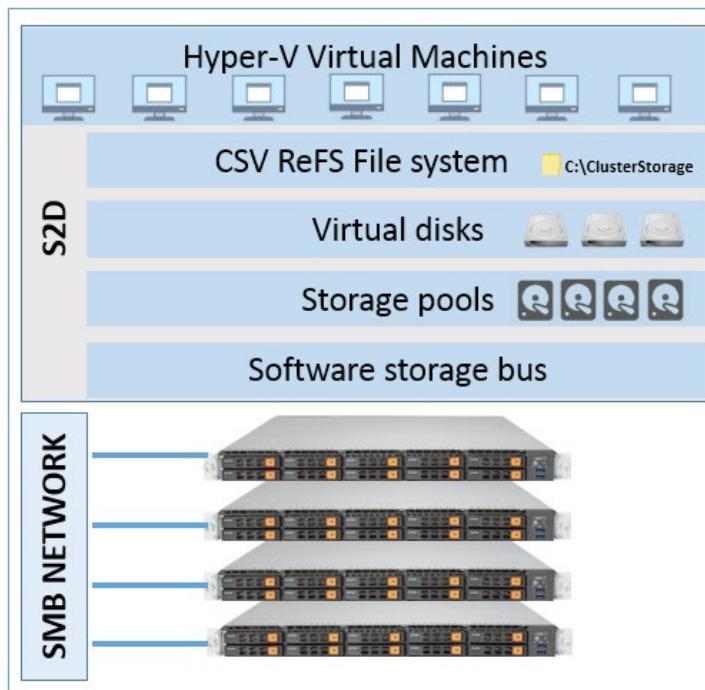


Figure 3. 1U Ultra 4-node hyper-converged infrastructure with S2D

KEY SOFTWARE-DEFINED FEATURES

The 1U Ultra HCI is a virtualization-oriented solution. Some of the key software-defined features include software-defined compute (Hyper-V), software-defined storage (S2D), and software-defined networking (SDN), as well as security and manageability, making it a comprehensive software-defined data center in a rack. Below, Table 1 lists the main software-defined features the Supermicro HCI solution offers.

CATEGORY	FEATURE
Compute	Hyper-V
Storage	Storage Spaces Direct (S2D) software-defined storage
Network	Software-defined networking capabilities: <ul style="list-style-type: none">• NIC teaming• vSwitches• SMB Direct
Security	BitLocker Shielded Virtual Machines Windows Defender Credential Guard Device Guard
Management	Virtual Machine Manager Operations Manager

Table 1. *Supermicro 1U Ultra HCI solution software-defined features*

MICROSOFT SDDC CERTIFICATION AND PERFORMANCE TEST

The 1U Ultra HCI 4-node reference architecture has been certified for SDDC Premium to run as a premium-level software-defined datacenter. The certification not only validates the 4-node cluster under normal stress, but it also requires all the key components, including NIC, drives, HBA, and switch, to pass the SDDC Premium tests.

The 4-node solution is validated to deliver over **1 million IOPS** for 4K blocks (100% random reads) using the Microsoft diskspd/VM Fleet test tool.

HARDWARE COMPONENTS

Supermicro SYS-1028U-S2D HCI solution reference architecture includes four 1U rackmount SYS-1028U-TN10RT+ servers, along with a Mellanox 100GbE SN2700 switch (or Supermicro SSE-C3632S 100G switch), and consumes only 5 units of rack space. Figure 1 illustrates the rack view of the 4-node HCI solution in a standard 16U rack.

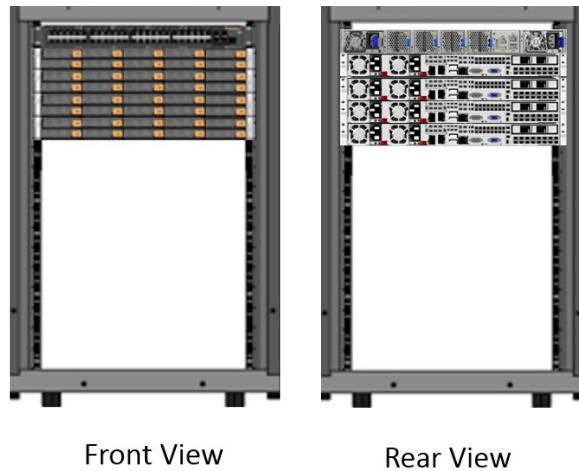


Figure 1. Supermicro SYS-1028U-S2D 4-node hyper-converged solution in a rack

Servers

The 1U SYS-1028U-TN10RT+ is a member of the Ultra SuperServer family. Supermicro Ultra systems deliver best-in-class performance with rich features for virtualization, and support highest performance dual Intel® Xeon® E5-2600 v4 CPUs with up to 3TB of DDR4 memory. SYS-1028U-TN10RT+ is the only 1U server on the market that can accommodate up to 10 hot-swap U.2 NVMe SSD drives. Figure 4 below shows the front view of SYS-1028U-TN10RT+.

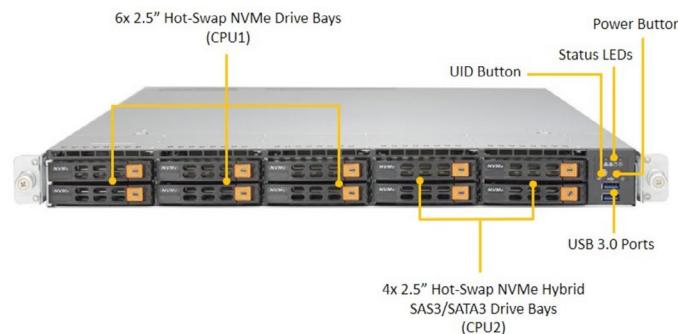


Figure 4. SMCI 1U Ultra SYS-1028U-TN10RT+: 10 2.5" Hot-swappable drive bays, 6 NVMe ports, and 4 NVMe/SAS hybrid ports

Processor

The Intel® Xeon® processor E5-2600 v4 product family is designed to better utilize resources and enable optimal performance across compute, storage, and network workloads. It is the best fit for optimized cloud deployment for software-defined enterprise data centers.

The reference architecture recommends two Intel® Xeon® processor E5-2650 v4 with 12 cores, at 2.2GHz, 30MB Smart Cache, and 9.6GT/s QPI links.

Memory

A total of 128 GB of DDR4-2133MHz memory is used on each node, which meets the requirements for a hyper-converged deployment and delivers the best memory performance for both the operating system as well as for the S2D caching needs.

Storage

Windows Server 2016 S2D recommends a combination of two tiers of storage devices, one optimized for performance (caching), and one for capacity. This 1U Ultra reference architecture employs an all-flash NVMe configuration for maximized performance, with write-optimized NVMe drives as the caching device to deliver higher write performance, and to reduce the wear of the capacity and read-optimized NVMe drives at the lower tier.

- **Caching devices**

Two 400GB Intel SSD DC P3700 U.2 NVMe drives as caching devices to provide better endurance and higher write performance.

- **Capacity devices**

Six 2TB Intel SSD DC P3600 U.2 NVMe drives as capacity devices to obtain high capacity and great read performance.

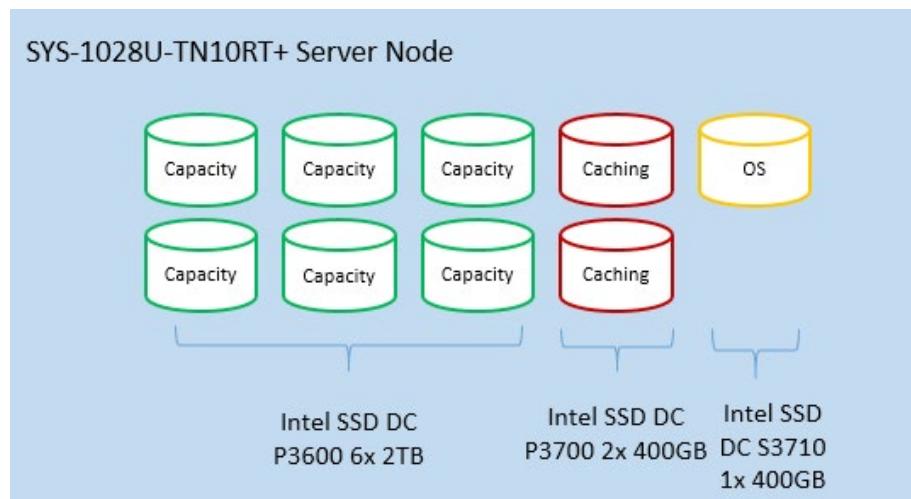


Figure 5. SMCI 1U Ultra SYS-1028U-TN10RT+ server node storage layout

In addition to the caching and capacity devices, a 400GB Intel SATA SSD DC S3710 is installed on each node as the boot device. The capacity is adequate for kernel memory dumps when issues arise. Figure 5 shows the overall storage layout on each Ultra server.

Networking

Networking is a key element in the S2D architecture, as a highly efficient and reliable networking fabric is critical to the overall performance of the S2D HCI. We have selected 25G Ethernet network interface cards supporting Remote Direct Memory Access (RDMA), along with a high-performance 100G Mellanox switch (or Supermicro SSE-C3632S 100G switch) for the 1U Ultra HCI reference architecture.

- **Network adapters**

The Supermicro AOC-S25G-M2S 25GbE is a high-performance network adapter that is built on the Mellanox ConnectX®-4 Lx EN chipset. It comes with dual SPF28 ports and supports the RoCE v2 specification with CPU offload, delivering high performance and low latency over Ethernet. In the reference architecture, one dual port NIC is installed on each server node.

- **TOR Switch**

Mellanox SN2700 32-port 100G Ethernet switch (or Supermicro SSE-C3632S 100G switch) is a reliable, high-density top-of-rack (ToR) solution that enables full rack connectivity to servers at full speed. The switch provides Enhanced Traffic Selection (ETS) 802.1Qaz and Priority Based Flow Control (PFC) (802.1p/Q and 802.1Qbb), thus fully supporting RoCE-capable NICs thus is a great choice for a S2D HCI solution.

- **Switch-to-Node Cabling**

The 1U Ultra HCI solution adopts the network configuration of one switch connecting to multiple connections per server for greater network resiliency and performance. Below, Figure 4 shows the cabling diagram.

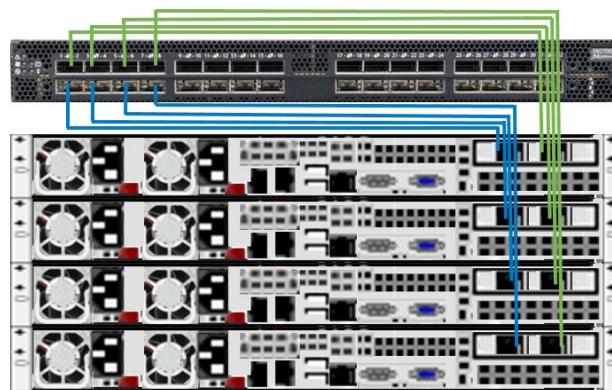


Figure 6. SMCI 1U Ultra SYS-1028U-TN10RT+ switch-to-node cabling diagram

Note: The diagram does not show the management network that aggregates the onboard dedicated IPMI LAN ports with a separate gigabit Ethernet switch.

For More Information

Supermicro Windows Server Software-Defined Solutions

www.supermicro.com/wssd

Supermicro 1U Ultra SuperServer 1028U-TN10RT+

www.supermicro.com/products/system/1U/1028/SYS-1028U-TN10RT.cfm

Microsoft Windows Server 2016

www.microsoft.com/en-us/cloud-platform/windows-server

Intel® Xeon® Processor E5-2600 v4 product family

www.intel.com/content/www/us/en/processors/xeon/xeon-e5-solutions.html

Intel® NVMe SSDs

www.intel.com/content/www/us/en/solid-state-drives/solid-state-drives-ssd.html

SECURITY

Windows Server 2016 introduced new security features, including Credential Guard and Device Guard, for protecting physical machines and virtual machines. In addition, Shielded VMs are used to protect virtual machines from a compromised host. To support the new features, each 1U Ultra server node is equipped with a TPM 2.0 module, with upgraded firmware that fully meets all SDDC requirements. The BIOS is pre-configured and optimized for deploying a secured hyper-converged solution out-of-the-box.

APPENDIX: SYS-1028U-S2D BILL OF MATERIALS

PART NUMBER	PART DESCRIPTION	QUANTITY	NOTE
SYS-1028U-TN10RT+	X10DRU-i+, 119UAC10-R1K02P-N10, AOC-URN6-i2XT	4	System name
P4X-DPE52650V4-5R2N3	BDW-EP 12C/24T E5-2650V4 2.2G 30M 9.6GT 105W R3 2011 M0	8	Dual Xeon processors per node
MEM-DR432L-SL01-ER24	32GB DDR4-2400 2Rx4 LP ECC REG DIMM	16	128GB DDR4 memory per node
HDS-2TM-SSDSC2BA400G4	Intel S3710 400GB, SATA 6Gb/s, HET MLC 2.5" 7.0mm, 20nm	4	1 SATA3 SSD for OS per node
HDS-2VM-SSDPE2MD400G4	Intel DC P3700 400B, NVMe PCI-E 3.0, HET, MLC 2.5" 20nm, HF, RoHS/REACH	8	2 NVMe SSDs for caching per node
HDS-2VM-SSDPE2ME020T4	Intel DC P3600 2TB, NVMe PCIe3.0, MLC 2.5" 20nm	24	6 HDDs drives for capacity per node
AOC-S25G-M2S-O	Standard low-profile 2-port 25GbE SFP28, based on Mellanox ConnectX-4 LX EN chipset, RoHS	4	1x dual port 25GbE NICs per node support RoCE
AOM-TPM-9655H	Horizontal TPM with Infineon 9655, RoHS/REACH, PBF	1	1 TPM 2.0 module per node

Note: The BOM above is for reference only and is a non-exhaustive list.

About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its "We Keep IT Green™" initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

www.supermicro.com

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, Twin, BigTwin, TwinPro, TwinPro², SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

All other brands names and trademarks are the property of their respective owners.

© Copyright 2017 Super Micro Computer, Inc. All rights reserved.

Printed in USA

Please Recycle

14_MS-HCI-1028U-Ref_170518_Rev05

